

Maximo Application Suite



Contents

- Welcome..... 1**

- Product overview..... 1**
 - Maximo Application Suite technical overview.....2
 - Prerequisite software..... 5
 - Compatibility matrix..... 39
 - Maximo Application Suite as a Service overview..... 41
 - What's new.....43
 - ... in Maximo Application Suite feature channel..... 43
 - What's new in Maximo Application Suite 9.1..... 44
 - ... in 9.0..... 49
 - ... in 8.11..... 52
 - ... earlier releases..... 57
 - ... in fix packs..... 63
 - ... in Maximo Application Suite as a Service..... 64
 - Applications, industry solutions, and add-ons..... 67
 - Applications..... 67
 - Industry solutions..... 72
 - Add-on and tools..... 74
 - Maximo Application Suite accelerators..... 76
 - Maximo Application Suite application URLs..... 77
 - Licensing in Maximo Application Suite..... 78
 - Licensing in Maximo Application Suite 9.1..... 78
 - Licensing in Maximo Application Suite 9.0 and earlier..... 105
 - Language and locale support..... 127
 - Accessibility..... 128
 - Federal requirements..... 129
 - Documentation conventions..... 130

- Getting started..... 130**
 - Application user..... 131
 - Application administrator..... 131
 - Application suite administrator..... 132
 - SaaS suite administrator..... 132

- Planning..... 132**
 - Planning for IBM Maximo Application Suite standard installation with CLI..... 132
 - Prerequisites for installing..... 133
 - ... to install in disconnected environment..... 135
 - ... to install foundation service..... 136
 - Planning to install on Amazon Web Services..... 138
 - Overview..... 138
 - Prerequisites for installing..... 143
 - Preparing to install..... 146
 - Configuring the installation permissions..... 154
 - Resizing Red Hat OpenShift cluster on Amazon Web Services..... 158
 - Security considerations..... 159
 - Planning to install on Microsoft Azure..... 160
 - Overview..... 160
 - Prerequisites for installing..... 163

Preparing to install.....	167
Configuring installation permissions.....	174
Security considerations.....	175
Planning for on premises.....	176
Preparing to install.....	176
Prerequisites for installing.....	177
Requirements and capacity planning.....	178
Supported software versions.....	178
Verify software entitlement.....	179
Networking considerations.....	182
System requirements.....	185
Installation topology.....	199
Planning for IBM Cloud	201
Creating your IBM Cloud account and configuring permissions.....	201
Obtaining your IBM Entitlement key from the IBM Entitled Registry.....	203
Installing IBM Cloud CLI.....	204
Installing Red Hat OpenShift Container Platform on IBM Cloud.....	205
Installing the Red Hat OpenShift Container Platform Command line Interface.....	211
Adding compute nodes to an existing Red Hat OpenShift Container Platform installation.....	212
Operational mode for installation.....	212
IBM Suite License Service.....	213
High availability.....	213
Logical architecture.....	213
Resilient architecture components.....	215
Installing.....	218
... with Maximo Application Suite CLI.....	218
... in disconnected environments.....	219
... with Amazon Web Services CloudFormation templates.....	222
Installing the Maximo Application Suite.....	222
Installing Cloud Pak for Data on an Amazon Web Services instance of Maximo Application Suite	232
Monitoring IBM Maximo Application Suite installation on Amazon Web Services.....	238
Accessing IBM Maximo Application Suite.....	239
Configuring Let's Encrypt.....	241
Configuring Maximo Application Suite on Amazon Web Services.....	245
Configuring Amazon Web Services DocumentDB.....	250
Configuring Amazon MSK.....	259
... with Microsoft Azure Resource Manager templates.....	262
Installing.....	262
... with Ansible collection.....	276
Maximo Application Suite Ansible collection examples.....	278
Setting up IBM Maximo Application Suite.....	281
Uninstalling.....	287
Uninstalling Maximo Application Suite.....	287
Deleting the Maximo Application Suite stack on Amazon Web Services.....	288
Uninstalling Maximo Application Suite on Microsoft Azure.....	290
Deploying.....	291
Applications.....	292
Maximo Collaborate.....	292
Maximo Health.....	296
Maximo Manage.....	297
Maximo Monitor.....	370
Maximo Predict.....	371
Maximo Real Estate and Facilities.....	373
Maximo Visual Inspection.....	391
IoT tool.....	403

Industry solutions.....	404
Maximo Health and Predict - Utilities.....	405
MRO Inventory Optimization.....	407
Add-ons.....	408
App Connect.....	408
IBM Maximo Visual Inspection Edge.....	408
IBM Parts Identifier.....	409
Deploying IBM Maximo Optimizer.....	410
Maximo AI Service and AI features in Maximo Manage.....	412
Accelerators.....	464
Activating and deactivating accelerators.....	464
Activating an accelerator from the details page.....	464
Activating an accelerator from the accelerators list page.....	464
Activating applications.....	465
Maximo Collaborate.....	465
Maximo Monitor.....	467
Maximo Predict.....	468
Maximo Visual Inspection.....	469
Maximo Optimizer.....	469
Maximo Health and Predict - Utilities.....	470
Deactivating and deleting applications.....	472
Upgrading.....	473
Upgrading Maximo Application Suite versions.....	473
Upgrade prerequisites for apps and add-ons.....	475
... using the channel subscription method.....	477
... using static catalog	478
... manually.....	480
Converting manual deployment to channel subscription.....	482
Upgrading to Maximo Application Suite 9.1.....	482
Updating applications.....	483
Upgrading from Maximo Asset Management to Maximo Manage.....	494
Overview.....	494
Planning.....	512
Preparing.....	515
Installing Maximo Application Suite.....	521
Creating a Maximo Manage database.....	522
Deploying Maximo Manage.....	546
Activating Maximo Manage.....	552
Migrating.....	552
Verifying.....	566
Overview of migrating TRIRIGA to Maximo Real Estate and Facilities.....	572
What's changed in Maximo Real Estate and Facilities.....	573
Migrating the application database.....	575
Migrating AES reversible encryption.....	575
Deployment and activation settings.....	576
Migrating TRIRIGA user files.....	576
Migrating users and licenses.....	576
Post-migration tasks.....	579
Configuring.....	580
Setting up IBM Maximo Application Suite.....	581
Storing configuration values as secrets.....	586
Configuring the global image pull secret.....	587
Certificate management.....	588
Creating a ClusterIssuer.....	588
Disabling default certificate authorities.....	589

Configuring certificate authority certificates.....	590
Configuring the size of public certificate resources.....	591
Manual certificate management.....	591
Authentication options for Db2U.....	602
Storage.....	603
User authentication.....	604
Authentication methods.....	605
Configuring default identity providers.....	611
Streamlined login.....	612
Self-registration for users.....	613
Configuring user authentication sessions.....	615
Configuring single sign-on properties.....	616
Configuring multiple identity providers.....	617
User synchronization.....	619
LDAP user registry synchronization.....	619
User synchronization with SCIM 2.0.....	625
Configuring multiple LDAP user registry synchronizations.....	631
Mapping LDAP users from Microsoft Active Directory.....	632
Mapping groups from LDAP to display group descriptions.....	633
Setting up email notifications.....	634
SMTP configuration.....	634
Configuring emails as optional in Maximo Application Suite 9.0.14.....	636
Creating custom email templates.....	637
Disabling email notifications.....	640
Changing the language of email notifications.....	642
Configuring external launchers.....	644
Configuring the minimum password length.....	646
Changing privacy access for obtaining user data.....	646
Enabling special characters for user ID and username.....	647
Customizing workloads.....	648
Supported pods.....	648
Customizing workload scale.....	703
Customizing workload affinity.....	705
Customizing workload tolerations.....	706
Customizing hostAliases in podTemplates.....	706
Enabling defaultJMS for Maximo Manage.....	707
Configuring the user interface.....	708
Updating the user interface.....	708
Enabling login notification.....	713
Disabling or hiding login options by using APIs.....	713
Disabling or hiding login options in custom resource.....	714
Hiding guided tours.....	716
Managing user profile.....	716
Disabling surveys.....	717
Configuring cross-origin resource sharing (CORS).....	717
Using single-node Red Hat OpenShift clusters.....	718
When to use single-node Red Hat OpenShift clusters.....	719
Installation prerequisites.....	719
Preparing a Docker container.....	719
Installing on bare metal or VMware vSphere.....	720
Installing logical volume manager storage.....	720
Installing Maximo Application Suite and Maximo Manage.....	721
Troubleshooting the image-registry-storage persistent volume claim.....	721
Administering	722
Performance optimization for IBM Maximo Application Suite.....	722
Key microservices and dependencies to scale for Maximo Application Suite core.....	722

Key MongoDB metrics.....	723
Scaling MongoDB community edition.....	723
Enabling Red Hat OpenShift Container PlatformCluster Insights Advisor.....	725
Configuring PID limits for Docker.....	726
Configuring the HAProxy router.....	726
Node considerations.....	727
Backing up and restoring IBM Maximo Application Suite.....	728
Overview of the components and processes for back up and restore procedures.....	728
Backing up Maximo Application Suite.....	734
Restoring and validating Maximo Application Suite.....	757
Backing up and restoring with IBM Storage Fusion.....	778
.....	781
User access and entitlements.....	782
Administering users.....	784
Administering security groups.....	789
User management APIs.....	793
Administering user sessions.....	794
Administering users and user access in Maximo Application Suite in 9.0 and earlier.....	796
Creating users.....	796
User entitlement and access.....	797
Importing users.....	801
Setting user account status.....	806
Setting language and time zone preferences for users.....	807
Deleting and anonymizing user data.....	808
Upgraded users from Maximo Asset Management.....	809
Managing users in Maximo Manage.....	809
Local user account settings.....	810
Configuring password settings.....	810
Setting password expiration.....	810
Enabling account lockout.....	811
Resetting login attempts.....	811
Administering licenses and AppPoints usage.....	813
Understand AppPoint allocation.....	813
View usage reports.....	813
Configure session idle timeout.....	814
Configure licenses.....	814
Generating and managing API keys.....	815
Audit logging in Maximo Application Suite.....	817
Importing data.....	818
Exporting data.....	819
Managing SaaS users.....	819
Requesting authentication and user sync.....	819
Adding users in SaaS.....	820
Monitoring AppPoint usage for Maximo Application Suite as a Service.....	822
Purchasing AppPoints for Maximo Application Suite as a Service.....	823
Scenario: Monitoring SaaS AppPoint usage for Maximo Manage.....	823
Monitoring.....	824
Configuring Red Hat OpenShift cluster monitoring.....	824
Installing Grafana.....	826
Using the serviceability dashboard.....	832
Installing Grafana, OpenTelemetry, and Prometheus.....	833
Installing the OpenTelemetry operator manually	833
Enabling the serviceability dashboard.....	833
Viewing metrics on the serviceability dashboard.....	834
Server-level monitoring metrics.....	835
Detailed metrics on the serviceability dashboard.....	839

Monitoring Maximo Manage data.....	844
Developing.....	844
Extending applications.....	844
Extensibility overview.....	845
Extension types.....	845
Characteristics of extensions.....	846
Extending Maximo Application Suite applications.....	847
Application Configuration.....	856
Application Configuration overview.....	856
Getting started for Application Configuration users.....	857
Application upgrade using Application Configuration	863
Application Configuration.....	870
Troubleshooting.....	881
Troubleshooting Amazon Web Services.....	882
Troubleshooting installation problems.....	882
Retrieving the installation source code version.....	888
Troubleshooting installation problems for Microsoft Azure.....	888
Retrieving installation source code version.....	894
Installation and configuration issues.....	894
Microsoft Active Directory user synchronization.....	894
LDAP user mapping from Microsoft Active Directory.....	895
Upgrade issues.....	896
General upgrade issues.....	897
Upgrade issues in Maximo Health , Maximo Predict, and Maximo Health and Predict - Utilities...	899
Upgrade issues in IoT tool.....	902
Upgrade issues in Maximo Manage.....	903
Upgrade issues in Maximo Monitor.....	904
Reference.....	905
Installing Red Hat OpenShift Container Platform and Maximo Application Suite on a Windows system.....	905
Installing Red Hat OpenShift Container Platform Local.....	906
Installing Maximo Application Suite and Maximo Manage.....	906
Maximo Application Suite core services.....	908
Maximo Application Suite pod details.....	917
Enabling access for identity provider administration by using APIs.....	922
Glossary.....	923
Notices.....	931
Index.....	935

Welcome

IBM Maximo Application Suite overview

IBM® Maximo® Application Suite provides powerful AI-driven asset and analytics solutions for your enterprise needs.

Maximo Application Suite is a powerful collection of:

- Applications
- Industry solutions
- Add-ons
- Tools

IBM Cloud Pak® for Data and Red Hat® OpenShift® Container Platform are included with Maximo Application Suite and used to manage the deployment of the containerized applications.

Maximo Application Suite can be an IBM-managed installation or a customer-managed installation. For customer-managed installations, when you purchase Maximo Application Suite, you also must purchase subscription licenses or perpetual licenses. For IBM-managed installations, you have the option to purchase Maximo Application Suite as a Service or Maximo Application Suite Dedicated.

You use Application Points (AppPoints) to manage application usage, runtime, and user access for the Maximo Application Suite applications, industry solutions, add-ons, and tools that you want to use. You do not need to deploy and use all of the applications, industry solutions, add-ons, and tools.

For more information about the features and benefits of Maximo Application Suite, see the [Applications, industry solutions, add-ons, and tools](#) documentation.

Installing Maximo Application Suite

Before installing Maximo Application Suite, review the “[Planning](#)” on page 132 documentation to ensure that you are aware of the dependencies and requirements that are necessary for your installation.

Then, review the “[Installing Maximo Application Suite](#)” on page 218 to learn more about the customer-managed scenarios for:

- Installing Maximo Application Suite on-premises
- Installing Maximo Application Suite in an AirGap environment
- Installing Maximo Application Suite on Amazon Web Services
- Installing Maximo Application Suite on Microsoft Azure

For IBM-managed installation, IBM performs the installation for you and you can proceed to complete post-installation tasks to set up Maximo Application Suite.

Using Maximo Application Suite

Customer-managed After Maximo Application Suite is installed, from the **Suite administration** page, you can use the product Catalog to configure and deploy applications, industry solutions, add-ons, and tools, administer your users, and more.

SaaS From the **Suite administration** page, you can manage users and their entitlement to Maximo Manage, request LDAP server authentication and user synchronization, and monitor application usage.

Use the Suite navigator and the app switcher in the menu bar to navigate between applications or to return to the Maximo Application Suite Suite navigator.

More information

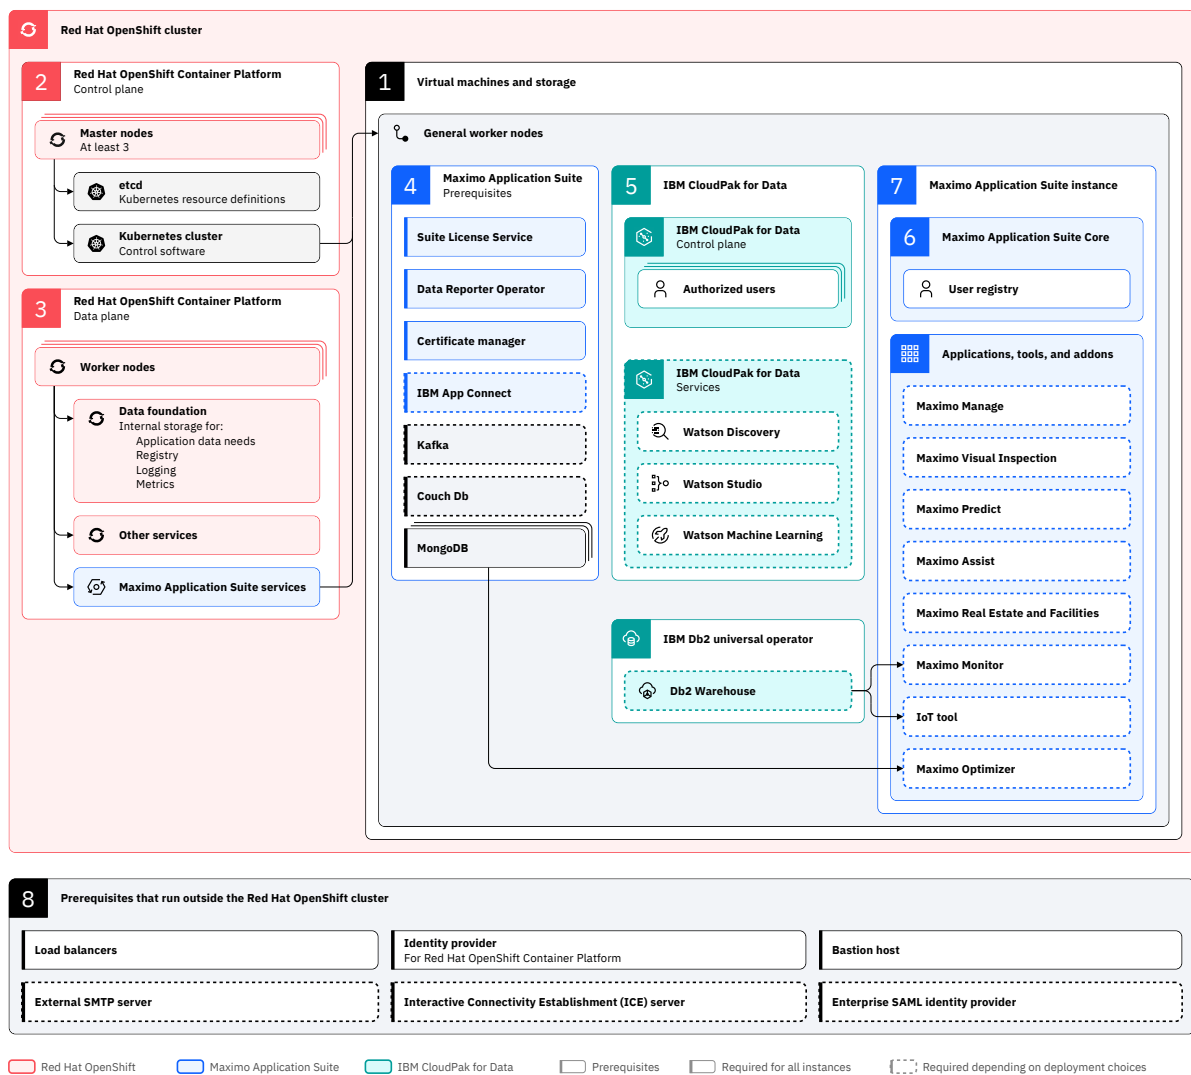
For more information about Maximo Application Suite, see [IBM Maximo Application Suite](#) on IBM.com.

Customer-managed **IBM Maximo Application Suite technical overview**

A Maximo Application Suite instance runs on Red Hat OpenShift Container Platform.

Red Hat OpenShift Container Platform enables portability between cloud and on-premises environments, containerization and container orchestration that supports robustness, resiliency, availability, and elasticity, repeatable deployments per the operator model, and automation and streamlining of the development process.

[Open image in new tab](#)



Refer to the description about each component in the following sections.

Virtual machines and storage

Most Maximo Application Suite applications and prerequisites run on a set of virtual machines that comprise a Red Hat OpenShift cluster. These virtual machines have separate IP addresses and appear as *nodes* in the Red Hat OpenShift cluster.

The virtual machines are provided and managed in two ways:

Infrastructure as a Service (IaaS) provider

For example, an on-premises provider, such as VMware, or a cloud provider, such as IBM Cloud®, AWS, or Azure.

You have direct access to the virtual machines, and you must install and manage the Red Hat OpenShift Container Platform software itself.

Managed Red Hat OpenShift Container Platform service provider

For example, Red Hat OpenShift on IBM Cloud.

Red Hat OpenShift Container Platform control plane

- A set of *master nodes* manages the Red Hat OpenShift cluster.

Note: To ensure continuous and high availability of the Red Hat OpenShift cluster, use a minimum of three master nodes.

- These master nodes run the Kubernetes cluster control software that manages what runs on the *worker nodes*, which are the other nodes in the cluster.
- The master nodes also maintain an internal database, an etcd that contains the Kubernetes resource definitions.
- If you are using a Managed Red Hat OpenShift Container Platform service provider, you don't see these master nodes because the service provider manages them.
- If you are running your own Red Hat OpenShift Container Platform service, you need to provision virtual machines. After you install the Red Hat OpenShift Container Platform, interaction with the virtual machines happens through Red Hat OpenShift Container Platform and Kubernetes APIs.

Red Hat OpenShift Container Platform infrastructure

Red Hat OpenShift Container Platform infrastructure components are installed as a part of the Red Hat OpenShift Container Platform installation.

- Red Hat OpenShift Container Platform infrastructure services, such as logging and monitoring.
- Red Hat OpenShift Container Platform provides a storage management mechanism, Red Hat OpenShift Container Platform Storage that runs in the Red Hat OpenShift Container Platform cluster itself.

Worker nodes have disk storage that is private to that node. The disk storage can be lost if the node malfunctions. Some Maximo Application Suite applications and prerequisites require storage that persists over node failures and that can be shared between nodes. However, if you are running in an external cloud environment, you can choose to use an external storage provider.

Some components must run on every worker node while other components run on three worker nodes. These components can run on any worker nodes in the cluster, but if possible, place the components on dedicated worker nodes. For more information, see [“Preparing to install Maximo Application Suite on premises” on page 176](#).

Maximo Application Suite prerequisites

Maximo Application Suite applications have several dependencies or prerequisites. Some prerequisites are necessary regardless of the selected applications. Other prerequisites are necessary only for specific applications. For most of these prerequisites, you can choose whether to deploy them in the Red Hat OpenShift cluster or run them externally, either in a separate Red Hat OpenShift cluster or by using an external service provider.

Cloud Pak for Data

One prerequisite that must run in the Red Hat OpenShift cluster is the Cloud Pak for Data.

Note: The Maximo Application Suite license entitles you to install and use several Cloud Pak for Data services, provided you are using them with Maximo Application Suite applications.

Cloud Pak for Data consists of a control plane, which has its own user interface and its own set of authorized users. After you install Cloud Pak for Data into Red Hat OpenShift cluster and you log in as an administrator, you can install one or more Cloud Pak for Data services into the Red Hat OpenShift cluster. You can use Cloud Pak for Data to install Db2® Warehouse, which is used by Maximo Monitor and IoT applications.

Maximo Application Suite applications, industry solutions, addons, and tools

Maximo Application Suite provides a suite of applications, tools, and add-ons. You can choose the applications to deploy if the appropriate prerequisites are in place.

Maximo Application Suite instance

Start by installing the base application, Maximo Application Suite core, into Red Hat OpenShift cluster.

- Use Maximo Application Suite core to install and manage the Maximo Application Suite applications, industry solutions, and add-ons that you want to use.
- Maximo Application Suite core maintains a registry of users. You can specify which users have access to which Maximo Application Suite applications.

Prerequisites that run outside the Red Hat OpenShift cluster

The following prerequisites run outside the Red Hat OpenShift cluster:

Load balancers

Used to allow access to the protocol endpoints that are used to communicate with the Red Hat OpenShift Container Platform and with the applications and services that it hosts.

Identity provider

Used to authenticate users when they log in to the Red Hat OpenShift cluster. Usually, the identity provider is an enterprise directory service that supports an LDAP interface.

Most Maximo Application Suite users do not need to log in to Red Hat OpenShift cluster.

You can use the same enterprise directory to manage the login credentials for Maximo Application Suite and Red Hat OpenShift cluster users.

The Maximo Application Suite core and Cloud Pak for Data control planes can both be configured to use this directory service for their user login.

Enterprise SAML identity provider

Allows users to share a single sign-on with other enterprise applications. A user who is signed in to another enterprise application can use the Maximo Application Suite application without reauthenticating.

External SMTP server

Needed to configure Maximo Application Suite core, Maximo Manage, and other applications to send emails to users.

Interactive Connectivity Establishment (ICE) server

Configured to use Voice over Internet Protocol (VoIP) connections to Maximo Collaborate.

Bastion host

A host that runs outside of the Red Hat OpenShift cluster. The bastion host has direct network access to the cluster nodes.

A bastion host is necessary if you are installing the Red Hat OpenShift cluster yourself. The bastion host is useful when you install Maximo Application Suite core, Cloud Pak for Data, and other prerequisites into Red Hat OpenShift cluster.

Other software

IBM Cognos® Analytics entitlement is included in Maximo Application Suite 8.10 and later.

Prerequisite software

Ensure that your Maximo Application Suite installation meets each application-specific prerequisite according to your chosen installation path.

The following application prerequisites must be met before you can deploy the applications.



Attention: IBM App Connect and Cloud Pak for Data do not support odd-numbered Red Hat OpenShift Container Platform versions. If you plan to deploy Maximo Collaborate, Maximo Health and Predict - Utilities, or Maximo Predict, you must use even-numbered Red Hat OpenShift Container Platform versions.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

		Maximo Application Suite core	IoT	Maximo Monitor	Maximo Manage	Maximo Health	Maximo Predict	Maximo Health and Predict - Utilities	Maximo Collaborate	Maximo Real Estate and Facilities	Maximo Visual Inspection
Red Hat OpenShift cluster services	IBM Certificate Manager	X	X	X	X	X	X	X	X	X	X
	Data Reporter Operator	X	X	X	X	X	X	X	X	X	X
	IBM Suite License Service	X	X	X	X	X	X	X	X	X	X
	CouchDB								X		
Cloud Pak for Data services	CPD control plane		X	X	O	O	X	X	X	O	
	Db2 Warehouse		X	X	O	O	O	O	O		
	Watson Studio						X				
	Watson Machine Learning							X			
	IBM Watson® Discovery								X		
Red Hat OpenShift cluster or external services	MongoDB	X	X	X							
	Kafka		X	X	O						
	Relational Database				X*	X*		X*			
	App Connect							O		O	
External services	SMTP Server	O			O			O		O	
	LDAP server	O	O	O	O	O	O	O	O	O	O
	SAML IdP	O	O	O	O	O	O	O	O	O	O
	OIDC Idp									O	
	Certificate issuer	O	O	O	O	O	O	O	O	O	O
	IBM Compliance Expert Server								O		

Legend:

X

You need all prerequisites that are shown with an X in the Maximo Application Suite core column and in the columns for all the applications that you plan to use. Some applications have dependencies between the applications.

X*

The requirement for a relational database can be met by the application that uses Db2 Warehouse.

O

Optional prerequisites are marked with an O. Maximo Application Suite can use them if they are present.

For more information about the supported applications and add-ons for each release, see [compatibility matrix](#).

Notes:

- IBM Certificate Manager controls certificate management in Maximo Application Suite 8.8.
- Service Binding Operator (SBO) is no longer a dependency and is not required in Maximo Application Suite 8.8.

Database configurations and scope

Some applications are not recommended to reuse the same database instance within Db2 Warehouse for production environments. For testing or development environments you could reuse the same instance. Suite applications also can use different scopes of database configuration such as workspace-application scope, application scope, workspace scope and system scope. For more information about the configuration scopes, see [“Configuring Maximo Application Suite”](#) on page 580.

The following table describes some recommendations you can follow to configure your Db2 Warehouse to support the Suite applications.

Note: Some applications and tools support other databases than Db2 Warehouse. For example, Maximo Manage also supports Microsoft SQL Server, Oracle Database and other versions of Db2. Maximo Real Estate and Facilities also supports Microsoft SQL Server, Oracle Database and other versions of Db2. See the specific application installation documentation for the details about the other database types that such application supports.

Application	Exclusive database recommended	Scope
IoT	No. Typically works fine sharing the database with Monitor, Predict, and Maximo Health and Predict - Utilities.	System
Monitor	No. Typically connects to the IoT database.	System
Maximo Manage + add-ons (Maximo Health add-on included)	Yes	Workspace-application
Maximo Health	Yes	Workspace-application
Predict	No. Typically uses the IoT or Maximo Monitor database.	System
Maximo Health and Predict - Utilities	No. Typically uses the same database as Predict.	System
Maximo Collaborate	Does not use Db2 database.	Does not use Db2 database.
Maximo Visual Inspection	Does not use Db2 database.	Does not use Db2 database.
Maximo Real Estate and Facilities	Yes	Workspace-application

Related information

[Software Product Compatibility Report](#)

Red Hat OpenShift cluster services

Maximo Application Suite runs on Red Hat OpenShift. Therefore, an Red Hat OpenShift cluster must be configured and running to install and configure Maximo Application Suite.

Red Hat OpenShift certificate manager

The Red Hat OpenShift certificate manager service helps you manage and deploy SSL/TLS certificates for your apps and services. It provides you with a security-rich repository for your certificates and their

associated private keys, and helps prevent outages by sending you notifications when your certificates are about to expire.

Red Hat OpenShift certificate manager operator is installed into the `cert-manager-operator` namespace and the operand is created in the `cert-manager` namespace.

Note: The `cert_manager` role supports migration from an existing IBM Certificate Manager installation to the Red Hat OpenShift Certificate Manager. The role configures the cluster resources namespace to `ibm-common-services` for compatibility with existing ClusterIssuers.

To run the `cert_manager` role, you must install the Red Hat OpenShift Operators CatalogSource.

For more information about installing the Red Hat OpenShift certificate manager, see [Ansible dev-ops role - cert_manager](#).

Related tasks

[Creating a ClusterIssuer](#)

Service Binding Operator

Maximo Application Suite uses Service Binding Operator (SBO) to bind applications together with operator-managed backing services.

Note: Maximo Application Suite uses the Service Binding Operator to bind applications together with operator-managed backing services. If you are upgrading from Maximo Application Suite 8.7.x, install a Service Binding Operator to ensure compatibility with your applications.

Suite License Service

IBM Suite License Service provides features for managing virtualized environments and measuring license utilization. Suite License Service discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and generates audit reports. Each report provides you with different information about your infrastructure, for example the computer groups, software installations, and the content of your software catalog.

The Suite License Service (SLS) stores and manages the Maximo Application Suite license. You upload the license file to the SLS server as part of initial setup.

When you install Suite License Service and the Federal Information Processing Standard (FIPS) is enabled, ensure that the following default ciphers for Java™ are supported.

```
sh-4.4$ java -Dsemeru.fips=true Ciphers
Default Cipher
* TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
* TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
* TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
* TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
* TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
* TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
* TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
* TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
* TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA
* TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
* TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
* TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
* TLS_ECDH_RSA_WITH_AES_128_CBC_SHA
* TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
* TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
* TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
```

Data Reporter Operator

The IBM Data Reporter Operator accepts events and transforms them into reports that are submitted to the Data Service of the IBM Metrics Operator.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

As a Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance.

Raw json event data is sent to an endpoint that is serviced by the IBM Data Reporter Operator. The event data is transformed into a report and is sent to the IBM Data Service. The IBM Data Service periodically uploads the reports to Red Hat Marketplace.

In non-airgap environments, event data is uploaded hourly. In airgap environments, IBM Data Reporter Operator provides a manual mechanism to obtain and upload events to IBM. For more information about uploading data, see [Export and import usage reports using Data collection CLI](#).

Related information

[IBM Maximo Application Suite - Migrate Maximo Application Suite from User Data Services \(UDS\) to Data Reporter Operator \(DRO\)](#)

Migrating Maximo Application Suite from User Data Services to Data Reporter Operator

As an IBM Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance. New and existing Maximo Application Suite users can install or migrate to DRO by using the IBM Maximo Application Suite command line interface (CLI), ansible role, or manually.

Procedure

- Installing with CLI

The Maximo Application Suite CLI provides a mechanism to choose DRO when you install the Maximo Application Suite.

- New DRO installation

Run the **mas install** command with **export UDS_ACTION=install-dro**. The CLI installs UDS by default. The installation and configuration is automated, and no additional steps are required.

```
docker run -ti --pull always quay.io/ibmmas/cli mas install export UDS_ACTION=install-dro
```

For more information, see [Maximo Application Suite CLI installation](#).

- Migrating from UDS to DRO

You can automate DRO migration with the CLI by running the **mas update** command. The utility runs the following validation tests to make sure that the cluster is compatible for DRO migration.

- All Maximo Application Suite instances must be either 8.10.x or 8.11.x
- Available Maximo Application Suite instances must be configured with a UDS local on the cluster.
- UDS is installed locally and running
- The **mas update** command must have the catalog version set to February 2024

If these criteria are met, UDS and DRO are installed. The BASCFG CR's containing DRO configurations are applied on the Red Hat OpenShift cluster to use DRO on all Maximo Application Suite instances.

For more information, see [Maximo Application Suite CLI update](#).

Note: Maximo Application Suite instances with remote UDS configurations are not supported in the automated DRO migration utility. In these cases, use ansible roles to install and configure DRO.

- Install DRO with ansible role

Run the ansible role `dro` to automate and configure DRO. DRO can co-exist along with UDS, or you can install and configure DRO and then uninstall UDS.

Install and configure DRO

```
export IBM_ENTITLEMENT_KEY=<Your IBM entitlement Key>
export DRO_CONTACT_EMAIL=xxx@xxx.com
export DRO_CONTACT_FIRSTNAME=xxx
export DRO_CONTACT_LASTNAME=xxx
export MAS_CONFIG_DIR=<path to masconfig dir>
export MAS_INSTANCE_ID=<your_instance_id>
```



```

export ROLE_NAME='dro'
export DRO_ACTION=install-dro
export DRO_STORAGE_CLASS=<your_ocp_storageclass> # optional field, can be used when using
custom storage classes
export DRO_MIGRATION=true # optional field, if set will also uninstall UDS.

ansible-playbook playbooks/run_role.yml

```

- MAS_CONFIG_DIR is an empty directory on your machine where the ansible role stores temporary files during the installation or migration process.
- DRO_STORAGE_CLASS is an optional field. If Red Hat OpenShift Container Platform is running on IBM Cloud , Amazon Web Services, or Microsoft Azure, the ansible playbook by default determines a suitable storage class and uses it. You can also use a custom storage class and update DRO_STORAGE_CLASS with a suitable name.
- When you set DRO_MIGRATION to True, the ansible role installs and configures DRO on multiple instances of Maximo Application Suite and uninstalls UDS.

For more information, see [Data Reporter Operator](#) role.

- Installing manually

Follow the instructions in the README file on the [Red Hat Marketplace GitHub repository](#).

After the DRO is installed, follow these steps to configure Maximo Application Suite to connect to the DRO instance.

1. Login to Red Hat OpenShift Container Platform as an administrator.
2. Search for BASCFG in **Administration > Custom Resource Definition**.
3. Open the BASCFG file and select the **Instances** tab.
4. Look for a BASCFG instance that is related to your environment. If you have one already, reuse the same name on the following sample `bascfg.yml` and run the following **oc** commands to apply the `bascfg` configuration on Maximo Application Suite.

```

oc login
oc project mas-{{ mas_instance_id }}-core
oc apply -f bascfg.yml

```

Sample `bascfg.yml` file. Insert valid values for all fields surrounded by double braces (`{{ }}`). For more information, see [Role Variables - BASCFG Generation](#).

```

-\\-\\-
apiVersion: v1
kind: Secret
type: opaque
metadata:
  name: dro-apikey
  namespace: "mas-{{ mas_instance_id }}-core"
stringData:
  api_key: "{{ dro_api_key }}"
-\\-\\-
apiVersion: config.mas.ibm.com/v1
kind: BasCfg
metadata:
  name: "{{ mas_instance_id }}-bas-system"
  namespace: "mas-{{ mas_instance_id }}-core"
  labels:
    mas.ibm.com/configScope: system
    mas.ibm.com/instanceId: "{{ mas_instance_id }}"
spec:
  displayName: DRO {{ mas_instance_id }}
  config:
    url: "{{ dro_endpoint_url }}"
    contact:
      email: "{{ dro_contact.email }}"
      firstName: "{{ dro_contact.first_name }}"
      lastName: "{{ dro_contact.last_name }}"
    credentials:
      secretName: dro-apikey
  certificates:
    - alias: {{ insert your Certificate alias name }}

```

```
cert: |
  {{ dro_certs }}
```

Note: If the cluster has multiple Maximo Application Suite instances, repeat the steps with each of the `mas_instance_ids`.

Related tasks

[Setting up Maximo Application Suite](#)

Related reference

[Connection details for installing Maximo Application Suite on Amazon Web Services](#)

Check connection details for existing Red Hat OpenShift cluster, network infrastructure, IBM Suite License Service instance, and IBM Data Reporter Operator instance.

[Connection details for installing Maximo Application Suite on Microsoft Azure](#)

Check connection details for existing Red Hat OpenShift cluster, network infrastructure, IBM Suite License Service instance, and IBM Data Reporter Operator instance.

Creating the Db2 instance by using the stand-alone Db2U operator

You can create an IBM Db2 instance that uses the stand-alone Db2U operator from the Red Hat OpenShift console or from the CLI.

Tip:

- This task maps to the following Ansible role: `db2`. For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.
- When you install Db2 on IBM Cloud and use IBM File Storage for IBM Cloud follow the steps that are outlined in the following topic to enable `no_root_squash`. For more information, see [Configuring IBM Cloud File Storage](#).

Before you begin

If your environment needs to be compliant with Federal Information Processing Standard (FIPS), in IBM Maximo Application Suite 8.10.1 or later, you must install a new instance of Db2 inside a Red Hat OpenShift cluster that is enabled for FIPS. The IPsec encryption must be enabled when you install the cluster, and the JDBC configuration must use the non-SSL connection.

For more information, see [IPsec encryption configuration](#).

If you are planning to install Maximo Application Suite, ensure the following default ciphers for Java are supported when you enable FIPS.

```
sh-4.4$ java -Dsemeru.fips=true Ciphers
Default Cipher
* TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
* TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
* TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
* TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
* TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
* TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
* TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
* TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
* TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA
* TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
* TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
* TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
* TLS_ECDH_RSA_WITH_AES_128_CBC_SHA
* TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
* TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
* TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
```

Installing by using the Red Hat OpenShift Container Platform web console

Procedure

1. In the banner, click **Import YAML** (). Enter the following YAML:

```

---
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: "db2u-ca-issuer"
  namespace: "db2u"
spec:
  selfSigned: {}

```

2. Click **Create** to provision the self-signed CA certificate issuer.

3. In the banner, click **Import YAML** (). Enter the following YAML:

```

---
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: "db2u-ca-certificate"
  namespace: "db2u"
spec:
  secretName: "db2u-ca"
  duration: "175200h0m0s"
  renewBefore: "2160h0m0s"
  issuerRef:
    name: "db2u-ca-issuer"
    kind: Issuer


  isCA: true
  keyAlgorithm: rsa
  keySize: 4096
  keyEncoding: pkcs8

  usages:
    - cert sign
    - digital signature
    - key encipherment
    - server auth

  commonName: "ca.db2u"
  organization:
    - "IBM Maximo Application Suite"
  subject:
    countries:
      - GB
    streetAddresses:
      - London
    localities:
      - London
    organizationalUnits:
      - IBM Maximo Application Suite DB2U

```

4. Click the **Create** button to provision the self-signed CA certificate for Db2.

5. In the banner, click **Import YAML** (). Enter the following YAML. This YAML creates the Db2 server certificate issuer. This issuer references the secret that is created from the CA certificate issuer in the previous step.

```

---
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: "db2u-issuer"
  namespace: "db2u"
spec:
  ca:
    secretName: "db2u-ca"

```

6. Verify that the secret db2u-ca exists, then click **Create**.

7. In the banner, click **Import YAML** (). Enter the following YAML:

```

---
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: "db2u-certificate"
  namespace: "db2u"
spec:
  secretName: "db2u-certificate"
  duration: "175200h0m0s"
  renewBefore: "2160h0m0s"
  issuerRef:
    name: "db2u-issuer"
    kind: Issuer
  usages:
    - cert sign
    - digital signature
    - key encipherment
    - server auth
  commonName: "db2u"
  dnsNames:
    - "db2u-manage-db2u.apps.cluster1.example-cluster.com"
    - "*.db2u-manage-db2u.apps.cluster1.example-cluster.com"
    - "c-db2u-manage-db2u-engn-svc.db2u.svc"
    - "*.c-db2u-manage-db2u-engn-svc.db2u.svc"
  organization:
    - "IBM Maximo Application Suite"
  subject:
    countries:
      - GB
    streetAddresses:
      - London
    localities:
      - London
    organizationalUnits:
      - IBM Maximo Application Suite DB2U

```

Note: For the DNS names, include the svc and routes that are used. Replace them with real environment information:

- `{{db2_instance_name}}-{{db2_namespace}}.{{cluster_subdomain}}`
- `*.{{db2_instance_name}}-{{db2_namespace}}.{{cluster_subdomain}}`
- `c-{{db2_instance_name}}-db2u-engn-svc.{{ db2_namespace }}.svc`
- `*.c-{{db2_instance_name}}-db2u-engn-svc.{{ db2_namespace }}.svc`

8. Click the **Create** button to create the Db2 server certificate.

9. In the banner, click **Import YAML** (). Enter the following YAML:

```

---
apiVersion: db2u.databases.ibm.com/v1
kind: Db2uCluster
metadata:
  name: "db2u-manage"
  namespace: "db2u"
spec:
  account:
    privileged: true
  add0ns:
    graph:
      enabled: false
    rest:
      enabled: false
  version: "11.5.7.0-cn4"
  size: 1
  environment:
    dbType: db2wh
    database:
      name: "BLUDB"
    settings:
      dftTableOrg: "ROW"
    ssl:
      secretName: "db2u-certificate"
      certLabel: "CN=db2u"
  instance:

```

```

registry:
  DB2_4K_DEVICE_SUPPORT: "ON"
  DB2AUTH: 'OSAUTHDB,ALLOW_LOCAL_FALLBACK,PLUGIN_AUTO_RELOAD'
  DB2_FMP_RUN_AS_CONNECTED_USER: 'NO'
  DB2_WORKLOAD: MAXIMO
mln:
  total: 1
license:
  accept: true
podConfig:
  db2u:
    resource:
      db2u:
        requests:
          cpu: "2"
          memory: "12Gi"
        limits:
          cpu: "6"
          memory: "18Gi"
storage:
- name: meta
  type: create
  spec:
    storageClassName: "ocs-storagecluster-cephfs"
    accessModes:
      - ReadWriteMany
    resources:
      requests:
        storage: "100Gi"
- name: data
  type: template
  spec:
    storageClassName: "ocs-storagecluster-ceph-ibrd"
    accessModes:
      - ReadWriteOnce
    resources:
      requests:
        storage: "500Gi"
- name: backup
  type: create
  spec:
    accessModes:
      - ReadWriteMany
    resources:
      requests:
        storage: "500Gi"
    storageClassName: "ocs-storagecluster-cephfs"
- name: activelogs
  spec:
    accessModes:
      - ReadWriteOnce
    resources:
      requests:
        storage: "100Gi"
    storageClassName: "ocs-storagecluster-ceph-ibrd"
  type: template
- name: tempts
  spec:
    accessModes:
      - ReadWriteOnce
    resources:
      requests:
        storage: "100Gi"
    storageClassName: "ocs-storagecluster-ceph-ibrd"
  type: template

```

10. Click the **Create** button to create the db2ucluster CR instance.

11. Verify that the db2ucluster is created successfully.

Click **Home** > **Search** page, select the db2u project and search for resource type Db2uCluster. Filter the service name keyword with db2u-manage. Then, click its details and wait until it is in Ready state. The process might take 10 minutes to complete.

12. Check the Db2 instance login, URL, password and ca crt information.

Check the internal service.

On the **Home** > **Search** page, select the db2u project and search for resource type Service. Filter the service name keyword with db2u-engn-svc. In the search results, it shows the service name similar to c-db2wh-xxxx-db2u-engn-svc.

Inside the cluster, the JDBC URL is:

```
jdbc:db2://c-db2u-manage-db2u-engn-svc.db2u.svc:50001/BLUDB:sslConnection=true;
```

Note: The name for the Db2 instance is the one provided in the db2ucluster YAML file. It is not a generated instance name.

When you install Maximo Application Suite 8.10 or later in FIPS enabled Red Hat OpenShift cluster, configure the JDBC URL to use a non-SSL connection.

- a. Edit the JdbcConfig file from **Administration** > **CustomResourceDefinitions** and search JdbcCfg.
- b. Change sslEnabled to false.
- c. Change url to use non-ssl port.

Check the Db2 instance password for default Db2 user db2inst1.

Click **Home** > **Search**, select the db2u project and search for resource type Secret. Filter the service name keyword with -instancepassword. In the search results, it shows the service name similar to c-db2u-manage-instancepassword. Then, reveal its password.

The JDBC user and password that is used to access the Db2 instance is:

```
user: db2inst1  
password: xxx
```

Note: Consider the [“Authentication options for Db2U”](#) on page 602 when you are working with Db2 user authentication.

Check the internal TLS that is used for JDBC SSL access:

Click **Home** > **Search**, select the db2u project and search for resource type Secret. Filter the service name keyword with db2u-certificate. Then, reveal its ca.crt.

Installing by using the Red Hat OpenShift command-line interface (CLI)

Procedure

1. Create a Db2U Operator catalog YAML file called db2catalog.yaml.

```
---  
apiVersion: operators.coreos.com/v1alpha1  
kind: CatalogSource  
metadata:  
  name: ibm-db2uoperator-catalog  
  namespace: openshift-marketplace  
spec:  
  sourceType: grpc  
  image: icr.io/cpopen/ibm-db2uoperator-catalog:latest  
  imagePullPolicy: Always  
  displayName: IBM Db2U Catalog  
  publisher: IBM  
  updateStrategy:  
    registryPoll:  
      interval: 45m
```

2. Apply the db2catalog.yaml file to the Red Hat OpenShift cluster.

```
oc apply -f db2catalog.yaml
```

3. Verify the catalog source status:

```
oc get catalogsource -n openshift-marketplace ibm-db2uoperator-catalog -o
jsonpath='{.status.connectionState.lastObservedState}' {"\n"}
```

4. Create a namespace called db2u:

```
oc new-project db2u
```

5. Create the OperandRequest YAML file db2u-operator.yaml to install the Db2u Operator:

```
---
apiVersion: operator.ibm.com/v1alpha1
kind: OperandRequest
metadata:
  name: db2u-request
  namespace: "db2u"
spec:
  requests:
  - operands:
    - name: ibm-db2u-operator
      registry: common-service
      registryNamespace: ibm-common-services
```

6. Apply the db2u-operator.yaml file to the Red Hat OpenShift cluster.

```
oc apply -f db2u-operator.yaml
```

7. Verify the Db2U operator is created and running successfully:

```
oc get sub -n ibm-common-services ibm-db2u-operator -o jsonpath='{.status.installedCSV}
{"\n"}
```

Sample output

```
db2u-operator.v1.1.13
```

```
oc get csv -n ibm-common-services db2u-operator.v1.1.13 -o jsonpath='{.status.phase } :
{.status.message}' {"\n"}
```

Sample output

```
Succeeded : install strategy completed with no errors
```

```
oc get deployments -n ibm-common-services db2u-operator-manager -o
jsonpath='{.status.availableReplicas}' {"'\n'"}
```

Sample output

```
1
```

8. To provision the self-signed CA certificate issuer, create a CA Issuer YAML file called caissuer.yaml.

```
---
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: "db2u-ca-issuer"
  namespace: "db2u"
```

```
spec:
  selfSigned: {}
```

9. Apply the `caissuer.yaml` file to the Red Hat OpenShift cluster.

```
oc apply -f caissuer.yaml
```

10. To provision the self-signed CA certificate for Db2, create a CA certificate YAML file called `cacert.yaml`.

```
---
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: "db2u-ca-certificate"
  namespace: "db2u"
spec:
  secretName: "db2u-ca"
  duration: "175200h0m0s"
  renewBefore: "2160h0m0s"
  issuerRef:
    name: "db2u-ca-issuer"
    kind: Issuer

  isCA: true
  keyAlgorithm: rsa
  keySize: 4096
  keyEncoding: pkcs8

  usages:
    - cert sign
    - digital signature
    - key encipherment
    - server auth

  commonName: "ca.db2u"
  organization:
    - "IBM Maximo Application Suite"
  subject:
    countries:
      - GB
    streetAddresses:
      - London
    localities:
      - London
    organizationalUnits:
      - IBM Maximo Application Suite DB2U
```

11. Verify that the CA certificate is created successfully.

```
oc get certificates -n db2u
NAME                                READY    SECRET    AGE    EXPIRATION
db2u-ca-certificate                 True     db2u-ca   26s    2042-05-04T08:26:02Z
```

12. Create the YAML file `issuer.yaml` to create the Db2 server certificate issuer. This issuer references the secret that is created from the CA certificate issuer in the preceding step.

```
---
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: "db2u-issuer"
  namespace: "db2u"
spec:
  ca:
    secretName: "db2u-ca"
```

13. Apply the `issuer.yaml` to the Red Hat OpenShift cluster.

```
oc apply -f issuer.yaml
```

14. Create the YAML file `certificate.yaml` to provision the Db2 server certificate:


```

---
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: "db2u-certificate"
  namespace: "db2u"
spec:
  secretName: "db2u-certificate"
  duration: "175200h0m0s"
  renewBefore: "2160h0m0s"
  issuerRef:
    name: "db2u-issuer"
    kind: Issuer

  usages:
    - cert sign
    - digital signature
    - key encipherment
    - server auth

  commonName: "db2u"
  dnsNames:
    - "db2u-manage-db2u.apps.cluster1.example-cluster.com"
    - "*.db2u-manage-db2u.apps.cluster1.example-cluster.com"
    - "c-db2u-manage-db2u-engn-svc.db2u.svc"
    - "*.c-db2u-manage-db2u-engn-svc.db2u.svc"
  organization:
    - "IBM Maximo Application Suite"
  subject:
    countries:
      - GB
    streetAddresses:
      - London
    localities:
      - London
    organizationalUnits:
      - IBM Maximo Application Suite DB2U

```

Note: For the DNS names, include the svc and routes that are used. Replace them with real environment information:

- {{db2_instance_name}}-{{db2_namespace}}.{{cluster_subdomain}}
- *.{{db2_instance_name}}-{{db2_namespace}}.{{cluster_subdomain}}
- c-{{db2_instance_name}}-db2u-engn-svc.{{ db2_namespace }}.svc
- *.c-{{db2_instance_name}}-db2u-engn-svc.{{ db2_namespace }}.svc

15. Apply the certificate.yaml to the Red Hat OpenShift cluster.

```
oc apply -f certificate.yaml
```

16. Create the YAML file db2cluster.yaml to create the Db2 cluster:

```

---
apiVersion: db2u.databases.ibm.com/v1
kind: Db2uCluster
metadata:
  name: "db2u-manage"
  namespace: "db2u"
spec:
  account:
    privileged: true
  addOns:
    graph:
      enabled: false
    rest:
      enabled: false
  version: "11.5.7.0-cn4"
  size: 1
  environment:
    dbType: db2wh
    database:
      name: "BLUDB"
    settings:

```

```

    dftTableOrg: "ROW"
    ssl:
      secretName: "db2u-certificate"
      certLabel: "CN=db2u"
  instance:
    registry:
      DB2_4K_DEVICE_SUPPORT: "ON"
      DB2AUTH: 'OSAUTHDB,ALLOW_LOCAL_FALLBACK,PLUGIN_AUTO_RELOAD'
      DB2_FMP_RUN_AS_CONNECTED_USER: 'NO'
      DB2_WORKLOAD: MAXIMO
  mln:
    total: 1
  license:
    accept: true
  podConfig:
    db2u:
      resource:
        db2u:
          requests:
            cpu: "2"
            memory: "12Gi"
          limits:
            cpu: "6"
            memory: "18Gi"
  storage:
  - name: meta
    type: create
    spec:
      storageClassName: "ocs-storagecluster-cephfs"
      accessModes:
      - ReadWriteMany
      resources:
        requests:
          storage: "100Gi"
  - name: data
    type: template
    spec:
      storageClassName: "ocs-storagecluster-ceph-ibd"
      accessModes:
      - ReadWriteOnce
      resources:
        requests:
          storage: "500Gi"
  - name: backup
    type: create
    spec:
      accessModes:
      - ReadWriteMany
      resources:
        requests:
          storage: "500Gi"
      storageClassName: "ocs-storagecluster-cephfs"
  - name: activelogs
    spec:
      accessModes:
      - ReadWriteOnce
      resources:
        requests:
          storage: "100Gi"
      storageClassName: "ocs-storagecluster-ceph-ibd"
    type: template
  - name: tempts
    spec:
      accessModes:
      - ReadWriteOnce
      resources:
        requests:
          storage: "100Gi"
      storageClassName: "ocs-storagecluster-ceph-ibd"
    type: template

```

17. Apply the `db2cluster.yaml` to the Red Hat OpenShift cluster.

```
oc apply -f db2cluster.yaml
```

18. Verify the Db2U cluster status:

```
oc get db2ucluster -n db2u db2u-manage -o jsonpath='{.status.state} {"\n"}'
```

19. Check the Db2 instance login, URL, and password.

Check the internal service:

```
oc get svc -n db2u | grep -i engn-svc
```

Sample output

c-db2u-manage-db2u-engn-svc	NodePort	172.30.120.206	<none>	50000:30601/
TCP,50001:30036/TCP				

Inside the cluster, the JDBC URL is:

```
jdbc:db2://c-db2u-manage-db2u-engn-svc.db2u.svc:50001/BLUDB:sslConnection=true;
```

Note: db2u-manage is the instance name for the Db2 instance.

When you install Maximo Application Suite 8.10 or later in FIPS enabled Red Hat OpenShift cluster, configure the JDBC URL to use a non-SSL connection.

- Edit the JdbcConfig file from **Administration > CustomResourceDefinitions** and search JdbcCfg.
- Change sslEnabled to false.
- Change url to use non-ssl port.

Check the Db2 instance password for default Db2 user db2inst1:

```
oc extract secret/c-db2u-manage-instancepassword -n db2u --keys=password --to=-
```

Sample output

```
# password  
PcJ1fKYfFdA5AtA
```

The JDBC user and password that is used to access the Db2 instance is:

```
user: db2inst1  
password: PcJ1fKYfFdA5AtA
```

Check the ca.crt used to connect the JDBC SSL port:

```
oc extract secret/db2u-certificate -n db2u --keys=ca.crt --to=-
```

Sample output

```
# ca.crt
-----BEGIN CERTIFICATE-----
MIIDuzCCAqOgAwIBAgIQf4KcTIk5y8EEZeiUtioEOTANBgkqhkiG9w0BAQsFADBT
MQswCQYDVQQGEwJHQjEPMA0GA1UEBxMGTG9uZG9uMQ8wDQYDVQQJEwZMb25kb24x
KjAoBgNVBAsTIU1CTSBNYXhpbnw8gQXBwbGljYXRpb24gU3VpdGUGREIyVTEQMA4G
A1UEAxMHY2EuZGIydTAeFw0yMjA1MDkxOTA2MjZaFw00MjA1MDQxOTA2MjZaMG0x
CzAJBgNVBAYTAkdCMQ8wDQYDVQOHEwZMb25kb24xZDZANBgNVBAkTBkxvbmRvbGJEq
MCgGA1UECxMhSUJNIE1heG1tbyBBcHBSaWVhbiBtdWl0ZSBEQjJVMRAwDgYD
VQQDEwdjYS5kYjJ1MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEApzw
Gz6FFCc1dsPAYJuxx7KnNsNJqyHCabo+VgQG4QkNzyMCwP2YmZZbvi7iajC+U20/
dMME06LeHdcBFsmkRE7dYcGw8YewsQ8mRTwYveP92h/yLUGzQ8IuhRZ70HZ5ozgt
4Cs5K0p0zqnkk0BbFltJZkWdnGerEnu025LUCwEfJ1sV3LmBTuOodKNLQ6VW5MWF
6HrLK4I2jPFPfNo1v/9V+rtRUIfXZSEwsm02imQTgVw9yM+oLZx4be05hY1fWE3c
nwVfHkygZKVxIsbd4zm/U7k/oHaEhIPt9gyWXL3pjdYo2jGXX2btp8xS8UcuI1w
Qew0QtIK6Kc56CgobwIDAQABo1cwVTAOBgNVHQ8BAf8EBAMCAqQwEwYDVR01BAww
CgyIKwYBBQUHAwEwDwYDVR0TAQH/BAUwAwEB/zAdBgNVHQ4EFgQUmk4PmrKaFUW0
18/P9Zgs6J9v9IwDQYJKoZIhvcNAQELBQADggEBAEmgttI/N5/91SDRH8AVTMyr
pr0QYvt9Wsm5hYvHIkuldtFMpKg6SZ7HzUnRiIGBz1XqC7TvAiQLauacp0JciBoq
Wgkh50gBB1d1/KZS8CuNk/KEym+DPw+cb8Lr1lpPNW/nKuc/0I8mDNsEv8zwYki
yymaTSr+MmNjztz+iqtmionwLVome721x11bEmcWUo6oxQxU9MbPmJiJOPzkcDx7
dZi087RuJ9aQxdkhZDwhZEUssGFDbq1+t1XZxy1DtE2spVaxXBai6wqScUceK8sE
geUcQR04VjMEb4RTiOb8QjJ0V0eOqE71nARyWTEjcvShstgRipcGnSlAVfLE7w=
-----END CERTIFICATE-----
```

What to do next

[“Configuring IBM Db2 Warehouse” on page 306](#)

Cloud Pak for Data services

IBM Cloud Pak for Data platform helps improve productivity and reduce complexity. Build a data fabric connecting siloed data distributed across a hybrid cloud landscape. This product offers a wide selection of IBM and third-party services spanning the entire data lifecycle. Deployment options include an on-premises software version that is built on the Red Hat OpenShift Container Platform or a fully managed version that is built on IBM Cloud.

Maximo Application Suite includes an entitlement to use Cloud Pak for Data.

For Maximo Application Suite users that require Maximo Predict or Maximo Collaborate applications Cloud Pak for Data is required to install the Watson Studio or Watson Discovery dependencies.

Important: For information about the Cloud Pak for Data components that are available to the licensee of Maximo Application Suite for restricted use, see the [License Information documents](#) for Maximo Application Suite.

Watson Studio

Watson Studio provides the environment and tools for you to collaboratively work on data to solve your business problems. You can choose the tools that you need to analyze and visualize data, to cleanse and modify data, to obtain streaming data, or to create and train machine learning models.

Required by the IBM Maximo Predict application.

Watson Machine Learning

Watson Machine Learning provides a full range of tools and services so that you can build, train, and deploy Machine Learning models. Choose the tool with the level of automation or autonomy that matches your needs, from a fully automated process to writing your own code.

For more information about Watson Machine Learning, see [Watson Machine Learning on Cloud Pak for Data](#).

Watson OpenScale

Watson OpenScale is an enterprise-grade environment for AI applications that provides your enterprise visibility into how your AI is built, is used, and delivers return on investment.

For more information, see [Watson OpenScale on Cloud Pak for Data](#).

Required by:

- Predict

Spark

You can optionally use Analytics Engine powered by Apache Spark to extend your jupyter notebooks capabilities while dealing with large data sets in Maximo Predict application.

With Analytics Engine powered by Apache Spark, you can:

- Run Jupyter notebooks and jobs from other tools in Watson Studio analytics projects by selecting a Spark environment runtime. Install this service either before or after you install the Watson Studio service.
- Run Spark SQL or jobs for data transformation, data science, or machine learning using Spark job APIs. The Spark job APIs do not require the Watson Studio service.

For more information, see [Analytics Engine Powered by Apache Spark on Cloud Pak for Data](#).

Optional for:

- Predict

Related tasks**Deploying IBM Maximo Predict**

Maximo Predict can use historical and recent asset performance data to correlate performance factors that predict asset degradation or failure. Other types of data that can be correlated include maintenance records, inspection reports, and environmental data. Maximo Predict uses artificial intelligence to optimize predictive model accuracy.

Db2 Warehouse

IBM Db2 Warehouse is an analytics data warehouse that features in-memory data processing and in-database analytics. The Cloud Pak for Data control plane is not required to install IBM Db2; alternatively, the Db2 operator can be installed standalone.

For more information about the Db2 Warehouse, see [Cloud Pak for Data : Db2 Warehouse on Cloud Pak for Data](#).

Red Hat OpenShift cluster or external services

Configure and deploy Red Hat OpenShift cluster or external services to install the IBM Maximo Application Suite.

MongoDB

Maximo Application Suite and Suite License Service use MongoDB for user management. The MongoDB instance can run in the Red Hat OpenShift cluster or external to it.

Prerequisites

To complete the setup of Maximo Application Suite, note the following MongoDB requirements.

- The hostnames and ports of the MongoDB servers.
- The `config db` name the database name that is used by MongoDB to store its user IDs.
- The user credentials for an admin user who has table creation authority.
- The MongoDB certificate authority (CA) certificates. If your MongoDB cluster uses self-signed CA certificates, you must retrieve them automatically or add them manually when you complete the setting up the Maximo Application Suite.

Note:

- Starting in 9.0.5 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x architecture, MongoDB is supported as an external service.

- Starting in 9.0.12 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM Power (ppc64le) architecture, MongoDB is supported as an external service.

For more information, see [Software Product Compatibility Reports \(SPCR\)](#).

To work with Maximo Application Suite, your MongoDB instance must support the transport layer security (TLS) communication protocol. For more information, see [TLS/SSL \(Transport Encryption\)](#) in the MongoDB documentation.

Federal Information Processing Standard (FIPS)

Starting with Maximo Application Suite 8.10.1, FIPS mode can be enabled. You can continue to use the MongoDB Community Edition for Maximo Application Suite in FIPS mode. However, if you want to have a FIPS compliant MongoDB instance as well, you must use MongoDB Enterprise Edition with your own license.

If FIPS is enabled for Maximo Application Suite, DocumentDB cannot be used as a dependency for Maximo Application Suite because DocumentDB does not support SCRAM-SHA-256 authentication, which is required by Maximo Application Suite in FIPS mode.

Certain components in Maximo Application Suite use the IBM Java Semeru Runtime, which has limited cipher support in FIPS Mode.

If the Maximo Application Suite is running in FIPS mode and configured to use the MongoDB instance in FIPS mode, ensure that the subset of ciphers that are supported by IBM Java Semeru are enabled for FIPS mode.

Ciphers	Name (OpenSSL)	Cipher suite Name (IANA)
0xC024	ECDHE-ECDSA-AES256-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
0xC028	ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
0xC023	ECDHE-ECDSA-AES128-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
0xC027	ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC026	ECDH-ECDSA-AES256-SHA384	TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
0xC02A	ECDH-RSA-AES256-SHA384	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
0xC025	ECDH-ECDSA-AES128-SHA256	TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
0xC029	ECDH-RSA-AES128-SHA256	TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
0xC00A	ECDHE-ECDSA-AES256-SHA	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
0xC014	ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
0xC009	ECDHE-ECDSA-AES128-SHA	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
0xC013	ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA

<i>Table 2. Supported ciphers (continued)</i>		
Ciphers	Name (OpenSSL)	Cipher suite Name (IANA)
0xC005	ECDH-ECDH-AES256-SHA	TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
0xC00F	ECDH-RSA-AES256-SHA	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
0xC004	ECDH-ECDH-AES128-SHA	TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA
0xC00E	ECDH-RSA-AES128-SHA	TLS_ECDH_RSA_WITH_AES_128_CBC_SHA

Maximo Application Suite 8.9 or earlier

For Maximo Application Suite 8.9 or earlier, you can use the Community Edition of MongoDB. To set up MongoDB Community Edition on the Red Hat OpenShift cluster, use the [Maximo Application Suite ansible devops collection](#).

Apache Kafka

Apache Kafka provides a buffer for messages that are sent to and received from external interfaces. Apache Kafka is not required if the IBM Maximo Manage software is not interfacing with external systems.

About this task

Apache Kafka is required by IoT and is optional for Manage


Note: When you are updating an operator, there might be delays applying changes to the Kafka configuration.

Procedure

- Install Kafka from Red Hat OpenShift web console or command line.
- [Red Hat AMQ Streams operator](#), which is based on the [Strimzi operator](#), can be used to install Kafka for on-premises installations. It can also be used to install Kafka in cloud-based Maximo Application Suite installations, when a managed Kafka service by the cloud provider is not desirable.

Note: Starting in 8.10.1, to deploy Maximo Application Suite in a FIPS enabled environment, it is recommended to install Kafka by using Strimzi Operator 0.33.2.

To install by using the Red Hat OpenShift web console:

1. From **Home > Projects**, click the **Create Project** button, enter the name `kafka`, and click **Create** to provision the new namespace for Kafka.
2. In the banner, click **Import YAML** (). Enter the following YAML.

```
---
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: "kafka"
  namespace: "kafka"
spec:
  targetNamespaces:
    - "kafka"
```

3. Click **Create** to create the operator group in the `kafka` namespace.

4. In the banner, click **Import YAML** (). Enter the following YAML.

```
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-streams
  namespace: "kafka"
spec:
  channel: amq-streams-1.8.x
  installPlanApproval: Automatic
  name: amq-streams
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

Tip: For Maximo Application Suite on AWS (BYOL) version 8.7, change `amq-streams-1.8.x` to `amq-streams-1.7.x` to match the version of AMQ streams that is installed in the BAS namespace.

5. Click **Create** to create the subscription resources.
6. From **Operators > Installed Operators**, search for AMQ Streams and verify that the operator Status is set to **Succeeded**.

7. In the banner, click **Import YAML** (). Enter the following YAML.

```
---
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: "maskafka"
  namespace: "kafka"
spec:
  # -----
  kafka:
    version: 2.7.0
    replicas: 3
    resources:
      requests:
        memory: 4Gi
        cpu: "1"
      limits:
        memory: 4Gi
        cpu: "2"
    jvmOptions:
      -Xms: 3072m
      -Xmx: 3072m
    config:
      offsets.topic.replication.factor: 3
      transaction.state.log.replication.factor: 3
      transaction.state.log.min.isr: 2
      log.message.format.version: "2.7"
      log.retention.hours: 24
      log.retention.bytes: 1073741824
      log.segment.bytes: 268435456
      log.cleaner.enable: true
      log.cleanup.policy: delete
      auto.create.topics.enable: false
    storage:
      type: jbod
      volumes:
        - id: 0
          type: persistent-claim
          class: "ocs-storagecluster-ceph-ibb"
          size: 100Gi
          deleteClaim: true
    authorization:
      type: simple
    listeners:
      - name: tls
        port: 9094
        type: route
        tls: true
        authentication:
```



```

        type: scram-sha-512
# -----
zookeeper:
  replicas: 3
  resources:
    requests:
      memory: 1Gi
      cpu: "0.5"
    limits:
      memory: 1Gi
      cpu: "1"
  jvmOptions:
    -Xms: 768m
    -Xmx: 768m
  storage:
    type: persistent-claim
    class: "ocs-storagecluster-ceph-rbd"
    size: 10Gi
    deleteClaim: true
# -----
entityOperator:
  userOperator: {}
  topicOperator: {}

```

Ensure that you modify the specified storage class `ocs-storagecluster-ceph-rbd` to use a supported storage class for your cluster.

8. Click **Create** to create the Kafka cluster.

9. From **Workloads > StatefulSets**, switch to the `kafka` project. You should see two StatefulSets: `maskafka-kafka` (the Kafka brokers) and `maskafka-zookeeper` (the Kafka ZooKeepers). Select each statefulset and verify that each has three pods, which are in Ready state.

10. In the banner, click **Import YAML** (). Enter the following YAML.

```

---
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaUser
metadata:
  name: "maskafkauser"
  labels:
    strimzi.io/cluster: "maskafka"
  namespace: "kafka"
spec:
  authentication:
    type: scram-sha-512
  authorization:
    type: simple
  acls:
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: topic
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: group
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: cluster
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: transactionalId

```

11. Click **Create** to create a Kafka user, which is used by Maximo Application Suite to authenticate connections to Kafka.
12. From **Workloads > Secrets**, switch to the kafka project. Verify that the maskafkauser secret was created by the user entity operator.
13. From **Networking > Routes**, switch to the kafka project. Verify that the maskafka-kafka-tls-bootstrap route was created.
14. Get the Kafka information.

To get the Kafka host and port:

```
oc get Kafka.kafka.strimzi.io maskafka -o
jsonpath="{.status.listeners[0].addresses[0]}"
```

Sample output

```
{"host": "maskafka-kafka-tls-bootstrap-kafka.apps.cluster1.example-
cluster.com", "port": 443}
```

To get the Kafka ca crt:

```
oc get Kafka.kafka.strimzi.io maskafka -o
jsonpath="{.status.listeners[0].certificates[0]}"
```

Sample output

```
-----BEGIN CERTIFICATE-----
MIIFLTCCAwwAwIBAgIUExU12XrdIPy6vZAtk9toGh2jbEwdQYJKoZIhvcNAQEN
BQAwLTETMBEGA1UECgwKaw8uc3RyYWw16aTEWMBQGA1UEAwwNY2x1c3R1ci1jYSB2
MDAeFw0yMjA1MTEyMTAyMzFaFw0yMjA1MTEyMTAyMzFaMC0xEzARBGNVBAoMcm1v
LnN0cm1temkxXjAUBG9NVBAMMDWNsdXN0ZXItY2EgdjAwggIiMA0GCSqGSIb3DQEB
AQUAA4ICDWAwggIKAoICAQDh6bYIudhZQ1/rR9IgsB7pzqTvtRiN0vzmnZPdtVtT
q7lNlytPqpR6uuC1rhpuR0CPb++RvjP2QrWgXr5VWBktT1MLk8WzDfX3+qxds5x8
B00EKneBZkhohxBdb0c081pxDpQAFty+SeXhuR0d5vWLEuh30JeZMEUfTcNfUbo
J/IHUIGeDmhK//DumQE79z3vflc2EcQgenMo0VoBy4ooQ2o4B7Y3p1XHustvtn6h
lam30rSA+p3nKskrMDDpNKadHtmCwI/rZZBFYb7DTdUpi69New3TEMRXGG3dMdk
YYTdkN0zkB5BTvRx5FC6GX+cz/Uq3Snx1SmWB1DT+2n1n1wzVAgbNdsW4HiDUIdI
FBjyQDqWTH9e7aUv3Rz1rT4c995YBTfH1Jdvq5mzneMf6lab7iZoW1hGYQLRRC5y
v8iTycwHd7EEGf/tjGrJ/s5nWPgGv/DE0g95/UvTRz9dZUWRwHCFANd0LaFW/HdF
qkhuivZ0KNXqfr7zxnCw/F+0408+vcR43HKUTwId7vq1+F+Egjt69U5pDF4sh6ep
SgLTHoCGd/bekq5HHkry1C0ty+ZU9EEWp4fQD+wN3RzGxJ080AA3RjkqsXmHbd5e
aXlnhDB68mWpohFuJ6YciNBX1C/2HhDer7PiMD9Zj0/7A3UHZj4hHXcS0cNsw7
mwIDAQABo0UwQzAdBgNVHQ4EFgQU6yQK1Z+FEJyMkjsPxmHERps1vgwEgYDVR0T
AQH/BAGwBgEB/wIBADAoBgNVHQ8BAf8EBAMCAQYwDQYJKoZIhvcNAQENBQADggIB
AEfcrS4I2xsbTuULmTh10GLgv7Mo+aJ80s+vCE+MvSMVrsVs1VnigzE6a5vi7Ys
TTpstmAhIf0cEEq1dRa5GcG6Az6NW1bskZXfftojWtjnZevkuRnn/xICdizX+mj4
A3wL/G0VpTAWVUa5+1Uh1AzFwhBw5kDvMxHyQhmpg9t98ptxNpj5n9cHSWwJpjX1
boNi1+Y5kA4raWGa6gE0E0lwmLyS5pj0WCTCTD2MvldNakYPMq0bVPE4DNia4qa1
hux0yxdri51KNBc7yVgQ1Fa7ZD+rF1a6aa6GwvAKYNoxd7VW7fmZBSckpuWer9+R
YCVvgE2a4vLnc5zLFw0fhjqazSiIx0PMEmkHx1ZTriVg0GVZ8beU+I9BxUQsJyJU
S4z9UaHexmYu/YRAQXKODw1xhqqR6oW2+CXYrtUvzN6kamFh8jN3AKf4PKA+TmjL
maW0M7FVp+0Erne59hBcZhKG0Yx4AkjCwKc1RwDBxXcBTcmXduDFeGzLub0napJ
Uczo2zURQ7L6qPew9Guh001dnGp+kgi8T8kt/DniMvQBWDK3GvFi0A5mVjLQqMHQ
HvAPzshx7S11045hepGK4fxQMCAHW6c1V3j10R8RHh7bck1d5mJ5Nh/BjZhk/LK
N5K1fwoek0QSVAXQfnX1YtJfrHfz5+TYx0NnYTCgX6fE
-----END CERTIFICATE-----
```

To get the Kafka username and password:

```
oc extract secret/maskafkauser -n kafka --keys=sasl.jaas.config --to=-
```

Sample output:

```
# sasl.jaas.config
org.apache.kafka.common.security.scram.ScramLoginModule required
username="maskafkauser" password="KbpatTNjUu5N";
```

Where the username is maskafkauser and the password is KbpatTNjUu5N.

To get the SASL Mechanism:

```
oc get Kafka.kafka.strimzi.io maskafka -n kafka -o
jsonpath='{.spec.kafka.listeners[0].authentication}' | jq -r .
```

Sample output:

```
{
  "type": "scram-sha-512"
}
```

Where the SASL Mechanism is SCRAM-SHA-512.

- To install by using the Red Hat OpenShift command-line interface (CLI):
 - From the bastion host, create the YAML file kafka-sub.yaml, containing the Namespace, OperatorGroup, and Subscription resources that are used to install Kafka:

```
---
apiVersion: v1
kind: Namespace
metadata:
  name: "kafka"
---
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: "kafka"
  namespace: "kafka"
spec:
  targetNamespaces:
    - "kafka"
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-streams
  namespace: "kafka"
spec:
  channel: amq-streams-1.8.x
  installPlanApproval: Automatic
  name: amq-streams
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

- Apply the kafka-sub.yaml file to the Red Hat OpenShift Container Platform cluster:

```
oc apply -f kafka-sub.yaml
```

- Verify that the AMQ Streams operator was successfully deployed:

```
oc get csv -n kafka -l operators.coreos.com/amq-streams.kafka
```

Sample output

NAME	DISPLAY	VERSION	REPLACES
amqstreams.v1.8.4 Succeeded	Red Hat Integration - AMQ Streams	1.8.4	amqstreams.v1.8.3

4. From the bastion host, create the YAML file `kafka-cluster.yaml`, containing the Kafka resource that describes the configuration of the Kafka cluster:

```
---
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: "maskafka"
  namespace: "kafka"
spec:
# -----
  kafka:
    version: 2.7.0
    replicas: 3
    resources:
      requests:
        memory: 4Gi
        cpu: "1"
      limits:
        memory: 4Gi
        cpu: "2"
    jvmOptions:
      -Xms: 3072m
      -Xmx: 3072m
    config:
      offsets.topic.replication.factor: 3
      transaction.state.log.replication.factor: 3
      transaction.state.log.min.isr: 2
      log.message.format.version: "2.7"
      log.retention.hours: 24
      log.retention.bytes: 1073741824
      log.segment.bytes: 268435456
      log.cleaner.enable: true
      log.cleanup.policy: delete
      auto.create.topics.enable: false
    storage:
      type: jbod
      volumes:
        - id: 0
          type: persistent-claim
          class: "ocs-storagecluster-ceph-rbd"
          size: 100Gi
          deleteClaim: true
    authorization:
      type: simple
    listeners:
      - name: tls
        port: 9094
        type: route
        tls: true
        authentication:
          type: scram-sha-512
# -----
    zookeeper:
      replicas: 3
      resources:
        requests:
          memory: 1Gi
          cpu: "0.5"
        limits:
          memory: 1Gi
          cpu: "1"
      jvmOptions:
        -Xms: 768m
        -Xmx: 768m
      storage:
        type: persistent-claim
        class: "ocs-storagecluster-ceph-rbd"
        size: 10Gi
        deleteClaim: true
# -----
    entityOperator:
      userOperator: {}
      topicOperator: {}
```

Ensure that you modify the specified storage class `ocs-storagecluster-ceph-rbd` to use a supported storage class for your cluster.

5. Apply the `kafka-cluster.yaml` file to the OCP cluster:

```
oc apply -f kafka-cluster.yaml
```

6. Verify that the Kafka cluster was successfully deployed. The Kafka CR should be in Ready state.

The Kafka CR specified in the following command is fully qualified with its API group name `kafkas.kafka.strimzi.io` to avoid ambiguity with the Kafka CR provided by `kafkas.ibmevents.ibm.com`.

```
oc get kafkas.kafka.strimzi.io -n kafka
```

Sample output

NAME	DESIRED KAFKA REPLICAS	DESIRED ZK REPLICAS	READY	WARNINGS
maskafka	3	3	True	

7. From the bastion host, create the YAML file `kafka-user.yaml`, containing the `KafkaUser` resource describing the configuration of the Kafka user which will be used by Maximo Application Suite to authenticate connections to Kafka:

```
---
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaUser
metadata:
  name: "maskafkauser"
  labels:
    strimzi.io/cluster: "maskafka"
  namespace: "kafka"
spec:
  authentication:
    type: scram-sha-512
  authorization:
    type: simple
  acls:
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: topic
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: group
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: cluster
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: transactionalId
```

8. Apply the `kafka-user.yaml` file to the OCP cluster:

```
oc apply -f kafka-user.yaml
```

9. Verify that the `maskafkauser` secret was created by the user entity operator:

```
oc get secret maskafkauser -n kafka
```

Sample output

NAME	TYPE	DATA	AGE
maskafkouser	Opaque	2	2m14s

10. Get the Kafka information.

To get the Kafka host and port:

```
oc get Kafka.kafka.strimzi.io maskafka -o  
jsonpath="{.status.listeners[0].addresses[0]}"
```

Sample output:

```
{"host": "maskafka-kafka-tls-bootstrap-kafka.apps.cluster1.example-  
cluster.com", "port": 443}
```

To get the Kafka ca crt:

```
oc get Kafka.kafka.strimzi.io maskafka -o  
jsonpath="{.status.listeners[0].certificates[0]}"
```

Sample output:

```
-----BEGIN CERTIFICATE-----  
MIIFLTCCAxWgAwIBAgIUExU12XrdIPy6vZAtk9toGh2jbEwDQYJKoZIhvcNAQEN  
BQAwLETETMBEGA1UECgwKaw8uc3RyaWw16aTEWMBQGA1UEAwwNY2x1c3Rlc1ci1jYSB2  
MDAeFw0yMjA1MTEyMTAyMzFaFw0yMzA1MTEyMTAyMzFaMC0xEzARBgNVBAoMcm1v  
LnNoOcm1temkxZjAUBGNVBAWMDWNSdXN0ZXItY2EgdjAwggIiMA0GCsQGSIB3DQEB  
AQUAA4ICDwAwggIKAoICAQDh6bYIudhZQ1/rR9Igsb7pzqTvtRiNOvzmnZPdtVtT  
q71NLyTqPqR6uuCIrhpuR0CPb++Rvjp2QrWgXr5VWBktT1MLk8WzDfX3+qxd5x8  
B00EKneBZkhohxBdb0co8ipxDpQAFTy+SeXhuR0d5vwLEuh30JeZMEUfTcNfUbo  
J/IHUIGeDmhK//DumQE79z3vfLc2EcQgenMo0VoBy4ooQ2o4B7Y3p1XHustvtn6h  
lam30rSA+p3nKskrMDDpNKadHtmCrwI/rZZBFYb7DTdUpi69New3TEMRXGG3dMdk  
YYTdKN0zkB5BTvRx5FC6GX+cz/Uq3Snx1SmWB1DT+2n1nlwzVAgbNdsW4HiDUIdI  
FBjyQDqWTH9e7aUv3RzlrT4c995YBTfh1Jdvq5mzneMf6lab7iZow1hGYQLRRC5y  
v8iTycwHd7EEGf/tjGrJ/s5nWPgGv/DEOg95/UvTRz9dZUWRwHCFAND0LaFW/HdF  
qkhuivZOKNXqfz7zxnCw/F+0408+vcR43HKUTwId7vq1+F+Egjt69U5pDF4sh6ep  
SgLTHoCGd/bekq5HHkrylCoty+ZU9EEWp4fQD+wN3RzGxJ080AA3RjkqsXmHbd5e  
aXlnhDB68mWpohFuJ6YciNBBX1C/2HhDeR7PiMD9Zj0/7A3UHJz4hHXcSgoCnSW7  
mwIDAQABo0UwQzAdBgNVHQ4EFgQU6yQK1Z+FEJyMkjsPxmHERpslvgeYDVR0T  
AQH/BAGwBgEB/wIBADA0BgNVHQ8BAf8EBAMCAQYwDQYJKoZIhvcNAQENBQADggIB  
AEfcrS4I2xsbTuULMtH10GLgv7Mo+aJ80s+vCE+MvSMVrsSvs1VnigzE6aSvi7Ys  
TTpstmAhIf0cEEqldRa5GcG6Az6Nw1bskZXfftojWtjnZevkuRnn/xICdizX+mj4  
A3WL/GOVpTAWVUa5+1Uh1AzFWhBw5kDvMxHyQhmpet98ptxNpj5n9cHSwwJpjX1  
boNi1+Y5kA4rawGa6gEOE0lwmLyS5pj0WCTCTD2MvldNakYPMq0bVPE4DNia4qa1  
hux0yxdz51KNBc7yVgQ1Fa7ZD+rF1a6aa6GvwwAKYN0xd7VW7fmZBSckpuWer9+R  
YCVgE2a4vLnc5zLfw0fhjqazSiIx0PMEmkHx1ZTriVg0GVZ8beU+I9BxUQsJyJU  
S4z9UaHexmYu/YRAQXK0Dw1xhqQR6ow2+CXyrtUvzN6kamFh8jN3AKf4PKA+TmjL  
maW0M7FVp+0Ene59hBcZhKG0QYx4AkjCwKc1RwDBXcBTcmXduDfEgZLub0napJ  
Uczo2zURQ7L6qPew9Guh001dnGp+kgi8T8kt/DniMvQBWDK3GvFi0A5mVjLQqMHQ  
HvAPzshx7Si1045hepGK4fxQMcmAHw6c1V3j10R8RHh7bckld5mJ5Nh/BjZhk/LK  
N5K1fwoek0QSVAXQfnX1YtJfRhfz5+TYx0NnYtCgX6fE  
-----END CERTIFICATE-----
```

To get the Kafka username and password:

```
oc extract secret/maskafkouser -n kafka --keys=sasl.jaas.config --to=-
```

Sample output:

```
# sasl.jaas.config  
org.apache.kafka.common.security.scram.ScramLoginModule required  
username="maskafkouser" password="KbpatTNjUu5N";
```

Where the username is maskafkauser and the password is KbpattNjUu5N.

To get the SASL Mechanism:

```
oc get Kafka.kafka.strimzi.io maskafka -n kafka -o  
jsonpath='{.spec.kafka.listeners[0].authentication}' | jq -r .
```

Sample output:

```
{  
  "type": "scram-sha-512"  
}
```

Where the SASL Mechanism is SCRAM-SHA-512.

- Installing Kafka by using the AMQ Streams Operator UI
 1. Installing the AMQ Streams and using the AMQ Streams UI to create a Kafka cluster and user.
 - Log into the OCP cluster with your username and password
 - Create a new project: kafka
 - Navigate to Operator Hub, and search for "AMQ Streams", then select the tile for Red Hat Integration - AMQ Streams
 - Click "Install"
 - On the next page, select the radio button for amq-streams-1.7.x.
 - Select a specific namespace on the cluster: kafka
 - click the "install" button
 - The operator is ready for use.
 2. Create a Kafka cluster.
 - In the kafka namespace, click "Installed Operators"
 - Navigate to the Kafka tab, and click "Create Kafka"
 - Enter the name "kafka" for the cluster
 - Expand the Kafka configuration and enter the following values:
 - Kafka Brokers: 3
 - Expand Storage:
Kafka Storage: jbod
Expand volumes and add a new volume:
 - id: 0
 - type: persistent-claim
 - Size: 100Gi
 - Storage class. ocs-storagecluster-ceph-rbd (replace with a supported block storage class)
 - Delete claim: true (checked)
 - Expand Listeners, and scroll down to the listener section named "tls"
 - port: 9093
 - type: route
 - tls: true (checked)
 - Expand "Authentication":
 - Type: scram-sha-512
 - Expand Authorization

- Type: Simple
- Expand the Zookeeper configuration and enter the following values:
 - e. Zookeeper Nodes: 3
 - f. Expand Storage:
 - Zookeeper Storage: persistent-claim
 - Size: 10Gi
 - class: ocs-storagecluster-ceph-rbd (replace with a supported block storage class)
 - Delete claim: true (checked)

Click Create.

For more information about available and tested storage classes, and for deployment size guidance, see the Monitor and IoT section of the [Maximo Application Suite system requirements](#) document.

3. Create the Kafka user.

- In the AMQ Streams Operator, navigate to Kafka User tab. Click "Create KafkaUser".
- Set name: "masuser".
- Expand authentication:

```
type: scram-sha-512
Expand authorization:
type: simple
```

- Expand acls and add 4 ACLs as shown:

```
a: host: '*'
    operation: All
    > Expand Resource:
name: '*'
  patternType: prefix
  type: topic
Type: Allow
b: host: '*'
    operation: All
    > Expand Resource:
    name: '*'
    patternType: prefix
    type: group
Type: Allow
  c: host: '*'
    operation: All
    > Expand Resource:
    name: '*'
    patternType: literal
type: topic
Type: Allow
  c: host: '*'
    operation: All
    > Expand Resource:
    name: '*'
    patternType: literal
type: group
Type: Allow
```

- Click Remove acl for any extra acls shown that are not configured.
- Click Create.

4. Create the Kafka topics.

- a. In the AMQ Streams Operator, navigate to Kafka Topic tab. Click Create KafkaTopic.

```
Name: cqin
Labels: strimzi.io/cluster=kafka
Partitions: 1
```



```
Replication factor: 3
Topic Name: cqin
```

b. Click Create.

c. Create another topic using the these values:

```
Name: cqinerr
Labels: strimzi.io/cluster=kafka
Partitions: 1
Replication factor: 3
Topic Name: cqinerr
```

d. Click Create.

e. Create another topic using the these values:

```
Name: sqin
Labels: strimzi.io/cluster=kafka
Partitions: 1
Replication factor: 3
Topic Name: sqin
```

f. Click Create.

g. Create another topic using the these values:

```
Name: sqout
Labels: strimzi.io/cluster=kafka
Partitions: 1
Replication factor: 3
Topic Name: sqout
```

h. Click Create

i. When you complete this section, a Kafka cluster, user and topics are created. You will need to collect the following details to complete the Apache Kafka Configuration in the Maximo Application Suite Administration UI.

- Kafka bootstrap hosts and ports
- username and password
- CA certificate

j. Next steps, configure the Suite parameters for Kafka in the Maximo Application Suite UI.

What to do next

Configure Maximo Application Suite parameters

Now you are ready to configure Apache Kafka details.

1. In the Maximo Application Suite instance, login to the **Administration** dashboard.
2. in **Other > Configurations**, select Apache Kafka. The following information is needed to configure the Apache Kafka details: Hosts/HostnamesUsername/passwordCertificates
Hosts - to obtain the bootstrap hosts, in the Red Hat OpenShift console.
3. In the Kafka project, go to **Networking > Routes** and search for the route kafka-kafka-tls-bootstrap.
4. Copy the value in the host field.
For example, kafka-kafka-tls-bootstrap-kafka.<yourdomain.com>.
5. The port number in the external route is 443. Enter the host name and port values in the hosts section.

6. To obtain the Kafka user's password, in the Red Hat OpenShift Kafka project, go to **Workloads > Secrets**. Search for your Kafka user, i.e. masuser. The data section will contain the user's password.
7. To obtain the certificates, during configuring the Apache Kafka parameters, click the Retrieve option to automatically retrieve the certificates from the Kafka bootstrap host.
Alternatively, you can obtain the certificate details prior to configuring the Maximo Application Suite parameters to enter the information manually.
 - a. To do this, in the Red Hat OpenShift console, switch to the Kafka project, and then navigate to Custom Resource Definitions. Search for Kafka.
 - b. Click the Instances tab and select your instance in the Kafka namespace.
 - c. Click YAML view. From this view, you can copy the certificate. The certificate will have BEGIN CERTIFICATE and END CERTIFICATE tags which must be included.

```
-----BEGIN CERTIFICATE-----
MIIDLTCcAhWgAwIBAgIJANfi6SPHo4cIM...
-----END CERTIFICATE-----
```

8. Copy the certificate text to be added in the Maximo Application Suite UI.
9. Log in to the Maximo Application Suite Admin User Interface and go to Administration.
10. Click Configurations
11. Click Apache Kafka
12. Add the tls bootstrap host name and port:
For example, Host Portxxx.xxx.xxx.xx.com 443
13. Enter the username and password
14. Enter an alias name. For example, strimzi
15. Add the copied certificate(s), and set the alias name then click confirm.
16. Click Save.
17. To confirm that the configuration is successful, in the OpenShift console, navigate to **Administration > Custom Resource Definitions**. Search for kafkacfg. Click the Instances tab. Click your instance, then view the YAML for any success or failure messages.

IBM Event Streams

Event Streams is an alternative for **AMQ Streams** for Kafka dependency which is available in IBM Cloud catalog.

1. To install an Event Streams instance in IBM Cloud, login to your **IBM Cloud account**, go to **Catalog** and search for **"Event Streams"**. Once you Click the "Event Streams" tile, go to **Create** tab and you will get to the provisioning details page where you will have to enter information regarding your **Event Streams instance**.
2. **Location** - It is recommended to choose a location that is close to the server/cluster location of your Maximo Application Suite instance for improved network performance.
3. **Pricing Plan** - Choose the plan that best fits your expected Kafka usage.
4. **Resource details** - Enter a Service name (it can be any unique name), and the optionally enter more details such as IBM Cloud resource group, and tags.
5. **Review the summary of your Events Streams instance**, review and accept the license agreement terms and click **Create**.

The Events Streams instance will be provisioned and you be redirected to the **Event Streams Home page**. Click **Service Credentials** from the menu. Click **"New credential"** to create a new **service credential** that contains the details that will be used to integrate **Event Streams** into your **Maximo Application Suite instance**.

- **Name:** Unique name for your service credential

- **Example:** Service credentials-1
- **Role:** Defined the level of permissions for your **Event Streams** instance
- **Example:** Manager (default)

When the **Event Streams service credentials is created**, expand it to see all the credential details. You will need the following information from **Event Streams service credentials** to configure it properly in Maximo Application Suite:

- **kafka_brokers_sasl** - Contains 6 hostnames for available Kafka brokers of your Event Streams instance.
- **user** - Kafka username, default is **token**
- **password** - Kafka password

Suite Configuration Parameters for event streams

Now you are ready to configure Event Streams into Maximo Application Suite.

1. Login to the Suite Administration dashboard of your Maximo Application Suite instance, go to **Other configurations > Configurations**.
2. Select Apache Kafka.
3. Enter the following information to configure Events Streams as a Kafka service for Maximo Application Suite:

- Hosts/Hostnames - Add a row for each of the six Kafka broker hostnames provided in the Event Streams service credential.

Note: Make sure you do not copy the port. Copy the Kafka broker hostname.

For example,

```
broker-0-<your-event-streams-broker-id>.kafka.svc07.us-south.eventstreams.cloud.ibm.com
broker-1...
....
broker-6-<your-event-streams-broker-id>.kafka.svc07.us-south.eventstreams.cloud.ibm.com
```

- Port - Enter the port associated to the kafka broker hostnames provided in the Event Streams service credential.

For example, 9093

- SASL Mechanism - Select plain. This is the default authentication mechanism for Event Streams.
- Username - Enter the user provided on Event Streams service credential.
- Password - Enter the password provided on Event Streams service credential.
- Certificates - Enter the chain of SSL certificates for your Event Streams instance.
- Click Add to add the intermediate of the certificate chain.
- Enter an alias.

For example, kafkacertpart1

- Enter the Certificate content. Here you will include the Let's Encrypt R3 intermediate certificate, issued to **US, Let's Encrypt, R3**. For more information about certificate content , see [here](#).

For example,

```
-----BEGIN CERTIFICATE-----
MIIF5jCCBM6gAwIBAgISA0Y...
-----END CERTIFICATE-----
```

4. Click Confirm. The first part of this certificate should have valid dates and look like the following example:

```
Issued to: US, Let's Encrypt, R3
Issued by: US, Internet Security Research Group, ISRG Root X1
```

```
Valid from: Thu Sep 01 2022
Valid to: Mon Sep 15 2025
```

This is the intermediate certificate which is required for the SSL connection to Event Streams endpoint.

5. Click Add to add the root of the certificate chain.
6. Enter an alias.

For example, `kafkacertpart2`

7. Enter the Certificate content. Here you will include the ISRG Root X1 cross-signed certificate, issued to **US, Internet Security Research Group, ISRG Root X1**. For more information about certificate content, see [here](#).

```
-----BEGIN CERTIFICATE-----
MIIFazCCA10gAw...
-----END CERTIFICATE-----
```

8. Click Confirm. The second part of this certificate should have valid dates and look like the following example:

```
Issued to: US, Internet Security Research Group, ISRG Root X1
Issued by: US, Internet Security Research Group, ISRG Root X1
Valid from: Thu Jun 04 2015
Valid to: Mon Jun 04 2035
```

This is the root certificate which is required for the SSL connection to Event Streams endpoint.

9. Save the Apache Kafka configuration.

Now, wait for the Apache Kafka configuration to reconcile, this process might take up to 10 minutes. The configuration will be successfully completed when the configuration status is set to Ready.

Configuration Ready - Kafka configuration was successfully verified

IBM® App Connect Enterprise

Use App Connect to connect your different applications and make your business more efficient. Set up flows that define how data moves from one application to one or more other applications. App Connect supports a range of skill levels and interfaces, giving you the flexibility to create integrations without writing a single line of code. You can use a web user interface or drop resources into a toolkit that gives a broader range of configuration options. Your entire organization can make informed business decisions by providing rapid access, visibility, and control over data as it flows through your business applications and systems.

Maximo Application Suite includes an entitlement to use IBM® App Connect Enterprise 12.0.1.0. Do not use an instance of IBM® App Connect Enterprise that is on IBM Cloud. You must install IBM® App Connect Enterprise in a Cloud Pak for Data cluster.

If you plan to use the data loader for Maximo Health and Predict - Utilities, you must configure IBM® App Connect Enterprise 12.0.2.0-r2, and your IBM App Connect dashboard must use an AppConnectEnterpriseProduction license.

Related tasks

Deploying IBM App Connect

With App Connect, you can connect applications and data from existing systems and modern technologies across all their environments.

Object Storage

Object storage is a computer data storage that manages data as objects, as opposed to other storage architectures like file systems which manages data as a file hierarchy, and block storage which manages data as blocks within sectors and tracks. This is commonly used by services and applications as persistent storage system.

Required by:

- Assist

Optional for:

- Manage (might be needed for persistent storage and attachments)

Cloud Object Storage is an alternative for the Object Storage dependency, which is available in the IBM Cloud catalog.

Related concepts

[Requirements and capacity planning](#)

[Provisioning storage](#)

Maximo Manage supports both ephemeral storage and persistent storage.

Related tasks

[Selecting S3-compatible object storage for IBM Maximo Collaborate](#)

Application database

To deploy Maximo Health, Maximo Real Estate and Facilities, or Maximo Manage, a database instance must be configured and running. The applications support Db2, Db2 Warehouse, Microsoft® SQL Server, or Oracle Database. If Maximo Health is deployed as part of Maximo Manage, the two applications share a database.

If you deploy IBM Maximo Health and Predict - Utilities, the industry solution uses the Maximo Health database.

Note: If you want to use Db2 Warehouse on Cloud Pak for Data for Maximo Health and Maximo Manage, you must configure a database at the workspace-application scope. You cannot use the Db2 Warehouse instance that is configured at the system scope.

Note: If you want to use Db2 Warehouse on Cloud Pak for Data for IBM Maximo Real Estate and Facilities, you must configure a database at the workspace-application scope. You cannot use the Db2 Warehouse instance that is configured at the system scope.

Related concepts

[Deploying Maximo Real Estate and Facilities in Maximo Application Suite](#)

Related tasks

[Deploying IBM Maximo Health](#)

Improve the reliability of your assets by proactively monitoring and managing asset health by using Maximo Health.

[Deploying IBM Maximo Manage](#)

To prepare for deployment, you must complete several tasks, such as configuring the database for Maximo Manage. Verify the compatibility of the industry solutions and add-ons if you want to deploy them and complete other required and optional preparation steps as needed. For example, if you are deploying in multiple languages, review the support languages that are available.

Related information

[Installing Db2 Warehouse](#)

External services

Configure external services such as SMTP server connection, certificate issuer, IBM Cloud Internet Services, Interactive Connectivity Establishment (ICE) protocol, and identity providers when you install the IBM Maximo Application Suite.

Simple Mail Transfer Protocol

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

You can configure SMTP as part of setup or later.

Related concepts

Simple Mail Transfer Protocol configuration

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

Related tasks

Setting up IBM Maximo Application Suite

Cloud Internet Services

IBM Cloud Internet Services, powered by Cloudflare, provides a fast, highly performant, reliable, and secure internet service for customers running their business on IBM Cloud. It is an alternative for a service that enables Domain Name Service (DNS) management which is required if you are planning to use custom cluster issuers signed by Let's Encrypt certificate authority for your Maximo Application Suite instance installation, instead of using the default self-signed certificates.

Interactive Connectivity Establishment (ICE) server

The ICE protocol is used to generate media traversal candidates which can be used in WebRTC applications, the two most prominent protocols for ICE servers are STUN and TURN. To ensure your collaboration works in different network, you need to configure at the least one STUN server and one TURN ICE server.

You can prepare and setup an open-source ICE server, which uses the STUN and TURN.

You can search and request some trial ICE services from the existing ICE service providers for temporary use.

For enterprise-level use, you can buy the ICE server/service from the company which provides the ICE services such as Twilio and Xirsys.

To use VoIP connections in Maximo Collaborate collaboration sessions, an Interactive Connectivity Establishment (ICE) server must be configured for Maximo Collaborate.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

The following types of ICE servers are supported:

- STUN
- TURN
- Twilio
- Xirsys

Required by

- Maximo Collaborate

Related tasks

Deploying IBM Maximo Collaborate

Identity providers

Identity provider prerequisites for Maximo Application Suite.

LDAP server

To use LDAP user registry with Maximo Application Suite, you need the following LDAP server information.

Important: To use the LDAP server for user and group synchronization, the server must support the secure LDAP (LDAPS) protocol. Non-TLS connections are not supported.

The following parameters are configurable:

- URL of your LDAP instance
- Bind DN, and Bind password.
- Base DN
- UserID Map

For more information, see [“Configuring LDAP authentication” on page 608](#)

SAML server

Configuring SAML user authentication for use with Maximo Application Suite is a multistep process:

1. Create SAML service provider information.

Your Maximo Application Suite server acts as service provider for the SAML identify provider (IdP). You need to provide a preferred service provider name and select a name identifier format, or use the defaults. The information is written to a service provider metadata file that you use to configure your SAML provider.

2. Register with the SAML provider.

Configure your SAML IdP to recognize Maximo Application Suite. Use the downloaded SP file and follow the information for your SAML provider to complete this step.

3. From your SAML IdP, download the SAML IdP metadata XML file to Maximo Application Suite.

The following parameters are configurable:

Service provider name

Use the default provided name or provide one of your own. This is the name that is used to register the Maximo Application Suite service provider.

Name identifier format

This is the format of the username identifier that is used with the SAML server.

For more information, see [“Configuring SAML authentication” on page 607](#)

OpenID Connect (OIDC)

Starting in Maximo Application Suite 9.1, you can configure OpenID Connect (OIDC) authentication with Maximo Application Suite for user authentication. For more information, see [“Configuring OIDC authentication” on page 609](#).

Related concepts

Authentication methods

Maximo Application Suite supports MongoDB, Lightweight Directory Access Protocol (LDAP) authentication, and Security Assertion Markup Language (SAML) authentication methods for local user authentication.

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Compatibility matrix

Information about the supported applications and add ons for each release of IBM Maximo Application Suite is available in the compatibility matrix.

Supported Maximo Application Suite releases

The following table shows the supported Maximo Application Suite releases for applications and add-ons.

Applications and add-ons	9.1	9.0	8.11	8.10
Maximo Collaborate*	9.0.x - 9.1	8.8 - 9.0	8.7 - 8.8	8.6 - 8.7
Maximo Real Estate and Facilities	9.1	N/A	N/A	N/A
Maximo Health and Predict - Utilities	9.0.x - 9.1	N/A	N/A	8.5 - 8.6
Maximo Manage	9.0.x - 9.1	8.7 - 9.0	8.6 - 8.7	8.5 - 8.6
Maximo Monitor	9.0.x - 9.1	8.11 - 9.0	8.10 - 8.11	8.9 - 8.10
Maximo Predict	9.0.x - 9.1	8.9 - 9.0	8.8 - 8.9	8.7 - 8.8
Maximo Safety	9.0.x - 9.1	N/A	N/A	N/A
Maximo Visual Inspection	9.0.x - 9.1	8.8 - 9.0	8.8 - 8.9	8.7 - 8.8
Maximo Optimizer	9.0.x - 9.1	8.5 - 9.0	8.4 - 8.5	8.3 - 8.4

*Starting in Maximo Application Suite 9.1, Maximo Assist is now called Maximo Collaborate.

For information about supported hardware and software dependencies for Maximo Application Suite, applications, and add-ons, follow these steps.

1. Go to the [Software Product Compatibility Reports \(SPCR\)](#) page.
2. From the menu, select the type of report to create. For example, **Detailed system requirements**.
3. Search for Maximo Application Suite, then select the product version for which to create a report.
4. Click **Submit** to create the report.
5. To view supported dependencies, select a tab such as **Containers, Prerequisites, Software, or Hardware**.

Note: For IBM Maximo Manage components, see [“Application-specific requirements for Maximo Manage”](#) on page 187.

Note: For the Maximo Real Estate and Facilities application compatibility matrix, see [“Application-specific requirements for Maximo Real Estate and Facilities”](#) on page 187.

Seamless upgrades

Maximo Application Suite administrators can independently upgrade the components of Maximo Application Suite. Maximo Application Suite application uses the n-1 compatibility with an earlier version model so that Maximo Application Suite administrators can independently upgrade applications in the suite. for more information, see [“Upgrading IBM Maximo Application Suite”](#) on page 473.

Unsupported application versions

When an application version becomes unsupported by a Maximo Application Suite release:

- Maximo Application Suite core upgrades are prevented if the unsupported version of the application is already installed.
- New application releases are no longer required to maintain compatibility with the unsupported version of the application.

For example, following the application compatibility with an earlier version model implement within Maximo Application Suite (n-1), assuming Assist 8.3 is released as part of Maximo Application Suite 8.6, when Maximo Application Suite 8.7 is released, Assist 8.1 will be unsupported.

SaaS **Maximo Application Suite as a Service overview**

IBM Maximo Application Suite as a Service is a cloud-based product that combines Maximo Application Suite functions on Red Hat OpenShift on Amazon Web Services.

Maximo Application Suite as a Service is available in the following editions:

- IBM Maximo Application Suite as a Service Essentials Edition
- IBM Maximo Application Suite as a Service Standard Edition
- IBM Maximo Application Suite as a Service Premium Edition

Each edition provides different options and combinations of capabilities in Maximo Application Suite to meet the business needs of your organization. For example, Maximo Application Suite as a Service Essentials Edition provides predefined configuration and applications to facilitate getting started quickly. Maximo Application Suite as a Service Standard Edition includes standardized configurations and a choice of applications to adapt to your business needs. Maximo Application Suite as a Service Premium Edition provides the most configurable options and flexible scheduling of upgrades to meet custom requirements.

For more information about choosing an edition, contact your IBM representative.

Note: Migrating from existing Maximo offerings, such as IBM Maximo Application Suite Dedicated or Maximo Asset Management to a Maximo Application Suite as a Service edition incurs a migration cost. For more information, contact your IBM representative.

Maximo Application Suite as a Service Essentials Edition

Maximo Application Suite as a Service Essentials Edition is available in two options that provide predefined configurations and applications to access enterprise asset management and mobile features. With Maximo Application Suite as a Service Essentials Edition, you can choose either Maximo Maintenance Essentials or Maximo Inspection Essentials.

With Maximo Maintenance Essentials, Maximo Application Suite includes the following suite applications and capabilities.

- IBM Maximo Manage
- IBM Maximo Health
- IBM Maximo Mobile

With Maximo Inspection Essentials, Maximo Application Suite includes the following suite application.

- IBM Maximo Visual Inspection

Maximo Application Suite as a Service Essentials Edition (9.1)

Maximo Application Suite as a Service Essentials Edition is available in five options, which provide predefined configurations and applications to access enterprise asset management, real estate and facilities management, and mobile features.

With Maximo Maintenance Essentials, Maximo Application Suite includes the following suite applications and capabilities.

- IBM Maximo Manage
- IBM Maximo Health
- IBM Maximo Mobile

With Maximo Inspection Essentials, Maximo Application Suite includes the following suite application.

- IBM Maximo Visual Inspection

With Maximo Lease Management Essentials, Maximo Application Suite includes the following suite application.

- IBM Maximo Real Estate and Facilities, with Lease Administration and Accounting.

With Maximo Space Management Essentials , Maximo Application Suite includes the following suite application.

- IBM Maximo Real Estate and Facilities, with Space and Reservation Management.

With Maximo Capital Planning Essentials, Maximo Application Suite includes the following suite application.

- IBM Maximo Real Estate and Facilities, Capital Projects, and Facility Condition Assessment

Maximo Application Suite as a Service Standard Edition

Note: If you are an existing Maximo Application Suite as a Service customer, your current instance is the same as Maximo Application Suite as a Service Standard Edition.

Maximo Application Suite as a Service Standard Edition includes all options from the essential editions and provides standardized configurations and availability of a wider range of capabilities.

Maximo Application Suite as a Service Standard Edition includes Maximo Application Suite with the choice of the following applications, industry solutions, and add-ons.

Applications

- IBM Maximo Manage
- IBM Maximo Collaborate
- IBM Maximo Health
- IBM Maximo Monitor
- IBM Maximo Predict
- IBM Maximo Visual Inspection

Industry solutions

- IBM Maximo Utilities
- IBM Maximo Oil & Gas
- IBM Maximo Nuclear
- IBM Maximo Transportation
- IBM Maximo Aviation
- IBM Maximo Civil Infrastructure

Add-ons

- IBM Maximo Service Provider
- IBM Maximo Health, Safety and Environment
- IBM Maximo Asset Configuration Manager
- IBM Maximo Spatial
- IBM Maximo Connector for Oracle applications
- IBM Maximo Connector for SAP Applications
- IBM Maximo Optimizer Limited
- IBM Maximo IT
- IBM Maximo Vegetation Management

In addition, the following add-ons that are specific to SaaS deployment and operations are available.

- Read-only database replicas

- Maximo Models for Electrical Distribution

Maximo Application Suite as a Service Premium Edition

Maximo Application Suite as a Service Premium Edition includes all options from the standard edition and also the availability of a wider range of capabilities. Maximo Application Suite as a Service Premium Edition provides the most configurable options and flexible scheduling of upgrades for custom and complex deployments and includes assigned support personnel.

Additional options can be purchased.

Switching editions

You can change to a different Maximo Application Suite as a Service edition, such as updating from Maximo Application Suite as a Service Essentials Edition to Maximo Application Suite as a Service Standard Edition or Maximo Application Suite as a Service Premium Edition. You can also change from Maximo Application Suite as a Service Premium Edition to Maximo Application Suite as a Service Standard Edition. Depending on your configuration, some migration paths might not be possible.

Changing from Maximo Application Suite as a Service Essentials Edition requires the purchase of additional AppPoints. If you change from Maximo Application Suite as a Service Standard Edition to the premium edition, you might need to purchase additional AppPoints.

For more information about changing Maximo Application Suite as a Service editions, contact your IBM representative.

Related concepts

[Getting started as a SaaS suite administrator](#)

Get started in Maximo Application Suite as a Service applications by adding users and specifying their entitlements, requesting server authentication and user synchronization, and monitoring application usage.

[Applications, industry solutions, add-ons, accelerators, and tools](#)

IBM Maximo Application Suite includes many applications, industry solutions, add-ons, and tools.

[What's new in Maximo Application Suite as a Service](#)

Related information

[IBM Maximo Application Suite as a Service operations documentation](#)

What's new

Learn about new features and capabilities.

Related information

[What's new in IBM Maximo Manage](#)

[What's new in IBM Maximo Health, IBM Maximo Predict, IBM Maximo Health and Predict](#)

[What's new in IBM Maximo Monitor](#)

[What's new in IBM Maximo Visual Inspection](#)

[What's new in IBM Maximo Assist](#)

What's new in the Maximo Application Suite feature channel

Learn more about what's new in the feature channel for nonproduction instances of Maximo Application Suite and for production instances of Maximo Application Suite as a Service.

Maximo Application Suite incrementally delivers updates to provide new feature capabilities to the feature channel.

As a customer-managed user, you can use the feature channel to update your nonproduction instances to preview new features.

As a SaaS user, you can use new features in your Maximo Application Suite as a Service environment.

You can subscribe to the feature channel by using a channel subscription. For more information, see [“Upgrading IBM Maximo Application Suite by using the channel subscription method”](#) on page 477. Information that provides bug fixes and security updates is available in the release notes. For more information, see [Maximo Application Suite feature channel release notes](#).

September 2025

Multiple architecture support for industry solutions and add-ons

The following industry solutions and add-ons now support architecture IBM Power® (ppc64le) to use existing Red Hat OpenShift capabilities through the Red Hat OpenShift Container Platform for scalability and resilience.

- IBM Maximo Asset Configuration Manager
- IBM Maximo Reliability Strategies
- IBM Maximo IT

SaaS

IBM Maximo Vegetation Management

Maximo Vegetation Management is a new SaaS add-on that you can purchase if you already have Maximo Manage. Use the add-on to understand and manage the state of vegetation across your entire service area. For more information, see [Managing vegetation](#).

To learn more about what’s new in the feature channel for suite applications, see [“Applications, industry solutions, and add-ons”](#) on page 44.

Applications, industry solutions, and add-ons

What's new in applications, industry solutions, and add-ons

Learn more about what’s new in the feature channel for the following applications, industry solutions, and add-ons.

- [What's new in the feature channel for Maximo Manage](#)
- [What's new in the feature channel for Maximo Mobile](#)

Customer-managed

What's new in Maximo Application Suite 9.1

Learn more about what's new and changed in IBM Maximo Application Suite 9.1.

Licensing in Maximo Application Suite

The licensing guidance provides information and updates to entitlements in Maximo Application Suite 9.1. For more information, see [“Licensing in Maximo Application Suite 9.1”](#) on page 78

Prerequisite software changes

Multiple architecture support for Maximo Application Suite core and Maximo Manage base

Maximo Application Suite core and Maximo Manage base now support architecture for IBM System/390x (S390x) and IBM Power (ppc64le) to use existing Red Hat OpenShift capabilities through the Red Hat OpenShift Container Platform for scalability and resilience.

You can choose to install or migrate your existing Maximo Application Suite core and Maximo Manage base on the IBM Z processing platform.

If Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x or IBM Power (ppc64le) architecture, consider the following conditions for prerequisite software.

- MongoDB is supported only as an external service. For more information, see [“MongoDB”](#) on page 21.

- IBM Db2 can be used as an external service. For more information, see [“Configuring IBM Db2” on page 302.](#)
- Grafana operators are not supported. For more information, see [“Installing Grafana” on page 826.](#)

For more information, see [Software Product Compatibility Reports \(SPCR\).](#)

Installation

Foundation service

The foundation service consists of selected functions that unify setting up, accessing, and viewing capabilities that were previously separate and specific to each suite application. By consolidating these common administrative tasks and operations, administrators can centrally administer users, security, and data and also configure or customize applications. For more information, see [“Foundation service” on page 136.](#)

Upgrade

Upgrading to Maximo Application Suite 9.1

When you upgrade to Maximo Application Suite 9.1, either Maximo Manage or the foundation service is required. The upgrade path to Maximo Application Suite 9.1 depends on your current installation and configuration. For more information, see [“Upgrading to Maximo Application Suite 9.1” on page 482.](#)

Configuration

Configuring applications with Application configuration

You can configure Maximo Application Suite applications with Application configuration. For more information, see [“Application Configuration” on page 856.](#)

Configuring cross-origin resource sharing (CORS)

You can configure CORS in Maximo Application Suite for WebSphere Liberty server by updating the custom resource file in Red Hat OpenShift Container Platform. For more information, see [“Configuring cross-origin resource sharing \(CORS\)” on page 717.](#)

Customizing hostAliases in podTemplates

Map a hostname to an IP address and then you can effectively bypass DNS resolution and create direct mapping between the hostname and the IP address. For more information, see [“Customizing hostAliases in podTemplates” on page 706.](#)

Disabling or hiding login options

A system administrator or an identity provider (IdP) administrator can hide or disable identity providers so that they are not accessible to users. Administrators can continue to maintain the existing IdP configuration and account links. For more information, see [“Disabling or hiding login options by using APIs” on page 713](#)

User management

Configuring multiple LDAP user registry synchronizations

You can configure multiple LDAP user registry synchronizations to synchronize users and groups from more than one LDAP server to your local Maximo Application Suite user registry. For more information, see [“Configuring multiple LDAP user registry synchronizations” on page 631.](#)

OpenID Connect

You can configure single sign-on with OpenID Connect (OIDC) for authentication and authorization to enable users to log in and access suite applications through a centralized identity provider. For more information, see [“Configuring OIDC authentication” on page 609.](#)

IDP administrator access

Administrators who have IDP administrator can configure the integrations, such as LDAP, SAML, SMTP, User registry synchronization, Certificate management, and User management.

For more information, see [“User access and entitlements in Maximo Application Suite 9.1”](#) on page 782

Password expiration for local users

You can enable password expiration to define when a user is required to change their password. You can specify how long the password is valid before a user must change their password. You can include a grace period during which users can still log in after their password expires. You can also send an email to users to notify when a password is scheduled to expire. For more information, see [“Setting password expiration for local users”](#) on page 810

Administering user sessions

Force log out

You can configure user sessions to automatically log out a user within a defined time period. For more information, see [“Configuring force user log out”](#) on page 794.

Enabling maintenance mode

To prevent all users from accessing the system during critical maintenance or when changes are being made, you can enable maintenance mode by using APIs. After you enable maintenance mode, active users are logged out, and login access is restricted for all users. Only system administrators, identity provider (IdP) management administrators, and super users can log in during maintenance mode. For more information, see [“Enabling maintenance mode by using APIs”](#) on page 795.

Using translated versions of email notifications

You can configure the Smtpcfg custom resource to use a translated version of Maximo Application Suite emails.

For more information, see [“Changing the language of email notifications”](#) on page 642.

Disabling Maximo Application Suite email notifications

You can configure the Smtpcfg custom resource to disable Maximo Application Suite emails from being sent to your users.

For more information, see [“Disabling email notifications”](#) on page 640.

Authorization services

Administering users

You can now centrally manage user records, their authorizations, and also security privileges in Maximo Application Suite. You can specify whether the user is granted concurrent or authorized access and also assign the user to security groups that gives them specific access to the applications and capabilities.

For more information, see [Administering users](#).

Note: Existing Maximo Manage 9.0.x users can continue to manage users in Maximo Manage by selecting the Users (Manage) application.

For more information, see [Managing users and groups](#).

Creating a user administrator

To set up a user as an administrator who can grant other users access to the suite applications, you assign them user management permissions. For more information, see [“Creating administrators for user management”](#) on page 787.

Administering security groups

You can provide authorization to suite applications, actions, and data to a group of users by using security groups. Security groups are used to control user access to applications, data, and actions by grouping users and assigning permissions. For more information, see [Administering security groups](#).

Note: If Maximo Manage is deployed, administrators can continue to use security groups in Maximo Manage, by selecting the Security groups (Manage) application. For more information, see [Managing users and groups](#).

User management APIs

New user management APIs replace the existing user management APIs. The new user management APIs provide greater granularity for user authorization to allow for a more flexible and secure way to manage access within the application. For more information, see [“User management APIs” on page 793](#)

Common navigation

New suite navigation menu

To unify and improve access to suite applications, you can navigate to the suite applications that you have access to by using the side navigation menu. This new navigation menu replaces the **Application switcher** and the administration icon on the header menu. A Read permission can be granted for each application so that it is visible in the side navigation menu.

Streamlined content access

To consolidate access to content in the user interface, the content that was available from the **Take a tour** button is now available in the **Help** menu. To access content, such as interactive tours, links to videos, and what's new information, select **Help > Show me how**.

Suite administration

Application-specific language support

You can now select languages such as Arabic and Hebrew from the user profile page. These languages are application-specific and supported only in IBM Maximo Manage and IBM Maximo Real Estate and Facilities.

For more information, see [“Language and locale support” on page 127](#) and [“Setting language and time zone preferences for users” on page 807](#).

Importing data

You can import data from a .csv file to update the database and simultaneously create multiple records, such as user records. You can use this file to add the data and ensure that the format adheres to the import processing rules. For more information, see [“Importing data” on page 818](#). For an example about how to use the importing data action to create user records, see [“Importing users in Maximo Application Suite 9.1” on page 787](#).

Exporting data

You can also export data from Maximo Application Suite as a .csv or JSON file to view or update the data. For example, if you want to view or edit user profiles, you can export and update the information in the file and import the updated file. For more information, see [“Exporting data” on page 819](#).

Disabling survey

The customer satisfaction survey gathers feedback about your experience with Maximo Application Suite. When enabled, the survey is available to all users to submit feedback. You can disable surveys. For more information, see [“Disabling surveys” on page 717](#).

Applications, industry solutions, and add-ons

New IBM Maximo Real Estate and Facilities suite application

IBM Maximo Real Estate and Facilities is now available in Maximo Application Suite. With Maximo Real Estate and Facilities, you can manage real estate portfolios and facility assets throughout their lifecycle with space management, reservations, capital projects, facility condition assessment, lease management, operations, and maintenance. For more information, see [“IBM Maximo Real Estate and Facilities” on page 71](#).

You can deploy and activate IBM Maximo Real Estate and Facilities as a suite application, create users, and assign user access. For more information, see [“Deploying Maximo Real Estate and Facilities in Maximo Application Suite ” on page 373](#) and [“Administering users and user access in Maximo Application Suite in 9.1” on page 781](#).

New Maximo Asset Investment Planning add-on in Maximo Manage

Maximo Asset Investment Planning is now available in Maximo Application Suite. Maximo Asset Investment Planning is a strategic planning tool that is designed to help asset-intensive organizations make better decisions about where and when to invest in their assets, while improving planning processes. For more information, see [Maximo Asset Investment Planning](#).

You deploy and activate Maximo Asset Investment Planning as an add-on to Maximo Manage. For more information, see [Deploying Maximo Manage](#).

Stand-alone Maximo Health no longer a suite application

Maximo Health is no longer available as a stand-alone suite application but remains an add-on in Maximo Manage.

If Maximo Health is deployed as a stand-alone suite application and you are upgrading to 9.1, you must add Maximo Health as an add-on in Maximo Manage 9.0. Then, you upgrade to 9.1. For more information, see [Upgrading Maximo Health 9.0 stand-alone to Maximo Manage 9.1 with Health](#)

IBM Maximo Assist renamed IBM Maximo Collaborate

Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate. If Maximo Assist is deployed in Maximo Application Suite 9.0 or earlier and you upgrade to Maximo Application Suite 9.1, the name is automatically changed in the user interface.

Note: In the Maximo Application Suite documentation, Maximo Assist is now referred to as Maximo Collaborate.

IBM Maximo Optimizer

For optimization models, disable the **Allow model customization artifacts** configuration option. For more information, see [“Activating IBM Maximo Optimizer”](#) on page 469.

On a single-node Red Hat OpenShift cluster, the Maximo Application Suite CLI utility installs the Maximo Optimizer Limited plan. For more information, see [“IBM Maximo Optimizer”](#) on page 75.

What's new in accelerators

The Red Hat Marketplace is deprecated. Maximo Application Suite accelerators are now hosted on the [Maximo Application Suite accelerators solutions page](#). For more information, see [“Maximo Application Suite accelerators”](#) on page 76.

What's new in Maximo Application Suite applications, industry solutions, and add-ons

Learn more about what's new and changed in the following applications, industry solutions, and add-ons in Maximo Application Suite 9.1.

- [What's new in Maximo Civil Infrastructure 9.1](#)
- [What's new in Maximo Collaborate 9.1](#)
- [What's new in Maximo Health 9.1](#)
- [What's new in Maximo Health, Safety and Environment 9.1](#)
- [What's new in Maximo IT 9.1](#)
- [What's new in Maximo Manage 9.1](#)
- [What's new in Maximo Mobile 9.1](#)
- [What's new in Maximo Monitor 9.1](#)
- [What's new in Maximo Oil & Gas 9.1](#)
- [What's new in Maximo Predict 9.1](#)
- [What's new in Maximo Real Estate and Facilities 9.1](#)
- [What's new in Maximo Visual Inspection 9.1](#)
- [What's new in Maximo Visual Inspection Edge 9.1](#)

What's new in Maximo Application Suite 9.0

Learn more about what's new and changed in IBM Maximo Application Suite 9.0.

Version alignment

The version for Maximo Application Suite, including all applications, industry solutions, and add-ons in Maximo Application Suite is now aligned, starting on version 9.0.

Version 9.0 is compatible with version 8.11, and you can upgrade to version 9.0 without the need to complete separate migration tasks.

Prerequisite software changes

Support for Internet Protocol version 6

Install, manage, and update Maximo Application Suite in an Internet Protocol version 6 (IPv6) environment.

Support for MongoDB 5.0 and 6.0

MongoDB 5.0 and 6.0 are supported in Maximo Application Suite. For more information, see [Update for MongoDB 5.0 and 6.0](#).

IBM Data Reporter Operator

As a Maximo Application Suite administrator, you configure IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance.

The User Data Services is now deprecated. Data Reporter Operator replaces User Data Services.

Data Reporter Operator has a reduced operational footprint and cost than IBM User Data Services to help you reduce the Maximo Application Suite resource usage.

If you are an existing Maximo Application Suite user, migrate to Data Reporter Operator because it includes tools that assist with migration. If you do not migrate to Data Reporter Operator, you must manually report metrics to IBM.

For more information, see [“Data Reporter Operator”](#) on page 7.

Upgrade

Rolling back Maximo Assist

Starting in Maximo Application Suite 9.0 after you upgrade to a later version of Maximo Assist, you can roll back to a previous version by updating the version properties.

For example, you can roll back Maximo Assist from 9.0.1 to 9.0.0. However, you cannot roll back Maximo Assist from 9.0.0 to 8.11.

For more information, see [Rolling back Maximo Collaborate](#).

User synchronization

Custom mapping with LDAP user synchronization

Custom mapping for user data that is synchronized between your LDAP server and the user registry is available in the user interface. By using user registry synchronization, you can specify custom map values for user and group data to synchronize with the LDAP server. Alternatively, you can use the default map values that are defined by the system.

For more information, see [“User and group registry mapping”](#) on page 622.

SCIM 2.0 protocol support for user and group synchronization

You can synchronize users and groups from an external identity provider (IdP) by using the System for Cross-domain Identity Management (SCIM) 2.0 protocol. The SCIM specification defines a common

schema for user and group resources with a protocol that defines how operations are performed on those resources. SCIM facilitates the exchange of these resources across different services.

For more information, see [“User synchronization with SCIM 2.0” on page 625](#) and [“Configuring Maximo Application Suite to synchronize user and groups with SCIM 2.0” on page 630](#).

User authentication

Initiated logout for the SAML service provider

When you configure SAML authentication, you can enable an initiated logout for the SAML service provider so that current user sessions are logged out before another user logs in with the same credentials.

For more information, see [“Configuring SAML authentication” on page 607](#).

Support for multiple IDPs for the same authentication type

You can now configure multiple identity providers (IDPs) for the same authentication type, such as SAML or LDAP, in Maximo Application Suite for user authentication. For more information, see [“Configuring multiple identity providers for same authentication type” on page 617](#).

Self-registration for user authentication

Users can self-register to create their own login accounts and use the applications that they have access to. Before users can self-register, an administrator must enable and configure access options that are associated with each identity provider that is configured.

For more information, see [“Self-registration for users” on page 613](#).

Configuration

Supported pods for IBM Data Dictionary and the IoT tool

You can customize the workload in the IBM Data Dictionary and the IoT tool by configuring the pods that are handled by custom resource objects.

For more information, see [“Supported pods for IBM Data Dictionary” on page 702](#) and [“Supported pods for IoT tool” on page 682](#).

Configuring the size of public certificate resources

You can change the private key size of public certificates that are provided by Maximo Application Suite.

For more information, see [“Configuring the size of public certificate resources” on page 591](#).

Db2 text search

Configure the Maximo Manage operator to support text search for Db2. When the operator starts the **maxinst** process, the **-q** parameter is used to enable text search for all columns that are flagged as searchable.

For more information, see [“Configuring and activating Maximo Manage” on page 342](#).

Configuring single sign-on properties

You can change the default single sign-on (SSO) token name **LTPAtoken2** to avoid conflicts with the same cookie name that is generated by other software.

For more information, see [“Configuring single sign-on properties” on page 616](#).

Changing privacy access for obtaining user data

You can configure the Suite custom resource (CR) file to control whether user information is available to all users.

For more information, see [“Changing privacy access for obtaining user data” on page 646](#).

Suite administration

User interface customization

As a suite administrator, you can change the appearance of the user interface to include your own logos, company branding, and custom visual styles. You can add CSS styling and globally update these changes across Maximo Application Suite.

For more information, see [“Updating the user interface” on page 708](#).

Account lockout

To increase application and user security for local authentication, you can define the conditions that prevent users from logging in after consecutive unsuccessful login attempts. You can define the number of consecutive password attempts before the users account is locked. You can also choose to lock the account by duration or until an administrator unlocks the account.

For more information, see [“Enabling account lockout” on page 811](#).

Anonymizing personal information

You can enable the option to anonymize personal information, such as username, emails, and display name, before you delete users. As an administrator, you can also anonymize the personal information as a global option by updating the custom resource in the Red Hat OpenShift web console. Then, when you delete users, this personal information is anonymized by default.

For more information, see [“Deleting and anonymizing user data” on page 808](#).

Configuring guided tours

Maximo Application Suite provides in-app guidance, such as guided tours, to help users learn more about different tasks and updates in the product. You can configure the user interface to hide guided tours.

For more information, see [“Hiding guided tours” on page 716](#).

User assistance

Overview of Licensing and AppPoints

Watch a video to understand Maximo Application Suite's customer-managed licensing model, and how to manage usage with AppPoints.

Applications, industry solutions, accelerators, and add-ons

Removal of MRO Inventory Optimization

Starting in Maximo Application Suite 9.0, MRO Inventory Optimization is no longer available to be launched from the Suite navigator as an externally configured application and must be accessed by the dedicated URL.

If MRO Inventory Optimization is configured as an external launcher and you are upgrading to Maximo Application Suite 9.0, you must remove MRO Inventory Optimization before you can complete the upgrade. To remove MRO Inventory Optimization as an external launcher, you delete the solution portal URL for MRO Inventory Optimization on the **External launcher** page. When the product URL is removed, users can no longer access it in the suite navigator.

For more information, see [“Configuring external launchers” on page 644](#).

Removal of Watson Discovery dependency in Maximo Assist

Starting in Maximo Application Suite 9.0, Watson Discovery, which is used to support the query and diagnose functions, is no longer available as a dependency in Maximo Assist. If Maximo Assist is already deployed and activated with Watson Discovery, and you are upgrading to Maximo Application Suite 9.0, before you can complete the upgrade, you must contact IBM Support to help with the manual removal of Watson Discovery.

Removal of voice inspections in Maximo Assist

Starting in Maximo Application Suite 9.0, voice inspections are no longer available in Maximo Assist. If voice inspections are enabled and you are upgrading to Maximo Application Suite 9.0, this feature is automatically removed during the upgrade.

Removal of Search in Maximo Assist

Starting in Maximo Application Suite 9.0, the search function is no longer available in Maximo Assist. If search is enabled and you are upgrading to Maximo Application Suite 9.0, the search function is automatically removed during the upgrade.

What's new in Maximo Application Suite applications, industry solutions, and add-ons

Learn more about what's new and changed in the following applications, industry solutions, and add-ons in Maximo Application Suite 9.0.

- [What's new in Maximo Assist 9.0](#)
- [What's new in Maximo Civil Infrastructure 9.0](#)
- [What's new in Maximo Health 9.0](#)
- [What's new in Maximo Health, Safety and Environment 9.0](#)
- [What's new in Maximo IT 9.0](#)
- [What's new in Maximo Manage 9.0](#)
- [What's new in Maximo Mobile 9.0](#)
- [What's new in Maximo Monitor 9.0](#)
- [What's new in Maximo Oil & Gas 9.0](#)
- [What's new in Maximo Predict 9.0](#)
- [What's new in Maximo Visual Inspection 9.0](#)
- [What's new in Maximo Visual Inspection Edge 9.0](#)

Related concepts

[What's new in Maximo Application Suite fix packs](#)

Learn more about updates that are provided in Maximo Application Suite fix packs to support prerequisite software, such as Red Hat OpenShift cluster services and Cloud Pak for Data services.

Customer-managed

What's new in Maximo Application Suite 8.11

Learn more about what's new and changed in IBM Maximo Application Suite 8.11.

Configuration

Manage trusted certificate authorities

Maximo Application Suite comes with a built-in set of certificate authority (CA) certificates and automatically trusts a certificate by default if the certificate is issued by one of these certificate authorities.

To disable the trust in the default certificate authority (CA) certificates, you can update the custom resource (CR) file for Maximo Application Suite. For more information, see [“Disabling default certificate authorities”](#) on page 589.

If you disable the default trust, then you need to specifically configure certificates and certificate authorities for all external systems that Maximo Application Suite connects to. For more information, see [Configuring certificate authority certificates](#).

Customizing the workload

As an administrator, you can manually configure workloads so that Maximo Application Suite can scale them to match demand. You can modify the supported pod's specifications, such as replicas,

container resources, affinity, anti-affinity, and tolerations. For more information, see [“Customizing workloads”](#) on page 648.

Horizontal and vertical pod scaling

You can scale horizontally or vertically by using podTemplates by setting their resources and replicas. Depending on the requirements of your workloads, you can set the values for these replicas to less than or more than the default value.

For more information, see [“Customizing workload scale”](#) on page 703.

Pod scheduling with affinity and anti-affinity

Affinity is one of the key features available in Kubernetes to customize and improve control of the pod scheduling process. The Kubernetes pod and node affinity and anti-affinity rules enable administrators to control where pods are scheduled. Specifying multiple rules helps facilitate a wide range of scheduling configurations.

With affinity and anti-affinity, administrators can:

- Define rules, including conditions with logical operators.
- Create preferred and required rules for a greater variety of matching conditions.
- Match the labels of pods that are running within nodes and determine the scheduling location of new pods.

For more information, see [“Customizing workload affinity”](#) on page 705.

Pod scheduling with tolerations

By using toleration provided by Kubernetes, you can schedule pods on a node that has a matching a key value pair that is assigned to a node.

For more information, see [“Customizing workload tolerations”](#) on page 706.

Suite administration

Configuring multiple login options

By configuring local, SAML, and LDAP authentication as identity providers (IdP) and associating each identity provider with user records, you can provide multiple login options for user authentication. If you configure more than one identity provider, you can also specify a default identity provider to be the primary login option for users on the suite login page.

If you specify SAML as the default identity provider, you can enable seamless login so that users authenticate to Maximo Application Suite by using the login page that uses the SAML identity provider.



Attention: If you enable seamless login, the Maximo Application Suite login page is not shown. If you need to display a security message to comply with federal regulations, ensure that seamless login is disabled. Otherwise, users do not see any system notification that might be shown on the Maximo Application Suite login page.

To configure login options, on the **Suite administration** page, from the side navigation menu, select **Users** and then click the **Authentication** tab.

For more information, see [User authentication](#).

Login notification

You can create and display a system message on the login page to provide security and privacy information to users. To enable the message, on the **Suite administration** page, from the side navigation menu, select **Users** and then click the **Authentication** tab. In the Login notification section, enable the message and enter the information. After you save, the message is shown on the login page.

For more information, see [Enabling login notification](#).

Activate and deactivate users

When system access for a user account is no longer needed, you can now deactivate that particular user account instead of deleting it. Deactivated users cannot log in to Maximo Application Suite. The suite administrator can activate or deactivate multiple user accounts in the user interface or by importing user details. You can deactivate a user account automatically by a date or a specified number of days of inactivity. If the user was an authorized user with permanently reserved AppPoints, then the allocated AppPoints are returned to the pool of AppPoints. When you activate an authorized user, AppPoints are reserved from the pool. If not enough AppPoint are available, then the user cannot be activated. For more information, see [Setting user account status](#).

Enhancement to delete users

Starting in Maximo Application Suite 8.11, the system retains the details of the deleted users.

The user deletion behavior is unchanged. The deleted users do not appear in the UI and are not included in any API responses. Any reserved AppPoints are relinquished, and the users cannot login to Maximo Application Suite. The only noticeable difference is that the user IDs, usernames, and email addresses of deleted users are not available for reuse by existing or new users.

Support for roll back and version lock

Starting in 8.11, you can roll back Maximo Application Suite to an earlier version. Maximo Application Suite can be locked to use a specific version even when the operator is updated or upgraded.

For more information, see [Setting version lock in Maximo Application Suite](#).

API keys for user management

You can generate and manage application programming interface (API) keys to use in automation processes for user management in IBM Maximo Application Suite.

For more information, see [“Generating and managing API keys” on page 815](#).

Audit log generation to support the Federal Information Security Management Act (FISMA)

Maximo Application Suite and its applications now support the generation of audit logs for some controls to support the Federal Information Security Management Act (FISMA).

As a Suite administrator, you can forward all logging from Red Hat OpenShift into an external system so that logs can be aggregated and securely stored. For more information, see [“Audit logging in Maximo Application Suite” on page 817](#).

Configuring passwords settings page moved to local authentication

Configuring passwords for local authentication was previously available by clicking **Users** from the side navigation menu and then clicking the **Password settings** tab. You can now configure passwords for local authentication on by clicking **Users** and then clicking the **Authentication** tab. In the Identity providers section, in the row for the local identity provider, click the **More actions** icon.

For more information, see [Configuring password settings](#).

Maximo Application Suite enhancements on Microsoft Azure

Support for new databases

Starting in Maximo Application Suite 8.11, configure databases such as Microsoft SQL Server or Oracle Database for using IBM Maximo Manage. These databases can be hosted on private subnet of another VPC.

For more information see, [“Maximo Application Suite offering type” on page 168](#).

Maximo Application Suite enhancements on Amazon Web Services

Support for new databases

Starting in Maximo Application Suite 8.11, configure databases such as Microsoft SQL Server or Oracle Database for using IBM Maximo Manage. These databases can be hosted on private subnet of another VPC. VPC peering is used to establish connection between the VPC of Amazon Web Services stack to the VPC of a database to establish database connection when the stack is deployed.

For more information, see [“Maximo Application Suite offering type for Amazon Web Services”](#) on page 147.

Support for selecting EBS volume type

Starting in Maximo Application Suite, support for selecting either gp3 or io1 as EBS volume type for worker nodes.

For more information, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

FIPS and FISMA support for Maximo Application Suite on Amazon Web Services Marketplace in US GovCloud regions

A Federal Information Processing Standard (FIPS) and Federal Information Security Management Act (FISMA) supported Maximo Application Suite can now be installed in the Amazon Web Services US GovCloud regions. The BYOL option is available for installing Maximo Application Suite in private hosted zone for existing Red Hat OpenShift cluster and New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI).

For more information, see [“Installing BYOL IBM Maximo Application Suite”](#) on page 222.

User assistance improvements

Upgrading from IBM Maximo Asset Management to IBM Maximo Manage

Watch a video to understand the changes in the new documentation for upgrading from Maximo Asset Management to Maximo Manage.

Managing users video

Watch this new video that shows how to manage users in Maximo Application Suite 8.11.

Applications, industry solutions, accelerators, and add-ons

IBM Maximo IT

IBM Maximo IT is now available in Maximo Application Suite. Maximo IT provides a single point of user support and enterprise service management of information technology (IT) and operational technology (OT) assets and processes.

For more information, see [Maximo IT in IBM Documentation](#).

You can deploy and activate Maximo IT as an add-on to Maximo Manage. Before you deploy Maximo IT, accept the license terms in Passport Advantage®. For more information, see [Deploying Maximo Manage](#).

Reliability Strategies

Reliability Strategies is now available in Maximo Application Suite. By using Reliability Strategies, you can access a library of maintenance strategies that are based on reliability-centered maintenance and include both failure details and mitigation activities for each failure. The library was developed by industry and domain experts and contains hundreds of assets and tens of thousands of possible failures across all known operating contexts. For more information, see [Reliability Strategies](#).

You can deploy and activate Reliability Strategies as an add-on to Maximo Manage. For more information, see [Deploying Maximo Manage](#).

Before you begin using Reliability Strategies, you must complete some configuration steps. For more information, see [Configuring Reliability Strategies](#).

New accelerators in the Maximo Application Suite catalog

Accelerators are solutions that are provided by IBM and IBM partners to complement or extend Maximo Application Suite capabilities. These accelerators are hosted on the [Red Hat Marketplace](#), which you can integrate with Maximo Application Suite and access from the Maximo Application Suite catalog.

For more information, see [“Maximo Application Suite accelerators”](#) on page 76.

IBM Maximo Models for Electrical Distribution

IBM Maximo Models for Electrical Distribution provides model templates to enhance the reliability of energy and utility assets to anticipate downtime, degradation, and failures. This accelerator includes prebuilt notebooks and configuration files for several transmission and distribution asset classes, incorporating features available in IBM Maximo Health and IBM Maximo Predict. Maximo Models for Electrical Distribution replaces IBM Maximo Health and Predict - Utilities in Maximo Application Suite 8.11.

For more information, see [Overview of Maximo Models for Electrical Distribution](#)

Maximo Manage has new function to reduce system downtime when you do an update

After you upgrade to IBM Maximo Application Suite operator version 8.11, configure a database state to reduce Maximo Manage system downtime as you upgrade to Maximo Manage 8.7. and subsequent Maximo Manage upgrades. For more information, see [“Updating IBM Maximo Manage”](#) on page 484 and [“Reducing system downtime”](#) on page 323.

Removal of Maximo Health and Predict - Utilities and Parts Identifier

Starting in Maximo Application Suite 8.11, the following industry solution and add-on are no longer available:

- Maximo Health and Predict - Utilities
- Parts Identifier

If either of these are deployed and active in your environment and you are upgrading to Maximo Application Suite 8.11, you must deactivate and delete them before you can complete the upgrade.

The Maximo Models for Electrical Distribution accelerator replaces Maximo Health and Predict - Utilities in Maximo Application Suite 8.11.

IBM Maximo Optimizer

You can now collect and publish monitoring metrics for Maximo Optimizer. For more information, see [“Server-level monitoring metrics”](#) on page 835 .

Maximo Optimizer includes a new integration with the new Dispatching dashboard. Also, some optimization models now use the new what-if analysis function. For more information, see [What's new in Maximo Manage 8.7](#).

Serviceability dashboard

By using the serviceability dashboard, you can monitor the health and performance of IBM Maximo Assist, IBM Maximo Health, IBM Maximo Optimizer 8.5.0 and later, and IBM Maximo Predict. The serviceability dashboard is implemented with Grafana, OpenTelemetry, and Prometheus. For more information, see [“Using the serviceability dashboard”](#) on page 832.

What's new in Maximo Application Suite applications, industry solutions, and add-ons

Learn more about what's new and changed in the following applications, industry solutions, and add-ons in Maximo Application Suite 8.11.

- [What's new in Maximo Manage 8.7](#)
- [What's new in Maximo Monitor 8.11](#)
- [What's new in Maximo Health 8.9](#)
- [What's new in Maximo Predict 8.9](#)
- [What's new in Maximo Visual Inspection 8.9](#)
- [What's new in Maximo Assist 8.8](#)
- [What's new in IBM Maximo Civil Infrastructure 8.6](#)
- [What's new in Maximo Mobile 8.11](#)

Related concepts

[What's new in Maximo Application Suite fix packs](#)

Learn more about updates that are provided in Maximo Application Suite fix packs to support prerequisite software, such as Red Hat OpenShift cluster services and Cloud Pak for Data services.

What's new in earlier releases

Learn more about what's new in earlier releases.

Customer-managed **What's new in Maximo Application Suite 8.10**

Learn more about what's new and changed in IBM Maximo Application Suite 8.10.

Installation

Support for update approval method

From Maximo Application Suite 8.10, the manual deployment for applications and that use installation script for Maximo Application Suite are discontinued. To upgrade Maximo Application Suite and its applications, you must now run a conversion script to use a subscription method and subscribe to the latest channel.

For more information, see [“Upgrading IBM Maximo Application Suite” on page 473](#) and [“Converting IBM Maximo Application Suite from manual deployment to channel subscription” on page 482](#).

Support for IBM Cloud Pak for Data

You can now install IBM Cloud Pak for Data 4.6 with IBM Maximo Application Suite 8.10.

For more information, see [Software Product Compatibility Reports \(SPCR\)](#).

New default IBM Db2 integration

A default Db2 instance is now configured when you install the Maximo Application Suite with IBM Maximo Manage on Amazon Web Services or Microsoft Azure environment.

For more information, see [“Installing BYOL IBM Maximo Application Suite” on page 222](#), [“Installing client managed IBM Maximo Application Suite for public paid offer” on page 225](#), and [“Installing Maximo Application Suite” on page 263](#).

Support for configuring Amazon MSK Kafka

You can now configure an Amazon Managed Streaming for Kafka by using the CloudFormation template when you install the Maximo Application Suite on an Amazon Web Services environment.

The Amazon MSK is configured to process data streaming for applications such as IoT and IBM Maximo Monitor from Maximo Application Suite. Previously, you were able to configure the Amazon MSK manually.

For more information, see [“Installing BYOL IBM Maximo Application Suite” on page 222](#) and [“Installing client managed IBM Maximo Application Suite for public paid offer” on page 225](#).

Support for configuring DocumentDB on Amazon Web Services

You can now configure a DocumentDB by using the CloudFormation template when you install the Maximo Application Suite. You can select either a new MongoDB or Amazon DocumentDB, or use an existing MongoDB or Amazon DocumentDB.

For more information, see [“Installing BYOL IBM Maximo Application Suite” on page 222](#) and [“Installing client managed IBM Maximo Application Suite for public paid offer” on page 225](#).

Support for Microsoft Azure Private DNS zone

You can now configure Private DNS zone to resolve host names in your public domain when you install the Maximo Application Suite in a new Red Hat OpenShift cluster.

For more information, see [“Microsoft Azure DNS zones” on page 170](#).

New support certificates from Let's Encrypt

You can use certificates from a well-known certificate authority (CA) when you deploy the Maximo Application Suite with the Bring Your Own License (BYOL) option in Amazon Web Services and Microsoft Azure.

For more information, see [“Configuring Let's Encrypt for Maximo Application Suite on Amazon Web Services”](#) on page 241 and [“Configuring Let's Encrypt for Maximo Application Suite on Microsoft Azure”](#) on page 272.

Suite administration

Backing up and restoring Maximo Application Suite core and Maximo Manage

As an administrator, you can plan and implement backup and restore strategies for Maximo Application Suite core and Maximo Manage.

For more information, see [“Backing up and restoring IBM Maximo Application Suite”](#) on page 728.

Importing user data

To create multiple user records simultaneously, you can import your users' information by using a template that contains the information. You can download the template, which is a comma-separated values file, on the **User** page in the Suite administration user interface. In the .csv template, you enter the information for each user, such as identity details, contact information, and access entitlements that the user might need for applications and administration tasks. When you upload the completed .csv file, the data that you provided is processed, and a record is created for each user in the file.

For more information, see [Importing users](#).

Updates to expiration times for user authentication sessions

Starting in Maximo Application Suite 8.10, the default expiration time for access token and refresh token has changed. The default expiration time for the access token is now 30 minutes. The default expiration time for the refresh token is now 12 hours.

For more information, see [Configuring user authentication sessions](#).

IBM Maximo Connector for Envizi

The catalog includes a new tile for Maximo Connector for Envizi that can be used for automated Environmental, Social, and Governance (ESG) reporting on locations and meters in Maximo Manage from Envizi ESG Suite.

For more information, see [Integrating with Maximo Connector for Envizi](#).

IBM Maximo Connector for TRIRIGA®

The catalog includes a new tile for Maximo Connector for TRIRIGA . Maximo Connector for TRIRIGA integrates Maximo Manage with TRIRIGA Application Suite for streamlined operations, reporting, and workflows across your enterprise assets and facilities.

For more information, see [Integrating with Maximo Connector for TRIRIGA](#) .

Support for Federal Information Processing Standard 140-2

Cryptographic modules, data in motion, and data at rest that are used in the following applications, dependencies, and industry solutions support the Federal Information Processing Standard (FIPS) 140-2:

- Maximo Application Suite core and its dependencies such as IBM Suite License Service 3.7.0.
- IBM Maximo Manage including add-ons, industry solutions, and connectors.

For more information, see [“Creating the Db2 instance by using the stand-alone Db2U operator”](#) on page 10, [“Suite License Service”](#) on page 7, and [“Apache Kafka”](#) on page 23.

The following add-ons, industry solutions, and connectors do not support FIPS 140-2: IBM Maximo Civil Infrastructure, IBM Maximo Connector for Envizi, IBM Maximo Spatial, IBM Maximo Connector for TRIRIGA , IBM Maximo Connector for Workday Applications, IBM Maximo Connector for SAP Applications, and IBM Maximo Connector for Oracle applications.

Applications

What's new in Maximo Application Suite applications

Learn more about what's new and changed for the following applications, industry solutions, and add-ons in Maximo Application Suite 8.10.

- [What's new in Maximo Manage 8.6](#)
- [What's new in Maximo Monitor 8.10](#)
- [What's new in Maximo Health 8.8](#)
- [What's new in Maximo Predict 8.8](#)
- [What's new in Maximo Visual Inspection 8.8](#)
- [What's new in Maximo Assist 8.7](#)
- [What's new in Maximo Health and Predict - Utilities 8.6](#)
- [What's new in IBM Maximo Civil Infrastructure 8.5](#)
- [What's new in Maximo Mobile 8.10](#)

Related concepts

[What's new in Maximo Application Suite fix packs](#)

Learn more about updates that are provided in Maximo Application Suite fix packs to support prerequisite software, such as Red Hat OpenShift cluster services and Cloud Pak for Data services.

Customer-managed

What's new in Maximo Application Suite 8.9

Learn more about what's new and changed in IBM Maximo Application Suite 8.9.

Installation

Support for Production or Non-production mode in installation

When you install IBM Maximo Application Suite, you can now enter a new parameter **Operational Mode** to configure the instance in Production or Non-production mode for IBM AppPoint optimization. You can use the non-production mode to install and deploy the IBM Maximo Application Suite for internal development and testing with no AppPoint costs incurred for installing applications, add-ons, solutions or other capabilities. However, you continue to be charged based on your access type. For more information, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier”](#) on page 796.

Note: The **Operational Mode** parameter is available only for installation. It is not supported for upgrading the IBM Maximo Application Suite.

For more information, see [“Operational mode for installation”](#) on page 212.

Suite administration

Manual certificate management

When you configure the suite, you can manually upload your public transport layer security (TLS) certificates that are used to establish secure connections with Maximo Application Suite. You can use certificate management to ensure that your certificates are valid and replace expired certificates.

By default, Maximo Application Suite uses IBM Certificate Manager service to control certificate management. To upload your own public TLS certificates, you must first enable manual certificate management.

For more information, see [Manual certificate management](#).

Configuring user sessions

To increase application security and ensure that AppPoints are promptly returned when users close their browsers or tabs without first logging out, you can configure user authentication sessions. You configure user authentication sessions by changing the expiration time for the access and refresh token in the custom resource (CR) file for Maximo Application Suite.

For more information, see [Configuring user sessions](#).

Configuring password settings

As a suite administrator, you can configure password settings to meet the requirements of your organization's security policy. On the **User** page in the Suite administration user interface, you can set password requirements for users, such as minimum password length, minimum number of uppercase or lowercase characters, and you can choose the placement of password characters. You can also choose whether you want users to change their password on first login.

These password rules are enforced when you create a password for new users or when users change their password by selecting **Profile > Manage profile > Change password**.

For more information, see [Configuring password settings](#).

Multiple customization archives for Maximo Manage configuration

When you configure Maximo Manage for activation, you can specify the customization archive file location in the configuration. You can now include multiple customization archives during the activation process.

If you used a single customization archive in a previous configuration, that archive information is automatically transferred to the Multiple customization archive table when you update the next configuration of Maximo Manage.

For more information, see [Adding customizations to Maximo Manage](#).

Field validation for Maximo Manage configuration

Error messages identify configuration issues before you start the activation process. Explanations and suggestions for corrective action are provided in messages that might be shown when you configure Maximo Manage for activation.

Applications

IBM Maximo Safety

From Maximo Application Suite 8.9, IBM Maximo Safety is no longer available. If Maximo Safety is installed and you are upgrading to Maximo Application Suite 8.9, you must deactivate and delete Maximo Safety before you can complete the upgrade.

What's new in Maximo Application Suite applications

Learn more about what's new and changed for the following applications, industry solutions, and add-ons in Maximo Application Suite 8.9.

- [What's new in Maximo Manage 8.5](#)
- [What's new in Maximo Monitor 8.9](#)
- [What's new in Maximo Health 8.7](#)
- [What's new in Maximo Predict 8.7](#)
- [What's new in Maximo Visual Inspection 8.7](#)
- [What's new in Maximo Assist 8.6](#)
- [What's new in Maximo Health and Predict - Utilities 8.5](#)
- [What's new in IBM Maximo Civil Infrastructure 8.4](#)
- [What's new in IBM Maximo Spatial 8.5](#)
- [What's new in Maximo Mobile 8.9](#)

IBM Maximo Application Suite 8.8 includes features and enhancements that improve the installation and getting started user experience.

Maximo Application Suite installation paths

Maximo Application Suite 8.8 supports the following installation paths and scenarios:

Customer-managed installation

Customer-managed installations consist of the following scenarios:

- Install Maximo Application Suite on-premises.
- Install Maximo Application Suite in an AirGap environment.
- Install Maximo Application Suite on Amazon Web Services.
- Install Maximo Application Suite on Microsoft Azure.
- Install Maximo Application Suite on IBM Cloud.

For more information, see [Customer-managed installations](#).

IBM-managed installation

Have IBM install IBM Maximo Application Suite for you. In this scenario, the installation is performed by IBM, and you can complete postinstallation tasks to set up IBM Maximo Application Suite.

Ansible® collection

To automate some of the manual steps that are involved with installing Maximo Application Suite and its components, use the Ansible collection roles that match your installation path or use case.

For more information, see [Ansible collection](#).

New installation options for Maximo Application Suite on Amazon Web Services

Introduced two options to install the Maximo Application Suite on an Amazon Web Services (AWS) environment:

Maximo Application Suite (BYOL)

The BYOL product can be installed from the AWS Marketplace by using any one of the following two CloudFormation templates:

1. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
2. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

Maximo Application Suite (client-managed)

You can install a new Red Hat OpenShift cluster before installing the Maximo Application Suite or you can use your existing Red Hat OpenShift cluster to install the Maximo Application Suite from the AWS Marketplace.

The following two CloudFormation templates are available for installing with a new Red Hat OpenShift cluster:

1. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
2. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

The following single CloudFormation template is available for installing with an existing Red Hat OpenShift cluster:

1. Existing Red Hat OpenShift cluster (OCP)

For more information, see [“Installing the Maximo Application Suite on Amazon Web Services” on page 222](#).

New installation options for Maximo Application Suite on Microsoft Azure

Maximo Application Suite (BYOL)

You can now install the BYOL product from the Microsoft Azure Marketplace by using any one of the following two CloudFormation templates:

1. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
2. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

For more information, see [“Installing IBM Maximo Application Suite on Microsoft Azure”](#) on page 262.

New product catalog support for IBM Maximo Application Suite installation in AirGap environment

New versions of product catalogs or applications are now added to support installation of the latest version of IBM Maximo Application Suite in an AirGap or disconnected environment.

For more information, see [Supported product catalogs for disconnected installation](#).

Certificate management with IBM Certificate Manager

IBM Certificate Manager is the supported service to control certificate management in Maximo Application Suite 8.8.

Service Binding Operator no longer a dependency

Service Binding Operator (SBO) is not required in Maximo Application Suite 8.8.

For more information, see [Service Binding Operator](#).

User Data Services replaces the Behavior Analytics Service

IBM User Data Services collect, transform, and transmit product usage data. The information includes overall AppPoint entitlement and usage, installed applications and capabilities, license and AppPoint usage for each application, and contact person's name and email address for IBM to reach for any licensing discrepancies. The name and email address is configured at installation time. You do not have to be an IBM Maximo Application Suite user. Details of individual Maximo Application Suite users and breakdown of their usage are not transmitted.

User Data Services replaces the Behavior Analytics Service (BAS) in Maximo Application Suite 8.8.

IBM Maximo Optimizer replaces IBM Maximo Scheduler Optimization

Starting in Maximo Application Suite 8.8, IBM Maximo Scheduler Optimization is replaced with IBM Maximo Optimizer.

You can deploy Maximo Optimizer by selecting one of the following purchasing plans:

Maximo Optimizer plan

No restrictions on the deployment of virtual processor cores (VPCs), parallel execution of jobs, or number of custom models.

Maximo Optimizer Limited plan

Limits the deployment to two VPCs and a serialized execution of jobs for only one model. The model is a default model or an extension of the default model that is provided by IBM.

For more information, see [Deploying Maximo Optimizer](#).

If Maximo Scheduler Optimization is installed and you are upgrading to Maximo Application Suite 8.8, you must deactivate and delete Maximo Scheduler Optimization before you can complete the upgrade.

For more information, see [Upgrading to Maximo Optimizer](#).

Improved access to administration tasks

Improvements to the **Suite administration** page provide direct access to administration and configuration tasks from the side navigation menu and separates these tasks from the catalog.

New Applications page for deployed applications

After you initially deploy and activate applications from the catalog, you can complete the following administrative tasks on the **Applications** page:

- Activate applications, if the application is not already activated.
- Administer versions.
- Deactivate applications.
- Delete applications.

On the **Suite administration** page, from the side navigation menu, select **Applications**.

New Workspace page to configure activated applications

After applications are activated, you can now manage the following configuration settings for those applications on the **Workspace** page:

- Update configuration.
- Download debug information.
- Deactivate.

On the **Suite administration** page, from the side navigation menu, select **Workspace**.

Configuring external launchers moved to the default workspace

Configuration of products as external launchers was previously available by selecting **Configurations** from the side navigation menu. You now can configure external products on the **Suite administration** page, by selecting **Workspace** from the side navigation menu and then clicking the **External launchers** tab.

Locale and time zone for users

When you create users, you can specify the locale and time zone for the user instead of updating these preferences from the users' profile. Users can still change their locale settings from their profiles, which are automatically applied to their user records.

For more information, see [Setting language and time zone preferences for users](#).

New user ID to identify users

When you add users, you create a unique user ID that identifies users internally. By default, the username that is used to access Maximo Application Suite is the same as the user ID, but you still have the flexibility to change the username. For example, if you need to change the name of the user, you can update the username, and user ID remains unchanged.

Documentation improvements

To optimize and improve content retrievability, the documentation is now organized and streamlined to identify information by product and version.

For more information, see [Documentation conventions](#).

Customer-managed

What's new in Maximo Application Suite fix packs

Learn more about updates that are provided in Maximo Application Suite fix packs to support prerequisite software, such as Red Hat OpenShift cluster services and Cloud Pak for Data services.

About fix packs

Maximo Application Suite incrementally delivers regular product fix packs to provide bug fixes, security updates, and support for prerequisite software.

Fix pack information for Maximo Application Suite and each application is provided in the Maximo Application Suite release notes. For more information, see [Maximo Application Suite release notes](#).

You can install fix packs by using a channel subscription. For more information, see [“Upgrading IBM Maximo Application Suite by using the channel subscription method” on page 477](#).

Prerequisite software

Learn more about updates to Maximo Application Suite to support prerequisite software changes that are delivered in Maximo Application Suite fix packs.

June 2025

Support for IBM System/390x (s390x) and IBM Power (ppc64le) in Maximo Application Suite core and Maximo Manage base 9.0.12

Maximo Application Suite core and Maximo Manage base 9.0.12 now support IBM System/390x (s390x) and IBM Power/ppc64le to leverage existing Red Hat OpenShift capabilities through the Red Hat OpenShift Container Platform for scalability and resilience.

You can choose to install or migrate your existing Maximo Application Suite core and Maximo Manage base on IBM Z processing platform.

If Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x or IBM Power (ppc64le) architecture, consider the following for prerequisite software.

- MongoDB is supported as an external service only. For more information, see [“MongoDB” on page 21](#).
- IBM Db2 can be used as an external service. For more information, see [“Configuring IBM Db2” on page 302](#).
- Grafana operators are not supported. For more information, see [“Installing Grafana” on page 826](#).

New features and updates

Learn more about updates to Maximo Application Suite to support new features and updates that are delivered in Maximo Application Suite fix packs.

Related concepts

[What's new in Maximo Application Suite 9.0](#)

Learn more about what's new and changed in IBM Maximo Application Suite 9.0.

[What's new in Maximo Application Suite 8.11](#)

Learn more about what's new and changed in IBM Maximo Application Suite 8.11.

[What's new in Maximo Application Suite 8.10](#)

Learn more about what's new and changed in IBM Maximo Application Suite 8.10.

SaaS

What's new in Maximo Application Suite as a Service

Learn more about what's new and changed in IBM Maximo Application Suite as a Service.

May 2025

IBM Maximo AI Service in Maximo Application Suite as a Service

IBM Maximo AI Service is now available in IBM Maximo Application Suite as a Service Standard Edition and IBM Maximo Application Suite as a Service Premium Edition. Maximo AI Service enables

access to the Maximo Application Suite AI framework and IBM watsonx™, which you can use to configure AI features in Maximo Manage.

For more information, see [Deploying Maximo AI Service](#).

For more information about IBM Maximo Application Suite as a Service Standard Edition and IBM Maximo Application Suite as a Service Premium Edition, see [“Maximo Application Suite as a Service overview”](#) on page 41.

December 2024

IBM Maximo IT in Maximo Application Suite as a Service

IBM Maximo IT is now available in IBM Maximo Application Suite as a Service Standard Edition and IBM Maximo Application Suite as a Service Premium Edition. Maximo IT provides a single point of user support and enterprise service management of information technology (IT) and operational technology (OT) assets and processes.

For more information, see [Maximo IT IBM Documentation](#).

For more information about IBM Maximo Application Suite as a Service Standard Edition and IBM Maximo Application Suite as a Service Premium Edition, see [“Maximo Application Suite as a Service overview”](#) on page 41.

August 2024

This release includes the following user synchronization and user authentication updates. For information about enabling these updates in Maximo Application Suite as a Service, you can open a case with IBM Support.

User synchronization

- Configuration for custom mapping of user data that is synchronized between your LDAP server and the user registry is available in the user interface. For more information, see [“User and group registry mapping”](#) on page 622.
- Synchronization of users and groups from an identity provider (IdP) by using the System for Cross-domain Identity Management (SCIM) 2.0 protocol is available. For more information, see [“User synchronization with SCIM 2.0”](#) on page 625 and [“Configuring Maximo Application Suite to synchronize user and groups with SCIM 2.0”](#) on page 630.

User authentication

- For SAML authentication, administrators can enable an initiated logout for the SAML service provider so that current user sessions are logged out before another user logs in with the same credentials. For more information, see [“Configuring SAML authentication”](#) on page 607.
- Configuration of multiple identity providers (IDPs) for the same authentication type, such as SAML or LDAP is available. For more information, see [“Configuring multiple identity providers for same authentication type”](#) on page 617.
- Users can self-register to create their own login accounts and use the applications that they have access to. Before users can self-register, an administrator must enable and configure access options that are associated with each identity provider that is configured. For more information, see [“Self-registration for users”](#) on page 613.

Additional features are provided to align with IBM Maximo Application Suite 9.0. For more information, see [What's new in Maximo Application Suite 9.0](#).

December 2023

Maximo Application Suite as a Service editions

Maximo Application Suite as a Service is now available in three editions.

- IBM Maximo Application Suite as a Service Essentials Edition

- IBM Maximo Application Suite as a Service Standard Edition
- IBM Maximo Application Suite as a Service Premium Edition

For more information, see [“Maximo Application Suite as a Service overview”](#) on page 41.

May 2023

Importing user data

To create multiple user records simultaneously, you can import your users' information by using a template that contains the information. You can download the template, which is a comma-separated values file, on the **User** page in the Suite administration user interface. In the .csv template, you enter the information for each user, such as identity details, contact information, and access entitlements that the user might need for applications and administration tasks. When you upload the completed .csv file, the data that you provided is processed and creates a record for each user in the file.

For more information, see [Adding users](#) and [“Importing users in Maximo Application Suite in 9.0 and earlier”](#) on page 801.

December 2022

The following applications, industry solutions, and add-ons are now available in Maximo Application Suite as a Service:

Applications

- IBM Maximo Assist
- IBM Maximo Health
- IBM Maximo Monitor
- IBM Maximo Predict
- IBM Maximo Visual Inspection

Industry solutions

- IBM Maximo Health and Predict - Utilities

Add-ons

- IBM Maximo Optimizer

For more information, see [Applications, industry solutions, add-ons, and tools](#).

July 2022

The first release of IBM Maximo Application Suite as a Service provides an operational environment that is deployed on Amazon Web Services (AWS).

This initial release includes IBM Maximo Manage and the following Maximo Manage industry solutions and add-ons:

Industry solutions

- IBM Maximo Utilities
- IBM Maximo Oil and Gas
- IBM Maximo Nuclear
- IBM Maximo Transportation
- IBM Maximo Aviation
- IBM Maximo Civil Infrastructure

Add-ons

- IBM Maximo Service Provider
- IBM Maximo Health, Safety, and Environment
- IBM Maximo Asset Configuration Manager
- IBM Maximo Spatial
- IBM Maximo Connector for Oracle Applications
- IBM Maximo Connector for SAP Applications
- IBM Maximo Mobile
- IBM Maximo Anywhere

For more information, see [Maximo Manage](#)

In this initial release, SaaS suite administrators can manage users and their entitlement to Maximo Manage, request LDAP server authentication and user synchronization, and monitor application usage.

[Learn more about getting started as a SaaS suite administrator.](#)

Applications, industry solutions, add-ons, accelerators, and tools

IBM Maximo Application Suite includes many applications, industry solutions, add-ons, and tools.

The Maximo Application Suite catalog is available on the suite administration page on the side navigation menu and includes the key applications, industry solutions, add-ons, and tools.

The following applications, industry solutions, add-ons, and tools are available for Maximo Application Suite:

- IBM Maximo Manage and its many components
- IBM Maximo Real Estate and Facilities and its many components, starting in Maximo Application Suite 9.1.
- IBM Maximo Collaborate

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

- IBM Maximo Health
- IBM Maximo Monitor
- IBM Maximo Predict
- IBM Maximo Visual Inspection
- IBM Maximo IT
- IBM MRO Inventory Optimization
- IBM Maximo Visual Inspection Edge
- IBM Maximo Optimizer
- IBM App Connect
- The IoT tool

The following industry solutions and tools are not available in 8.11:

- IBM Maximo Health and Predict - Utilities. For more information, see [IBM Maximo Health and Predict - Utilities](#).
- IBM Parts Identifier. For more information, see [IBM Parts Identifier](#).

Maximo Application Suite applications

Applications add functionality to the core Maximo Application Suite.

IBM Maximo Collaborate

Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate. If Maximo Assist is deployed in Maximo Application Suite 9.0 or earlier and you upgrade to Maximo Application Suite 9.1, the name is automatically changed in the user interface.

Note: In the Maximo Application Suite documentation, Maximo Assist is now referred to as Maximo Collaborate.

IBM Maximo Collaborate, is an application in Maximo Application Suite. By using Maximo Collaborate, you can use AI-powered guidance and a knowledge base of equipment maintenance data to reduce the time that is required to diagnose and repair equipment problems, improve first-time fix rates, improve diagnosis accuracy, and drive higher levels of technician productivity. Using an intuitive mobile interface, you can diagnose equipment problems, find recommended solutions, and collaborate with experts to resolve problems.

Maximo Collaborate is available from the Suite navigator or directly at the following URL.

```
https://<workspace_id>.collaborate.<mas_domain>
```

Related tasks

[Deploying IBM Maximo Collaborate](#)

[Activating IBM Maximo Collaborate](#)

[Updating Maximo Collaborate](#)

Related information

[Getting started](#)

[Maximo Assist](#)

IBM Maximo Health

IBM Maximo Health is an application in Maximo Application Suite. By using Maximo Health, you can improve your asset's reliability by understanding asset health and taking action. You can review your assets' performance and condition indicators, such as the last failure date and the maintenance-to-replacement ratio (MRR), and take action by creating work orders and service requests. You can use work queues to improve the quality of your asset's details and related data. You can also configure scoring for assets' health, criticality, and risk.

Maximo Health is available from the Suite navigator.

Access the application from the following URLs:

Application	URL
Maximo Manage + Health	<code>https://<workspace_id>.manage.<mas_domain>/maximo/oslc/graphite/reengineer/index.html</code>
Maximo Health stand-alone	<code>https://<workspace_id>.health.<mas_domain>/maximo/oslc/graphite/reengineer/index.html</code>

Related tasks

[Deploying IBM Maximo Health](#)

Improve the reliability of your assets by proactively monitoring and managing asset health by using Maximo Health.

[Updating IBM Maximo Health](#)

When new versions of IBM Maximo Health become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

Related information

[Getting started with Maximo Health and Predict - Utilities](#)

IBM Maximo Manage

IBM Maximo Manage is an application in Maximo Application Suite. By using Maximo Manage, you can get a comprehensive view of all of your asset types, their conditions and locations, and the work processes that support them, to support optimal planning, control, audit, and compliance capability.

The database provides critical information about asset resources, including key attributes, their configuration, and their physical and logical relationships to other resources.

From the hierarchical navigator, users can drill down through layers from a system-wide view to individual assets. Analytic functions are applied to input data, and the output is displayed on value cards, tables, images, line graphs, and alert tables.

Maximo Manage can be deployed and accessed through Maximo Application Suite. You can access the suite from the following URLs:

- https://home.<mas_domain> - The Maximo Application Suite Suite navigator.
- https://admin.<mas_domain> - The administration dashboard.

Maximo Manage add-ons and industry solutions

With Maximo Manage, you can optionally include additional components, such as add-ons and industry solutions during the deployment and activation process, and IBM Maximo Mobile. You can integrate with SAP, Oracle, and Workday applications by using a connector.

The following add-ons are available for Maximo Manage:

- IBM Maximo Anywhere
Available only in Maximo Application Suite 8.8 and earlier. For more information, see [Maximo Anywhere documentation](#).
- IBM Maximo Asset Configuration Manager
For more information, see [Maximo Asset Configuration Manager documentation](#).
- Maximo Asset Investment Planning
For more information, see [Maximo Asset Investment Planning](#).
- IBM Maximo Connector for SAP Applications
For more information, see [Maximo Enterprise Adapter for SAP Applications documentation](#).
- IBM Maximo Connector for Oracle applications
For more information, see [Maximo Enterprise Adapter for Oracle Applications documentation](#).
- **Customer-managed** IBM Maximo Connector for Workday Applications
For more information, see [Integrating with Workday Connector](#).
- IBM Maximo Health, Safety and Environment
For more information, see [Maximo Health, Safety, and Environment Manager documentation](#).
- IBM Maximo Service Provider
For more information, see [Maximo for Service Providers documentation](#).
- IBM Maximo Spatial
For more information, see [Maximo Spatial Asset Management documentation](#).

- **Customer-managed** IBM Maximo Connector for Envizi
[Integrating with Maximo Connector for Envizi.](#)
- **Customer-managed** IBM Maximo Connector for TRIRIGA
For more information, see [Integrating with Maximo Connector for TRIRIGA.](#)
- IBM Maximo IT
For more information, see [IBM Maximo IT.](#)
- Reliability Strategies
For more information, see [Reliability Strategies.](#)

The following industry solutions are available for Maximo Manage

- IBM Maximo Aviation
For more information, see [Maximo for Aviation documentation.](#)
- IBM Maximo Civil Infrastructure
For more information, see [Maximo Civil Infrastructure documentation.](#)
- IBM Maximo Nuclear
For more information, see [Maximo for Nuclear Power documentation.](#)
- IBM Maximo Oil & Gas
For more information, see [Maximo for Oil and Gas documentation.](#)
- IBM Maximo Transportation
For more information, see [Maximo for Transportation documentation.](#)
- IBM Maximo Utilities
For more information, see [Maximo for Utilities.](#)

IBM Maximo Mobile

When Maximo Manage is deployed and activated, you can download the Maximo Mobile application from the Apple App Store or Google Play store. For more information, see [Configuring Maximo Mobile and What's new in Maximo Mobile.](#)

Related tasks

[Deploying IBM Maximo Manage](#)

To prepare for deployment, you must complete several tasks, such as configuring the database for Maximo Manage. Verify the compatibility of the industry solutions and add-ons if you want to deploy them and complete other required and optional preparation steps as needed. For example, if you are deploying in multiple languages, review the support languages that are available.

Related information

[Getting started with IBM Maximo Manage](#)

IBM Maximo Monitor

IBM Maximo Monitor is an application in Maximo Application Suite. By using Maximo Monitor, you can visualize current and historical trend data for devices and assets in customizable dashboards.

Analytic functions are applied to input data, and the output is displayed on value cards, tables, images, line graphs, and alert tables.

Anomaly detectors run on the input data to detect outliers, gaps, and flat lines in the data and fire alerts. The anomalous data points are highlighted on line graphs.

Maximo Monitor is available from the Suite navigator or directly at the following URL:

[https://<workspace_id>.monitor.<mas_domain>/](https://<workspace_id>.monitor.<mas_domain>)

Related tasks

[Deploying IBM Maximo Monitor](#)

By using Maximo Monitor, you can visualize current and historical trend data for your devices and assets on customizable dashboards.

Related information

[Maximo Monitor](#)

IBM Maximo Real Estate and Facilities

IBM Maximo Real Estate and Facilities is an application in IBM Maximo Application Suite. It is an integrated workplace management system that enables access to a full set of real estate and facilities applications. By using Maximo Real Estate and Facilities, you have the flexibility to start with any real estate or facilities discipline and expand into other areas.

Maximo Real Estate and Facilities offers:

- Increased visibility into under-performing facilities, resources, and process.
- Improved control of facility occupancy and operating costs.
- Engaging workplace services for building occupants.
- Automated activities that are designed to increase the efficiency and organizational effectiveness of real estate.
- Facility management, and environmental sustainability functions within mid- and large-sized commercial and public enterprises.
- For more information about Maximo Real Estate and Facilities, see the [Maximo Real Estate and Facilities overview](#).
- For information about deploying Maximo Real Estate and Facilities, see “[Deploying Maximo Real Estate and Facilities in Maximo Application Suite](#)” on page 373.

When Maximo Real Estate and Facilities is deployed, it is available from the Suite navigator.

In the common navigation menu, selecting Maximo Real Estate and Facilities opens the Maximo Real Estate and Facilities navigation separately.

Maximo Real Estate and Facilities is also available from the following URL:

[https://<workspace_id>.facilities.<mas_domain>/](https://<workspace_id>.facilities.<mas_domain>)

IBM Maximo Predict

IBM Maximo Predict is an application in Maximo Application Suite. By using Maximo Predict, you can leverage your historical and near real-time asset performance data, maintenance records, inspection reports, and environmental data to correlate performance factors that predict asset degradation or failure. Maximo Predict also uses artificial intelligence to optimize predictive model accuracy.

Maximo Predict can calculate predicted values, such as estimated time to failure for an asset, the probability of a failure occurring in a selected prediction window, the historical trend of failure probability scores for an asset, the probability of end of life failure for an asset, anomalies that have occurred, and the trend of anomaly scores for an asset over time.

Maximo Predict is available from the Suite navigator.

If Maximo Health is deployed as part of Maximo Manage, the Maximo Predict application is available at the following URL: https://<workspace_id>.manage.<mas_domain>/maximo/oslc/graphite/reengineer/index.html

If Maximo Health is not deployed as part of Maximo Manage, the Maximo Predict application is available at the following URL: https://<workspace_id>.health.<mas_domain>/maximo/oslc/graphite/reengineer/index.html

For more information, see the [Maximo Predict documentation](#).

Related tasks

[Deploying IBM Maximo Predict](#)

Maximo Predict can use historical and recent asset performance data to correlate performance factors that predict asset degradation or failure. Other types of data that can be correlated include maintenance records, inspection reports, and environmental data. Maximo Predict uses artificial intelligence to optimize predictive model accuracy.

[Activating IBM Maximo Predict](#)

Before you can grant users access and start working with Maximo Predict, you must activate the application. You can activate Maximo Predict after the deployment is complete.

[Updating Maximo Predict](#)

Related information

[Getting started with Maximo Predict](#)

IBM Maximo Visual Inspection

IBM Maximo Visual Inspection is a machine-learning application for video and image analysis in Maximo Application Suite. By using Maximo Visual Inspection, you can use built-in deep learning models to analyze images and video streams for classification and object detection.

Maximo Visual Inspection includes tools and interfaces for anyone with limited skills in deep learning technologies. You can use Maximo Visual Inspection to label images and videos that can be used to train and manage a model. The model can then be validated and deployed in customized solutions that demand image classification and object detection.

Maximo Visual Inspection is available from the suite catalog or directly at the following URL:

`https://<workspace_id>.visualinspection.<mas_domain>`

Related concepts

[IBM Maximo Visual Inspection Edge](#)

IBM Maximo Visual Inspection Edge is an add-on in Maximo Application Suite. By using IBM Maximo Visual Inspection Edge, you can perform inference operations from edge devices by calling or deploying multiple models that were trained in IBM Maximo Visual Inspection.

Related information

[Getting started with IBM Maximo Visual Inspection](#)

[IBM Maximo Visual Inspection](#)

Maximo Application Suite Industry solutions

On the Industry solutions tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove industry solutions.

Application suite administrators can manage the Maximo Application Suite industry solutions collection.

Related tasks

[Configuring external launchers](#)

[Deploying industry solutions](#)

The industry solutions that are available for your installed Maximo Application Suite version can be configured for use with your environment.

[Deploying applications, add-ons and industry solutions](#)

By using the **Applications** tab in the Maximo Application Suite catalog, administrators who have system configuration privileges can add and remove applications.

[Deploying IBM Maximo Health and Predict - Utilities](#)

Maximo Health and Predict - Utilities supports maintenance, operations, and performance of assets and networks for energy and utility companies.

Maximo Health and Predict - Utilities

Note:

Starting in Maximo Application Suite 8.11, Maximo Health and Predict - Utilities is no longer available as a separate industry solution. The information that is provided is applicable only to Maximo Application Suite 8.10 and earlier versions. For more information, see [“Upgrading IBM Maximo Application Suite” on page 473](#). Before you upgrade to Maximo Application Suite 8.11, deactivate and delete Maximo Health and Predict - Utilities.

IBM Maximo Health and Predict - Utilities is an industry solution in Maximo Application Suite. By using Maximo Health and Predict - Utilities, you can configure scoring and predictive models to support the maintenance, operations, and performance of energy and utility assets and networks.

Maximo Health and Predict - Utilities is available from the Suite navigator.

If Maximo Health is deployed as part of Maximo Manage, the Maximo Health and Predict - Utilities industry solution is available at the following URL: https://<workspace_id>.manage.<mas_domain>/maximo/oslc/graphite/reengineer/index.html

If Maximo Health is not deployed as part of Maximo Manage, the Maximo Health and Predict - Utilities industry solution is available at the following URL: https://<workspace_id>.health.<mas_domain>/maximo/oslc/graphite/reengineer/index.html

For more information, see the [Maximo Health and Predict - Utilities documentation](#).

Related tasks

[Deploying IBM Maximo Health](#)

Improve the reliability of your assets by proactively monitoring and managing asset health by using Maximo Health.

[Updating IBM Maximo Health](#)

When new versions of IBM Maximo Health become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

Related information

[Getting started with Maximo Health and Predict - Utilities](#)

Customer-managed

IBM MRO Inventory Optimization

Note:

Starting in Maximo Application Suite 9.0, MRO Inventory Optimization is no longer available to be added as an externally configured application and must be accessed by the dedicated URL. The information that is provided is applicable to Maximo Application Suite 8.11 and earlier versions. If MRO Inventory Optimization is configured as an external launcher and you are upgrading to Maximo Application Suite 9.0, you must remove MRO Inventory Optimization before you can complete the upgrade.

MRO Inventory Optimization is an industry solution in Maximo Application Suite.

For more information, see the [IBM MRO Inventory Optimization product page](#)

Related tasks

[Configuring IBM MRO Inventory Optimization in Maximo Application Suite](#)

In version 8.11 and earlier, MRO Inventory Optimization is a stand-alone but linked product that requires an externally purchased license. To give users access to MRO Inventory Optimization in Maximo Application Suite, you can add MRO Inventory Optimization as an external launcher.

Related information

[Getting started with MRO Inventory Optimization](#)

Maximo Application Suite add-ons and tools

On the Add-ons tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove add-ons.

Application suite administrators can manage the Maximo Application Suite application collection.

Customer-managed **IBM Maximo IT**

IBM Maximo IT provides a single point of user support and enterprise service management of information technology (IT) and operational technology (OT) assets and processes.

You can deploy and activate Maximo IT as an add-on to Maximo Manage. Before you deploy Maximo IT, accept the license terms in Passport Advantage.

Related tasks

Deploying IBM Maximo Manage

To prepare for deployment, you must complete several tasks, such as configuring the database for Maximo Manage. Verify the compatibility of the industry solutions and add-ons if you want to deploy them and complete other required and optional preparation steps as needed. For example, if you are deploying in multiple languages, review the support languages that are available.

Related information

Maximo IT IBM Documentation

Customer-managed **IBM Maximo Visual Inspection Edge**

IBM Maximo Visual Inspection Edge is an add-on in Maximo Application Suite. By using IBM Maximo Visual Inspection Edge, you can perform inference operations from edge devices by calling or deploying multiple models that were trained in IBM Maximo Visual Inspection.

You can use IBM Maximo Visual Inspection Edge to remotely call or locally deploy multiple AI models, which were trained in IBM Maximo Visual Inspection, on edge devices. For example, servers that are connected to CCTV or IP cameras that are pointed at assembly lines. The inferencing capabilities of IBM Maximo Visual Inspection Edge enable you to use computer vision to perform inspections, detect manufacturing anomalies, and prescribe corrective actions.

The prebuilt runtime environment that IBM Maximo Visual Inspection Edge provides includes deep learning libraries to deploy models quickly and spend more time customizing your workflow.

For more information, see the IBM Maximo Visual Inspection Edge documentation.

IBM Maximo Visual Inspection Edge

By using IBM Maximo Visual Inspection Edge, you can deploy multiple models that were trained in IBM Maximo Visual Inspection to edge devices.

Configuration parameters

The following parameters are configurable:

- Location and purpose of the edge deployment.

Required by

- Requires Maximo Visual Inspection at the Application scope.

Related concepts

IBM Maximo Visual Inspection

IBM Maximo Visual Inspection is a machine-learning application for video and image analysis in Maximo Application Suite. By using Maximo Visual Inspection, you can use built-in deep learning models to analyze images and video streams for classification and object detection.

IBM Maximo Optimizer

IBM Maximo Optimizer is an add-on to Maximo Application Suite. By using Maximo Optimizer, you can automate efficient decisions for long-range planning, scheduling, and dispatching of resources for asset maintenance while balancing competing objectives and constraints.

Maximo Optimizer includes IBM Maximo Optimization Framework for data and application management of optimization jobs and embeds IBM ILOG® CPLEX® Optimization Studio for solving optimization models. Maximo Optimizer is a Maximo Application Suite add-on and can be used with Maximo Manage.

You can deploy Maximo Optimizer with one of the following plans:

Maximo Optimizer plan

The Maximo Optimizer plan with no restrictions on the deployment of virtual processor cores (VPCs), parallel execution of jobs, or number of custom models.

Maximo Optimizer Limited plan

The Maximo Optimizer Limited plan limits the deployment to two VPCs and a serialized execution of jobs for only one model. The model is a default model or a customized version of the default model that is provided by IBM.

If you purchase the Maximo Optimizer plan instead of the Maximo Optimizer Limited plan, then the Maximo Optimizer plan is your default installation. However, only the Maximo Optimizer Limited plan is supported on a single-node Red Hat OpenShift cluster. At installation time, the Maximo Application Suite CLI utility automatically detects if the cluster has more than 1 node or not. If the CLI utility detects a single node cluster at installation time, then the Maximo Optimizer Limited plan gets installed. A message displays in the CLI utility about this behavior.

After deployment and activation, Maximo Optimizer is available from the Suite navigator or directly at the following URL:

`https://<workspace_id>.optimizer.<mas_domain>/`

Related tasks

Deploying IBM Maximo Optimizer

By using Maximo Optimizer, you can automate efficient decisions for long-range plans, schedules, and the dispatch of resources for asset maintenance, helping to balance competing objectives and constraints.

Activating IBM Maximo Optimizer

Before Maximo Optimizer is available for use, you must activate Maximo Optimizer. Activating the application does not automatically grant your users access to the application.

Modifying optimization system properties

To run optimization, system properties for Maximo Optimizer are automatically enabled during deployment and activation. If needed, you can modify the system properties in Maximo Manage.

Upgrading from Maximo Scheduler Optimization to Maximo Optimizer

In Maximo Application Suite 8.8, you use Maximo Optimizer instead of Maximo Scheduler Optimization. If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade.

Related information

Maximo Optimization Framework optimization models

IBM Parts Identifier

IBM Parts Identifier is an IBM Cloud service that can be used as an add-on to IBM Maximo Manage and IBM Maximo Mobile in Maximo Application Suite. By using Parts Identifier, you can search for and identify industrial parts from a mobile device.

Note: Starting in Maximo Application Suite 8.11, Parts Identifier is no longer available. If Parts Identifier is deployed and active in your environment and you are upgrading to Maximo Application Suite 8.11, you must deactivate and delete Parts Identifier before you can complete the upgrade.

Related tasks

[Deploying IBM Parts Identifier](#)

Parts Identifier is an IBM Cloud service that can be used as an add-on with IBM Maximo Manage and IBM Maximo Mobile. It enables technicians to search for and identify industrial parts on a mobile device.

IBM App Connect

App Connect is an add-on to Maximo Application Suite. By using App Connect, you can connect applications and data from existing systems and modern technologies across all their environments.

Use App Connect to connect your different applications and make your business more efficient. Set up flows that define how data is moved from one application to one or more other applications. App Connect supports a range of skill levels and interfaces, giving you the flexibility to create integrations without writing a single line of code. You can use a web user interface or drop resources into a toolkit that gives a broader range of configuration options. Your entire organization can make smarter business decisions by providing rapid access, visibility, and control over data as it flows through your business applications and systems from a single place - App Connect.

For more information, see the [App Connect documentation](#).

Related tasks

[Deploying IBM App Connect](#)

With App Connect, you can connect applications and data from existing systems and modern technologies across all their environments.

Customer-managed **IoT tool**

The IoT tool is an add-on to Maximo Application Suite. By using the IoT tool, you can access device management tools and configure device connectivity, data filtering, and data mapping.

The IoT tool is available from the Suite navigator or directly from the following URL:
https://<workspace_id>.iot.<mas_domain>/ibmssologin

The IoT tool APIs are available at the following URLs:

- https://<workspace_id>.messaging.iot.<mas_domain> - IoT MQTT broker
- https://<workspace_id>.iot.<mas_domain> - IoT APIs
- https://<workspace_id>.iot.<mas_domain>/docs/index.html - IoT API docs

Related tasks

[Deploying IoT tool](#)

The IoT tool provides device connectivity, data filtering and mapping, and device management tools and is needed by Maximo Monitor.

Related information

[Getting started with IBM Maximo Monitor](#)

[About the IoT tool](#)

Customer-managed **Maximo Application Suite accelerators**

Accelerators are solutions that complement or extend Maximo Application Suite capabilities or experience or accelerate time to value.

IBM and IBM partners develop accelerators, and they are hosted on [Maximo Application Suite accelerators solutions page](#). IBM reviews and approves all accelerators through a verification program. Maximo Application Suite administrators who have workspace management access can add and remove accelerators from the **Accelerators** tab in the Maximo Application Suite catalog. Business partners interested in participating in the marketplace must have their accelerator verified and approved for listing.

Accelerators range from simple content to automation scripts to comprehensive solutions that are built on Maximo Application Suite. Maximo Application Suite includes the following types of accelerators:

Content

Content accelerators include configuration guidelines, application dashboards, record attachments, model templates, and data loaded by a user through the UI.

Connectors

Connectors enable stand-alone applications to integrate with Maximo Application Suite.

Applications

Applications are stand-alone applications that might include integration with Maximo Application Suite or provide data that can be loaded into Maximo Application Suite.

Scripting and Code Patterns

Scripting and code patterns accelerators consist of automation scripts, DBC scripts, or custom code to automate tasks, load data, and other actions.

Related concepts

Accelerators

Accelerators are solutions that are provided by IBM and partners of IBM that extend Maximo Application Suite capabilities. These accelerators are hosted on the [Maximo Application Suite accelerators solutions page](#).

Customer-managed

Maximo Application Suite application URLs

When you installed Maximo Application Suite, you defined the following environment variables, which are used to create the application access URLs.

- `<mas_domain>` - Domain that your instance of Maximo Application Suite runs under.
- `<workspace_id>` - Identifier for the unique workspace in which your applications are deployed.

Each Maximo Application Suite core component has a subdomain:

- `https://admin.<mas_domain>` - The Suite administration.
- `https://home.<mas_domain>` - The Suite navigator.

Maximo Collaborate

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

`https://<workspace_id>.collaborate.<mas_domain>/` - Maximo Collaborate application

Maximo Health and Maximo Predict

If Maximo Health is deployed as part of Maximo Manage, the applications are available at the following URL: `https://<workspace_id>.manage.<mas_domain>/maximo/oslc/graphite/reengineer/index.html`

If Maximo Health is not deployed as part of Maximo Manage, the applications are available at the following URL: `https://<workspace_id>.health.<mas_domain>/maximo/oslc/graphite/reengineer/index.html`

Maximo Monitor

- `https://<workspace_id>.monitor.<mas_domain>/` - IBM Maximo Monitor application

IoT tool

- `https://<workspace_id>.iot.<mas_domain>/ibmssologin` - IoT tool

In addition, URIs are available for special purposes as needed:

- `https://<workspace_id>.iot.<mas_domain>` - IoT tool APIs
- `https://<workspace_id>.iot.<mas_domain>/docs/index.html` - IoT API docs
- `https://<workspace_id>.messaging.iot.<mas_domain>` - IoT tool MQTT broker

IBM Maximo Visual Inspection

`https://<workspace_id>.visualinspection.<mas_domain>` - Maximo Visual Inspection

Licensing in Maximo Application Suite

The licensing guidance provides information about licensing and AppPoints entitlements for IBM Maximo Application Suite.

Overview of Licensing and AppPoints

Watch a video to understand Maximo Application Suite's customer-managed licensing model, and how to manage usage with AppPoints.

Customer-managed **Licensing in Maximo Application Suite 9.1**

The licensing guidance provides information about licensing and entitlements for IBM Maximo Application Suite 9.1.

Notice: This Licensing Guidance is intended to provide only supplementary information to assist you in deploying the Program(s) you have licensed from IBM within your purchased entitlement. Your license agreement, such as the IBM International Program License Agreement (IPLA) or equivalent and its transaction documents, including the License Information for IBM Maximo Application Suite is the sole and complete agreement between you and IBM regarding use of the Program.

- [“Listing of licenses by type” on page 78](#)
- [“What do you get with your purchase of IBM Maximo Application Suite, and what is your entitlement?” on page 79](#)
- [“License ratio & user type” on page 85](#)
- [Offering-specific licenses](#)
- [Non-production instances](#)
- [Additional terms for service providers](#)
- [Cloned or custom applications](#)
- [Back-up and restore](#)
- [Licensing requirements for indirect access scenarios](#)
- [“Appendix: Modules and applications in Maximo Application Suite 9.1” on page 95](#)

Note: The Licensing Guidance is also available in [PDF](#).

Maximo Application Suite uses:

AppPoints

AppPoints is a unit of measure by which the Program can be licensed. An AppPoint is a common unit of value. Sufficient AppPoint entitlements must be obtained to cover the total number of entitlements required for Licensee's Authorized Use of the Program.

Red Hat OpenShift

All deployments of Maximo Application Suite that are deployed on Red Hat OpenShift Container Platform must have sufficient entitlement for the Maximo Application Suite cores that are used.

Listing of licenses by type

See also

- [IBM Maximo Application Suite 9.1 \(L-ZMZT-HPUYXK\)](#)

The following license types are used for AppPoints in Maximo Application Suite:

- [“Concurrent user” on page 79](#)
- [“Authorized user ” on page 79](#)
- [“Access for Concurrent & Authorized users” on page 79](#)
- [“Install” on page 79](#)

Concurrent user

A unit of measure by which the Program can be licensed. A Concurrent user is a person who is accessing the Program at any particular point in time. Regardless of whether the person is simultaneously accessing the Program multiple times, the person counts only as a single Concurrent User. The Program may be installed on any number of computers or servers, but the Licensee must obtain entitlements (Proof of Entitlement(PoE)) for the maximum number of Concurrent Users simultaneously accessing the Program. The licensee must obtain an entitlement for each simultaneous Concurrent User accessing the Program in any manner directly or indirectly (for example: via a multiplexing program, device, or application server) through any means.

Authorized user

A unit of measure by which the Program can be licensed. An Authorized user is a unique person who is given access to the Program. The Program may be installed on any number of computers or servers and each Authorized User may have simultaneous access to any number of instances of the Program at one time. The licensee must obtain separate, dedicated entitlements for each Authorized user that is given access to the Program in any manner directly or indirectly (for example, through a multiplexing program, device, or application server) through any means. An entitlement for an Authorized user is unique to that Authorized User and may not be shared, nor may it be reassigned other than for the permanent transfer of the Authorized user entitlement to another person.

Access for Concurrent & Authorized users

A customer can set up any users access type as Authorized or Concurrent. Entitlement levels for Users include:

- Applications that do not require entitlement.
- Limited Users with entitlement to 3 modules of the Maximo Manage application, the Monitor application, and the Maximo Real Estate and Facilities application.
- Base Users with entitlement to many modules of the Maximo Manage application, and the Maximo Health, Maximo IT, and Maximo Real Estate and Facilities applications.
- Premium Users with access to the entire suite of applications.

Each level of access inherits the entitlements that are granted with every lower tier license.

Install

A software program instance. Licensee must obtain an entitlement for each Installation of the Program, except where outlined in the section of Applications that do not require entitlement.

In addition to the above, the following terms apply to Licensee's use of the Program.

Content token

Content token is a unit of measure by which the Program can be licensed. A Content token is a unit of input and output content such as individual words, and sub words of a sentence. Sufficient entitlements must be obtained to cover the maximum number of Content Tokens in any calendar month.

What do you get with your purchase of IBM Maximo Application Suite, and what is your entitlement?

- [Maximo Manage](#)
- [Maximo Manage industry solutions and add-ons](#)
- [Maximo Manage Mobile](#)
- [Maximo Manage integration solutions](#)
- [Maximo Monitor](#)
- [Maximo Health](#)
- [Maximo Predict](#)
- [Maximo Visual Inspection and Maximo Visual Inspection Edge](#)
- [Maximo Collaborate](#)

- [Maximo Real Estate and Facilities](#)
- [Maximo AI Service](#)

Maximo Manage

- Provides comprehensive asset lifecycle and maintenance management for all asset types.
- Licensing is Authorized and Concurrent User; AppPoints based on Tier level – Base or Limited.
- Maximo Manage now includes access to Maximo Spatial, Linear, and Calibration at both the Base and Limited tier. Maximo Spatial requires a separate installation.
- Maximo Manage now includes Scheduler at the Base tier level.

Maximo Spatial

- Embeds geographic information system (GIS) functionality in the Maximo User interface that enables users visualize assets, work orders, and locations on a **Map** tab. Uses technology from Esri ArcGIS.
- Enables map and navigation controls for dynamic visual display of work orders, work requests, assets, and locations in Maximo.
- Maximo Spatial must be used with Esri ArcGIS.
- Customers must purchase Esri ArcGIS and any required analyst desktop licenses separately from Esri or an authorized agent.
- Entitled as part of Maximo Manage and requires a separate installation.
- Customers must purchase AppPoints to install Maximo Spatial.

Linear

- Linear is used for managing assets such as roads, pipelines, rail lines, and transmission lines.
- Entitled and installed as part of Maximo Manage.

Calibration

- Enables traceability and reverse traceability, all calibration history data, calibration data sheets, and required reporting.
- Entitled and installed as part of Maximo Manage.

Scheduler

- Scheduler provides an innovative means to graphically plan, and schedule current and upcoming work based on available resources and to manage aspects of advanced work management that include intuitive graphical assignment, mobile workforce management, real-time communication with the field, and schedule optimization.
- An advanced work management tool to enable users to manage large projects, such as shutdowns, outages, turnarounds, and customer appointments, and all planned and unplanned maintenance across a broad geographic area or where weather plays an important factor in asset availability.
- Licensing is Authorized and Concurrent user; AppPoints based on 'Base' or if user is given access to IS/Add-on applications 'Premium' Tier level.
- Entitled and installed as part of Maximo Manage.

Maximo Manage industry solutions and add-ons

Maximo Manage includes the following industry solutions and add-ons:

Industry solutions

- Maximo Aviation
- Maximo Civil Infrastructure
- Maximo Nuclear
- Maximo Oil & Gas

- Maximo Transportation
- Maximo Utilities

Four add-ons

- Maximo Asset Configuration Manager
- Maximo Service Provider
- Maximo Health, Safety and Environment
- Maximo IT

Four supplemental add-ons

- Maximo Optimizer
- Maximo Reliability Strategies
- Maximo Asset Investment Planning
- Maximo Maintenance Cost Insights
- Licensing is Authorized and Concurrent user.
- For industry solutions and add-ons, the AppPoints are based on 'Premium' or 'Limited' Tier level. Limited users can select 3 modules, which include the modules that are unique to an industry solution. See the [Appendix for a list of Modules](#) that are specific to the Industry Solutions.
- For supplemental add-ons the AppPoints are based on user Tier level.
- Install AppPoints might also be required based on IS/Add-on installed.

Maximo Aviation

- Provides Maintenance, Repair, and Overhaul (MRO) organizations that maintain aircraft or aircraft components a solution to help reduce the turnaround time and improve the amount of time available for revenue generation.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Civil Infrastructure

- Provides asset lifecycle management capabilities to help organizations reduce the total cost of ownership and improve the monitoring of the civil infrastructures they manage. Civil Infrastructure drives inspections and records inspection results about concrete and steel structures and related assets. Visualizes anomalies and work in structure-appropriate views, including linear, spatial, asset-based, and building information management views; manages planned and unplanned work for work orders, contractors, and purchasing management and creates preventive maintenance plans and job plans to be incorporated into the work schedule.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.
- The Civil Infrastructure application also requires per Install licenses; AppPoints allocated is 50.

Maximo Oil & Gas

- Provides Oil & Gas companies with best practices to help improve the productivity and efficiency of their critical assets that are related to the oil and gas industry.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Transportation

- Provides organizations with best practices to help improve the productivity of critical transportation assets, including over-the-road vehicles, rail, aircraft, and marine vessels.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Nuclear

- Provides owners, operators, and maintainers of nuclear power plants best practices to help improve the productivity and efficiency of their critical assets.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Utilities

- Helps increase asset and resources effectiveness in water, gas, and electric utilities.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Asset Configuration Manager

- Maximo Asset Configuration Manager enables real-time calculation of both an asset's build and a component's life, with benefits such as improved compliance regulations and reduced operating costs.
- An asset is a complete parent asset, such as aircraft, locomotives, and passenger rail cars. An asset is not components, such as engines or axles, or groups of assets such as trains, which can be made up of multiple locomotives and passenger cars. In aviation, this is also referred to as tail numbers.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Health, Safety and Environment

- Maximo Health, Safety and Environment provides key extensions to Maximo Manage to provide strategic applications for audit management, risk assessment, safety reporting, management of change, condition reporting, corrective actions, and training.
- Facilitates meeting or exceeding regulations and legislated requirements that impact health, safety, and the environment. Best in class capabilities and best practices for improving safety, reliability, and compliance.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo IT

- Maximo IT is an integrated service management solution that helps you manage a comprehensive range of IT processes, services, and assets.
- Separate Part purchase with Maximo Application Suite being a prerequisite.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Limited' or 'Base' Tier level. 'Limited' or 'Premium' Tier level if Service Provider is installed.

Maximo Optimizer

- Optimizer installation is required if clients want to use optimization in Maximo Application Suite.
- Licensing is per installation. AppPoints allocated is 220 or Limited 60.

Maximo Service Provider

- Maximo Service Provider is intended for companies who deliver asset management services to their third-party user clients in a revenue-generating business model. Examples include: IBM ITD, EDS, CSC, Johnson Controls, UNICCO, Coor Service Management.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.
- IBM SP Rule: a customer that uses Maximo to manage assets for a 3rd-party company cannot own Maximo without also deploying Maximo Service Provider. A waiver can be obtained under certain circumstances with approval from WW Maximo Sales Leader, WW Price Leader, and Product Management.

Reliability Strategies

- Reliability Strategies helps maximize asset performance, reduce downtime, and boost revenues by helping clients migrate from reactive to proactive strategy-based maintenance.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' 'Base' or 'Limited' Tier level.

Maximo Asset Investment Planning

- Maximo Asset Investment Planning is an integrated planning solution with a focus on leveraging EAM and APM data to help users generate asset investment plans that factor in capex, opex, and average annual cost of the asset versus risk.
- Maximo Optimizer is a prerequisite.
- Licensing is Authorized and/or Concurrent User; AppPoints based on 'Premium' 'Base' or 'Limited' Tier level.

Maximo Maintenance Cost Insights

- Maximo Maintenance Cost Insights is designed for better visibility to the total cost of maintenance including work order, labor, services, materials, and tools.
- Separate Maximo Application Suite as a Service Part purchase with Maximo Application Suite being a prerequisite. 80 Maximo Application Suite as a Service AppPoints is required.
- Licensing is Authorized and/or Concurrent User; AppPoints based on 'Premium' 'Base' or 'Limited' Tier level.

Maximo Mobile

- The Maximo Mobile app is available on Google, Android, Windows and iOS stores.
- The Maximo Mobile app integrates with Maximo Collaborate to provide based assistance to technicians on their mobile device.
- Maximo Mobile integrates with Maximo Manage to provide Inspection, Technician work orders, Asset, and Inventory apps on their device.

Maximo Manage integration solutions

Maximo Connector for SAP Applications

- A business interface that enables real-time information exchange between base products and SAP or SAP PM ERP applications by providing bidirectional connectivity with pre-built integrations between them
- Used with base products including Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for Oracle Applications

- A business interface that enables real-time information exchange between base products and Oracle E-Business Suite applications by providing bidirectional connectivity by using pre-built integrations between them.
- Used with base products that include Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for Workday Applications

- A business interface that enables real-time information exchange between base products and Workday applications by providing bidirectional connectivity by using pre-built integrations between them.
- Used with base products that include Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for TRIRIGA

- Enables bidirectional synchronization of portfolio data about people, places, and assets for consistent operations. Near real-time service request and work order routing for coordinated operations and near real-time project tasks to work-order synchronizations to enable streamlined operations.
- Licensing is per installation. Allocated AppPoints is 0.

Maximo Connector for Envizi

- Enables automatic synchronization of locations and meter readings from Maximo Manage to Maximo Connector for Envizi to automate tracking of energy usage and calculating related scope 1 and 2 emissions of electricity, natural gas, and water in Envizi.
- Licensing is per installation. Allocated AppPoints is 0.

Maximo Spatial

- Maximo Spatial installation is required if clients want to use Maximo Spatial in Maximo Application Suite.
- Licensing is per installation. Allocated AppPoints is 20.

Maximo Monitor

- A monitoring solution for real-time visibility, root-cause troubleshooting, and AI driven alerts at scale.
- Through monitoring dashboards, clients have a consolidated view of operations with data integration from multiple IoT and non-IoT sources.
- Licensing is Authorized and Concurrent user. AppPoints is based on Tier level Limited.

Maximo Health

- Maximo Health provides a first step in managing the performance of critical assets. By capturing information from the asset, the age of the asset and maintenance history, organizations can quickly understand which assets are in poor health and need attention to prevent disruption.
- Maximo Health uses ready-to-use common scoring elements and knowledge from equipment and maintenance engineers or maintenance professionals. Data sources include any information in the asset record such as maintenance and age, weather information in which the asset has been operating, and operational data, such as supervisory control and data acquisition (SCADA), historians, MQ Telemetry Transport (MQTT), Open Platform Communications (OPC), Open Systems Interconnection (OSI), enterprise asset management (EAM), geographic information system (GIS), and IoT.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Base' Tier level.

Maximo Predict

- Maximo Predict helps identify and manage asset reliability risks that could adversely affect plant or business operations. The solution enables organizations to apply machine learning and analytics to improve maintenance strategies. The cost of maintenance management is minimized by automating the steps to predict failure based on readily available operational data.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' Tier level.

Maximo Visual Inspection and Maximo Visual Inspection Edge

- Maximo Visual Inspection enables AI computer vision models to be built efficiently and easily while maintaining high levels of accuracy. Subject matter experts are provided with the power of AI. Whether it is inspecting an asset or monitoring quality on a production line, clients can experience a return on investment.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' Tier level.
- The Maximo Visual Inspection application also requires per install licenses. Allocated AppPoints is 45.

- Maximo Visual Inspection Edge is a connected device at the edge, for example, a mobile device, camera, data stream/storage device, used to perform an AI workflow resulting in an inference.
- Maximo Visual Inspection Edge also requires per install licenses. Allocated AppPoints is 1 per Maximo Visual Inspection Edge that is installed.

Maximo Collaborate

- Maximo Collaborate helps organization reduce meantime to repair, improve first-time fix rates, and improve overall field technician productivity by enabling technicians to contact remote experts for assistance through a remote collaboration session.
- Maximo Collaborate helps organizations scale the knowledge of their experts and improve the training experience for their new workforce.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Limited' Tier level.

Maximo Real Estate and Facilities

- Provides comprehensive real estate and facilities management that helps organizations manage and optimize their workplace experience, facility resource scheduling, facility strategic planning, lease accounting, and space management.
- Licensing is Authorized and/or Concurrent User; AppPoints based on Tier level – Self- Service, Limited, Base or Premium.
- Reserve requires per Install licenses; AppPoints allocated based on tier level; Limited Tier: 20 AppPoints < 10K reserve users; Base Tier: 120 AppPoints – 10K – 100K reserve users; Advanced Tier: 200 AppPoints > 100K reserve users

Maximo AI Service

- Provides a central AI engine for generative AI use cases by supporting numerous use cases through the service back-end service that hosts embedding models, processing and built with watsonx to support foundation model related use cases.
- License usage is measured by the number of content tokens consumed by features enabled by the AI Service, such as the Maximo AI assistant. Token consumption is converted to AppPoints and billed accordingly based on actual usage.
- 1 Billion Content Tokens per month / 10 AppPoints

License ratio & user type

- [License ratios](#)
- [Application user type](#)
- [Noncharged entitlement for application users](#)
- [Administrator user type](#)
- [Install type](#)

License ratios

IBM Administrative Base user

Entitlement conversion ratio: 1 Authorized User is 10 AppPoints.

IBM Administrative Premium User

Entitlement conversion ratio: 1 Authorized User is 15 AppPoints.

IBM Application Limited user

Entitlement conversion ratio: 1 Concurrent User is 5 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 2 AppPoints.

IBM Application Base user

Entitlement conversion ratio: 1 Concurrent User is 10 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 3 AppPoints.

IBM Application Premium user

Entitlement conversion ratio: 1 Concurrent User is 15 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 5 AppPoints.

IBM Maximo Connector for Oracle applications install

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Connector for SAP Applications

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Connector for Workday Applications

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Spatial install

Entitlement conversion ratio: 1 installation is 20 AppPoints.

IBM Maximo Civil Infrastructure

Entitlement conversion ratio: 1 installation is 50 AppPoints.

IBM Maximo Optimizer install

Entitlement conversion ratio: 1 installation is 220 AppPoints.

IBM Maximo Optimizer Limited install

Entitlement conversion ratio: 1 installation is 60 AppPoints.

Note: The use of IBM Maximo Optimizer Limited has the following additional restrictions: (a) deployment is limited to 2 virtual private clouds (VPCs), (b) use of a single optimizer model as provided by or IBM its extension, (c) job execution must be performed serially.

IBM Maximo Visual Inspection install

Entitlement conversion ratio: 1 installation is 45 AppPoints.

IBM Maximo Visual Inspection Edge install

Entitlement conversion Ratio: 1 installation is 1 AppPoints.

IBM Maximo Real Estate and Facilities Limited install

Entitlement conversion ratio: 1 install/ 20 AppPoints.

IBM Maximo Real Estate and Facilities Base install

Entitlement conversion ratio: 1 install/ 120 AppPoints

IBM Maximo Real Estate and Facilities Reservation Advanced install

Entitlement conversion ratio: 1 install/ 200 AppPoints

IBM Maximo AI Service Content Tokens per month

Entitlement Conversion Ratio: 1 Billion Content Tokens per month/ 10 AppPoints.

Application user types**Application limited user**

An Application Limited user accessing Maximo Manage functions of the Program can complete the following tasks:

- Use all modules for the limited purpose of running and viewing reports, viewing records as read only, changing the status of records, and updating Work Orders and conducting Inspections assigned to the Maximo Manage Application Limited User
- Use without restriction a maximum of three Maximo modules but excludes Planning and Scheduling, Application Administration, Administration, Integration, and Security and System Configuration.

An Application Limited user can also be granted access to any Maximo Mobile app that falls within the three-module requirement, as well as Maximo Monitor.

Note:

Additionally, a Limited User may update work orders assigned to them using the Work Order Tracking, Quick Reporting, Labor Reporting, or Activity and Tasks applications. For example, a user can have access to applications in the Assets, Planning and Purchasing modules as their 3-modules along with access to the Work Order Tracking application to update work orders and labor reporting and still be classified as Limited.

Note:

If an IS/Add-on is installed in a MAS instance then users given access to these applications are either Self-Service, Limited or Premium – there are no Base users if given access to Manage IS/Add-on applications.

Application base user

An Application Base User can access Maximo Manage including access to Planning and Scheduling, Administration, Integration, Security, and System Configuration modules but excludes the industry solution and add-on applications or modules. An Application Base user also has access to Maximo Health.

Application premium user

An Application Premium user can be granted access to all applications in Maximo Application Suite except the Suite administration application for Deployment, User Management, and Configuration.

Application user type - AppPoint allocation

<i>Table 4. AppPoint allocation for application users</i>				
Application user	Limited	Base	Premium	Description
Maximo Manage Limited	✓	✓	✓	<p>Any 3 Maximo Manage modules but cannot be Planning and Scheduling, Application Administration, Administration, Integration, System Config, Security.</p> <p>Ability to view reports, change status, and update work orders in any module (former Express capability). Manage now includes Linear, Calibration, and Spatial users. Spatial requires an installation.</p> <p>For more information, see Maximo Manage modules and applications table in the Appendix.</p>

Application user	Limited	Base	Premium	Description
Maximo Manage industry solution and add-on Limited	✓	-	✓	Includes Maximo Manage industry solutions (IS) and add-ons along with Maximo IT. If installed and more than 3 modules are selected and access is granted to at least one IS or Add-on application, then the user must be a Premium user and cannot be a Base user. The 3 modules that are selected cannot be Planning and Scheduling, Application Administration, Administration, Integration, System Config, or Security. For more information, see Maximo Manage modules and applications table in the Appendix
Maximo Mobile	✓	✓	✓	Includes Maximo Mobile (3-modules for Limited use), Maximo Collaborate or third-party mobile.
Maximo Monitor	✓	✓	✓	
Maximo Real Estate and Facilities	✓	✓	✓	Maximo Real Estate and Facilities can be limited, base or premium users. For more information, see “Maximo Real Estate and Facilities AppPoints consumption” on page 102.
Maximo Manage		✓	✓	Includes Scheduler, Linear, Calibration, and Spatial. Spatial requires an installation. Includes Maximo IT if Service Provider (SP) is not installed.
Maximo Health		✓	✓	
Industry solutions			✓	The Maximo Manage IS are: <ul style="list-style-type: none"> • Maximo Aviation • Maximo Civil Infrastructure • Maximo Oil & Gas • Maximo Transportation • Maximo Nuclear • Maximo Utilities

<i>Table 4. AppPoint allocation for application users (continued)</i>				
Application user	Limited	Base	Premium	Description
Add-ons			✓	Maximo Manage add-on - <ul style="list-style-type: none"> • Maximo Asset Configuration Manager • Maximo Service Provider • Maximo Health, Safety and Environment
Maximo Predict			✓	
Maximo Visual Inspection			✓	Install AppPoints are required.

Noncharged entitlements for application users

Self-service applications do not require entitlement - 0 AppPoints.

- Entering service requests and viewing the status of their service requests by using **Create Service Request** and **View Service Request**.
- Creating and viewing requisitions, viewing templates, and viewing drafts in the Desktop Requisitions application.
- Mobile Service Request application.
- Create Incident in Maximo Oil & Gas.
- Review Incidents in Maximo Oil & Gas.
- Service Requests in Maximo Oil & Gas.
- Create a Management of Change (MOC) request, with view-only access to related artifacts, in the MOC in Maximo Oil & Gas.
- Create Incident in Maximo Health, Safety and Environment.
- Review Incidents in Maximo Health, Safety and Environment.
- Service Requests in Maximo Health, Safety and Environment.
- Create a Management of Change request (with view-only access to related artifacts) in MOC in Maximo Health, Safety and Environment.
- Graphical Appointment Book in Maximo Manage.
- Vehicle Requests in Maximo Transportation.
- Access the Bill Review applications in Maximo Service Provider.
- Creating and viewing condition reports in Create Condition Report in Maximo Nuclear.
- Users who are created by default by the Program for performance improvement and not used for any other purposes. For example, Maximo Mobile Template User.
- Users who are created solely for purpose of performance testing in the Licensee's internal test environment will not consume AppPoints when enforcement is tuned off. For example, User System testing and performance load testing of the Maximo Application Suite environment.
- Program Installs of a component that is used as part of the Licensee's internal development and test environment with at least one installation of the same component in the production environment.
- Use of IBM App Connect Enterprise for solution integration where one end of the integration pattern is Maximo Application Suite, and the other end of the integration pattern is IBM TRIRIGA Application Suite, IBM Environmental Intelligence Suite, IBM TRIRIGA Portfolio Data Manager, IBM TRIRIGA Application Platform, or IBM Facilities and Real Estate Management on Cloud (TRIRIGA).
- Install and use of IBM Maximo Models for Electrical Distribution.
- Create Service requests in Maximo Real Estate and Facilities.

- View news in Maximo Real Estate and Facilities.
- Create Real Estate requests in Maximo Real Estate and Facilities.
- Create Project requests in Maximo Real Estate and Facilities.
- Create Move requests in Maximo Real Estate and Facilities.
- Create Product requests in Maximo Real Estate and Facilities.
- Search locations, people and assets in the Locate perceptive App in Maximo Real Estate and Facilities.

Administrative user types

Maximo Application Suite administrative entitlements are available for Base and Premium levels of access.

Administrative User Type - AppPoint allocation

Table 5. AppPoint allocation for administrative user

Administrative user	Base (10 AppPoints)	Premium (15 AppPoints)	Description
Application	✓	✓	Application administrators administrate one or more applications.
Suite		✓	Suite administrator manages overarching system configuration settings from the Suite Administration page.

- Administrators are defined users in the system.
- AppPoints used for administration entitlement are reserved continuously from the pool of available AppPoints in the same way authorized users are, so that administrators always have access.
- Administrators can also hold application entitlement of any tier, enabling them to be granted access to all application functionality without additional AppPoints requirements. For example, a Base administrator withdraws 10 AppPoints is given Premium application entitlement and access to Maximo Predict, which would usually require 15 AppPoints.

Install types

Install AppPoint allocation

Table 6. AppPoint allocation for specific application installation

Application Installs	AppPoints
Maximo Connector for SAP Applications	80
Maximo Connector for Oracle Applications	80
Maximo Connector for Workday Applications	80
Maximo Spatial	20
Maximo Civil Infrastructure	50
Maximo Optimizer Limited	60
Maximo Optimizer	220
Maximo Visual Inspection	45

<i>Table 6. AppPoint allocation for specific application installation (continued)</i>	
Application Installs	AppPoints
Maximo Visual Inspection Edge	1
Reserve Limited	20
Reserve Base	120
Reserve Advanced	200
AI Service	10

- Program Installs of a component that is used as part of the Licensee's internal development and test environment with at least one (1) Install of the same component in the production environment, do not consume AppPoints.
- Install AppPoints are consumed when the environment is deployed and are unavailable for use by other users or installs for the duration of that environment's life.

Offering-specific licenses

- [Red Hat Entitlements](#)
- [App Connect Entitlements](#)
- [Cognos Entitlements](#)
- [DB2 Entitlements](#)

Red Hat entitlements

Red Hat products that are listed include software, maintenance, and support in the form of Red Hat Software Subscriptions governed by a separate Red Hat Enterprise Agreement set forth at <https://www.redhat.com/en/about/agreements> (or if applicable, a negotiated version of such agreement between Licensee and Red Hat). Under this License Information, Licensee acquires Red Hat Software Subscription entitlements for the listed Red Hat Products in the quantities and/or ratios set out in this LI. The Red Hat Software Subscription entitlements are applicable and valid only to the extent the Red Hat Products are used in support of the Program, and only while Licensee has IBM Software Subscription and Support (S&S) in effect for the Program. Red Hat Software Subscription entitlements do not include support of the Red Hat Products to run, deploy or otherwise support any software other than the Program, and Licensee acknowledges additional per unit fees at Red Hat standard rates will apply for any broader usage. Red Hat's standard Software Subscription support does not include Red Hat's remote access to Licensee's network and/or systems, but, if the parties mutually agree that remote access is needed for a support issue, such remote access is subject to the terms of the Remote Access Rider to the Red Hat agreement set forth at: <https://www.redhat.com/wapps/tnc/standalone/RemoteAccessRider> (Red Hat login required). If Licensee's Program entitlement is designated "Reserved", the following entitlements to the Red Hat Products are not included and do not apply.

Red Hat products

- Red Hat Enterprise Linux®
- Red Hat OpenShift Container Platform
 - Additional Entitlement Ratio: First AppPoint / 166 VPC, each additional 8 AppPoints / 1 VPC.
- Red Hat OpenShift AI
 - Additional Entitlement Ratio: First AppPoint / 45 VPC, each additional 30 AppPoints / 1 VPC
- Supporting programs, open source, and separately licensed code packages and their operators required to operate Maximo Application Suite are permitted to be deployed on the same Red Hat OpenShift Container Platform clusters as Maximo Application Suite.

- Provided sufficient additional Red Hat OpenShift entitlements have been purchased from Red Hat for software other than the Program (including third party or other IBM software), a client may deploy such software in same Red Hat OpenShift clusters as Maximo Application Suite.
- For deployment of permitted workloads on infrastructure nodes, including security agents, monitoring agents (that maybe custom or 3rd party), refer to the [Red Hat OpenShift subscription guide](#).

App Connect entitlements

- License covers inclusion of a restricted use entitlement to IBM App Connect Enterprise for solution integration where one end of the integration pattern is to IBM Maximo Application Suite.
- Use of IBM App Connect Enterprise for solution integrations with IBM Maximo Health, IBM Maximo Predict, or IBM Maximo Manage, for Health functions, is limited to 3 VPCs per production environment
- Use of IBM App Connect Enterprise for solution integrations with IBM Maximo Connector for Workday Applications is limited to 4 VPCs per production environment.

The number of VPCs used by IBM App Connect Enterprise in connection with the individual applications set forth above may be reallocated among those applications by the Customer at its discretion so long as the total number of VPCs used does not exceed the aggregate number of VPCs authorized under this section.

Cognos entitlements

Additional restriction specific to creation & editing of reports or dashboards:

- Requires IBM Administrative Base User license.
- Limited to up to three IBM Administrative Base Users who can perform creation/editing.

Db2 entitlements

Db2 Warehouse

Warm Standby: allocate 1 VPC (counted towards overall VPC entitlement through the Maximo Application Suite license).

Db2 Standard Edition

The Supporting Program may use a maximum of 16 processor cores and 128 GB of memory on each physical or virtual server.

Non-Production instances

Licensee may only use the following components as part of Licensee's internal development and test environment for internal non-production activities ("Non-Production").

Non-Production components

IBM provides non-production components exclusively through a continuous delivery (CD) stream as part of the Maximo Application Suite Program. This stream introduces new features, developed for the next release, allowing early access for non-production environment use. This stream is offered alongside and in parallel with our normal maintained streams.

For a list of non-production components or features being introduced in a monthly CD stream, see <https://ibm.biz/MASFC-whats-new>

Additional terms for Service Providers

Additional definitions

1. Licensee is a "Service Provider" if Licensee uses the Programs to provide services in support of or to benefit one or more of Licensee's Client(s).
2. "Client" means a third party that engages Licensee to provide services using the Program; a Client cannot, by more than 50%, own, be owned by, or be under common ownership with, Licensee. A Client cannot acquire Authorized User or Limited Use User entitlements to the Program.

License scope and restrictions

1. The Program is for use within Licensee's Enterprise only and may not be resold, rented, leased, or transferred to third parties. Any attempt to do so in violation of these provisions is void. In addition, Licensee may not use the Program to provide any timesharing, commercial hosting or service bureau usage. The foregoing notwithstanding:
 - Licensee may use the Programs to process data on behalf of Licensee's Client(s).
 - Licensee may use the Programs to support more than one Client at a time, up to the total number of users set forth in the applicable PoE.
 - While Client(s) may use or benefit from the use of the Programs provided to Licensee under this Agreement, no Client is a party to this Agreement.
2. Licensee will maintain and provide support to Licensee's Client(s) for any Program used as part of services Licensee provides to Licensee's Client(s), and ensure that each Client is notified to look to Licensee, not IBM, for any such maintenance and support.
3. Licensee is responsible for Licensee's Client's compliance with the terms of this Agreement. In this regard, Licensee must enter into Licensee's own written agreement with each Client that:
 - Prohibits any action by the Client that is not in compliance with the terms of this Agreement, including, without limitation, reverse assembly, reverse compilation, reverse engineering or other translation of the Program;
 - Notifies the Client that the Program is copyrighted and licensed (not sold) and that at no time is the license to the Program transferred to the Client;
 - Notifies the Client that "SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM MAKES NO WARRANTIES OR CONDITIONS, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THE USE OF THE PROGRAM, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT";
 - Notifies the Client that IBM is not liable for any lost profits, lost savings, or any incidental or other economic consequential damages resulting from the use of the Program, regardless of whether the Licensee, the Client or IBM have been advised of the possibility of such damages.
 - Prohibits the Client's use of the Program except when used as part of services Licensee provides to the Client, and obligates the Client to destroy or return all copies of the Program to Licensee when Licensee's services to the Client end.
4. Licensee will:
 - Immediately notify IBM if Licensee becomes aware of any act or failure to act by a Client which, if committed by Licensee, would immediately or with the passage of time become a breach of the Agreement;
 - Provide all information and assistance requested by IBM in preventing, investigating, and mitigating any such breach;
 - Pay all license fees owed and indemnify IBM against any and all claims, damages, losses, costs and expenses suffered or incurred by IBM arising out of any such breach or threatened breach or of a Client's use of or problems with the Programs.
5. General:
 - Licensee agrees to indemnify and defend IBM against any claims or lawsuits, including reasonable attorney's fees, that arise or result, directly or indirectly, from providing services using the Program to Licensee's Client or such Client's accessing or use of the Program.
 - Licensee is solely responsible for the payment of any duty, tax, levy, or fee that any authority may impose that results from Licensee providing services using the Program.

Cloned or custom applications

- Maximo Manage user-defined applications can either be cloned applications or custom applications. Cloned applications are based on a copy of an existing Maximo Manage application, which is also called

the originating application. Custom applications are built from scratch, have no originating application and are built based on either a custom object or an existing Maximo object.

- Users being granted access to a cloned application must have at least the same level of access in the originating application.
- Limited Use users may be granted access to custom applications when the main object is a custom object or an object the user is otherwise entitled to access. The custom application must be a part of one of the three modules selected by the user.
- If the user needs to access custom or cloned applications that are based on an application or object not within the three modules selected for Manage Limited, then Manage Base or Industry Solution/Add-on Premium is required.
- Custom or cloned applications shall not be utilized to circumvent licensing or grant users functionality or access they would otherwise not be entitled to access.

Back-up and restore

- Governed by IPLA level language.
- https://www.ibm.com/about/software-licensing/assets/guides_pdf/Backup.pdf

Standby scenarios for Production backup:

- Cold (entitlement is not required)
- Warm (depends on access to live production data)
- Hot (entitlement IS required)

Licensing requirements for indirect access scenarios

This section covers scenarios related to license requirements when integrating with external systems, external custom applications, websites, or 3rd party mobile solutions.

Licensee must obtain separate, dedicated entitlements for Concurrent Users or Authorized Users given access to the Program in any manner directly or indirectly (for example: via a multiplexing program, device, or application server) through any means.

License Circumvention Licensee will not attempt to aggregate users or otherwise circumvent IBM's licensing restrictions via technical means, including without limitation the use of any interface between the Program and other software that performs functionality substantially similar to that contained in any IBM program then offered by IBM when Licensee acquires the Program.

Third-party mobile licensing

For 3rd party mobile access and email approval, capabilities licensing is required.

Licensing is Authorized and Concurrent user. AppPoints based on Tier level.

Integrations and external applications in Maximo Manage

- There are instances when an organization licensed to use Maximo Manage will elect to use an external application, develop custom applications, websites, or use 3rd party mobile applications as a way to enable their users to interact with records, data, and information that is stored in Maximo Manage.
- If the applications are integrated with Maximo Manage, (importing/exporting data), the distinction of which system is the master of that data also determines if a license is required to create or update those records.
- If data that is considered to be created and housed as the 'source of truth' in Maximo Manage, then a license is required to insert or update that data. If a third-party system is considered the master of the data, any replication of that data into Maximo Manage for use within other processes does not require a license.

- Examples of when each user would require a Maximo Manage License include the use of a 3rd party application to update Work Orders; create or update Asset records; change status of a record, or any supply chain related functions if they are primarily initiated and managed in Maximo.
- Customers would need to purchase enough AppPoints to cover these users but is not registered as part of the Suite License server.

When is integration allowed without any additional end-user license required?

- Maximo Manage includes an extensive set of integration capabilities, referred to as the Maximo integration framework (MIF). There are also specific adapters available for integration with SAP or Oracle ERP systems. The MIF and adapters allow companies to share and exchange information between Maximo Manage and external systems in real time or in batch mode. The MIF is run using the Maximo Manage admin (max integration admin) user which does not require a specific license.
- Examples of a third-party system that are considered the master data repository and do not require a Maximo Manage license (except for established Maximo Application Suite connectors) would include things like a HR system with personnel information, pay rates and certification data that is synched into a labor record in Maximo Manage on a regular basis; an ERP system that manages all the Item/Inventory and purchasing transactions that pushes updates like new item records or receipts into Manage; or a Scada or other system with operational asset data being pushed into an asset meter reading record. In all of these cases the external system is the system of record and the synchronization of that data into Maximo Manage does not require an end user license.
- Maximo Manage Self Service applications/use cases are exempt from licensing.
 - For example, a college creates a web app that enables students to create service requests to report problems/incidents in the dorms. This is the same as users of the Manage self-service service request (create/view) applications, for which users do not require a license entitlement.
- Extracting Maximo Manage data is for reporting purposes only.
 - A user is emailed a copy of a report from Maximo Manage.
 - Maximo Manage data is copied into data warehouse (no updates back into Maximo Manage).
 - In these cases, there is no license required for users receiving data in this way.

Appendix: Modules and applications in Maximo Application Suite 9.1

- [Manage Modules](#)
- [Manage Industry Solution and Add-on Modules](#)
- [Maximo Real Estate and Facilities AppPoints consumption](#)

Maximo Manage modules

The following table lists the modules and applications that are included in Maximo Manage. The applications are organized and grouped by module.

<i>Table 7. Modules and applications in</i>	
Maximo Application Suite	
Module and sub-module	Application
Application Administration	Mobile Configuration
	Application Configuration
Security	User
	Security Group

Table 8. Modules and applications in Maximo Manage

Maximo Manage	
Module and sub-module	Application
Administration	Sets
	Organizations
	Calendars
	Bulletin Board
	Communication Templates
	Report Administration
	Conditional Expression Manager
	Classifications
	Work View
	Service Addresses
	Map Manager
	Record Release
	Time Zone Rules
	Scheduler Administration
	Push Notification Administration
	Synchronization ArcGIS
	Work Queue Manager
	Guest User
	AI Configuration
	Visual Inspection Models
Resources	Labor
	Qualifications
	People
	Person Group
	Crafts
	Crew Types
	Crews
KPI	KPI Manager
	KPI Templates
Scheduler Administration	Appointment Book Manager
	Scheduling Alternate Resource
	Configure ToolTips
	Scheduler Data Manager

Table 8. Modules and applications in Maximo Manage (continued)

Maximo Manage	
Module and sub-module	Application
	Optimizer Administrator
Analytics	Report Viewer
	KPI Viewer
	Design Data Sets
	Business Analyst
Assets	Assets
	Asset Templates
	Locations
	Features
	Meters
	Relationships
	Meter Groups
	Condition Monitoring
	Failure Codes
	Reliability Strategies
Role Based Applications	Asset Manager
Building Information Models	BIM Projects
Contracts	Purchase Contracts
	Lease Rental Contracts
	Labor Rate Contracts
	Master Contracts
	Warranty Contracts
	Terms and Conditions
Financial	Currency Codes
	Exchange Rates
	Chart of Accounts
	Cost Management
	Budget Monitoring
Integration	Object Structures
	Publish Channels
	Invocation Channels
	Enterprise Services
	Web Services Library

Table 8. Modules and applications in Maximo Manage (continued)

Maximo Manage	
Module and sub-module	Application
	End Points
	External Systems
	Logical Management Operations
	Integration Modules
	Launch in Context
	Message Tracking
	Message Reprocessing
Interactions	Create Interaction
	Interactions
	OSLC Providers
	OSLC Resources
	JSON Resources
	JSON Mapping
	Notifications
Inventory	Item Master
	Service Items
	Tools
	Stocked Tools
	Inventory
	Inventory Usage
	Shipment Receiving
	Condition Codes
	Storerooms
	Collections
	Count Book
Role Based Application	Inventory Counting
	Inventory Receiving
	Issues and Transfers
Planning and Scheduling	Graphical Assignment
	Graphical Work Week
	Graphical Scheduling
	Graphical Appointment Book
	Graphical Crew Management

Table 8. Modules and applications in Maximo Manage (continued)

Maximo Manage	
Module and sub-module	Application
	Graphical Resource View
	Scheduling Dashboard
	Dispatching Dashboard
	Planning Dashboard
	LBS Location History
Planning	Job Plans
	Routes
	Inspection Forms
Safety	Hazards
	Precautions
	Lock Out / Tag Out
	Safety Plans
	Data Sheet Template
Preventive Maintenance	Preventive Maintenance
	Master PM
Purchasing	Purchase Requisitions
	Purchase Orders
	Receiving
	Shipment Receiving
	Invoices
	Request for Quotations
	Companies
	Company Master
	Terms and Conditions
Security	Security Groups (Manage)
	Users (Manage)
Self Service	
Desktop Requisitions	Create Requisitions
	View Requisitions
	View Templates
	View Drafts
Service Requests	Create Service Request
	View Service Request

Table 8. Modules and applications in Maximo Manage (continued)

Maximo Manage	
Module and sub-module	Application
	Search Solutions
Role Based Applications	Service Request
Service Level	Service Level Agreements
	Service Groups
Service Desk	Activities and Tasks
	Service Requests
	Solutions
	Ticket Templates
System Configuration	
Platform Configuration	System Properties
	Logging
	Domains
	Database Configuration
	Application Designer
	Communication Templates
	Actions
	Roles
	Escalations
	Workflow Designer
	Workflow Administration
	Cron Task Setup
	E-mail Listeners
	Web Services Library
	Launch in Context
	Instant Messaging Configuration
	Automation Scripts
Migration	Object Structures
	Migration Manager
	Migration Groups
	Migration Collections
	Maximo Management Interface
Task Management	Activities and Tasks
Work Orders	Work Order Tracking

Table 8. Modules and applications in Maximo Manage (continued)

Maximo Manage	
Module and sub-module	Application
	Labor Reporting
	Quick Reporting
	Activities and Tasks
	Assignment Manager
	Service Requests
Work Centers	Conduct an Inspection
Role Based Applications	Inspections
	Technician
	Work Approval
	Service Request

Maximo Manage industry solution and Add-on modules

The following table lists the modules for Maximo Manage industry solutions and add-on modules.

Table 9. Industry solutions and add-on modules in Maximo Manage. Industry solutions and add-on modules in Maximo Manage

Industry /Add-on	Modules
Asset Configuration Manager (ACM)	Asset Configuration Manager
	Flight Logs
Service Provider (SP)	Service Provider (SP)
Health, Safety and Environment (HSE)	Change (HSE)
	Operations (HSE)
	Safety and Quality Management (HSE)
Oil & Gas	Change (Oil)
	Operations (Oil)
	Safety and Quality Management (Oil)
Transportation	Data Import (Tr)
	Motor Pool (Tr)
	Warranties (Tr)
Nuclear	Configuration Change (Nuc)
	Operation Management (Nuc)
	Permits (Nuc)
	Condition Report (Nuc)

Table 9. Industry solutions and add-on modules in Maximo Manage. Industry solutions and add-on modules in Maximo Manage (continued)

Industry /Add-on	Modules
Aviation	Aircraft and Equipment
	Configuration Management
	Consists
	Data Import
	Flight Log
	Exchange Orders
	Motor Pool
	Operations
	Safety and Quality Management
	Warranties

Maximo Real Estate and Facilities AppPoints consumption

User are entitled to use all applications and components in the assigned tier.

Table 10. Access type for Maximo Real Estate and Facilities users

Access type	Acquire and Dispose	Build-out	Populate portfolio data	Administer real estate	Experience	Maintain
Self service <ul style="list-style-type: none"> Authorized user 0 AppPoints Concurrent users 0 AppPoints 	<ul style="list-style-type: none"> Service requests News Real estate requests 	<ul style="list-style-type: none"> Service requests News Project requests 		<ul style="list-style-type: none"> Service requests News Real estate requests 	<ul style="list-style-type: none"> Service request Locate News Move requests 	<ul style="list-style-type: none"> Service request News Product requests

Table 10. Access type for Maximo Real Estate and Facilities users (continued)

Access type	Acquire and Dispose	Build-out	Populate portfolio data	Administer real estate	Experience	Maintain
Limited <ul style="list-style-type: none"> • Authorized user 2 AppPoints • Concurrent users 5 AppPoints 	<ul style="list-style-type: none"> • Self service access • Reports and metrics • Approvals • Contract requests • Offline forms 	<ul style="list-style-type: none"> • Self service access • Reports and metrics • Approvals • Vendor bid response • Change requests • Project tasks • Offline forms 	<ul style="list-style-type: none"> • Self service access • Floor plan reports 	<ul style="list-style-type: none"> • Self service access • Reports and metrics • Approvals • Contract requests • Offline forms 	<ul style="list-style-type: none"> • Self service access • Reports and metrics • Approvals • Move & reserve tasks • Space assessment • Offline forms 	<ul style="list-style-type: none"> • Self service access • Reports and metrics • Approvals • Work tasks exec • Contact center • Offline forms

Table 10. Access type for Maximo Real Estate and Facilities users (continued)

Access type	Acquire and Dispose	Build-out	Populate portfolio data	Administer real estate	Experience	Maintain
<p>Base</p> <ul style="list-style-type: none"> • Authorized user 3 AppPoints • Concurrent users 10 AppPoints 	<ul style="list-style-type: none"> • Limited access • Advanced scenarios • Transaction management • Portfolio planning 	<ul style="list-style-type: none"> • Limited access • Cost management • Schedule management • Purchase order approvals • Procurement management • Program management • Facility condition assessment 	<ul style="list-style-type: none"> • Limited access • Drawing and model management • Energy star benchmarking 	<ul style="list-style-type: none"> • Limited access • Real estate lease administration • AR tenant tracking • Payment processing • Generate journal entries 	<ul style="list-style-type: none"> • Limited access • Space planning • Move management • Reserve coordination • Facility project management • Strategic planning 	<ul style="list-style-type: none"> • Limited access • FCA inspections • O&M management • Model management • Asset lease administration • Facility project management • Facility condition assessment • Opportunity analysis • Work requests from SNMP alerts • Energy and waste log management • Utility invoice tracking
<p>Premium</p> <ul style="list-style-type: none"> • Authorized user 5 AppPoints • Concurrent users 15 AppPoints 			<ul style="list-style-type: none"> • Base • Configuration tools 			

Table 10. Access type for Maximo Real Estate and Facilities users (continued)

Access type	Acquire and Dispose	Build-out	Populate portfolio data	Administer real estate	Experience	Maintain
Reserve + room panel • Install user	• Tier 1 20 AppPoints <10K reserve users • Tier 2 120 AppPoints 10K-100K reserve users • Tier 3 200 AppPoints >100K reserve users				• Self service access • Room and desk reservations	

Customer-managed **Licensing in Maximo Application Suite 9.0 and earlier**

The Licensing Guidance provides information about licensing and entitlements for IBM Maximo Application Suite.

Notice: This Licensing Guidance is intended to provide only supplementary information to assist you in deploying the Program(s) you have licensed from IBM within your purchased entitlement. Your license agreement (such as the IBM International Program License Agreement (IPLA)) or equivalent and its transaction documents, including the License Information for IBM Maximo Application Suite is the sole and complete agreement between you and IBM regarding use of the Program.

- [Listing of licenses by type](#)
- [What do you get with your purchase of IBM Maximo Application Suite, and what is your entitlement?](#)
- [“License ratio & user type” on page 111](#)
- [Offering-specific licenses](#)
- [Non-production instances](#)
- [Additional terms for service providers](#)
- [Cloned or custom applications](#)
- [Back-up and restore](#)
- [Licensing requirements for indirect access scenarios](#)
- [Appendix: Modules & applications](#)

Note: The Licensing Guidance is also available in [PDF](#).

Maximo Application Suite uses:

AppPoints

AppPoints is a unit of measure by which the Program can be licensed. An AppPoint is a common unit of value. Sufficient AppPoint entitlements must be obtained to cover the total number of entitlements required for Licensee's Authorized Use of the Program.

Red Hat OpenShift

All deployments of Maximo Application Suite that are deployed on Red Hat OpenShift Container Platform must have sufficient entitlement for the Maximo Application Suite cores that are used.

Listing of licenses by type

See also

- [IBM Maximo Application Suite 8.11 \(L-XSNF-SHG8GG\)](#)
- [IBM Maximo Application Suite 9.0 \(L-GRJQ-AJY62V\)](#)

The following license types are used for AppPoints in Maximo Application Suite:

- [“Concurrent user” on page 106](#)
- [“Authorized user ” on page 106](#)

- [“Access for Concurrent & Authorized users” on page 106](#)
- [“Install” on page 106](#)

Concurrent user

A unit of measure by which the Program can be licensed. A Concurrent user is a person who is accessing the Program at any particular point in time. Regardless of whether the person is simultaneously accessing the Program multiple times, the person counts only as a single Concurrent User. The Program may be installed on any number of computers or servers, but the Licensee must obtain entitlements ([Proof of Entitlement\(PoE\)](#)) for the maximum number of Concurrent Users simultaneously accessing the Program. The licensee must obtain an entitlement for each simultaneous Concurrent User accessing the Program in any manner directly or indirectly (for example: via a multiplexing program, device, or application server) through any means.

Authorized user

A unit of measure by which the Program can be licensed. An Authorized user is a unique person who is given access to the Program. The Program may be installed on any number of computers or servers and each Authorized User may have simultaneous access to any number of instances of the Program at one time. The licensee must obtain separate, dedicated entitlements for each Authorized user that is given access to the Program in any manner directly or indirectly (for example, through a multiplexing program, device, or application server) through any means. An entitlement for an Authorized user is unique to that Authorized User and may not be shared, nor may it be reassigned other than for the permanent transfer of the Authorized user entitlement to another person.

Access for Concurrent & Authorized users

A customer can set up any users access type as Authorized or Concurrent. Entitlement levels for Users include:

- Applications that do not require entitlement.
- Limited Users with entitlement to 3 modules of the Maximo Manage application, and the Monitor application.
- Base Users with entitlement to many modules of the Maximo Manage application, as well as the Maximo Health application.
- Premium Users with access to the entire suite of applications.

Each level of access inherits the entitlements that are granted with every lower tier license.

Install

A software program instance. Licensee must obtain an entitlement for each Installation of the Program, except where outlined in the section of Applications that do not require entitlement.

What do you get with your purchase of IBM Maximo Application Suite, and what is your entitlement?

- [Maximo Manage](#)
- [Maximo Manage industry solutions and add-ons](#)
- [Maximo Manage Mobile](#)
- [Maximo Manage integration solutions](#)
- [Maximo Monitor](#)
- [Maximo Health](#)
- [Maximo Predict](#)
- [Maximo Visual Inspection and Maximo Visual Inspection Edge](#)
- [Maximo Assist](#)

Maximo Manage

- Provides comprehensive asset lifecycle and maintenance management for all asset types.
- Licensing is Authorized and Concurrent User; AppPoints based on Tier level – Base or Limited.

- Maximo Manage now includes access to Maximo Spatial, Linear, and Calibration at both the Base and Limited tier. Maximo Spatial requires a separate installation.
- Maximo Manage now includes Scheduler at the Base tier level.

Maximo Spatial

- Embeds geographic information system (GIS) functionality in the Maximo User interface that enables users visualize assets, work orders, and locations on a **Map** tab. Uses technology from Esri ArcGIS.
- Enables map and navigation controls for dynamic visual display of work orders, work requests, assets, and locations in Maximo.
- Maximo Spatial must be used with Esri ArcGIS.
- Customers must purchase Esri ArcGIS and any required analyst desktop licenses separately from Esri or an authorized agent.
- Entitled as part of Maximo Manage and requires a separate installation.
- Customers must purchase AppPoints to install Maximo Spatial.

Linear

- Linear is used for managing assets such as roads, pipelines, rail lines, and transmission lines.
- Entitled and installed as part of Maximo Manage.

Calibration

- Enables traceability and reverse traceability, all calibration history data, calibration data sheets, and required reporting.
- Entitled and installed as part of Maximo Manage.

Scheduler

- Scheduler provides an innovative means to graphically plan, and schedule current and upcoming work based on available resources and to manage aspects of advanced work management that include intuitive graphical assignment, mobile workforce management, real-time communication with the field, and schedule optimization.
- An advanced work management tool to enable users to manage large projects, such as shutdowns, outages, turnarounds, and customer appointments, and all planned and unplanned maintenance across a broad geographic area or where weather plays an important factor in asset availability.
- Licensing is Authorized and Concurrent user; AppPoints based on 'Base' or if user is given access to IS/Add-on applications 'Premium' Tier level.
- Entitled and installed as part of Maximo Manage.

Maximo Manage industry solutions and add-ons

Maximo Manage includes the following industry solutions and add-ons:

Industry solutions

- Maximo Aviation
- Maximo Civil Infrastructure
- Maximo Nuclear
- Maximo Oil & Gas
- Maximo Transportation
- Maximo Utilities

Add-ons

- Maximo Asset Configuration Manager
- Maximo Health, Safety and Environment

- Maximo IT
- Maximo Optimizer
- Maximo Service Provider
- Maximo Reliability Strategies
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level. Limited users can select 3 modules, which include the modules that are unique to an industry solution. See the [Appendix for a list of Modules](#) that are specific to the Industry Solutions.
- Install AppPoints might also be required based on IS/Add-on installed.

Maximo Aviation

- Provides Maintenance, Repair, and Overhaul (MRO) organizations that maintain aircraft or aircraft components a solution to help reduce the turnaround time and improve the amount of time available for revenue generation.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Civil Infrastructure

- Provides asset lifecycle management capabilities to help organizations reduce the total cost of ownership and improve the monitoring of the civil infrastructures they manage. Civil Infrastructure drives inspections and records inspection results about concrete and steel structures and related assets. Visualizes anomalies and work in structure-appropriate views, including linear, spatial, asset-based, and building information management views; manages planned and unplanned work for work orders, contractors, and purchasing management and creates preventive maintenance plans and job plans to be incorporated into the work schedule.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.
- The Civil Infrastructure application also requires per Install licenses; AppPoints allocated is 50.

Maximo Oil & Gas

- Provides Oil & Gas companies with best practices to help improve the productivity and efficiency of their critical assets that are related to the oil and gas industry.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Transportation

- Provides organizations with best practices to help improve the productivity of critical transportation assets, including over-the-road vehicles, rail, aircraft, and marine vessels.
- Licensing is Authorized and Concurrent User. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Nuclear

- Provides owners, operators, and maintainers of nuclear power plants best practices to help improve the productivity and efficiency of their critical assets.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Utilities

- Helps increase asset and resources effectiveness in water, gas, and electric utilities.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Asset Configuration Manager

- Maximo Asset Configuration Manager enables real-time calculation of both an asset's build and a component's life, with benefits such as improved compliance regulations and reduced operating costs.
- An asset is a complete parent asset, such as aircraft, locomotives, and passenger rail cars. An asset is not components, such as engines or axles, or groups of assets such as trains, which can be made up of multiple locomotives and passenger cars. In aviation, this is also referred to as tail numbers.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo Health, Safety and Environment

- Maximo Health, Safety and Environment provides key extensions to Maximo Manage to provide strategic applications for audit management, risk assessment, safety reporting, management of change, condition reporting, corrective actions, and training.
- Facilitates meeting or exceeding regulations and legislated requirements that impact health, safety, and the environment. Best in class capabilities and best practices for improving safety, reliability, and compliance.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.

Maximo IT

- Maximo IT is an integrated service management solution that helps you manage a comprehensive range of IT processes, services, and assets.
- Separate Part purchase with Maximo Application Suite being a prerequisite.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Limited' or 'Base' Tier level. 'Limited' or 'Premium' Tier level if Service Provider is installed.

Maximo Optimizer

- Optimizer installation is required if clients want to use optimization in Maximo Application Suite.
- Licensing is per installation. AppPoints allocated is 220 or Limited 60.

Maximo Service Provider

- Maximo Service Provider is intended for companies who deliver asset management services to their third-party user clients in a revenue-generating business model. Examples include: IBM ITD, EDS, CSC, Johnson Controls, UNICCO, Coor Service Management.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' or 'Limited' Tier level.
- IBM SP Rule: a customer that uses Maximo to manage assets for a 3rd-party company cannot own Maximo without also deploying Maximo Service Provider. A waiver can be obtained under certain circumstances with approval from WW Maximo Sales Leader, WW Price Leader, and Product Management.

Maximo Reliability Strategies

- Maximo Reliability Strategies helps maximize asset performance, reduce downtime, and boost revenues by helping clients migrate from reactive to proactive strategy-based maintenance.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' 'Base' or 'Limited' Tier level.

Maximo Mobile

- The Maximo Mobile app is available on Google, Android, and iOS stores.
- The Maximo Mobile app integrates with Maximo Assist to provide AI/AR based assistance to technicians on their mobile device.

- Maximo Mobile integrates with Maximo Manage to provide Inspection, Technician work orders, Asset, and Inventory apps on their device.

Maximo Manage integration solutions

Maximo Connector for SAP Applications

- A business interface that enables real-time information exchange between base products and SAP or SAP PM ERP applications by providing bidirectional connectivity with pre-built integrations between them
- Used with base products including Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for Oracle Applications

- A business interface that enables real-time information exchange between base products and Oracle E-Business Suite applications by providing bidirectional connectivity by using pre-built integrations between them.
- Used with base products that include Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for Workday Applications

- A business interface that enables real-time information exchange between base products and Workday applications by providing bidirectional connectivity by using pre-built integrations between them.
- Used with base products that include Maximo Manage and Maximo Industry Solutions.
- Licensing is per installation. Allocated AppPoints is 80.

Maximo Connector for TRIRIGA

- Enables bidirectional synchronization of portfolio data about people, places, and assets for consistent operations. Near real-time service request and work order routing for coordinated operations and near real-time project tasks to work-order synchronizations to enable streamlined operations.
- Licensing is per installation. Allocated AppPoints is 0.

Maximo Connector for Envizi

- Enables automatic synchronization of locations and meter readings from Maximo Manage to Maximo Connector for Envizi to automate tracking of energy usage and calculating related scope 1 and 2 emissions of electricity, natural gas, and water in Envizi.
- Licensing is per installation. Allocated AppPoints is 0.

Maximo Spatial

- Maximo Spatial installation is required if clients want to use Maximo Spatial in Maximo Application Suite.
- Licensing is per installation. Allocated AppPoints is 20.

Maximo Monitor

- A monitoring solution for real-time visibility, root-cause troubleshooting, and AI driven alerts at scale.
- Through monitoring dashboards, clients have a consolidated view of operations with data integration from multiple IoT and non-IoT sources.
- Licensing is Authorized and Concurrent user. AppPoints is based on Tier level Limited.

Maximo Health

- Maximo Health provides a first step in managing the performance of critical assets. By capturing information from the asset, the age of the asset and maintenance history, organizations can quickly understand which assets are in poor health and need attention to prevent disruption.
- Maximo Health uses ready-to-use common scoring elements and knowledge from equipment and maintenance engineers or maintenance professionals. Data sources include any information in the asset record such as maintenance and age, weather information in which the asset has been operating, and operational data, such as supervisory control and data acquisition (SCADA), historians, MQ Telemetry Transport (MQTT), Open Platform Communications (OPC), Open Systems Interconnection (OSI), enterprise asset management (EAM), geographic information system (GIS), and IoT.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Base' Tier level.

Maximo Predict

- Maximo Predict helps identify and manage asset reliability risks that could adversely affect plant or business operations. The solution enables organizations to apply machine learning and analytics to improve maintenance strategies. The cost of maintenance management is minimized by automating the steps to predict failure based on readily available operational data.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' Tier level.

Maximo Visual Inspection and Maximo Visual Inspection Edge

- Maximo Visual Inspection enables AI computer vision models to be built efficiently and easily while maintaining high levels of accuracy. Subject matter experts are provided with the power of AI. Whether it is inspecting an asset or monitoring quality on a production line, clients can experience a return on investment.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Premium' Tier level.
- The Maximo Visual Inspection application also requires per install licenses. Allocated AppPoints is 45.
- Maximo Visual Inspection Edge is a connected device at the edge, for example, a mobile device, camera, data stream/storage device, used to perform an AI workflow resulting in an inference.
- Maximo Visual Inspection Edge also requires per install licenses. Allocated AppPoints is 1 per Maximo Visual Inspection Edge that is installed.

Maximo Assist

- Maximo Assist helps organization reduce meantime to repair, improve first-time fix rates, and improve overall field technician productivity by enabling technicians to contact remote experts for assistance through a remote collaboration session.
- Maximo Assist helps organizations scale the knowledge of their experts and improve the training experience for their new workforce.
- Licensing is Authorized and Concurrent user. AppPoints based on 'Limited' Tier level.
- Maximo Assist 8.8 and earlier also requires per Install licenses. Allocated AppPoints is 150.

License ratio & user type

- [License ratios](#)
- [Application user type](#)
- [Noncharged entitlement for application users](#)
- [Administrator user type](#)
- [Install type](#)

License ratios

IBM Administrative Base user

Entitlement conversion ratio: 1 Authorized User is 10 AppPoints.

IBM Administrative Premium User

Entitlement conversion ratio: 1 Authorized User is 15 AppPoints.

IBM Application Limited user

Entitlement conversion ratio: 1 Concurrent User is 5 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 2 AppPoints.

IBM Application Base user

Entitlement conversion ratio: 1 Concurrent User is 10 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 3 AppPoints.

IBM Application Premium user

Entitlement conversion ratio: 1 Concurrent User is 15 AppPoints.

Entitlement conversion ratio: 1 Authorized User is 5 AppPoints.

IBM Maximo Connector for Oracle applications install

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Connector for SAP Applications

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Connector for Workday Applications

Entitlement conversion ratio: 1 installation is 80 AppPoints.

IBM Maximo Spatial install

Entitlement conversion ratio: 1 installation is 20 AppPoints.

IBM Maximo Civil Infrastructure

Entitlement conversion ratio: 1 installation is 50 AppPoints.

IBM Maximo Optimizer install

Entitlement conversion ratio: 1 installation is 220 AppPoints.

IBM Maximo Optimizer Limited install

Entitlement conversion ratio: 1 installation is 60 AppPoints.

Note: The use of IBM Maximo Optimizer Limited has the following additional restrictions: (a) deployment is limited to 2 virtual private clouds (VPCs), (b) use of a single optimizer model as provided by or IBM its extension, (c) job execution must be performed serially.

IBM Maximo Visual Inspection install

Entitlement conversion ratio: 1 installation is 45 AppPoints.

IBM Maximo Visual Inspection Edge install

Entitlement conversion Ratio: 1 installation is 1 AppPoints.

IBM Maximo Assist install

For IBM Maximo Assist 8.8 and earlier, Entitlement conversion ratio: 1 installation is 150 AppPoints.

Note: Starting in Maximo Application Suite 9.0, installation points are no longer required.

Application user types

Application limited user

An Application Limited user accessing Maximo Manage functions of the Program can complete the following tasks:

- Use all modules for the limited purpose of running and viewing reports, viewing records as read only, changing the status of records, and updating Work Orders and conducting Inspections assigned to the Maximo Manage Application Limited User
- Use without restriction a maximum of three Maximo Manage modules but excludes Planning and Scheduling, Administration, Integration, and Security and System Configuration.

An Application Limited user can also be granted access to any Maximo Mobile app that falls within the three-module requirement, as well as Maximo Monitor.

Application base user

An Application Base User can access Maximo Manage including access to Planning and Scheduling, Administration, Integration, Security, and System Configuration modules but excludes the industry solution and add-on applications or modules. An Application Base user also has access to Maximo Health.

Application premium user

An Application Premium user can be granted access to all applications in Maximo Application Suite except the Suite administration application for Deployment, User Management, and Configuration.

Application user type - AppPoint allocation

<i>Table 11. AppPoint allocation for application users</i>				
Application user	Limited	Base	Premium	Description
Maximo Manage Limited	✓	✓	✓	<p>Any 3 Maximo Manage modules but cannot be Planning and Scheduling, Administration, Integration, System Config, Security.</p> <p>Ability to view reports, change status, and update work orders in any module (former Express capability). Manage now includes Linear, Calibration, and Spatial users. Spatial requires an installation.</p> <p>Refer to the Maximo Manage modules and applications table in the Appendix.</p>

Table 11. AppPoint allocation for application users (continued)

Application user	Limited	Base	Premium	Description
Maximo Manage industry solution and add-on Limited	✓	-	✓	Includes Maximo Manage industry solutions (IS) and add-ons along with Maximo IT. If installed and more than 3 modules are selected and access is granted to at least one IS or Add-on application, then the user must be a Premium user and cannot be a Base user. The 3 modules that are selected cannot be Planning and Scheduling, Administration, Integration, System Config, or Security. Refer to the <u>Maximo Manage modules and applications table</u> in the Appendix
Maximo Mobile	✓	✓	✓	Includes Maximo Mobile (3-modules for Limited use), Maximo Assist (Assist requires Install) or third-party mobile.
Maximo Monitor	✓	✓	✓	
Maximo Manage		✓	✓	Includes Scheduler, Linear, Calibration, and Spatial. Spatial requires an installation. Includes Maximo IT if Service Provider (SP) is not installed.
Maximo Health		✓	✓	
Industry solutions			✓	The Maximo Manage IS are: <ul style="list-style-type: none"> • Maximo Aviation • Maximo Civil Infrastructure • Maximo Oil & Gas • Maximo Transportation • Maximo Nuclear • Maximo Utilities
Add-ons			✓	Maximo Manage add-on - <ul style="list-style-type: none"> • Maximo Asset Configuration Manager • Maximo Service Provider • Maximo Health, Safety and Environment
Maximo Predict			✓	
Maximo Visual Inspection			✓	Install AppPoints are required.

Noncharged entitlements for application users

Self-service applications do not require entitlement - 0 AppPoints.

- Entering service requests and viewing the status of their service requests by using **Create Service Request** and **View Service Request**.
- Creating and viewing requisitions, viewing templates, and viewing drafts in the Desktop Requisitions application.
- Mobile Service Request application.
- Create Incident in Maximo Oil & Gas.
- Review Incidents in Maximo Oil & Gas.
- Service Requests in Maximo Oil & Gas.
- Create a Management of Change (MOC) request, with view-only access to related artifacts, in the MOC in Maximo Oil & Gas.
- Create Incident in Maximo Health, Safety and Environment.
- Review Incidents in Maximo Health, Safety and Environment.
- Service Requests in Maximo Health, Safety and Environment.
- Create a Management of Change request (with view-only access to related artifacts) in MOC in Maximo Health, Safety and Environment.
- Graphical Appointment Book in Maximo Manage.
- Vehicle Requests in Maximo Transportation.
- Access the Bill Review applications in Maximo Service Provider.
- Creating and viewing condition reports in Create Condition Report in Maximo Nuclear.
- Users who are created by default by the Program for performance improvement and not used for any other purposes. For example, Maximo Mobile Template User.
- Users who are created solely for purpose of performance testing in the Licensee's internal test environment will not consume AppPoints when enforcement is tuned off. For example, User System testing and performance load testing of the Maximo Application Suite environment.
- Program Installs of a component that is used as part of the Licensee's internal development and test environment with at least one installation of the same component in the production environment.
- Use of IBM App Connect Enterprise for solution integration where one end of the integration pattern is Maximo Application Suite, and the other end of the integration pattern is IBM TRIRIGA Application Suite, IBM Environmental Intelligence Suite, IBM TRIRIGA Portfolio Data Manager, IBM TRIRIGA Application Platform, or IBM Facilities and Real Estate Management on Cloud (TRIRIGA).

Administrative user types

Maximo Application Suite administrative entitlements are available for Base and Premium levels of access.

Administrative User Type - AppPoint allocation

<i>Table 12. AppPoint allocation for administrative user</i>			
Administrative user	Base (10 AppPoints)	Premium (15 AppPoints)	Description
Application	✓	✓	Application administrators administer one or more applications, adds, and assigns users to these applications, and uses the application-specific user interfaces to manage further user privileges.
Suite		✓	Suite administrator manages overarching system configuration settings from the Suite Administration page.

- Administrators are defined users in the system.
- AppPoints used for administration entitlement are reserved continuously from the pool of available AppPoints in the same way authorized users are, so that administrators always have access.
- Administrators can also hold application entitlement of any tier, enabling them to be granted access to all application functionality without additional AppPoints requirements. For example, a Base administrator withdraws 10 AppPoints is given Premium application entitlement and access to Maximo Predict, which would usually require 15 AppPoints.

Install types

Install AppPoint allocation

<i>Table 13. AppPoint allocation for specific application installation</i>	
Application Installs	AppPoints
Maximo Connector for SAP Applications	80
Maximo Connector for Oracle Applications	80
Maximo Connector for Workday Applications	80
Maximo Spatial	20
Maximo Civil Infrastructure	50
Maximo Optimizer Limited	60
Maximo Optimizer	220
Maximo Visual Inspection	45
Maximo Visual Inspection Edge	1

- Program Installs of a component that is used as part of the Licensee's internal development and test environment with at least one (1) Install of the same component in the production environment, do not consume AppPoints.
- Install AppPoints are consumed when the environment is deployed and are unavailable for use by other users or installs for the duration of that environment's life.

Offering-specific licenses

- [Red Hat Entitlements](#)
- [App Connect Entitlements](#)
- [Cognos Entitlements](#)
- [DB2 Entitlements](#)

Red Hat entitlements

Red Hat products that are listed include software, maintenance, and support in the form of Red Hat Software Subscriptions governed by a separate Red Hat Enterprise Agreement set forth at <https://www.redhat.com/en/about/agreements> (or if applicable, a negotiated version of such agreement between Licensee and Red Hat). Under this License Information, Licensee acquires Red Hat Software Subscription entitlements for the listed Red Hat Products in the quantities and/or ratios set out in this LI. The Red Hat Software Subscription entitlements are applicable and valid only to the extent the Red Hat Products are used in support of the Program, and only while Licensee has IBM Software Subscription and Support (S&S) in effect for the Program. Red Hat Software Subscription entitlements do not include support of the Red Hat Products to run, deploy or otherwise support any software other than the Program, and Licensee acknowledges additional per unit fees at Red Hat standard rates will apply for any broader usage. Red Hat's standard Software Subscription support does not include Red Hat's remote access to Licensee's network and/or systems, but, if the parties mutually agree that remote access is needed for a support issue, such remote access is subject to the terms of the Remote Access Rider to the Red Hat agreement set forth at: <https://www.redhat.com/wapps/tnc/standalone/RemoteAccessRider> (Red Hat login required). If Licensee's Program entitlement is designated "Reserved", the following entitlements to the Red Hat Products are not included and do not apply.

Red Hat products

- Red Hat Enterprise Linux
- Red Hat OpenShift Container Platform
 - Additional Entitlement Ratio: First AppPoint / 166 VPC, each additional 8 AppPoints / 1 VPC.
- Supporting programs, open source, and separately licensed code packages and their operators required to operate Maximo Application Suite are permitted to be deployed on the same Red Hat OpenShift Container Platform clusters as Maximo Application Suite.
- Provided sufficient additional Red Hat OpenShift entitlements have been purchased from Red Hat for software other than the Program (including third party or other IBM software), a client may deploy such software in same Red Hat OpenShift clusters as Maximo Application Suite.
- For deployment of permitted workloads on infrastructure nodes, including security agents, monitoring agents (that maybe custom or 3rd party), refer to the [Red Hat OpenShift subscription guide](#).

App Connect entitlements

The number of VPCs used by IBM App Connect Enterprise in connection with the individual applications set forth above may be reallocated among those applications by the Customer at its discretion so long as the total number of VPCs used does not exceed the aggregate number of VPCs authorized under this section.

Cognos entitlements

Additional restriction specific to creation & editing of reports or dashboards:

- Requires IBM Administrative Base User license.
- Limited to up to three IBM Administrative Base Users who can perform creation/editing.

Db2 entitlements

Db2 Warehouse

Warm Standby: allocate 1 VPC (counted towards overall VPC entitlement through the Maximo Application Suite license).

Db2 Standard Edition

The Supporting Program may use a maximum of 16 processor cores and 128 GB of memory on each physical or virtual server.

Non-Production instances

Licensee may only use the following components as part of Licensee's internal development and test environment for internal non-production activities ("Non-Production").

Non-Production components

IBM provides non-production components exclusively through a continuous delivery (CD) stream as part of the Maximo Application Suite Program. This stream introduces new features, developed for the next release, allowing early access for non-production environment use. This stream is offered alongside and in parallel with our normal maintained streams.

For a list of non-production components or features being introduced in a monthly CD stream, see <https://ibm.biz/MASFC-whats-new>

Additional terms for Service Providers

Additional definitions

1. Licensee is a "Service Provider" if Licensee uses the Programs to provide services in support of or to benefit one or more of Licensee's Client(s).
2. "Client" means a third party that engages Licensee to provide services using the Program; a Client cannot, by more than 50%, own, be owned by, or be under common ownership with, Licensee. A Client cannot acquire Authorized User or Limited Use User entitlements to the Program.

License scope and restrictions

1. The Program is for use within Licensee's Enterprise only and may not be resold, rented, leased, or transferred to third parties. Any attempt to do so in violation of these provisions is void. In addition, Licensee may not use the Program to provide any timesharing, commercial hosting or service bureau usage. The foregoing notwithstanding:
 - Licensee may use the Programs to process data on behalf of Licensee's Client(s).
 - Licensee may use the Programs to support more than one Client at a time, up to the total number of users set forth in the applicable PoE.
 - While Client(s) may use or benefit from the use of the Programs provided to Licensee under this Agreement, no Client is a party to this Agreement.
2. Licensee will maintain and provide support to Licensee's Client(s) for any Program used as part of services Licensee provides to Licensee's Client(s), and ensure that each Client is notified to look to Licensee, not IBM, for any such maintenance and support.
3. Licensee is responsible for Licensee's Client's compliance with the terms of this Agreement. In this regard, Licensee must enter into Licensee's own written agreement with each Client that:
 - Prohibits any action by the Client that is not in compliance with the terms of this Agreement, including, without limitation, reverse assembly, reverse compilation, reverse engineering or other translation of the Program;
 - Notifies the Client that the Program is copyrighted and licensed (not sold) and that at no time is the license to the Program transferred to the Client;
 - Notifies the Client that "SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM MAKES NO WARRANTIES OR CONDITIONS, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THE USE OF THE PROGRAM, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT";

- Notifies the Client that IBM is not liable for any lost profits, lost savings, or any incidental or other economic consequential damages resulting from the use of the Program, regardless of whether the Licensee, the Client or IBM have been advised of the possibility of such damages.
- Prohibits the Client's use of the Program except when used as part of services Licensee provides to the Client, and obligates the Client to destroy or return all copies of the Program to Licensee when Licensee's services to the Client end.

4. Licensee will:

- Immediately notify IBM if Licensee becomes aware of any act or failure to act by a Client which, if committed by Licensee, would immediately or with the passage of time become a breach of the Agreement;
- Provide all information and assistance requested by IBM in preventing, investigating, and mitigating any such breach;
- Pay all license fees owed and indemnify IBM against any and all claims, damages, losses, costs and expenses suffered or incurred by IBM arising out of any such breach or threatened breach or of a Client's use of or problems with the Programs.

5. General:

- Licensee agrees to indemnify and defend IBM against any claims or lawsuits, including reasonable attorney's fees, that arise or result, directly or indirectly, from providing services using the Program to Licensee's Client or such Client's accessing or use of the Program.
- Licensee is solely responsible for the payment of any duty, tax, levy, or fee that any authority may impose that results from Licensee providing services using the Program.

Cloned or custom applications

- Maximo Manage user-defined applications can either be cloned applications or custom applications. Cloned applications are based on a copy of an existing Maximo Manage application, which is also called the originating application. Custom applications are built from scratch, have no originating application and are built based on either a custom object or an existing Maximo object.
- Users being granted access to a cloned application must have at least the same level of access in the originating application.
- Limited Use users may be granted access to custom applications when the main object is a custom object or an object the user is otherwise entitled to access. The custom application must be a part of one of the three modules selected by the user.
- If the user needs to access custom or cloned applications that are based on an application or object not within the three modules selected for Manage Limited, then Manage Base or Industry Solution/Add-on Premium is required.
- Custom or cloned applications shall not be utilized to circumvent licensing or grant users functionality or access they would otherwise not be entitled to access.

Back-up and restore

- Governed by IPLA level language.
- https://www.ibm.com/about/software-licensing/assets/guides_pdf/Backup.pdf

Standby scenarios for Production backup:

- Cold (entitlement is not required)
- Warm (depends on access to live production data)
- Hot (entitlement IS required)

Licensing requirements for indirect access scenarios

This section covers scenarios related to license requirements when integrating with external systems, external custom applications, websites, or 3rd party mobile solutions.

Licensee must obtain separate, dedicated entitlements for Concurrent Users or Authorized Users given access to the Program in any manner directly or indirectly (for example: via a multiplexing program, device, or application server) through any means.

License Circumvention Licensee will not attempt to aggregate users or otherwise circumvent IBM's licensing restrictions via technical means, including without limitation the use of any interface between the Program and other software that performs functionality substantially similar to that contained in any IBM program then offered by IBM when Licensee acquires the Program.

Third-party mobile licensing

For 3rd party mobile access and email approval, capabilities licensing is required.

Licensing is Authorized and Concurrent user. AppPoints based on Tier level.

Integrations and external applications in Maximo Manage

- There are instances when an organization licensed to use Maximo Manage will elect to use an external application, develop custom applications, websites, or use 3rd party mobile applications as a way to enable their users to interact with records, data, and information that is stored in Maximo Manage.
- If the applications are integrated with Maximo Manage, (importing/exporting data), the distinction of which system is the master of that data also determines if a license is required to create or update those records.
- If data that is considered to be created and housed as the 'source of truth' in Maximo Manage, then a license is required to insert or update that data. If a third-party system is considered the master of the data, any replication of that data into Maximo Manage for use within other processes does not require a license.
- Examples of when each user would require a Maximo Manage License include the use of a 3rd party application to update Work Orders; create or update Asset records; change status of a record, or any supply chain related functions if they are primarily initiated and managed in Maximo.
- Customers would need to purchase enough AppPoints to cover these users but is not registered as part of the Suite License server.

When is integration allowed without any additional end-user license required?

- Maximo Manage includes an extensive set of integration capabilities, referred to as the Maximo integration framework (MIF). There are also specific adapters available for integration with SAP or Oracle ERP systems. The MIF and adapters allow companies to share and exchange information between Maximo Manage and external systems in real time or in batch mode. The MIF is run using the Maximo Manage admin (max integration admin) user which does not require a specific license.
- Examples of a third-party system that are considered the master data repository and do not require a Maximo Manage license (except for established Maximo Application Suite connectors) would include things like a HR system with personnel information, pay rates and certification data that is synched into a labor record in Maximo Manage on a regular basis; an ERP system that manages all the Item/Inventory and purchasing transactions that pushes updates like new item records or receipts into Manage; or a Scada or other system with operational asset data being pushed into an asset meter reading record. In all of these cases the external system is the system of record and the synchronization of that data into Maximo Manage does not require an end user license.
- Maximo Manage Self Service applications/use cases are exempt from licensing.
 - For example, a college creates a web app that enables students to create service requests to report problems/incidents in the dorms. This is the same as users of the Manage self-service service request (create/view) applications, for which users do not require a license entitlement.
- Extracting Maximo Manage data is for reporting purposes only.

- A user is emailed a copy of a report from Maximo Manage.
- Maximo Manage data is copied into data warehouse (no updates back into Maximo Manage).
- In these cases, there is no license required for users receiving data in this way.

Appendix: Modules & applications

- [Manage Modules](#)
- [Manage Industry Solution and Add-on Modules](#)

Maximo Manage modules

The following table lists the modules and applications that are included in Maximo Manage. The applications are organized and grouped by module.

Table 14. Modules and applications in Maximo Manage

Module and sub-module	Application
Administration	Sets
	Organizations
	Calendars
	Bulletin Board
	Communication Templates
	Report Administration
	Conditional Expression Manager
	Classifications
	Work View
	Service Addresses
	Map Manager
	Record Release
	Time Zone Rules
	Scheduler Administration
	Push Notification Administration
	Synchronization ArcGIS
Resources	Labor
	Qualifications
	People
	Person Group
	Crafts
	Crew Types
	Crews
KPI	KPI Manager
	KPI Templates

Table 14. Modules and applications in Maximo Manage (continued)

Module and sub-module	Application
Analytics	Report Viewer
	KPI Viewer
	Design Data Sets
	Business Analyst
Assets	Assets
	Asset Templates
	Locations
	Features
	Meters
	Relationships
	Meter Groups
	Condition Monitoring
	Failure Codes
	Reliability Strategies
Role Based Applications	Asset Manager
Building Information Models	BIM Projects
Contracts	Purchase Contracts
	Lease Rental Contracts
	Labor Rate Contracts
	Master Contracts
	Warranty Contracts
	Terms and Conditions
Financial	Currency Codes
	Exchange Rates
	Chart of Accounts
	Cost Management
	Budget Monitoring
Integration	Object Structures
	Publish Channels
	Invocation Channels
	Enterprise Services
	Web Services Library
	End Points
	External Systems

Table 14. Modules and applications in Maximo Manage (continued)

Module and sub-module	Application
	Logical Management Operations
	Integration Modules
	Launch in Context
	Message Tracking
	Message Reprocessing
Interactions	Create Interaction
	Interactions
	OSLC Providers
	OSLC Resources
	JSON Resources
	JSON Mapping
	Notifications
Inventory	Item Master
	Service Items
	Tools
	Stocked Tools
	Inventory
	Inventory Usage
	Shipment Receiving
	Condition Codes
	Storerooms
	Collections
	Count Book
Role Based Application	Inventory Counting
	Inventory Receiving
	Issues and Transfers
Planning and Scheduling	Graphical Assignment
	Graphical Work Week
	Graphical Scheduling
	Graphical Appointment Book
	Graphical Crew Management
	Graphical Resource View
	Scheduling Dashboard
	Dispatching Dashboard

Table 14. Modules and applications in Maximo Manage (continued)

Module and sub-module	Application
Planning	Job Plans
	Routes
	Manage Inspections Forms
Safety	Hazards
	Precautions
	Lock Out / Tag Out
	Safety Plans
Preventive Maintenance	Preventive Maintenance
	Master PM
Purchasing	Purchase Requisitions
	Purchase Orders
	Receiving
	Shipment Receiving
	Invoices
	Request for Quotations
	Companies
	Company Master
	Terms and Conditions
Security	Security Groups
	Users
Self Service	
Desktop Requisitions	Create Requisitions
	View Requisitions
	View Templates
	View Drafts
Service Requests	Create Service Request
	View Service Request
	Search Solutions
Role Based Applications	Service Request
Service Level	Service Level Agreements
	Service Groups
Service Desk	Activities and Tasks
	Service Requests
	Solutions

Table 14. Modules and applications in Maximo Manage (continued)

Module and sub-module	Application
	Ticket Templates
System Configuration	
Platform Configuration	System Properties
	Logging
	Domains
	Database Configuration
	Application Designer
	Communication Templates
	Actions
	Roles
	Escalations
	Workflow Designer
	Workflow Administration
	Cron Task Setup
	E-mail Listeners
	Web Services Library
	Launch in Context
	Instant Messaging Configuration
	Automation Scripts
Migration	Object Structures
	Migration Manager
	Migration Groups
	Migration Collections
	Maximo Management Interface
Task Management	Activities and Tasks
Work Orders	Work Order Tracking
	Labor Reporting
	Quick Reporting
	Activities and Tasks
	Assignment Manager
	Service Requests
Work Centers	Conduct an Inspection
Role Based Applications	Inspections
	Technician

Table 14. Modules and applications in Maximo Manage (continued)

Module and sub-module	Application
	Work Approval
	Service Request

Maximo Manage industry solution and Add-on modules

The following table lists the modules for Maximo Manage industry solutions and add-on modules.

Table 15. Industry solutions and add-on modules in Maximo Manage. Industry solutions and add-on modules in Maximo Manage

Industry /Add-on	Modules
Asset Configuration Manager (ACM)	Asset Configuration Manager
	Flight Logs
Service Provider (SP)	Service Provider (SP)
Health, Safety and Environment (HSE)	Change (HSE)
	Operations (HSE)
	Safety and Quality Management (HSE)
Oil & Gas	Change (Oil)
	Operations (Oil)
	Safety and Quality Management (Oil)
Transportation	Data Import (Tr)
	Motor Pool (Tr)
	Warranties (Tr)
Nuclear	Configuration Change (Nuc)
	Operation Management (Nuc)
	Permits (Nuc)
	Condition Report (Nuc)
Aviation	Aircraft and Equipment
	Configuration Management
	Consists
	Data Import
	Flight Log
	Exchange Orders
	Motor Pool
	Operations
	Safety and Quality Management
Warranties	

Language and locale support

IBM Maximo Application Suite supports the use of preferred language and locale for the Maximo Application Suite user interfaces. The preferences that are applied override the language and locale settings for the browser that is used to access Maximo Application Suite.

For example, if a user's preferred language is set to German and Maximo Application Suite is accessed from a browser set to English, the user interface is displayed in German.

Note: The following information is applicable to locale management for customer-managed Maximo Application Suite.

The user interface language

The languages that are displayed in the user interface are controlled by your browser settings, your preferred language setting, and the application base language.

Browser language setting

When users access the Maximo Application Suite user interface, the displayed language is set by their browser language settings, if the language is available. If the language that is selected in the browser is not available, the application base language is displayed.

Preferred language setting

Users can overrule the browser setting by selecting a preferred language, which is then displayed in supported applications no matter what the current browser language setting is. If the application does not support the preferred language setting, then the displayed language is controlled by the browser language setting.

Application-specific support

For a breakdown of the application-specific language support, see the following list.

Maximo Application Suite home page and Administration pages

Base language is English (en_US). Preferred language is supported.

Maximo Collaborate

Base language is English (en_US). Preferred language is supported.

Maximo Health

Base language is [configurable](#). Preferred language is supported.

Maximo Manage

Base language is [configurable](#). Preferred language is not supported.

Maximo Predict

Base language is English (en_US). Preferred language is not supported.

Maximo Monitor

Base language is English (en_US). Preferred language is supported in 8.10 and later versions.

Maximo Visual Inspection

Base language is English (en_US). Preferred language is not supported.

Industry solutions

Base language is English (en_US). Preferred language is not supported.

The following languages and locale are supported in Maximo Application Suite:

- Arabic (ar)
Supported only in Maximo Manage and IBM Maximo Real Estate and Facilities.
- Brazilian Portuguese (pt-BR)
- Croatian (hr)
- Czech (cs)
- Danish (da)
- Dutch (nl)

- English (en_US)
- English (en_UK)
- Finnish (fi)
- French (fr)
- German (de)
- Hebrew (he_IS)

Supported only in Maximo Manage and IBM Maximo Real Estate and Facilities.

- Hungarian (hu)
- Italian (it)
- Japanese (ja)
- Korean (ko)
- Norwegian (no)
- Polish (pl)
- Simplified Chinese (zh_CN)
- Slovak (sk)
- Slovenian (sl)
- Spanish (es)
- Swedish (sv)
- Traditional Chinese (zh_TW)
- Turkish (tr)

Related tasks

[Setting language and time zone preferences for users](#)

[Managing user profile](#)

Starting in IBM Maximo Application Suite 9.1, edit your user information, change password, or update locale and region settings.

Accessibility

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology successfully.

Overview

IBM Maximo Application Suite includes the following accessibility features:

- In IBM Maximo Application Suite 9.1, 9.0 and 8.11, select areas of IBM Maximo Manage are accessible. For more information, see the Accessibility Conformance Reports.
 - [Accessibility Conformance Reports for IBM Maximo Application Suite 9.1](#)
 - [Accessibility Conformance Reports for IBM Maximo Application Suite 9.0](#)
 - [Accessibility Conformance Reports for IBM Maximo Application Suite 8.11](#)
- The IBM Maximo Application Suite online product documentation is available in [IBM Documentation](#), which is viewable in a standard web browser.

IBM Maximo Application Suite 8.10 and earlier does not conform to accessibility standards. For more information, see

- [Accessibility Conformance Reports for IBM Maximo Application Suite 8.9 and 8.10](#)
- [Accessibility Conformance Reports for IBM Maximo Application Suite 8.8](#)

IBM and accessibility

For more information about IBM's commitment to accessibility and to view Accessibility Conformance Reports for IBM products, see [IBM Accessibility](#).

Federal requirements

If you are using Maximo Application Suite as part of a federal program, you need to comply with various government regulations, including the Federal Information Security Modernization Act (FISMA).

In Maximo Application Suite features are provided that can be used in preparation for US Federal compliance assessments, including FISMA compliance. In Maximo Application Suite 9.1, the features meet FISMA high level compliance. In Maximo Application Suite 8.11 and 9.0, the features meet FISMA moderate level compliance.

Federal Information Processing Standard (FIPS) 140

Note: Starting in Maximo Application Suite 8.11, support for FIPS is applicable only to Maximo Application Suite core and the IBM Maximo Manage application.

Note: Starting in Maximo Application Suite 9.1, support for FIPS is available for IBM Maximo Real Estate and Facilities.

Cryptographic modules, data in motion, and data at rest that are used in the following applications, dependencies, and industry solutions support the Federal Information Processing Standard (FIPS) 140:

- Maximo Application Suite core and its dependencies such as IBM Suite License Service 3.7.0.
- Maximo Manage including add-ons, industry solutions, and connectors.
- Maximo Real Estate and Facilities.

For more information, see

- [“MongoDB” on page 21](#)
- [“Creating the Db2 instance by using the stand-alone Db2U operator” on page 10](#)
- [“Suite License Service” on page 7](#)
- [“Apache Kafka” on page 23](#)

Audit logging

You can forward all logging from Red Hat OpenShift into an external system so that logs can be aggregated and securely stored. The Maximo Application Suite logs can be aggregated from a Red Hat OpenShift cluster to third-party systems. For more information, see [“Audit logging in Maximo Application Suite” on page 817](#)

Internet Protocol version 6

Starting in Maximo Application Suite 8.11.6 and IBM Maximo Manage 8.7.4, Internet Protocol version 6 (IPv6) is supported in Maximo Application Suite core and the Maximo Manage application.

For more information about supported versions, see [Software Product Compatibility Reports \(SPCR\)](#).

Starting in Maximo Application Suite 9.1, support for IPv6 is available for IBM Maximo Real Estate and Facilities.

Section 508 accessibility

IBM is committed to accessibility. Accessibility Compliance Reports contain details on accessibility compliance with standards, including the Worldwide Consortium Web Content Accessibility Guidelines, European Standard EN 301 349, and US Section 508.

For more information, see [IBM Accessibility](#).

Software Bill of Materials (SBOM)

To request a Software Bill of Materials (SBOM) for Maximo Application Suite, contact your IBM representative.

Secure Software Development Framework (SSDF)

Maximo Application Suite follows the SPbD@IBM discipline. For more information, see [IBM Security and Privacy by Design \(SPbD@IBM\)](#).

FISMA controls documentation

Further information about the use of Maximo Application Suite in Federal Information Security Management Act (FISMA) certified environments is available on request. Contact your IBM representative.

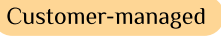
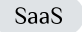
IBM Maximo operator catalog content verification

Starting in Maximo Application Suite 9.1, to ensure integrity and authenticity of IBM content, all IBM content is digitally signed. The signatures can be verified using the Maximo Application Suite public key.

For more information, see [Verifying IBM Maximo operator catalog content](#).

Documentation conventions

The following table describes the conventions that are used in this documentation.

Icon	Description
	Identifies IBM Maximo Application Suite that is offered as a customer-managed product on Red Hat OpenShift.
	Identifies IBM Maximo Application Suite as a Service that is offered as a cloud-based product on Amazon Web Services (AWS).
No icon	Applicable to both Maximo Application Suite and Maximo Application Suite as a Service.

Getting started

IBM Maximo Application Suite places a collection of powerful asset and data management tools at your fingertips.

Step 1: Log in to Maximo Application Suite

Log in to the Maximo Application Suite web console.

- Customer managed

Log in using the URL and credentials that your Maximo Application Suite administrator provides. Your Maximo Application Suite URL is of the format:

`https://<workspace_id>.home.<mas_domain>`

- IBM managed

For IBM Maximo Application Suite Dedicated, log in using the URL and credentials that are provided in the Welcome letter. Your Maximo Application Suite Dedicated URL is of the format:`https://home.<organization>.suite.maximo.com`

- SaaS

Log in using the URL and credentials that are provided in the Welcome letter. Your Maximo Application Suite as a Service URL is of the format: `https://home.<organization>.suite.maximo.com`

Step 2: Optionally set the user interface language and region

You can set the preferred locale and time zone for supported Maximo Application Suite applications. The preferences that you set here override the default settings in the browser that you use to access the suite.

For example, if your preferred language is set to German and you access Maximo Application Suite from a browser set to English, the user interface is displayed in German. For more information, see [Optionally set user preferred Maximo Application Suite language and region](#).

Tip: You can update the language and region settings at any time.

Step 3: Get started

Your next action depends on your role, which is defined by your entitlements and access settings. For more information, see [“Example of user roles configurations”](#) on page 800.

Getting started as an application user

Get started as an application user by selecting the applications that you have access to from the suite navigator.

In Maximo Application Suite 9.1, you can navigate to the suite applications that you have access to by using the side navigation menu. This navigation menu replaces the **Application switcher** and the administration icon on the header menu that is in earlier versions.

In Maximo Application Suite 9.0 and earlier, you can use the application switcher in the main toolbar to switch to other applications, or return to the home page. The following video shows you how you can navigate in Maximo Application Suite 9.0 and earlier.

Tip: If your user ID does not have access to the application that you want to use, contact your administrator to get access.

- [Getting started with IBM Maximo Collaborate](#)
- [Getting started with IBM Maximo Health, IBM Maximo Predict, or IBM Maximo Health and Predict - Utilities](#)
- [Getting started with IBM Maximo Manage](#)
- [Getting started with IBM Maximo Monitor](#)
- [Getting started with IBM Maximo Visual Inspection](#)

Customer-managed

Getting started as an application administrator

The application administrator manages Maximo Application Suite applications, including adding users, deploying and activating applications, and more.

Starting in Maximo Application Suite 9.1, administrators have access to the **Suite > Administration** page on the side navigation menu.

In Maximo Application Suite 9.0 and earlier, administrators can access the **Suite administration** page by selecting the admin icon in the menu bar.

Typical tasks:

- [Administer applications](#)
- [Administer users](#)

Getting started as an application suite administrator

The application suite administrator manages the underlying infrastructure resources, such as Red Hat OpenShift, and Db2.

Typical tasks:

- [Administer suite configurations](#)
- [Administer license and AppPoints](#)

Getting started as a SaaS suite administrator

Get started in Maximo Application Suite as a Service applications by adding users and specifying their entitlements, requesting server authentication and user synchronization, and monitoring application usage.

To access the **Administration** page, in the side navigation, select the **Suite** application.

Related concepts

[Monitoring AppPoint usage for Maximo Application Suite as a Service](#)

As a SaaS suite administrator, you can use the Usage dashboard to monitor the peak usage trends of how AppPoints are consumed in each application, such as when users use the Maximo Manage application. You can manage the daily usage of AppPoints for each application and ensure that your organization stays within their AppPoint capacity.

[Maximo Application Suite as a Service overview](#)

Related tasks

[Managing SaaS users](#)

You can create and manage SaaS users at the suite-level and also give users access and entitlement to the applications that they need. After the user is created, you can give users specific permissions to access individual applications. To connect an external server for authentication or user sync registry, you can submit a request with IBM Support.

Planning

IBM Maximo Application Suite uses Red Hat OpenShift run anywhere model. You can deploy on premises, or on a supported hyperscaler.

Planning for IBM Maximo Application Suite standard installation with CLI

Before you install IBM Maximo Application Suite using the Maximo Application Suite CLI utility, you must consider the dependencies that must be met for installation.

Prerequisites

You can use IBM's container image to install Maximo Application Suite. The container image provides an out of the box environment for managing Maximo Application Suite on Red Hat OpenShift, with numerous dependencies preinstalled.

In a scenario where you need to install the CLI utility, which is an open source tool, on your local system, you must configure the following software.

- Bash (v4)
- Red Hat OpenShift client
- IBM Cloud client with container plug-in enabled

Note: IBM Cloud is not required if you are deploying Maximo Application Suite on an organization's Red Hat OpenShift cluster.

- Ansible
- Python
- Network access to the Red Hat OpenShift cluster

For more information about installing the utility, see [Maximo Application Suite CLI Utility](#).

Related tasks

Standard installation with [IBM Maximo Application Suite CLI](#)

You can install the IBM Maximo Application Suite by using a command-line interface (CLI) utility.

Prerequisites for installing

Before you begin, ensure that your environment meets the prerequisites by downloading, and installing the software and interfaces that you use to install IBM Maximo Application Suite.

Prerequisites

You must have access to the following components:

- A private image registry setup and running in the restricted network, and secured with certificates. Configure one of the following options:
 - A bastion host with access to product images on the internet and the restricted network and has support for running docker containers. The docker image that contains the IBM Maximo Application Suite command line utility on the bastion host.
 - A host outside the restricted network with access to product images on the internet and with support for running docker containers. The docker image that contains the Maximo Application Suite command line utility on the host. Portable disk space sufficient to store the required images. A host inside the restricted network with support for running docker containers and can access the images downloaded to the portable disk space.
- Red Hat OpenShift cluster setup as an air gap cluster for disconnected installation.

Ensure you use the Maximo Application Suite sizing calculator to estimate your Red Hat OpenShift Worker Node configuration, storage, and memory requirements.

For more information, see [“Requirements and capacity planning” on page 178](#).

- IBM entitlement key.
- IBM Maximo Application Suite license file.

Red Hat OpenShift cluster

A Red Hat OpenShift cluster must be configured in disconnected environment or restricted network for IBM Maximo Application Suite installation.

To install the Red Hat OpenShift in your environment, see [Red Hat OpenShift Container Platform installation overview](#) and [Installing a Red Hat OpenShift cluster on any platform](#).

To install the Red Hat OpenShift in a disconnected or restricted network, see [Mirroring images for a disconnected installation](#).

To convert a connected Red Hat OpenShift cluster to a disconnected cluster, see [Converting a connected cluster to a disconnected cluster](#).

IBM entitlement key

Access the [Container Software Library](#) by using your IBM id to obtain your entitlement key.

IBM Maximo Application Suite license file

Access the [IBM License Key Center](#) to get the license key of IBM Maximo Application Suite. From **Get Keys > IBM App Point Suites**, select IBM MAXIMO APPLICATION SUITE AppPOINT LIC and enter the following information on the next page:

Field	Content
Number of Keys	Use this field to decide the number of App Points that are assigned in the license file.
Host ID Type	Enter Ethernet Address.
Host ID	Enter a host ID. The host ID can be any 12-digit hexadecimal string.
Hostname	Enter a hostname. The hostname is used to associate the license file to the Red Hat OpenShift Container Platform cluster instance.
Port	Enter the port number 27000.

Tip: You can retain the default values for other fields.

Generate and download the license file to your home directory in a file that is named `entitlement.lic`.

For more information about accessing the IBM License Key Center, see [Getting started guide](#).

Operator catalog

Select the catalog source for your installation. For more information, see [Catalog selection](#) or contact IBM Support for guidance.

Fusion Data Foundation

The Maximo Application Suite license includes IBM Storage Fusion Advanced. The Maximo Application Suite catalog has the Fusion operator which allows you to install Fusion Data Foundation.

1. Check installed catalogs.

```
oc get catalogsources.operators.coreos.com -n openshift-marketplace
```

2. Install the Maximo catalog.

```
export ROLE_NAME=ibm_catalogs
ansible-playbook ibm.mas_devops.run_role
```

3. Validate the updated catalogs.

```
oc get catalogsources.operators.coreos.com -n openshift-marketplace
```

4. Update the pull-secret.

```
export IBM_ENTITLEMENT_KEY=<your entitlement key>
export CNV_ENT_KEY=$(echo -n "cp:${IBM_ENTITLEMENT_KEY}" | openssl base64 -A)
oc extract secret/pull-secret -n openshift-config --confirm
sed -i 's/{"auths":{}/&"cp.icr.io":{"auth": "'"$CNV_ENT_KEY"'',"email": "not-used"}/' .dockerconfigjson
```

5. Install IBM Storage Fusion. For more information, see [Data Foundation](#).

Related tasks

[IBM Maximo Application Suite installation in disconnected environments](#)

You can install IBM Maximo Application Suite in an air gap environment, which is also known as disconnected, offline, or restricted network.

[Upgrading Maximo Application Suite by using static catalog](#)

Related information

[Sizing guidance](#)

Planning to install in disconnected environment

You can install IBM Maximo Application Suite in a disconnected environment such as air gap with a private mirror registry by using one of the static catalogs and automatic approval strategy. The installation is also known as offline or restricted network installation. In disconnected installations, you can use the policies that are configured by your organization to download images from the local docker registry.

For more information, see [Catalog selection](#).

Overview

If your cluster is not connected to the internet, install IBM Maximo Application Suite in your cluster by using the IBM Maximo Application Suite command line interface (CLI) utility on a bastion host.

For more information, see [“Standard installation with IBM Maximo Application Suite CLI” on page 218](#).

A bastion server is a device that has access to both the public internet and the local intranet where a local registry and Red Hat OpenShift Container Platform clusters are configured. You can store the product code and replicate your images through the bastion server directly to the local intranet registry in the local air gap network.

Installation overview

From a high level, air gap installation consists of the following steps.

1. Setup an air gap Red Hat OpenShift cluster by including private image registry and bastion host.
2. Select an IBM operator catalog as a source for installation.
3. Mirror product images based on your installation scenario.
4. Configure Red Hat OpenShift to use your private registry for Maximo Application Suite.
5. Install Maximo Application Suite.

Supported product applications for air gap installation

Installation in an air gap environment is supported from IBM Maximo Application Suite 8.8 and later for some product catalogs or applications. The following list of applications shows the supported and unsupported versions:

Application	First version to support air gap installation
IBM Maximo Application Suite Core	Supported from 8.8.0
IBM Maximo Assist or IBM Maximo Collaborate*	Supported from 8.7.1
IBM Maximo Health and Predict - Utilities	Supported from 8.6.1
IoT	Supported from 8.5.1
IBM Maximo Manage	Supported from 8.4.0
IBM Maximo Monitor	Supported from 8.10.0
IBM Maximo Optimizer	Supported from 8.2.0

Application	First version to support air gap installation
IBM Maximo Predict	Supported from 8.8.1
IBM Maximo Visual Inspection	Supported from 8.8.1

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

For more information, see [the list of static operator catalog versions](#).

Related tasks

[Upgrading Maximo Application Suite by using static catalog](#)

Foundation service

Starting in Maximo Application Suite 9.1, the foundation service is available and must be installed to use capabilities such as unified authorization services and the common navigation for all suite applications.

The foundation service consists of selected functions that unify setting up, accessing, and viewing capabilities that were previously separate and specific to each suite application. By consolidating these common administrative tasks and operations, administrators can centrally administer users, security, and data and also configure or customize applications. The foundation service is a subset of Maximo Manage that provides the back-end operations and data schema and excludes all other applications in Maximo Manage.

When you install or upgrade to Maximo Application Suite 9.1, you also install the foundation service bundle to use the following capabilities.

Uniform authorization

The security profile authorization is based on role-based access control (RBAC) and row-level access control (RLAC) for all Maximo Application Suite users. Roles are sets of permissions that you can use to grant or restrict access to specific operations. You can use roles to manage permissions for groups of users and applications.

Common entitlement calculator

The unified entitlement calculator calculates the user entitlement based on access to the suite applications, such as Maximo Application Suite, Maximo Monitor, and Maximo Visual Inspection. For more information, see [“Access entitlement” on page 782](#).

Common user management

When you create users, you can centrally manage their identifying information and authorizations and assign security privileges by using security groups. The foundation service also For more information, see [“Administering users and user access in Maximo Application Suite in 9.1” on page 781](#).

Common navigation

To unify the navigation across suite applications, you can switch between suite applications that you have access to by using the side navigation menu.

Common application configurator and repository

By using Maximo Application Framework Configuration, you can configure and customize applications, such as add queries to service requests or add custom fields to work orders. For more information, see [Application framework configuration overview](#) and [Getting started for Maximo application framework configuration application users](#).

Manage foundation service bundle

The foundation service consists of selected functions and a database that enables administrators to centrally manage and configure users, security, and data.

Foundation service bundle

The foundation service bundle is a license restricted version of Maximo Manage that includes only the following capabilities.

- User and security groups
- Application configuration
- Mobile configuration
- API keys
- APIs to access the security profile and navigation

Foundation database

The database for the foundation service bundle is the same as the Maximo Manage database.

Installing the foundation service

You install the foundation service by using the [command-line interface \(CLI\) utility](#). When you install the foundation service in Maximo Application Suite 9.1, administrators have a choice of the following options:

- Install only the foundation service
- Install Maximo Manage

If you are installing Maximo Application Suite and want to deploy suite applications that are not Maximo Manage, then the foundation service bundle is included during the installation. The Manage namespace is created. In the user interface, the Manage tile is added to the catalog to indicate that only the foundation service is installed. The Maximo Manage component is not installed.

If you are installing Maximo Application Suite and want to deploy the Maximo Manage suite application to access full Manage functionality, deployers use the **cron**, **mea**, **report**, and **ui** server bundles for Maximo Manage. You use these bundles for optimal load balancing. In this case, the foundation service is also added as an additional bundle. Alternatively, if you choose to deploy only the **all** server bundle, which combines these bundle types, then the foundation service is included in the **all** bundle.

In most deployment scenarios, the foundation service bundle is also the target bundle for user synchronization. The `isUserSyncTarget` flag that is in the CR of the server bundle is deprecated. If you set a value of true or false, the operator ignores the value and uses the internal rules that are defined for `isUserSyncTarget`.

The following internal rules are applied with the Manage operator to determine whether to include the foundation service bundle as an additional deployment.

If **all** bundle is selected

If the administrator selects the **all** bundle, the foundation service bundle is not deployed. The **all** bundle handles requests from the user interface. User synchronization or `isUserSyncTarget` is also targeted by the **all** bundle.

If **all** bundle and foundation service bundle are selected

The administrator might choose to override the Manage operator's logic and define a foundation bundle in the CR spec to coexist with the **all** bundle. In this case, the Manage operator applies the request and deploys the foundation service bundle along with the **all** bundle. The user synchronization is targeted for the foundation service bundle.

If **ui** bundle and any combination of **mea**, **reports**, **cron**, or **jms** bundles are selected

If the administrator selects the **ui** bundle with any combination of **cron**, **mea**, **report**, or **jms** bundles, a foundation service bundle is automatically deployed by the Manage operator. User synchronization is targeted for the foundation service bundle.

If Manage base is not selected

If Manage is not deployed but other suite applications, such as Maximo Monitor or Maximo Visual Inspection, are deployed, only the foundation service bundle is deployed automatically by the Manage operator. User synchronization is targeted for the foundation service bundle.

When you upgrade to Maximo Application Suite 9.1, either Maximo Manage or the foundation service is required. The upgrade path to Maximo Application Suite 9.1 depends on your current installation and configuration. For more information, see [“Upgrading to Maximo Application Suite 9.1” on page 482](#)

Related concepts

[Upgrading to Maximo Application Suite 9.1](#)

When you upgrade to Maximo Application Suite 9.1, either Maximo Manage or the foundation service is required. The upgrade path to Maximo Application Suite 9.1 depends on your current installation and configuration.

Related tasks

[Standard installation with IBM Maximo Application Suite CLI](#)

You can install the IBM Maximo Application Suite by using a command-line interface (CLI) utility.

Planning to install on Amazon Web Services

Before you can install IBM Maximo Application Suite on Amazon Web Services, you must consider your installation preferences, such as whether you want to create a Red Hat OpenShift cluster or reuse an existing cluster.

Related concepts

[Requirements and capacity planning](#)

Maximo Application Suite on Amazon Web Services overview

IBM Maximo Application Suite can be installed in the Amazon Web Services (AWS) cloud. You subscribe to this product, specify the installation parameters, and install Maximo Application Suite.

In Amazon Web Services Marketplace, Maximo Application Suite is available either as a bring-your-own-license (BYOL) product or as a paid offering.

• BYOL product

The BYOL product can be installed from the AWS Marketplace by using any one of the following CloudFormation templates:

- An existing Red Hat OpenShift cluster
- In Maximo Application Suite 8.8 and later, a new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- In Maximo Application Suite 8.8 and later, a new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

After you receive your Maximo Application Suite license, you can use it to install Maximo Application Suite in your Amazon Web Services cloud account.

Note: Starting in 8.11, to create prerequisite network resources such as VPC, private, and public subnet required in UPI mode of Maximo Application Suite deployment on AWS, unzip the [pre-req-vpc-subnets.zip](#) file and follow the steps given in the Readme file.

Note that a Terraform client and an AWS client must be installed on the machine from where the script is being executed.

• Client-managed product

In Maximo Application Suite 8.8 and later, you can install a new Red Hat OpenShift cluster as part of the Maximo Application Suite deployment from the Amazon Web Services Marketplace.

Maximo Application Suite client-managed with Red Hat OpenShift entitlement

This product includes the subscription for the Red Hat OpenShift cluster that is deployed during the installation process.

The following CloudFormation templates are available as fulfillment options:

- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

Maximo Application Suite client-managed without Red Hat OpenShift entitlement

This product does not include the subscription for the Red Hat OpenShift cluster that is deployed during the installation process or the existing one provided by the user. You must have your own Red Hat OpenShift subscription to deploy Maximo Application Suite by using this product.

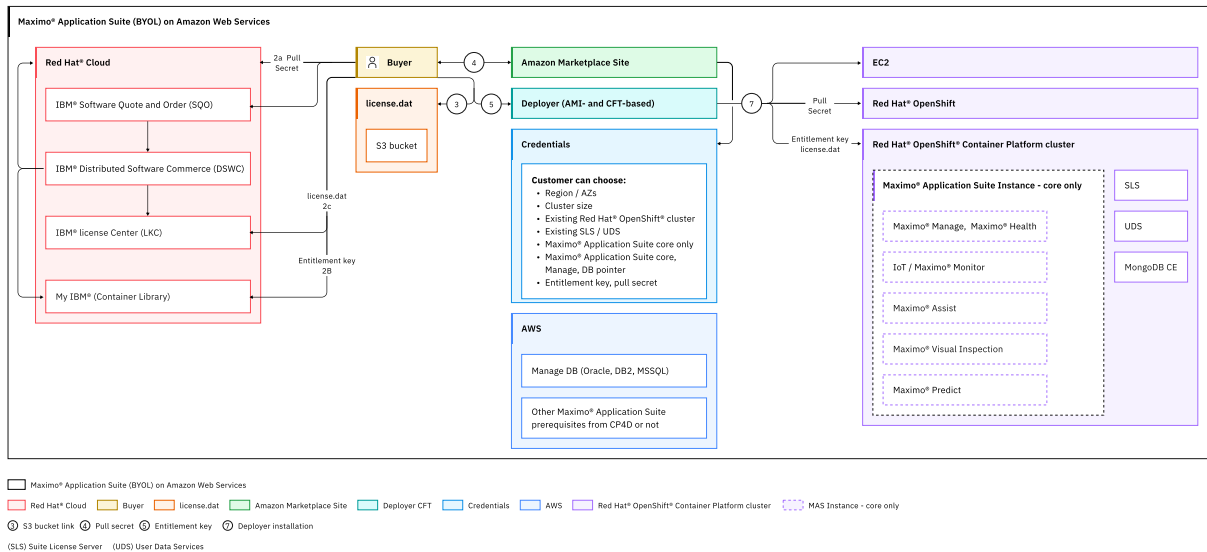
The following CloudFormation templates are available as fulfillment options:

- An existing Red Hat OpenShift cluster
- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

Purchase AppPoints on the AWS marketplace page. After the purchase, these AppPoints are available as licenses, which you can use to install Maximo Application Suite in your Amazon Web Services cloud account.

High-level architecture

The following diagram shows the high-level architecture for Maximo Application Suite BYOL on Amazon Web Services:



1. You buy IBM Maximo Application Suite from the Amazon Web Services Marketplace. Systems that are required to install Maximo Application Suite, such as Red Hat OpenShift, IBM license Center, and My IBM Container Library are enabled.

For procedure to link IBM purchase and Red Hat OpenShift account, see <https://www.ibm.com/docs/en/cloud-paks/1.0?topic=iocpc-accessing-red-hat-entitlements-from-your-cloud-paks>.

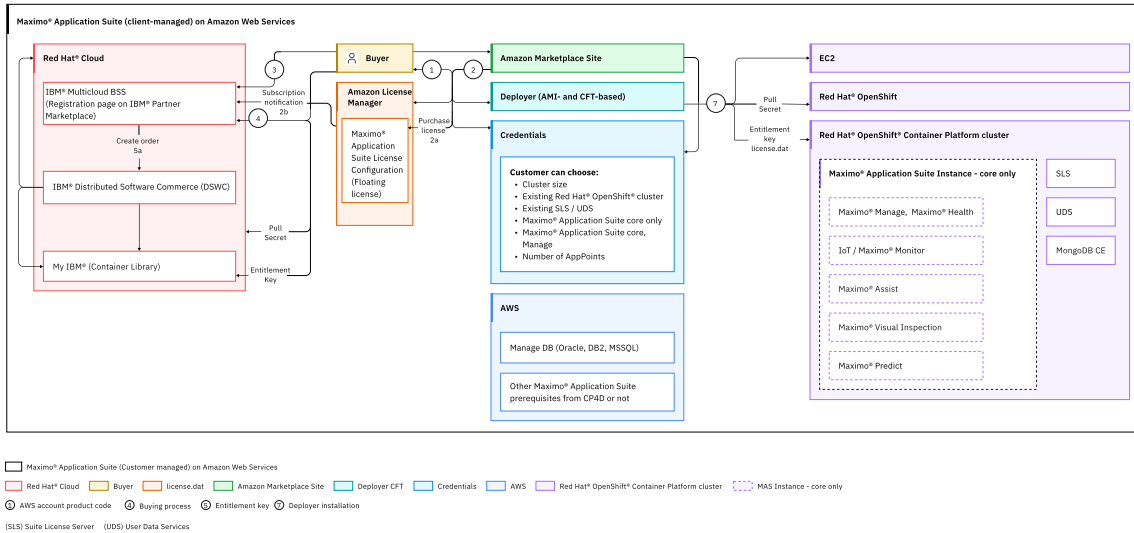
2. Retrieve the following:
 - a. The pull secret from the Red Hat OpenShift Cloud site.
 - b. Entitlement key from the My IBM site.
 - c. The license.dat file from IBM license Center (LKS) site.

A placeholder serverID is used.

3. The license . dat file is added in an S3 bucket that can be reached by using the URL.
4. You provide credentials and the S3 bucket link in the license . dat file.
5. The Cloud Formation Template is launched.
6. You can retrieve the license . dat file from the S3 bucket.
7. You install the Red Hat OpenShift Container Platform cluster and Maximo Application Suite in a way to match the placeholder serverID and then installs the license . dat file.

For more information, see [Cloud Pak for Data BYOL](#).

The following diagram shows the high-level architecture for Maximo Application Suite Paid on Amazon Web Services:



1. In Amazon Web Services Marketplace, you can subscribe to the public Maximo Application Suite. You might also find and buy a private offer for Maximo Application Suite. The Amazon Web Services account number and IBM product code are shown. You can then record the number and code.
2. a. The purchased license is provided in the License Manager of your account.
b. A notification is sent to the IBM Partner Marketplace that the client CustomerID subscribed to Product Code. For Private Offer, an offerID is included in the subscription notification.
3. IBM Partner Marketplace collects any missing information from the notification by using the daily reports. All information must be available within 72 hours. For a private offer, IBM Partner Marketplace might get information from sales qualified opportunities (SQO) to complete the order.
4. You can go to the IBM Partner Marketplace to complete the buying process and to get credentials. On the IBM Partner Marketplace, you get an IBM ID if you do not have one and must input the required information including the purchased product code and Amazon Web Services account number. You can now access instructions on how to get the entitled register key and pull secret. This flow might not be possible before 72 hours from the subscription in the public offering case.
5. The following process takes place:
 - a. If the subscription records are pending, the IBM Partner Marketplace creates the order on its internal customer portal now that the IBMid is known.
 - b. The order creation triggers the access to the IBM Container Library (Entitled Register) for Maximo Application Suite and IBM Cloud Pak for Data. Your Red Hat OpenShift account is enabled for linking.
6. You can now do the following:

- a. Access My IBM, retrieve the entitlement key, and link your IBM and Red Hat OpenShift accounts. For more information, see <https://www.ibm.com/docs/en/cloud-paks/1.0?topic=suocpc-accessing-red-hat-entitlements-from-your-cloud-paks>.
 - b. Register with Red Hat OpenShift and retrieves the pull secret.
7. You start the deployment and use the entitlement key and pull secret to install Red Hat OpenShift and Maximo Application Suite.
 8. Red Hat OpenShift Container Platform is installed with Cloud Credential Operator (CCO) in Mint mode and a credential request is created to create the secret that is used by IBM Suite License Service to allow interface Amazon Web Services License Manager Service.
 9. The Maximo Application Suite license server connects the Amazon License Manager to check out or check in AppPoints. This information is refreshed every 15 minutes.
- For more information, see [IBM Cloud Pak for Data Quick-start](#).

Next steps

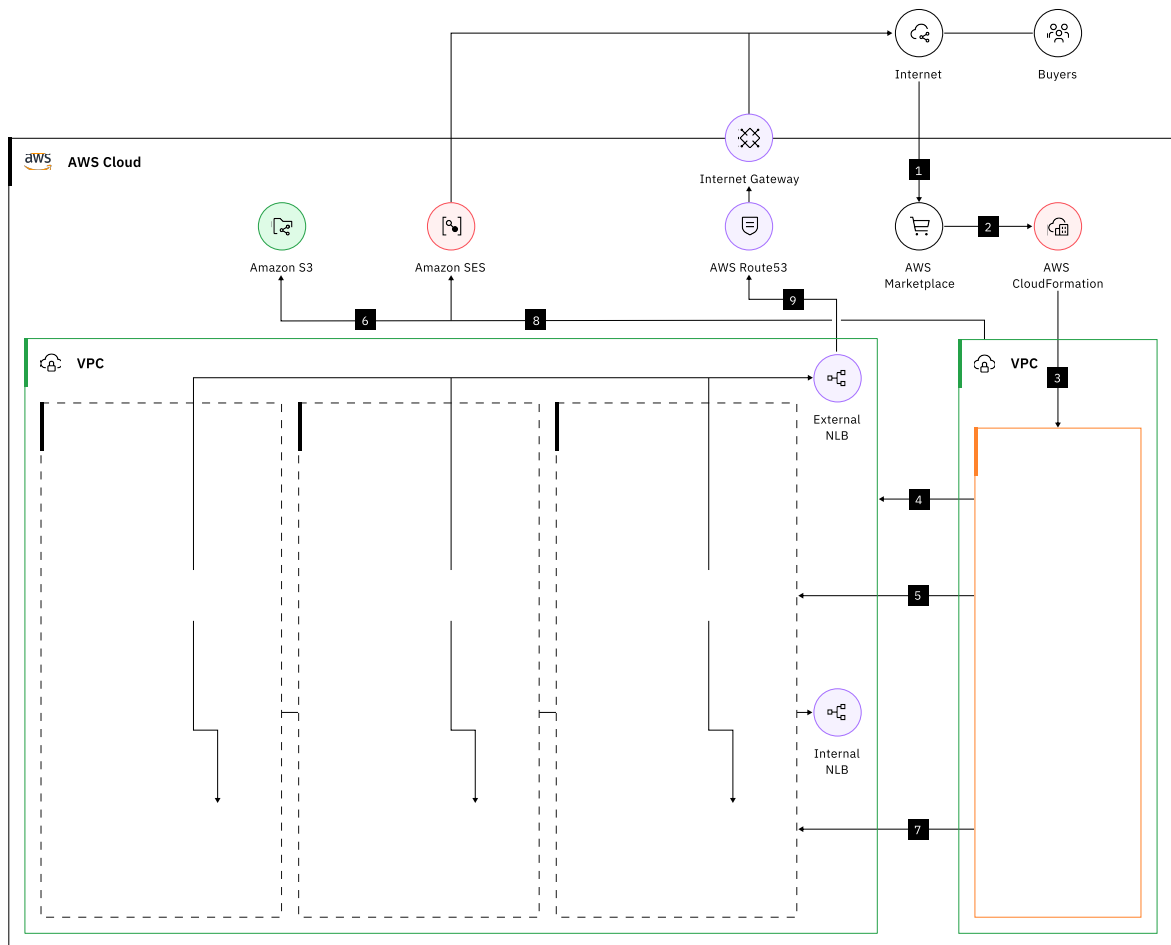
Before you can install IBM Maximo Application Suite on Amazon Web Services, you must consider your installation preferences, such as whether you want to create a Red Hat OpenShift cluster or reuse an existing cluster.

For more information, see [“Planning to install on Amazon Web Services”](#) on page 138.

Maximo Application Suite on Amazon Web Services installation topology

When you install IBM Maximo Application Suite in Amazon Web Services, the virtual network infrastructure, Red Hat OpenShift cluster, Maximo Application Suite prerequisites, and Maximo Application Suite are created in your Amazon Web Services cloud account.

The following figure and steps illustrate the sequence of events during an installation and the resulting topology:



- 1 In AWS Marketplace, you subscribe to the Maximo Application Suite product.
 - 2 In the AWS CloudFormation console, you configure a CloudFormation stack, specify the installation parameters, and start the installation by creating the stack.
 - 3 In your AWS cloud account, a Bootnode is created to complete the installation. The Bootnode contains all of the required information to install the Maximo Application Suite in an Red Hat OpenShift cluster.
- Note:** For the Red Hat OpenShift UPI option, the Bootnode is created in the same network infrastructure where the Red Hat OpenShift Container Platform is configured.
- 4 The Bootnode creates the required network infrastructure by using Terraform templates.
 - 5 The Bootnode launches an Red Hat OpenShift Container Platform (OCP) cluster deployment that creates the OCP bootstrap node, which uses the Red Hat OpenShift installer to create master and worker nodes.
 - 6 Information about the installation context and the Terraform state is saved to an Amazon S3 storage bucket in your AWS account.
 - 7 The Bootnode installs the Maximo Application Suite prerequisites and the Maximo Application Suite itself.
 - 8 If you configured a verified Amazon simple email service (SES) email address, you receive emails that contain the connection and authentication details for the Suite.
 - 9 The cluster is ready and you can access the Maximo Application Suite.

10 In the cluster, a bastion host is used to allow access to cluster nodes by using Secure Shell (SSH). The bastion host is kept in a shutdown state but, if required, you can restart it.

Note: The bastion host is not created for Red Hat OpenShift UPI and existing OCP cluster options.

11 After the installation is complete, the Bootnode is moved to a shutdown state. If required, you can restart the Bootnode and access it by using Secure Shell (SSH).

Maximo Application Suite unique identifiers for Amazon Web Services

During an IBM Maximo Application Suite installation, identifiers are created that uniquely identify the installation, the Red Hat OpenShift cluster, and the Maximo Application Suite instance.

When you install the Maximo Application Suite on Amazon Web Services, identifiers are created to uniquely identify components that are created during the installation. These identifiers are referred to throughout this documentation.

Installation identifier

When you install the Maximo Application Suite, a unique 6-character string, for example `drq2wd`, is generated to identify the installation itself. This identifier is also included in the names of several other components that are created during the installation, such as the Red Hat OpenShift cluster.

In this documentation, `<unique-string>` refers to the installation identifier.

Cluster identifier

When the Red Hat OpenShift cluster is created during the installation, it is identified by the following string: `masocp-<unique-string>`. For example, if the installation identifier is `drq2wd`, the cluster identifier is `masocp-drq2wd`. This identifier is passed to the cluster installer and is included in the cluster URL that you can access after the installation is complete.

In this documentation, `<cluster-name>` refers to the cluster identifier.

Instance identifier

When the Maximo Application Suite instance is created during the installation, it is identified by the following string: `mas-<unique-string>`. For example, if the installation identifier is `drq2wd`, the instance identifier is `mas-drq2wd`.

In this documentation, `<instance-id>` refers to the instance identifier.

Prerequisites for installing Maximo Application Suite on Amazon Web Services

Before you install IBM Maximo Application Suite, you must set up several components and gather the information that you need to specify the installation parameters.

The following components are required to install Maximo Application Suite in Amazon Web Services (AWS):

Amazon Web Services account

To install the Maximo Application Suite on Amazon Web Services (AWS), you must have an AWS account. During a Maximo Application Suite installation, a new virtual private cloud (VPC) is created in the AWS account. The Red Hat OpenShift cluster is then deployed in the VPC, and the Maximo Application Suite is deployed in the cluster.

For more information, see [How do I create and activate a new AWS account?](#)

Consider deploying into a new AWS account because the permissions that are defined in the Maximo Application Suite deployment automation allows you to create administrative roles, users, or groups.

In your AWS account, create an identity and access management (IAM) user and assign a managed policy to this user that includes the required permissions to install the Maximo Application Suite. You can assign the existing `AdministratorAccess` managed policy or create and assign your own managed policy. For more information, see [Configuring the installation permissions](#).

By default, your AWS account has access to all of the Amazon services that you need to install the Maximo Application Suite. These services include Amazon Route 53, CloudFormation, EC2 instances, S3 storage buckets, and the simple email service (SES). When you configure these services, ensure that you select the same geographical region where you want to install the Maximo Application Suite. For the list of supported regions, see the [installation considerations](#) topic.

Your AWS account includes default quotas for each AWS service. In the future, you might need to increase these quotas, for example if your Red Hat OpenShift cluster requires more resources. For more information, see [AWS service quotas](#).

SSH key pair

The Amazon EC2 service uses virtual computing environments, which are also known as instances, to provide computing capacity in the AWS cloud. In an installed Maximo Application Suite on Amazon Web Services, the Red Hat OpenShift cluster nodes are EC2 instances.

For more information, see [What is Amazon EC2?](#)

When an EC2 instance is created, a key pair, which consists of a public key and a private key, must be provided. Because the Maximo Application Suite installation creates many EC2 instances, you must generate a key pair and upload it in the Amazon EC2 service before you install the Maximo Application Suite. Your Amazon Web Services account has access to the EC2 service. For more information, see [Create a key pair](#).

When you specify the parameters for a Maximo Application Suite installation on Amazon Web Services, in the `SSHKey` parameter, select the public key that you generated. Keep the private key in a safe place and store it in the PEM format.

After the installation is complete, you can use the private key to connect to the cluster nodes by using Secure Shell (SSH). For more information, see [Accessing the Bootnode and Red Hat OpenShift cluster](#).

Bootnode CIDR IP address range

When you begin a Maximo Application Suite installation, a Bootnode is created that contains all of the required information, including the parameters that you specify, to complete the installation. The Bootnode creates and installs the VPC, Red Hat OpenShift cluster, Maximo Application Suite prerequisites, and Maximo Application Suite itself.

To connect to the Bootnode, you must specify a range of IPv4 addresses that can access it. In the `BootnodeSGIngressCidrIp` installation parameter, enter the IP address range by using the classless inter-domain routing (CIDR) notation. For example, to allow all IP addresses to access the Bootnode, enter `0.0.0.0/0`. To allow one specific IP address to access the Bootnode, enter `x.x.x.x/32`, for example `192.12.33.3/32`. For more information about CIDR notation, see https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing#CIDR_notation.

Maximo Application Suite license for BYOL product

If you do not already have your Maximo Application Suite license key file, you can create and download it in the [IBM License Key Center](#).

For instructions, see the steps in the `Upload your license key file` section of the [Setting up Maximo Application Suite](#) topic. When you create the license, provide the server parameters that are indicated in the following table:

Parameter	Value
Configuration	Single License Server
Host ID Type	Ethernet address

Parameter	Value
Host ID	A unique 12-character hexadecimal value, such as `0abcac110f02`. You can generate this value by using one of the following methods:

- An online hexadecimal number generator, such as [Online Hex Tools](#).
- A command-line utility, such as the Linux hexdump utility. For example, the following command generates a lowercase 12-character hexadecimal string:

```
hexdump -n 6 -e '8/8 "%08X" 1 "\n" /dev/urandom | tr '[:upper:]' '[:lower:]'
```

After you download the Maximo Application Suite license key file, upload it to an Amazon S3 storage bucket. Your AWS account has access to the Amazon S3 storage service. If any storage buckets are not created, first sign in to the AWS Management Console and open the Amazon S3 console. Then, click **Create bucket** and follow the steps in the bucket wizard. If you already have storage buckets created, in the **Buckets** list, choose the name of the bucket that you want to upload your Maximo Application Suite license key file to. Then, click **Upload** and follow the steps in the upload wizard.

For more information about using Amazon S3 buckets, see [Buckets overview](#) in the AWS documentation.

After you upload the license to the S3 bucket, record either the HTTP or the S3 location of the license, for example `s3://masocp-license/entitlement.lic`. You can find the S3 location in the Amazon S3 console by viewing the bucket name's properties. You enter this location when you specify the `MASLicenseUrl` installation parameter.

IBM Entitled Registry

The IBM Entitled Registry key is used during the installation to download the container images for the Maximo Application Suite and its applications from the IBM Entitled Registry.

Download this key from the [IBM Container Library](#).

Pull secret

To create the Red Hat OpenShift cluster, you must provide a pull secret. To access the pull secret, complete the following steps:

1. Create a Red Hat OpenShift account, if you do not have one already.

[Create a Red Hat Login](#).

2. To access your Red Hat OpenShift entitlements, in the Passport Advantage website, link your purchased Maximo Application Suite software to your Red Hat OpenShift account.

For more information about how to link purchased IBM software to a Red Hat OpenShift account, see [Accessing Red Hat entitlements from your IBM Cloud Paks](#).

3. Log in to the Red Hat Hybrid Cloud Console by using your Red Hat OpenShift account credentials. <https://cloud.redhat.com/>
4. In the **Clusters** page, click **Download Pull Secret**.

The pull secret is a JSON-formatted text file. When you configure the Maximo Application Suite installation parameters, copy the JSON text from the file into the `OpenShiftPullSecret` parameter.

Red Hat OpenShift subscription

If you are deploying the Maximo Application Suite using IBM Maximo Application Suite (client-managed, without Red Hat OpenShift entitlement) marketplace product, you should own Red Hat OpenShift Container Platform (OCP) subscription, which is valid and active throughout the term of Maximo Application Suite subscription.

For more information, see [Managing clusters](#).

Preparing to install Maximo Application Suite on Amazon Web Services

Before you install IBM Maximo Application Suite, consider your installation preferences, such as the type of Maximo Application Suite offering that you prefer and whether you want to reuse existing Maximo Application Suite compatible components in your new Maximo Application Suite instance.

You must also configure several prerequisite components before you install the Maximo Application Suite. For more information, see [Installation requirements](#).

When you install the Maximo Application Suite, you must specify options for the following items:

- [“Amazon Web Services region for installing Maximo Application Suite” on page 147](#)
- [“Maximo Application Suite offering type for Amazon Web Services” on page 147](#)

Note: Starting in 8.11, for US GovCloud regions, you can install Maximo Application Suite in private hosted zones for existing Red Hat OpenShift cluster and New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI).

If this installation is your first Maximo Application Suite installation on Amazon Web Services, a new Red Hat OpenShift cluster must be created. If you already installed an instance of the Maximo Application Suite on Amazon Web Services, you can reuse the existing cluster.

If you want the installation process to create a new cluster, you must configure and consider the following items:

- [“Public hosted zone” on page 148](#) or [Private hosted zone](#)
- [“Red Hat OpenShift cluster size” on page 148](#)

When you reuse an existing Red Hat OpenShift cluster, a prerequisites check is run in the cluster for following components:

- Cloud Pak for Data, if you selected the Maximo Application Suite core and Cloud Pak for Data offering type.
- SLS, if you do not specify an existing SLS instance to reuse.
- Data Reporter Operator , if you do not specify an existing User Data Services instance to reuse.
- MongoDB
- Cert-manager

The components are reused if they are installed in the cluster with supported version. Else, new instances of the components are created in the cluster.

The deployment is stopped if an unsupported version is preinstalled for any of the components or if multiple instances of a component are running in existing cluster.

The prerequisites check also verifies whether the following storage classes are available in the cluster:

- gp2
- ocs-storagecluster-cephfs, if you selected Maximo Application Suite core and Cloud Pak for Data.

The following components are required to ensure that the cluster has enough resources to accommodate the new Maximo Application Suite instance.

- Minimum number of worker nodes: 3
- Minimum CPUs per node: 8 cores
- Minimum memory per node: 32 GB

EBSVolumeType - Starting in Maximo Application Suite 8.11, select your preferred EBS Volume type which could be either gp3 or io1 for worker nodes based on your performance requirements and cost of the volume types.

For more information on volume types, see [Amazon EBS volume types](#)

Amazon Web Services region for installing Maximo Application Suite

In Amazon Web Services Marketplace, after you subscribe to the IBM Maximo Application Suite product, you must select the geographical region where you want to install the Maximo Application Suite.

To ensure that the installation succeeds, select one of the following supported regions:

Region name	Region endpoint identifier
Africa (Cape Town)	af-south-1
Asia Pacific (Hong Kong S.A.R. of the PRC)	ap-east-1
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Osaka)	ap-northeast-3
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Canada (Central)	ca-central-1
Europe (Frankfurt)	eu-central-1
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2
Europe (Milan)	eu-south-1
Europe (Paris)	eu-west-3
Europe (Stockholm)	eu-north-1
Middle East (Bahrain)	me-south-1
South America (São Paulo)	sa-east-1
US East (N. Virginia)	us-east-1
US East (Ohio)	us-east-2
US East (Oregon)	us-west-2
US GovCloud (East)	us-gov-east-1
US GovCloud (West)	us-gov-west-1

Maximo Application Suite offering type for Amazon Web Services

When you specify the parameters for a IBM Maximo Application Suite installation on Amazon Web Services, you can choose different Maximo Application Suite offering types.

- Maximo Application Suite core and Cloud Pak for Data
 - Db2 Warehouse and Db2 Data Management Console are the only Cloud Pak for Data services that are installed.
 - You install the Maximo Application Suite applications manually after installation.
 - You configure the prerequisites for any Maximo Application Suite applications that you install.
 - Starting in 8.11, you can configure Maximo Application Suite core with Cloud Pak for Data for Red Hat OpenShift on Amazon Web Services (ROSA).
 - Starting in 8.11, for US GovCloud regions, Cloud Pak for Data is not available for configuration.

- Maximo Application Suite core and Maximo Manage
 - Cloud Pak for Data is not included in this offering type.
 - For IBM Maximo Manage, you configure databases such as IBM Db2, Microsoft SQL Server, or Oracle Database.

Starting in 8.11, the databases can be hosted on private subnet of another VPC. Ensure that the VPC ID resides in the current deployment region, and does not have matching or overlapping IPv4 CIDR blocks 10.0.0.0/16.

Note: Starting in 8.11, for US GovCloud regions, IBM Db2 is not available for configuration.

Amazon Web Services zones for Red Hat OpenShift cluster

To create a Red Hat OpenShift cluster, you can select either a public hosted or private hosted zone in Amazon Web Services.

Public hosted zone

For IBM Maximo Application Suite installations, if you want to create a new Red Hat OpenShift cluster, you must configure a public hosted zone. A public hosted zone is a container that holds information about how you want to route internet traffic for a specific domain.

When the Maximo Application Suite cluster is created, the public hosted zone is used to allow internet traffic to access the cluster from outside the cluster's virtual private cloud (VPC). For more information, see [Working with public hosted zones](#).

You configure the public hosted zone in the Amazon Route 53 service. Amazon Route 53 is a domain name system (DNS) service that provides domain registration, routing, and health-checking functions. You can register any valid and available domain name. Your Amazon Web Services account has access to the Amazon Route 53 service. For the instructions to create a public hosted zone in Amazon Route 53, see [Creating a public hosted zone](#).

Private hosted zone

For Maximo Application Suite installations, if you want to create a new private Red Hat OpenShift cluster, you must configure a private hosted zone. A private hosted zone is a container that holds information about how you want to route internet traffic for a specific domain.

When the Maximo Application Suite cluster is created, the private hosted zone will not allow internet traffic to access the cluster from outside the cluster's virtual private cloud (VPC) unless accessed from Jumphost, which is present in the same VPC. For more information, see [Working with private hosted zones](#).

You configure the private hosted zone in the Amazon Route 53 service. Amazon Route 53 is a domain name system (DNS) service that provides domain registration, routing, and health-checking functions. You can register any valid and available domain name. Your AWS account has access to the Amazon Route 53 service. For the instructions to create a private hosted zone in Amazon Route 53, see [Creating a private hosted zone](#).

Red Hat OpenShift cluster size

When you specify the IBM Maximo Application Suite installation parameters, if you want to create a new Red Hat OpenShift cluster, you must choose the size of the cluster that you want to create.

Cluster size

Choose a cluster size that is appropriate for the scale of your Maximo Application Suite installation. Consider the Maximo Application Suite applications that you might want to deploy after the installation is complete.

Tip: You can also [resize your Red Hat OpenShift cluster](#) after it is provisioned.

For more information, see [“Requirements and capacity planning”](#) on page 178.

You can choose to create a small, medium, or large cluster. By default, the cluster that you create spans multiple availability zones. For the node, CPU, and memory dimensions for each cluster size, see the following table:

Cluster size	Node type	Number of nodes	EC2 instance size	CPUs per node	Total CPUs	Memory per node (GB)	Total memory (GB)
Small	Primary	3	m5.2xlarge	8	24	32	96
	Worker	3	m5.4xlarge	16	48	64	192
	Bootnode and Bastion host	1	t3.small	2	4	2	2
	Infra	3	m5.4xlarge	16	48	64	192
Medium	Primary	3	m5.2xlarge	8	24	32	96
	Worker	5	m5.4xlarge	16	80	64	320
	Bootnode and Bastion host	1	t3.small	2	4	2	2
	Infra	3	m5.4xlarge	16	48	64	192
Large	Primary	5	m5.2xlarge	8	40	32	160
	Worker	7	m5.4xlarge	16	112	64	448
	Bootnode and Bastion host	1	t3.small	2	4	2	2
	Infra	3	m5.4xlarge	16	48	64	192

Note: The Bootnode is used only during the installation and shutdown after the installation is complete. It is a workstation from which the Red Hat OpenShift cluster deployment, Maximo Application Suite is deployed. Hence, it can be used later to access the deployment files, configurations, logs, or to run the ad hoc actions on the cluster as needed.

Connection details

Check connection details for existing Red Hat OpenShift cluster, network infrastructure, IBM Suite License Service instance, and IBM Data Reporter Operator instance.

Connection details for an existing Red Hat OpenShift cluster

If you want to reuse an existing Red Hat OpenShift cluster, verify that it has enough resources to accommodate the new Suite instance. You can resize your Red Hat OpenShift cluster by completing the following task: [“Resizing Red Hat OpenShift cluster on Amazon Web Services”](#) on page 158.

Note: For details on Maximo Application Suite offerings and supported versions of Red Hat OpenShift, see [Software Product Compatibility Reports \(SPCR\)](#).

To reuse the cluster, enter the following installation parameters:

- The cluster's API URL in the format `https://api.<cluster-name>.<domain_name>`. Do not specify the port number. For example, `https://api.masocp-joalae.mas4aws.com`

For more information on the cluster name's format, see [Unique identifiers](#).

- `BootNodeVPCId`. The bootnode is added to this VPC.

Note: Ensure the existing Red Hat OpenShift cluster that is specified is reachable from this VPC.

- `BootNodeSubnetId`. Select the subnet ID from the bootnode VPC to be used for the bootnode.

Note: Ensure that the bootnode is in public subnet of VPC for the bootnode to be accessible.

- Username
- Password

– If the cluster was created during a previous Suite installation on AWS and you received the connection credentials in an email, retrieve the username and password from the email.

– If you did not receive an email, the Red Hat OpenShift cluster credentials can be obtained from the AWS Secret Manager service with the name `maximo-ocp-secret-<unique-string>`.

When you reuse an existing Red Hat OpenShift cluster, a prerequisites check is run in the cluster for following components:

- Cloud Pak for Data, if you selected the Maximo Application Suite core and Cloud Pak for Data offering type.
- Suite License Service, if you do not specify an existing SLS instance to reuse.
- IBM Data Reporter Operator

Tip: Starting in Maximo Application Suite 9.0, if you are using IBM User Data Services, you must migrate to Data Reporter Operator . For information, see [“Migrating Maximo Application Suite from User Data Services to Data Reporter Operator ”](#) on page 8.

- MongoDB

Starting in Maximo Application Suite 8.11, for US GovCloud regions, the external MongoDB cluster that is configured by you must follow the Federal Information Processing Standard (FIPS).

- Cert-manager

The components are reused if they are installed in the cluster with supported version. Else, new instances of the components are created in the cluster.

The deployment is stopped if an unsupported version is preinstalled for any of the components or if multiple instances of a component are running in existing cluster.

The prerequisites check also verifies whether the following storage classes are available in the cluster:

- `gp2`
- `ocs-storagecluster-cephfs`, if you selected Maximo Application Suite core and Cloud Pak for Data.

The following components are required to ensure the cluster has enough resources to accommodate the new Maximo Application Suite instance.

- Minimum number of worker nodes: 3
- Minimum CPUs per node: Eight cores
- Minimum memory per node (GB): 32

`ExocpProvisionedVPCId`

Starting in Maximo Application Suite 8.11, enter the VPC ID where your existing OCP cluster is provisioned. This is required to establish VPC Peering from this VPC to the VPC of database to establish database connection.

Connection details for using an existing network infrastructure

If you want to reuse an existing network infrastructure present in your region, verify that it has all the available network resources. It must have sufficient resources to accommodate the new Suite instance.

To reuse the existing network infrastructure, enter the following installation parameters:

1. **HostedZone**. Select a public or private hosted zone that you created in the AWS Route 53 service from the drop-down values.

For example, `mas4aws.myorg.com`

2. Select the values for an existing network.

- a. **Private cluster**. Select **true** to create a private Red Hat OpenShift cluster.

The default value is set to **false** and creates a public Red Hat OpenShift cluster.

Tip: When the **hostedzone** is private, **Private cluster** value must be set to true and when the **hostedzone** is public, the **Private cluster** value must be set to false.

Starting in Maximo Application Suite 8.11, for US GovCloud regions, by default the **Private cluster** value is set to true.

- b. Values of VPC ID, public subnet, and private subnet must be selected from the parameter drop-down menu.

For 3 availability zone, there are 3 public and 3 private subnet.

Connection details for an existing Suite License Service instance

If you previously installed Suite-compatible components in your AWS account, you can reuse them when you install your new Suite instance.

To reuse an existing Suite License Service (SLS) instance, you can enter the HTTP or S3 location of the service's public certificate, registration key, and endpoint URL installation parameters. To get these values:

1. Log in to the Red Hat OpenShift web console.
2. Go to **Workloads > ConfigMaps**.
3. Select the **sls** project.
4. Search for **suite-registration**.
5. Go to the **Data** section and copy the values for **ca**, **registrationKey** and **url**.

If you leave these parameters empty, a new SLS instance is created during the installation.

Data Reporter Operator connection details

Starting in Maximo Application Suite 9.0, to use a Data Reporter Operator instance, you can enter the following installation parameters:

Endpoint URL

On the **Routes details** page in the `redhat-marketplace` project, you can find the location of **ibm-data-reporter**.

API key

On the **Workloads > Secrets** page of the `mas-<instanceId>-core` project, you can find the **dro-apikey**.

HTTP or S3 location of the service's public certificate

To find the HTTP or S3 location of the service's public certificate such as `s3://masocp-license/dro-certificate.crt`, follow the steps.

- In the existing Maximo Application Suite instance that uses the Data Reporter Operator, in the `mas-<instance-id>-core` namespace, in the `<instance-id>-dro-cfg` secret, retrieve the certificate from the `ca-bundle.pem` file.

- Upload the certificate to your Amazon S3 bucket and record the certificate's HTTP or S3 location.

If you leave these parameters empty, a new Data Reporter Operator instance is created during the installation.

Related tasks

Migrating Maximo Application Suite from User Data Services to Data Reporter Operator

As an IBM Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance. New and existing Maximo Application Suite users can install or migrate to DRO by using the IBM Maximo Application Suite command line interface (CLI), ansible role, or manually.

Verified Amazon SES email address

Create and verify your email address in the Amazon simple email service (SES) and set an installation parameter to configure the IBM Maximo Application Suite. This email is used to confirm that the installation succeeded and provide you with the access details.

Verified Amazon SES email address

Your AWS account has access to Amazon SES. Create and verify at least one email address. For more information, see [Creating and verifying an email address identity](#) in the AWS documentation.

When you specify the installation parameters for the Suite, set the `EmailNotification` parameter to `true`. After the installation is complete, each verified SES address receives two emails that contain the details to access the Suite.

The first email contains connection URLs and usernames for the cluster and Suite in the following format:

- Maximo Application Suite provisioning status: SUCCESS
- Region: <region-code>
- Unique String: <unique-string>
- Red Hat OpenShift cluster URL: <ocp-cluster-url>
- Red Hat OpenShift API URL: <ocp-api-url>
- Red Hat OpenShift User: <ocp-username>
- IBM Suite License Service Endpoint URL: <sls-url>
- DRO Endpoint URL: <dro-url>
- Maximo Application Suite Initial setup URL: <setup-url>
- Maximo Application Suite Admin URL: <admin-url>
- Maximo Application Suite workspace URL: <workspace-url>
- Maximo Application Suite User: <mas-username>

The second email contains password information in the following format:

- Maximo Application Suite provisioning status: SUCCESS
- Region: <region-code>
- Unique String: <unique-string>
- Red Hat OpenShift Password: <ocp-password>
- Maximo Application Suite Password: <mas-password>

Database configuration details

If your IBM Maximo Application Suite offering includes the IBM Maximo Manage application, you must configure a database that this application can use.

Database configuration details for Maximo Manage

Maximo Manage supports IBM Db2, IBM Db2 Warehouse, Microsoft SQL Server, and Oracle Database.

For more information, see [Preparing your database for deployment](#).

After you configure the database that you want Maximo Manage to use, you can specify the following configuration details when you enter the Suite installation parameters:

- Username
- Password
- Java database connectivity (JDBC) URL.

For example, use the following parameters for the database that is configured.

- For IBM Db2 database:

```
jdbc:db2://1.2.3.4:50051/FTMDB:sslConnection=true
```

- Starting in 8.11, for Microsoft SQL Server database:

```
jdbc:sqlserver://;serverName=10.5.0.4;portNumber=1433;databaseName=maxdb80;  
sendStringParametersAsUnicode=false;selectMethod=cursor;encrypt=true;trustServerCertificate=  
true;
```

- Starting in 8.11, for Oracle Database:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=63.246.112.120)  
(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=maxdb)))
```

Note: Only secure sockets layer (SSL) connections to the database are supported.

Specify the JDBC connection URL, the username, and password to connect to the database that is configured for Maximo Manage. If you use SSL, ensure that the Port field contains your SSL port, as shown in the following examples:

IBM Db2

For Db2 SSL database connections in Maximo Manage, you must specify `sslConnection=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB:sslConnection=true;
```

Note: Starting in 8.11, for US GovCloud regions, Db2 is not available for configuration.

- HTTP or S3 location of the database's public certificate, for example `s3://masocp-license/db-certificate.crt`.

Ensure that you can retrieve your certificate depending on your database type. For example, for IBM Db2 Warehouse, go to the **Details** view of the Db2 instance in Cloud Pak for Data, and click the **Download SSL certificate** button.

In your AWS account, upload the certificate to an Amazon S3 bucket. Record the certificate's HTTP or S3 location.

Your AWS account has access to the Amazon S3 storage service. If you do not yet have any storage buckets created, first sign in to the AWS Management Console and open the Amazon S3 console. Then, click **Create bucket** and follow the steps in the bucket wizard. If you already have storage buckets created, in the **Buckets** list, choose the name of the bucket that you want to upload your Maximo Application Suite license key file to. Then, click **Upload** and follow the steps in the upload wizard.

- Whether you want to import demo data into the database. This data might be useful for development or test environments.

- `TablespaceNames` - Starting in 8.11, if you do not use the default names `maxdata` and `maxindex`, enter the tablespace and indexespace names. Each name is separated by a delimiter colon (:). For example, `<tablespace_name>:<indexespace_name>`

Note: This is not applicable to Microsoft SQL Server.

- `DBProvisionedVPCId` - Starting in 8.11, enter the VPC ID where your existing database such as IBM Db2, Microsoft SQL Server, or Oracle Database is provisioned on a private subnet of another VPC to establish the VPC peering from this VPC and the VPC created during current stack deployment to establish database connection. Ensure that this VPC ID resides in the current deployment region and does not have matching or overlapping IPv4 CIDR blocks 10.0.0.0/16.

Configuring the installation permissions

To ensure that you can install Maximo Application Suite on AWS, create an identity and access management (IAM) user in your AWS account and grant the required permissions to this user.

During a Maximo Application Suite installation on Amazon Web Services, identity and access management (IAM) resources are created by the Bootnode and the Red Hat OpenShift installer in the following order:

- The Bootnode creates the IAM resources that it needs to create the Red Hat OpenShift cluster.
- In the cluster, the Red Hat OpenShift installer creates the IAM resources that it needs to create master and worker nodes.

Procedure

In your AWS account, you create an IAM user that you can use to perform the Maximo Application Suite installation. This user must have the permissions to create the IAM resources that the Bootnode and Red Hat OpenShift installer require.

To grant the required permissions to this user, choose one of the following options:

- Assign the `AdministratorAccess` managed policy to the user. For more information, see Creating an administrator IAM user and user group in the AWS documentation.
- Create a customer-managed policy that includes the permissions in the following table and assign this policy to the user:

Permission
<code>autoscaling:DescribeAutoScalingGroups</code>
<code>ec2:AllocateAddress</code>
<code>ec2:AssociateAddress</code>
<code>ec2:AssociateDhcpOptions</code>
<code>ec2:AssociateRouteTable</code>
<code>ec2:AttachInternetGateway</code>
<code>ec2:AttachNetworkInterface</code>
<code>ec2:AuthorizeSecurityGroupEgress</code>
<code>ec2:AuthorizeSecurityGroupIngress</code>
<code>ec2:CopyImage</code>
<code>ec2>CreateDhcpOptions</code>
<code>ec2>CreateInternetGateway</code>
<code>ec2>CreateNatGateway</code>
<code>ec2>CreateNetworkInterface</code>

Permission
ec2:CreateRoute
ec2:CreateRouteTable
ec2:CreateSecurityGroup
ec2:CreateSubnet
ec2:CreateTags
ec2:CreateVolume
ec2:CreateVpc
ec2:CreateVpcEndpoint
ec2>DeleteDhcpOptions
ec2>DeleteInternetGateway
ec2>DeleteNatGateway
ec2>DeleteNetworkInterface
ec2>DeleteRoute
ec2>DeleteRouteTable
ec2>DeleteSecurityGroup
ec2>DeleteSnapshot
ec2>DeleteSubnet
ec2>DeleteTags
ec2>DeleteVolume
ec2>DeleteVpc
ec2>DeleteVpcEndpoints
ec2:DeregisterImage
ec2:Describe
ec2:DescribeAccountAttributes
ec2:DescribeAddresses
ec2:DescribeAvailabilityZones
ec2:DescribeDhcpOptions
ec2:DescribeImages
ec2:DescribeInstanceAttribute
ec2:DescribeInstanceCreditSpecifications
ec2:DescribeInstanceTypeOfferings
ec2:DescribeInstanceTypes
ec2:DescribeInstances
ec2:DescribeInternetGateways
ec2:DescribeKeyPairs

Permission
ec2:DescribeNatGateways
ec2:DescribeNetworkAcls
ec2:DescribeNetworkInterfaces
ec2:DescribePrefixLists
ec2:DescribeRegions
ec2:DescribeRouteTables
ec2:DescribeSecurityGroups
ec2:DescribeSubnets
ec2:DescribeTags
ec2:DescribeVolumes
ec2:DescribeVpcAttribute
ec2:DescribeVpcClassicLink
ec2:DescribeVpcClassicLinkDnsSupport
ec2:DescribeVpcEndpoints
ec2:DescribeVpcs
ec2:DetachInternetGateway
ec2:DisassociateRouteTable
ec2:GetEbsDefaultKmsKeyId
ec2:ModifyInstanceAttribute
ec2:ModifyNetworkInterfaceAttribute
ec2:ModifySubnetAttribute
ec2:ModifyVpcAttribute
ec2:ReleaseAddress
ec2:ReplaceRouteTableAssociation
ec2:RevokeSecurityGroupEgress
ec2:RevokeSecurityGroupIngress
ec2:RunInstances
ec2:TerminateInstances
elasticloadbalancing:*
iam:AddRoleToInstanceProfile
iam:AttachUserPolicy
iam:CreateAccessKey
iam:CreateInstanceProfile
iam:CreatePolicy
iam:CreateServiceLinkedRole

Permission
iam:CreateRole
iam:CreateUser
iam>DeleteAccessKey
iam>DeleteInstanceProfile
iam>DeleteRole
iam>DeleteRolePolicy
iam>DeleteUser
iam>DeleteUserPolicy
iam:DetachUserPolicy
iam:GetInstanceProfile
iam:GetRole
iam:GetRolePolicy
iam:GetUser
iam:GetUserPolicy
iam:ListAccessKeys
iam:ListAttachedRolePolicies
iam:ListAttachedUserPolicies
iam:ListInstanceProfiles
iam:ListInstanceProfilesForRole
iam:ListRolePolicies
iam:ListRoles
iam:ListUserPolicies
iam:ListUsers
iam:PassRole
iam:PutRolePolicy
iam:PutUserPolicy
iam:RemoveRoleFromInstanceProfile
iam:SimulatePrincipalPolicy
iam:TagRole
iam:TagUser
iam:UntagRole
kms:CreateKey
kms:DescribeKey
kms:GetKeyPolicy
kms:GetKeyRotationStatus

Permission
kms:ListResourceTags
kms:ScheduleKeyDeletion
logs:CreateLogGroup
logs:CreateLogStream
logs:PutLogEvents
route53:ChangeResourceRecordSets
route53:ChangeTagsForResource
route53:CreateHostedZone
route53>DeleteHostedZone
route53:GetChange
route53:GetHostedZone
route53:ListHostedZones
route53:ListHostedZonesByName
route53:ListResourceRecordSets
route53:ListTagsForResource
route53:UpdateHostedZoneComment
s3:*
servicequotas:ListAWSDefaultServiceQuotas
servicequotas:ListServiceQuotas
ses:GetIdentityVerificationAttributes
ses:ListIdentities
ses:SendEmail
ses:SendRawEmail
sts:GetCallerIdentity
tag:GetResources
tag:TagResources
tag:UntagResources
secretsmanager:CreateSecret

Resizing Red Hat OpenShift cluster on Amazon Web Services

You can change Red Hat OpenShift primary and secondary nodes on Amazon Web Services according to your sizing needs. This process involves modifying machine sets, and for Amazon Web Services clusters, changing the instance type.

For more information, see [Modifying a machine set](#) and [Red Hat OpenShiftResizing nodes in Red Hat OpenShift Container Platform 4.6 or later](#).

In AWS clusters, you change nodes in the Amazon Web Services console. The **EC2 Services** page displays information about various EC2 instance types. To go to this page, in the Amazon Web Services console, click expand **Instances** then click **Instance Types**.

Procedure

- Changing the instance type on Amazon Web Services nodes.
 - a) In the AWS management console, go to EC2 Services page for the region where your Maximo Application Suite instance is installed.
If the Maximo Application Suite instance is installed in us-east-1, the URL for EC2 Services is <https://us-east-1.console.aws.amazon.com/ec2/v2/home?region=us-east-1#Home>:
 - b) Click **Instances** or **Instances (running)**.
 - c) Find the EC2 instance (the Red Hat OpenShift cluster node) you want to change.
Tip: Filter by the `ClusterUniqueString` of your Maximo Application Suite instance. This process is especially useful if other EC2 instances in the region that are not part of the Maximo Application Suite instance you are working on.
 - d) Check the checkbox for the EC2 instance that you want to modify. Click **Instance state**, then select **Stop instance**.
 - e) When the instance is stopped, click **Actions** then select **Instance settings > Change instance type**.
 - f) On the **Change instance type** page, select the instance type.
 - g) Restart the instance. Wait 10 minutes for the instance to stabilize.
- Adding a node to an Amazon Web Services
Adding machine sets in AWS clusters can be done in similar way as the procedure provided for adding a GPU node to an AWS Red Hat OpenShift cluster. Use a different EC2 instance type.

For more information, see [“Adding a GPU worker node to a Red Hat OpenShift cluster on AWS” on page 394](#).

What to do next

Install the Maximo Application Suite on Amazon Web Services

You can install IBM Maximo Application Suite in the Amazon Web Services (AWS) cloud by using the Amazon Web Services CloudFormation templates. Maximo Application Suite is available as a bring-your-own-license (BYOL) and contract pricing product in AWS Marketplace. After you configure the installation requirements and consider your installation preferences, you subscribe to the product, specify the installation parameters, and start the installation.

Security considerations

Before you can install IBM Maximo Application Suite, you must ensure that your environment meets all security requirements.

Maximo Application Suite deployment complies with the following security considerations:

- Communication to the IBM Maximo Manage database uses JDBC with SSL enabled.
- SSH keys used for the connection to the bootnode and Red Hat OpenShift cluster nodes (master/worker/infra).
- Bootnode runs within the customer's Amazon Web Services account and does not have connectivity to the external network during and post deployment.
- Product images are pulled from authenticated IBM entitled registries.
- Credentials are kept in Red Hat OpenShift secrets.
- Access to the Red Hat OpenShift cluster nodes is only through the bastion host using private SSH key.
- The Amazon Web Services portal uses HTTPS (SSL/TLS certificates) for encryption.

What to do next

[“IBM Maximo Application Suite installation with Amazon Web Services CloudFormation templates” on page 222](#)

You can install IBM Maximo Application Suite in the Amazon Web Services (AWS) cloud by using the Amazon Web Services CloudFormation templates. Maximo Application Suite is available as a bring-your-own-license (BYOL) and contract pricing product in AWS Marketplace. After you configure the installation requirements and consider your installation preferences, you subscribe to the product, specify the installation parameters, and start the installation.

Planning to install on Microsoft Azure

Before you can install Maximo Application Suite, you must configure the installation requirements and consider your installation preferences, such as whether you want to create an Red Hat OpenShift cluster or reuse an existing one.

Related concepts

[Requirements and capacity planning](#)

Maximo Application Suite on Microsoft Azure overview

IBM Maximo Application Suite can be installed in the Microsoft Azure cloud. You subscribe to this product, specify the installation parameters, and install the Maximo Application Suite.

Maximo Application Suite is available either as a BYOL product or a paid offering in Microsoft Azure Marketplace.

After you configure the installation requirements and consider your installation preferences, you are ready to install the Maximo Application Suite. Connect to Microsoft Azure Marketplace, find Maximo Application Suite, and then deploy the product by providing necessary details in Microsoft Azure Resource Manager (ARM) template.

From Maximo Application Suite 8.8 or later, the Installer-provisioned Infrastructure (IPI) and User-provisioned Infrastructure (UPI) methods of Red Hat OpenShift cluster deployment are supported. These deployment methods are available as two different plans in the Marketplace product. The **New OpenShift cluster (IPI)** is IPI mode and **New OpenShift cluster, existing network (UPI)** is the UPI mode.

In your Microsoft Azure subscription, the installation process creates a new resource group (called the *boot node resource group*). The boot node and other necessary resources are created in the boot node resource group. The installation process then passes all of the installation details, including the parameters that you specified, to the boot node. The boot node creates the virtual network infrastructure in a different resource group that is created by the Red Hat OpenShift installer itself called the Red Hat OpenShift Container Platform cluster resource group. The boot node then installs Red Hat OpenShift cluster, prerequisites, and the Maximo Application Suite.

The Microsoft Azure Bastion service is used to provide Secure Shell (SSH) access to the Red Hat OpenShift cluster nodes. You can connect to the Red Hat OpenShift Container Platform cluster nodes over SSH from Microsoft Azure portal by selecting the *Bastion* option from the *Connect* list on the virtual machine details page.

After the installation is complete, you can access the Maximo Application Suite. If you provide the SMTP configuration during the deployment, you receive emails that contain the information that you need to access the Maximo Application Suite. After you log in to the Maximo Application Suite as the administrator, you can deploy the Maximo Application Suite applications, create users, and administer Maximo Application Suite configurations.

BYOL product

The BYOL product can be installed from the Microsoft Azure Marketplace by using any one of the following three plans:

1. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
2. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

3. Existing Red Hat OpenShift cluster

After you receive your Maximo Application Suite license, you can use it to install the Maximo Application Suite in your Microsoft Azure cloud account.

Paid product

You can install a new Red Hat OpenShift cluster as part of the Maximo Application Suite deployment from the Microsoft Azure Marketplace.

Deploy Maximo Application Suite from the Microsoft Azure Marketplace by subscribing to two different products.

1. IBM Maximo Application Suite (Paid) - You must subscribe first to complete the transactional aspect of the purchase process. The following plans are available for this product.

- Fixed offer for IBM Maximo Application Suite
- Custom offer for IBM Maximo Application Suite

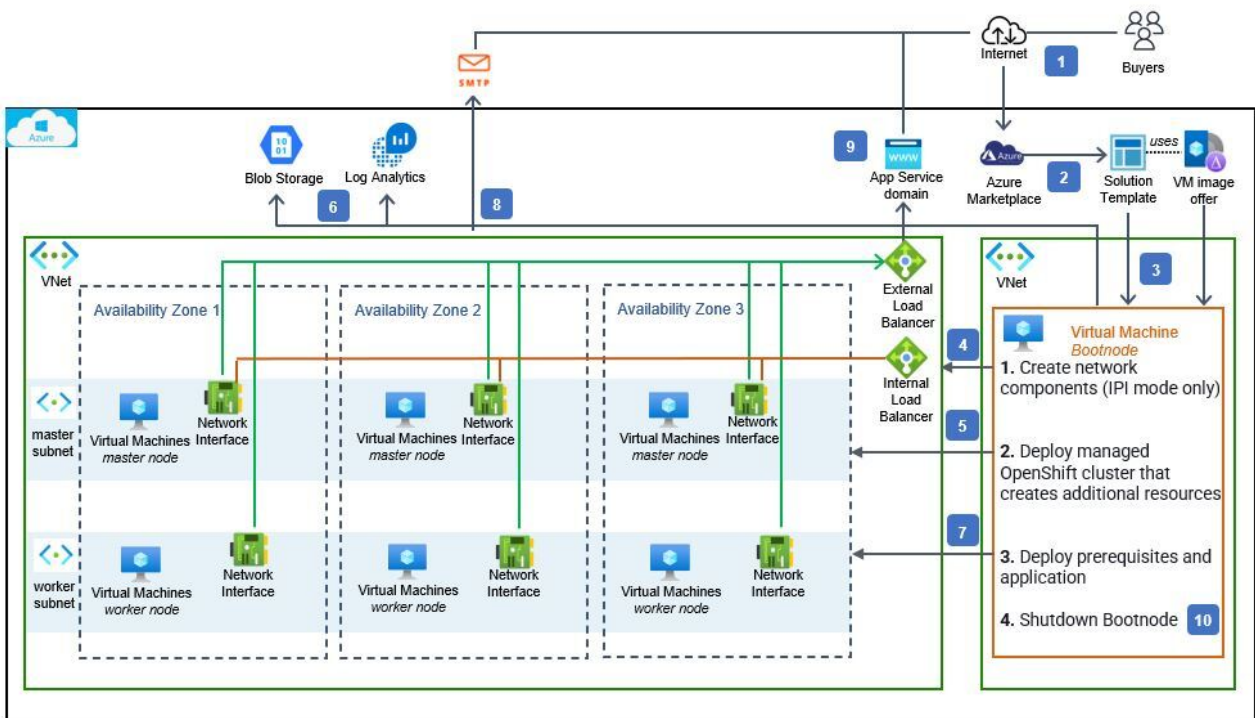
2. IBM Maximo Application Suite (BYOL) - You must subscribe to this product as a next step to complete the actual product deployment process. Refer to the BYOL section for plan details.

Maximo Application Suite on Azure installation topology

When you install Maximo Application Suite in Azure, the virtual network infrastructure, Red Hat OpenShift cluster, Suite prerequisites, and Suite are created in your Azure account.

The virtual network infrastructure configured with the **New OpenShift cluster** option, Red Hat OpenShift cluster, Suite prerequisites, and Suite are created in your Azure account when the Maximo Application Suite is installed in Azure.

The following figure and steps illustrate the sequence of events during an installation and the resulting topology:



1 In Azure Marketplace, you subscribe to the Maximo Application Suite product.

2 In the user interface section of the product deployment process, you specify the installation parameters, and start the installation by creating the product instance.

3

In your Azure subscription, a boot node is created in a new resource group (called the *boot node resource group*) to complete the installation. The boot node contains all of the required information to install the Suite in an Red Hat OpenShift cluster.

4

The boot node starts a Red Hat OpenShift Container Platform (OCP) cluster deployment that creates the OCP bootstrap node and network infrastructure. The OCP bootstrap node uses the Red Hat OpenShift installer to create controller and worker nodes. The Red Hat OpenShift cluster resources are created in a different resource group that is created by Red Hat OpenShift installer itself, called the *OCP cluster resource group*.

When you use the **New OpenShift cluster** option, you can select New OpenShift cluster (IPI), New OpenShift cluster, existing network (UPI), and Existing OpenShift cluster.

5

Information about the installation context and the Terraform state is saved to an Azure Blob storage in your Azure subscription.

6

The boot node installs the Suite prerequisites and the Suite itself.

7

If you provide the SMTP configuration during the deployment, you receive emails that contain the information that you need to access the Suite.

8

The cluster is ready, and you can access the Suite.

9

After the installation is complete, the boot node is moved to a shutdown state. If required, you can restart the boot node and access it by using Secure Shell (SSH).

1. Subscribe to the Maximo Application Suite in the Microsoft Azure Marketplace.
2. Specify the installation parameters, and start the installation by creating the product instance in the user interface section of the product deployment process.
3. A boot node is created in a new resource group (called the boot node resource group) to complete the installation in your Microsoft Azure subscription. The boot node contains all the required information to install the Maximo Application Suite in Red Hat OpenShift cluster.
4. Bootnode creates required network components in Microsoft Azure using Solution template (only in IPI mode).
5. Bootnode starts Red Hat OpenShift cluster deployment that creates the bootstrap node. The bootstrap node uses the Red Hat OpenShift installer to create master and worker nodes. In case of IPI mode, the Red Hat OpenShift cluster resources are created in a new resource group that is created by Red Hat OpenShift installer itself. In case of UPI mode, the Red Hat OpenShift cluster resources are created in the same resource group where existing network infrastructure resides. The resource group where Red Hat OpenShift cluster resources are created is called as OCP cluster resource group.
6. Information about the installation context and the Terraform® state is saved to an Microsoft Azure Blob storage in your Microsoft Azure subscription. Bootstrap logs are streamed to Log Analytics workspace.
7. The boot node installs the Maximo Application Suite prerequisites and the application itself.
8. If you provide the SMTP configuration during the deployment, you receive emails that contain the information that you need to access Maximo Application Suite.
9. The cluster is ready, and you can access the Maximo Application Suite.
10. After the installation is complete, the boot node is moved to a shutdown state. If required, you can restart the boot node and access it by using Secure Shell (SSH).

Maximo Application Suite unique identifiers for Microsoft Azure

During a Maximo Application Suite installation, identifiers are created that uniquely identify the installation, the Red Hat OpenShift cluster, and the Maximo Application Suite instance.

When you install the Maximo Application Suite on Microsoft Azure, identifiers are created to uniquely identify components that are created during the installation. These identifiers are referred to throughout this documentation.

Installation identifier

When you install Maximo Application Suite, a unique 6-character string, for example `drq2wd`, is generated to identify the installation itself. This identifier is also included in the names of several other components that are created during the installation, such as the Red Hat OpenShift cluster.

In this documentation, `<unique-string>` refers to the installation identifier.

Cluster identifier

When the Red Hat OpenShift cluster is created during the installation, it is identified by the following string: `masocp-<unique-string>`. For example, if the installation identifier is `drq2wd`, the cluster identifier is `masocp-drq2wd`. This identifier is passed to the cluster installer and is included in the cluster URL that you can access after the installation is complete.

In this documentation, `<cluster-name>` refers to the cluster identifier.

Instance identifier

- When the Maximo Application Suite instance is created during the installation, it is identified by the `mas-<unique-string>` string.

For example, if the installation identifier is `drq2wd`, the instance identifier is `mas-drq2wd`.

- When the Maximo Application Suite instance is created during the installation, it is identified by the same unique string that is used by installation identifier `<unique-string>`.

For example, if the installation identifier is `drq2wd`, the Maximo Application Suite instance identifier is the same `drq2wd`.

In this documentation, `<instance-id>` refers to the instance identifier.

Prerequisites for installing Maximo Application Suite on Microsoft Azure

Before you install the IBM Maximo Application Suite, you must set up several components, and gather the information that you need to specify in the installation parameters.

The following components are required to install Maximo Application Suite in Microsoft Azure.

Microsoft Azure account and subscription

To install the Maximo Application Suite on Microsoft Azure, you must have a Microsoft Azure account and a subscription created in it. During Maximo Application Suite installation, a new virtual private cloud VNet is created in the Microsoft Azure subscription. The Red Hat OpenShift cluster is then deployed in the VNet, and the Maximo Application Suite is deployed in the cluster.

During IBM Maximo Application Suite installation, when you use the 'New Red Hat OpenShift cluster' option, a new virtual private cloud VNet is created in the Microsoft Azure subscription.

For more information, see [How do I create and activate a new Azure account?](#).

In your Microsoft Azure account, create active directory (AD) user and assign AD roles to this user that includes the required permissions to install the Maximo Application Suite. You need to assign the existing AD roles `Contributor` and `User Access Administrator` to the user. For more information, see [“Configuring installation permissions on Microsoft Azure” on page 174](#).

By default, your Microsoft Azure account has access to all of the Microsoft Azure services that you need to install the Maximo Application Suite. These services include Microsoft Azure DNS, Resource Manager Templates, virtual machines, Blob storage, Log Analytics, and other services. When you configure these services, ensure that you select the same geographical region where you want to install the Maximo Application Suite. For the list of supported regions, see [“Preparing to installing Maximo Application Suite on Microsoft Azure” on page 167](#).

Your Microsoft Azure subscription includes default service limits, quotas, and constraints for each Microsoft Azure service. In the future, you might need to increase these quotas, for example if your Red Hat OpenShift cluster requires more resources. For more information, see [Azure service limits, quotas, and constraints](#) in the Microsoft Azure documentation.

SSH key pair

The Microsoft Azure virtual machines service uses virtual computing environments to provide computing capacity in the Microsoft Azure cloud. In an installed Maximo Application Suite on Microsoft Azure, the Red Hat OpenShift cluster nodes are Microsoft Azure virtual machines.

For more information, see [Virtual machines in Azure](#).

When a virtual machine is created, a key pair, which consists of a public key and a private key, must be provided. Because the Maximo Application Suite installation creates many virtual machines, you must generate a key pair before you install the Maximo Application Suite. For more information, see [Create a key pair using a third-party tool](#) in the Microsoft Azure documentation.

When you specify the parameters for a Maximo Application Suite installation on Microsoft Azure, in the SSHKey parameter, contain the public key that you generated. Keep the private key in a safe place and store it in the PEM format.

After the installation is complete, you can use the private key to connect to the cluster nodes by using Secure Shell (SSH). For more information, see [“Accessing the boot node and Red Hat OpenShift cluster” on page 274](#).

Boot node CIDR IP address range

When you begin a Maximo Application Suite installation, a boot node is created that contains all of the required information, including the parameters that you specify, to complete the installation. The boot node creates and installs the VNet, Red Hat OpenShift cluster, Maximo Application Suite prerequisites, and Maximo Application Suite itself.

To connect to the boot node, you must specify a range of IPv4 addresses that can access it. In the bootnodeSGIngressCidrIp installation parameter, enter the IP address range by using the classless inter-domain routing (CIDR) notation. For example, to allow all IP addresses to access the boot node, enter 0.0.0.0/0. To allow one specific IP address to access the boot node, enter x.x.x.x/32, for example 192.12.33.3/32. For more information, see [Classless_Inter-Domain_Routing#CIDR_notation](#).

Maximo Application Suite license

The Maximo Application Suite license needs to be retrieved from the IBM License Key Center.

If you want to purchase the Maximo Application Suite license through Microsoft Azure Marketplace, subscribe to the IBM Maximo Application Suite (Paid) product first. The subscription to this process initiates the purchase process with IBM through Microsoft Azure Marketplace. You can select either public offer or a private offer with negotiated pricing.

If you want to purchase the Maximo Application Suite Maximo Application Suite license outside of the Microsoft Azure Marketplace, contact your IBM sales representative.

If you do not already have your Maximo Application Suite license key file, you can create and download it in the [IBM License Key Center](#). For instructions, see the steps in the [Upload your license key file](#) section of the [Setting up Maximo Application Suite](#) topic. When you create the license, provide the server parameters that are indicated in the following table:

Parameter	Value
Configuration	Single license Server

Parameter	Value
Host ID Type	Ethernet address
Host ID	A unique 12-character hexadecimal value, such as `0abcac110f02`. You can generate this value by using one of the following methods:

- An online hexadecimal number generator, such as [Online Hex Tools](#).
- A command-line utility, such as the Linux hexdump utility. For example, the following command generates a lowercase 12-character hexadecimal string:

```
hexdump -n 6 -e '8/8 "%08X" 1 "\n" /dev/urandom | tr '[:upper:]' '[:lower:]'
```

After you download the Maximo Application Suite license key file, upload it to a Blob storage. Your Microsoft Azure account has access to the Microsoft Azure Blob storage. For more information about using Microsoft Azure Blob storage, see [Introduction to Azure Blob storage](#).

After you upload the license to the Blob storage container, create a shared access token in the container with appropriate expiry time. This step provides you a URL to the storage container that has a SAS token as part of it. The URL works only while it is valid. Complete the deployment before the SAS token expires. Record the HTTPS location of the license, for example `https://masocpstgacnt.blob.core.windows.net/masocpfiles/entitlement.lic?sp=r&st=2022-04-06T04:02:45Z&se=2022-06-30T12:02:45Z&spr=https&sv=2020-08-04&sr=c&sig=CN27jhRfxHDmDgz%2FYgkyGY7h%2BEZdp9H5PVAoaxP%2FURY%3D`. You enter this location when you specify the `masLicenseUrl` installation parameter.

IBM Entitled Registry

The IBM Entitled Registry key is used during the installation to download the container images for the Maximo Application Suite and its applications from the IBM Entitled Registry. Download this key from the [IBM Container Library](#).

You must purchase a Maximo Application Suite license to receive a valid entitled registry key to download the container images for the Maximo Application Suite.

Pull secret

To create the Red Hat OpenShift cluster, you must provide a pull secret. To access the pull secret, complete the following steps:

1. Ensure that you have a Red Hat account. If you do not have a Red Hat account, [create one](#).
2. To access your Red Hat entitlements, in the Passport Advantage website, link your purchased Maximo Application Suite software to your Red Hat account. For more information about how to link purchased IBM software to a Red Hat account, see [this Cloud Paks topic](#).
3. Log in to the Red Hat Hybrid Cloud Console by using your Red Hat account credentials.
4. In the **OpenShift > Clusters** page, click **Download Pull Secret**.

The pull secret is a JSON-formatted text file. When you configure the Maximo Application Suite installation parameters, copy the JSON text from the file into the `OpenShiftPullSecret` parameter.

Storage account with Blob storage container

You must provide certain files when you deploy the Maximo Application Suite

1. Maximo Application Suite license file for the BYOL product.
2. Maximo Application Suite Manage Db2 certificate if Maximo Application Suite and Manage are deployed.
3. IBM Suite License Service (SLS) public certificate if you are using existing SLS.
4. Data Reporter Operator (DRO) public certificate if you are using existing DRO.

The files must be kept in a Blob storage container that is created in a storage account. The shared access token used to generate the shared access signature (SAS) URI with at the minimum read access to these files must be provided.

The SAS URI of each file must be provided during the deployment as applicable.

Existing network infrastructure by using Red Hat OpenShift UPI mode

To reuse the existing network infrastructure, select the New OpenShift cluster, existing network (UPI) option when you deploy the Maximo Application Suite. This mode requires certain Microsoft Azure resources to be created in advance.

1. Generate an infrastructure ID.

Define an infrastructure ID to be used as the base name for the resources. It can be any unique string with alphanumeric characters.

Important: The infrastructure ID must use some standard for this string and the first character should be a letter. It can be a random alphanumeric string of 6 characters.

You can use the following command in bash shell to generate a string.

```
INFRA_ID="$(cat /dev/urandom | tr -dc 'a-z' | fold -w 3 | head -n 1)$(cat /dev/urandom | tr -dc '0-9' | fold -w 3 | head -n 1)"
echo $INFRA_ID
```

2. Provide specifications for existing network resources.

Reuse the existing network resources such as VNet, Subnet, and Network security group.

You must follow specific naming convention for these resources. The resource name must begin with INFRA_ID. If the naming conventions are not followed, the deployment might be unsuccessful.

- a. In a browser window, open the `json` file, right-click the page and save the file to your local machine by using the name `01_vnet.json`.

Starting in Maximo Application Suite 8.10 or later, open the `01_vnet.json` file.

Tip: You can edit the `json` file with your own Classless Inter-Domain Routing (CIDR) values for VNet `addressPrefix`, master subnet `masterSubnetPrefix`, and worker subnet `nodeSubnetPrefix`. However, ensure that you do not change the values for master subnet `masterSubnetName` and worker subnet `nodeSubnetName`.

If you are on Maximo Application Suite 8.9 or earlier, open the `01_vnet.json` file.

- b. Create the resource group that contains the network resources. In the command, the resource group name is defined as a variable.

```
RESOURCE_GROUP=$INFRA_ID-vnet-rg
az group create -l eastus2 -n $RESOURCE_GROUP
```

Note: You can use other naming conventions that are appropriate to your environment. You can also change the region as needed.

- c. To create the resource by using a JSON, run the following command.

```
az deployment group create -g $RESOURCE_GROUP --template-file "01_vnet.json" --
parameters baseName="$INFRA_ID"
```

Next steps

Consider your installation preferences, such as capacity planning and the type of Maximo Application Suite offering that you want. For more information, see [“Preparing to installing Maximo Application Suite on Microsoft Azure” on page 167](#).

Preparing to installing Maximo Application Suite on Microsoft Azure

Before you install IBM Maximo Application Suite, consider your installation preferences, such as the type of Maximo Application Suite offering that you prefer and whether you want to reuse existing Maximo Application Suite compatible components in your new application instance. You can consider reusing your existing network infrastructure before you install Maximo Application Suite.

You must also configure several prerequisite components before you install the Suite. For more information, see [Installation requirements](#).

When you install the Suite, you must specify options for the following items:

- [Installation region](#)
- [Maximo Application Suite offering type](#)

If this installation is your first Maximo Application Suite installation on Azure, a new Red Hat OpenShift cluster must be created. If you already installed an instance of the Suite on Azure, you can reuse the existing cluster.

Note: The deployment is supported on existing Red Hat OpenShift cluster that is either created by using the automation that is provided by IBM for the Maximo Application Suite Marketplace product or any pre-provisioned Red Hat OpenShift cluster. In case the pre-provisioned Red Hat OpenShift cluster does not meet the requirements, the deployment might fail and you are recommended to create a new cluster by using the automation provided by IBM for the Maximo Application Suite Marketplace product.

If you want the installation process to create a new cluster, you must configure and consider the following items:

- [App Service Domain](#)
- [Cluster size](#)
- [Public DNS](#)
- [Private DNS](#)

The private DNS configuration is available from Maximo Application Suite 8.10 for an existing private infrastructure.

Alternatively, if you want to reuse an existing Red Hat OpenShift cluster, or reuse cluster components, such as the Suite License Service (SLS) or Data Reporter Operator (DRO), gather the following details:

- [Connection details for an existing Red Hat OpenShift cluster](#)
- [Connection details for an existing Suite License Service instance](#)
- [Connection details for a DRO instance](#)

Depending on your preferences, you can also set up and gather the following items of information:

- [Email notifications](#)
- [Database configuration details for Maximo Manage](#)

Microsoft Azure region for installing Maximo Application Suite

In Microsoft Azure Marketplace, you must select the geographical region where you want to install the IBM Maximo Application Suite.

To ensure that the installation succeeds, select one of the following supported regions:

Region name	Region endpoint identifier
(US) East US	eastus
(US) East US 2	eastus2
(US) South Central US	southcentralus

Region name	Region endpoint identifier
(US) West US 2	westus2
(US) West US 3	westus3
(Asia Pacific) Australia East	australiaeast
(Asia Pacific) Southeast Asia	southeastasia
(Europe) North Europe	northeurope
(Europe) Sweden Central	swedencentral
(Europe) UK South	uksouth
(Europe) West Europe	westeurope
(US) Central US	centralus
(Africa) South Africa North	southafricanorth
(Asia Pacific) Central India	centralindia
(Asia Pacific) East Asia	eastasia
(Asia Pacific) Japan East	japaneast
(Asia Pacific) Korea South	koreacentral
(Canada) Canada East	canadacentral
(Europe) France Central	francecentral
(Europe) Germany West Central	germanywestcentral
(Europe) Norway East	norwayeast
(South America) Brazil South	brazilsouth

When you set up the Azure services that you need to configure the installation prerequisites, ensure that you use the same region where you want to install the Suite.

Maximo Application Suite offering type

When you specify the parameters for a Maximo Application Suite installation on Microsoft Azure, you can choose from the following Maximo Application Suite offering types:

1. Maximo Application Suite core and Cloud Pak for Data
 - Db2 Warehouse and Db2 Data Management Console are the only Cloud Pak for Data services that are installed.
 - You install Maximo Application Suite applications manually after installation.
 - You configure the prerequisites for any Maximo Application Suite applications that you install.
2. Maximo Application Suite core and Maximo Manage
 - Cloud Pak for Data is not included in this offering type.
 - For IBM Maximo Manage, you configure databases such as IBM Db2, Microsoft SQL Server, or Oracle Database.

Starting in 8.11, these databases can be hosted on private IPs of different Vnet.

App service domain

For IBM Maximo Application Suite installations, if you want to create a new Red Hat OpenShift cluster, you must configure an App Service Domain in Microsoft Azure. App Service Domain is a container that holds information about how you want to route internet traffic for a specific domain.

When the Maximo Application Suite cluster is created, the App Service Domain is used to allow internet traffic to access the cluster from outside the cluster's virtual network (VNet). For more information, see [App Service](#).

You configure the public domain in the Microsoft Azure App Service Domain service. The zone is created in Microsoft Azure DNS zones service. Microsoft Azure App Service, along with Microsoft Azure DNS zones, is a domain name system (DNS) service that provides domain registration, routing, and health-checking functions. You can register any valid and available domain name. Your Microsoft Azure account has access to the App Service Domain service. For the instructions to, create a domain in Microsoft Azure App Service, see [Buy a custom domain name for Azure App Service](#).

Red Hat OpenShift cluster size

When you specify the installation parameters for IBM Maximo Application Suite, if you want to create a new Red Hat OpenShift cluster, you must choose the size of the cluster that you want to create. Choose a cluster size that is appropriate for the scale of your Suite installation. Consider the Suite applications that you might want to deploy after the installation is complete.

You can choose to create a small, medium, or large cluster. By default, the cluster that you create spans multiple availability zones. For the node, CPU, and memory dimensions for each cluster size, see the following table:

Cluster size	Node type	Number of nodes	Virtual machine size	CPUs per node	Total CPUs	Memory per node (GB)	Total memory (GB)	VM OS disk per node (GB)	Disk storage per node (GB)
Small	Master	3	Standard_D8s_v3	8	24	32	96	256	0
	Worker	3	Standard_D16s_v3	16	48	64	192	256	100 (Approx.)
	Boot node	1	Standard_D2s_v3	2	2	8	8	64	64
Medium	Master	3	Standard_D8s_v3	8	24	32	96	256	0
	Worker	5	Standard_D16s_v3	16	80	64	320	256	100 (Approx.)
	Boot node	1	Standard_D2s_v3	2	2	8	8	64	64
Large	Master	5	Standard_D8s_v3	8	40	32	160	256	0

Cluster size	Node type	Number of nodes	Virtual machine size	CPUs per node	Total CPUs	Memory per node (GB)	Total memory (GB)	VM OS disk per node (GB)	Disk storage per node (GB)
	Worker	7	Standard_D16s_v3	16	112	64	448	256	100 (Approx.)
	Boot node	1	Standard_D2s_v3	2	2	8	8	64	64

The boot node is used only during the installation and shutdown after the installation is complete. It is a workstation from which the Red Hat OpenShift cluster deployment, Maximo Application Suite is deployed. Hence, it can be used later to access the deployment files, configurations, logs, or to run the ad hoc actions on the cluster as needed.

For more information, see [“Requirements and capacity planning” on page 178](#).

Microsoft Azure DNS zones

To create a Red Hat OpenShift cluster, you can select either a public hosted or private hosted zone in Microsoft Azure.

Public hosted zone

For IBM Maximo Application Suite installations, if you want to create a new Red Hat OpenShift cluster, you must configure a public hosted zone. Microsoft Azure DNS is a hosting service for DNS domains that provides name resolution by using Microsoft Azure infrastructure.

When the Maximo Application Suite cluster is created, the public hosted zone is used to allow internet traffic to access the cluster from outside the cluster's virtual private cloud (VPC).

Private hosted zone

For Maximo Application Suite installations, if you want to create a new private Red Hat OpenShift cluster, you must configure a private hosted zone. Microsoft Azure private DNS is a hosting service for DNS domains that provides name resolution by using Microsoft Azure infrastructure.

When the Maximo Application Suite cluster is created, the private DNS zone does not allow internet traffic to access the cluster from outside unless accessed from Jumphost. The Jumphost must be present in the same Vnet of Red Hat OpenShift cluster.

For more information about creating public or private hosted zones, see [How to create a DNS zone](#).

Connection details

Check connection details for existing Red Hat OpenShift cluster, network infrastructure, IBM Suite License Service instance, and IBM Data Reporter Operator instance.

Connection details for an existing Red Hat OpenShift cluster

If you want to reuse an existing Red Hat OpenShift or Microsoft Azure Red Hat OpenShift cluster, verify that it has enough resources to accommodate the new IBM Maximo Application Suite instance.

Note: For details on Maximo Application Suite offerings and supported Red Hat OpenShift versions, see [Software Product Compatibility Reports \(SPCR\)](#).

To reuse the cluster, enter the following installation parameters:

- The cluster's API URL in the format `https://api.<cluster-name>.<domain_name>`. Do not specify the port number. For example, `https://api.masocp-joalae.mas4multucloud.com`
 - For more information on the cluster name's format, see [Unique identifiers](#).
- Username
- Password
 - If the cluster was created during a previous Maximo Application Suite installation on Microsoft Azure and you received the connection credentials in an email, retrieve the username and password from the email.
 - If you did not receive an email, the Red Hat OpenShift Container Platform cluster connection details can be obtained from the **Outputs** section of the deployment that is created in the boot node resource group. The Red Hat OpenShift Container Platform cluster credentials can be obtained from the Microsoft Azure Vault with name `masocp-<unique-string>`.
- Starting in 8.11, VNetId of DB Provisioned
 - Provide the VnetId of database provisioned
- Starting in 8.11, VnetId of existing Red Hat OpenShift cluster
 - Provide the existing Red Hat OpenShift cluster's Vnet ID to establish the Vnet peering from this Vnet and the Vnet where the existing database virtual machine is created to establish database connection.

If you leave these parameters empty, a new Red Hat OpenShift cluster is created during the installation.

When you reuse an existing Red Hat OpenShift cluster, a prerequisites check is executed in the cluster for following components:

- Cloud Pak for Data, if you selected the Maximo Application Suite core and Cloud Pak for Data offering type.
- SLS, if you do not specify an existing SLS instance to reuse.
- Starting in Maximo Application Suite 9.0, IBM Data Reporter Operator (DRO)

Tip: If you are using IBM User Data Services, you must migrate to Data Reporter Operator . For information, see [“Migrating Maximo Application Suite from User Data Services to Data Reporter Operator ” on page 8.](#)

- MongoDB
- cert-manager

The components are reused if they are installed in the cluster with supported version. Else, new instances of the components are created in the cluster.

The deployment is terminated if an unsupported version is pre-installed for any of the components or if there are multiple instances of a component running in existing cluster.

The prerequisites check will also verify if the following storage classes are available in the cluster:

- managed-premium
- azurefiles-premium, if you selected Maximo Application Suite core and Cloud Pak for Data.

The following components are required to ensure the cluster has enough resources to accommodate the new Maximo Application Suite instance.

- Minimum number of worker nodes: 3
- Minimum CPUs per node: 8 cores
- Minimum memory per node (GB): 32

Connection details for an existing Suite License Service instance

If you previously installed Suite-compatible components in your Microsoft Azure account, you can reuse them when you install your new Suite instance.

To reuse an existing Suite License Service (SLS) instance, you can enter the following installation parameters.

- The endpoint URL.
 - In Red Hat OpenShift, in the namespace where SLS is installed, retrieve this URL from the `sls-*` route. This URL must be accessible from the new Suite instance.
- The registration key.
 - In the Red Hat OpenShift web console, search for resources that have a resource type of `LicenseService`.
 - In the result list, click the SLS instance that you want to reuse. SLS instances that are created by Suite installations include `<unique-string>` in their instance and namespace names.
 - In the **YAML** page, retrieve the registration key from the `status.registrationKey` field.
- HTTP location of the service's public certificate along with SAS token, for example `https://masocpstgacnt.blob.core.windows.net/masocpfiles/sls-certificate.crt?sp=r&st=2022-04-06T04:02:45Z&se=2022-06-30T12:02:45Z&spr=https&sv=2020-08-04&sr=c&sig=CN27jhRfxHDmDgz%2FYgkyGY7h%2BEZdp9H5PVAoaxP%2FURY%3D`.
 - In the existing Suite instance that uses the SLS, in the `mas-<instance-id>-core` namespace, in the `<instance-id>-sls-cfg` secret, retrieve the certificate from the `ca.crt` file.
 - Upload the certificate to your Blob storage container and record the certificate's HTTP location.

If you leave these parameters empty, a new SLS instance is created during the installation.

Data Reporter Operator connection details

To use a Data Reporter Operator instance, you can enter the following installation parameters:

Endpoint URL

On the **Routes details** page in the `redhat-marketplace` project, you can find the location of **ibm-data-reporter**.

API key

On the **Workloads > Secrets** page of the `mas-<instanceId>-core` project, you can find the **dro-apikey**.

HTTP location of the service's public certificate along with SAS token

To find the HTTP location of the service's public certificate with the SAS token such as `https://masocpstgacnt.blob.core.windows.net/masocpfiles/dro-certificate.crt?sp=r&st=2022-04-06T04:02:45Z&se=2022-06-30T12:02:45Z&spr=https&sv=2020-08-04&sr=c&sig=CN27jhRfxHDmDgz%2FYgkyGY7h%2BEZdp9H5PVAoaxP%2FURY%3D`, follow the steps.

- In the existing Maximo Application Suite instance that uses the Data Reporter Operator, in the `mas-<instance-id>-core` namespace, in the `<instance-id>-dro-cfg` secret, retrieve the certificate from the `ca-bundle.pem` file.
- Upload the certificate to your Blob storage container and record the certificate's HTTP location.

If you leave these parameters empty, a new Data Reporter Operator instance is created during the installation.

Related tasks

[Migrating Maximo Application Suite from User Data Services to Data Reporter Operator](#)

As an IBM Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance. New and existing Maximo Application Suite users can install or migrate to DRO by using the IBM Maximo Application Suite command line interface (CLI), ansible role, or manually.

Email notifications

Maximo Application Suite deployment on Azure uses the external SMTP server details that are provided by the user to send the email notifications. The setup, configuration, security of the external SMTP server must be managed by the user.

When you specify the installation parameters for the Suite, set the `EmailNotification` parameter to `true`. After the installation is complete, each verified SES address receives two emails that contain the details to access the Suite.

The first email contains connection URLs and usernames for the cluster and Suite in the following format:

- Maximo Application Suite provisioning status: SUCCESS
- Region: <region-code>
- Unique String: <unique-string>
- Red Hat OpenShift cluster URL: <ocp-cluster-url>
- Red Hat OpenShift API URL: <ocp-api-url>
- Red Hat OpenShift User: <ocp-username>
- SLS Endpoint URL: <sls-url>
- DRO Endpoint URL: <dro-url>
- Maximo Application Suite Initial Setup URL: <setup-url>
- Maximo Application Suite Admin URL: <admin-url>
- Maximo Application Suite Workspace URL: <workspace-url>
- Maximo Application Suite User: <mas-username>

The second email contains password information in the following format:

- Maximo Application Suite provisioning status: SUCCESS
- Region: <region-code>
- Unique String: <unique-string>
- Red Hat OpenShift Password: <ocp-password>
- Maximo Application Suite Password: <mas-password>

Database configuration details for Maximo Manage

If your Maximo Application Suite offering includes the Maximo Manage application, you must configure a database that this application can use. Maximo Manage supports IBM Db2, IBM Db2 Warehouse, Microsoft SQL Server, and Oracle Database.

Note: If you choose to install Maximo Application Suite with Manage, you can use the default IBM Db2 instance that is provisioned by IBM instead of configuring your own external Db2 instance. To configure the default IBM Db2, do not add information in the username, password, JDBC URL, certificate URL, and demo data fields.

The internal Db2 configuration is available from Maximo Application Suite 8.10 or later.

Restriction: The default IBM Db2 instance for configuring Internal Db2 is unavailable for Microsoft Azure Red Hat OpenShift.

For more information, see [“Preparing your database for deployment” on page 301](#).

After you configure the database that you want Maximo Manage to use, you can specify the following configuration details when you enter the Maximo Application Suite installation parameters:

- Username
- Password
- Java database connectivity (JDBC) URL,

For example, use the following parameters for the database that is configured.

– For IBM Db2 database:

```
jdbc:db2://1.2.3.4:50051/FTMDB:sslConnection=true
```

– Starting in 8.11, for Microsoft SQL Server database:

```
jdbc:sqlserver://;serverName=10.5.0.4;portNumber=1433;databaseName=maxdb80;  
sendStringParametersAsUnicode=false;selectMethod=cursor;encrypt=true;trustServerCertificate=  
true;
```

– Starting in 8.11, for Oracle Database:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=63.246.112.120)  
(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=maxdb)))
```

Note: Only secure sockets layer (SSL) connections to the database are supported.

- HTTP location of the database's public certificate with SAS token, for example: `https://masocpstgacnt.blob.core.windows.net/masocpfiles/db2-certificate.crt?sp=r&st=2022-04-06T04:02:45Z&se=2022-06-30T12:02:45Z&spr=https&sv=2020-08-04&sr=c&sig=CN27jhRfxHDmDgz%2FugkyFT7h%2BEZdp9H5fVgoaxP%2FURY%3D`.

Ensure that you can retrieve your certificate based on your database type. For example, for IBM Db2 Warehouse, go to the **Details** view of the IBM Db2 instance in Cloud Pak for Data, and click **Download SSL certificate**.

In your Microsoft Azure account, upload the certificate to a storage account Blob container. Create the SAS token and form the HTTP location for the certificate file.

Your Microsoft Azure account has access to the Microsoft Azure storage account service. If you do not yet have any storage accounts created, first sign in to the Microsoft Azure portal and open the **Storage accounts** service. Then, click **Create** and follow the steps in the wizard. If you already created a storage account, in the storage accounts list, choose the name of the storage account that you want to upload your Maximo Application Suite license key file to, create a container, and then upload the files to the container.

- Whether you want to import demo data into the database. This data might be useful for development or test environments.

Ensure that you are aware of the location site for the database, depending on where the database was provisioned compared to where the Cloud Pak for Data server is located. If the locations are far apart, the maxinst/updatedb operations might take several days or fail due to connection issues, especially in cases where you install several Manage extensions.

If you deploy the Maximo Manage application, you can add or [update the database configuration information](#).

Configuring installation permissions on Microsoft Azure

To ensure that you can install Maximo Application Suite on Azure, create an Active Directory user in your Azure account and grant the required permissions to this user. Then, this user can be used to create a service principal and you can deploy Maximo Application Suite.

Procedure

1. Create an Active Directory user with following Azure roles at the subscription level.
 - Contributor
 - User Access Administrator
2. Create the Azure Service principal that is used by Red Hat OpenShift installer.
 - a. Log in to the Azure that uses az cli:

```
az login -u "<username>" -p "<password>"
```

b. Create service principal with Contributor role:

```
az ad sp create-for-ibac --role="Contributor" --name="<any-name>" --scopes="/subscriptions/<subscription-id>"
```

Note the value of `appId`, `password` and `tenant` attributes.

c. Get the service principal's object ID:

```
az ad sp list --filter "appId eq '<appId-from-previous-output>'"
```

Note the value of `objectId` attribute.

d. Assign User Access Administrator role to the service principal.

```
az role assignment create --role "User Access Administrator" --assignee-object-id "<objectId-from-previous-output>" --assignee-principal-type ServicePrincipal
```

You now have the following details that are used during the deployment.

Subscription ID

The subscription ID is used in the command to create the service principal.

Tenant ID

The value of `tenant` attribute in the service principal command output.

Client ID

The value of `appId` attribute in the service principal command output.

Client secret

The value of `password` attribute in the service principal command output.

Results

During a Suite installation on Azure, the AD user who is logged in to the marketplace must have the same permissions as the user that created the service principal. At a minimum, ensure that user has the following roles at a subscription level.

- Contributor
- User Access Administrator

What to do next

The deployment from the marketplace happens in following two steps.

1. Boot node and few other resources are created by using the ARM template with current user's permissions.
2. Red Hat OpenShift cluster IPI infrastructure and resources are created with the service principal's permissions.

The ARM template is instantiated under this user's permissions.

Security considerations for Microsoft Azure

Before you can install IBM Maximo Application Suite, you must ensure that your environment meets all security requirements.

Maximo Application Suite deployment complies with the following security considerations:

- Communication to the IBM Maximo Manage database uses JDBC with SSL enabled.

- SSH keys used for the connection to the bootnode and Red Hat OpenShift cluster nodes (master/worker/infra).
- Bootnode runs within the customer's Microsoft Azure account and does not have connectivity to the external network during and post deployment.
- Product images are pulled from authenticated IBM entitled registries.
For more information, see [“Networking considerations” on page 182](#).
- Credentials are kept in Red Hat OpenShift secrets.
- Access to the Red Hat OpenShift cluster nodes is only through the bastion host using private SSH key.
- Microsoft Azure portal uses HTTPS (SSL/TLS certificates) for encryption.

What to do next

“IBM Maximo Application Suite installation with Microsoft Azure Resource Manager templates” on page 262

You can install Maximo Application Suite in the Microsoft Azure cloud by using the Microsoft Azure Resource Manager templates. In Microsoft Azure Marketplace, you subscribe to Maximo Application Suite, configure the installation parameters, and install the application. The network infrastructure, Red Hat OpenShift cluster, and Maximo Application Suite components are created in your Microsoft Azure cloud account.

Planning for on premises

Before you install Maximo Application Suite, configure installation requirements and consider installation preferences, such as whether you want to create a Red Hat OpenShift cluster or reuse an existing one.

Preparing to install Maximo Application Suite on premises

Review the minimal resource requirements for Red Hat OpenShift Container Platform nodes and size the Red Hat OpenShift Container Platform compute nodes before you install IBM Maximo Application Suite on-premises.

Sizing Red Hat OpenShift Container Platform compute nodes

For most Maximo Application Suite workloads, size Red Hat OpenShift Container Platform (OCP) compute nodes at 4 GB memory per CPU. The number of CPU per compute node varies based on the following factors:

- The number of Maximo Application Suite applications and add-ons deployed.
- The size of the Maximo Application Suite workload.
- The degree of availability during a compute node outage.

Smaller workloads favor smaller compute nodes, which can still provide availability during of a compute node outage. Take, for example, the following OCP clusters.

1. cluster1 with 3 compute nodes, each with 16 CPU and 64 GB of memory
2. cluster2 with 6 compute nodes, each with 8 CPU and 32 GB memory

Both clusters have the same compute and memory resources allocations. During a compute node outage, cluster1 loses a third of its capacity, while cluster2 loses a sixth of its capacity. Conversely, if the workload is so large that it requires many tens of 8x32 compute nodes in cluster2, then provision fewer larger-sized computes nodes and minimize the inter-node I/O requirements. To determine the optimal compute node size, consider both factors.

Another rule to apply when you size worker nodes is to allocate 15 GB - 25 GB of disk storage per CPU allocated to the compute nodes. If insufficient disk space is allocated to the compute nodes, then pod evictions due to disk pressure are observed as the pod density per compute node increases.

The following table illustrates some examples.

Table 17. Sizing OCP compute nodes

Workload	CPU cores	Memory (GB)	Disk (GB)
Small	4	16	100
Medium	8	32	200
Large	16	64	400
Extra Large	32	128	800

Prerequisites for installing Maximo Application Suite on premises

Configure your IBM Cloud account permissions, install some tools in a client machine to run command-line instructions, and install the IBM Maximo Application Suite prerequisites and dependencies.

Enabling the Red Hat OpenShift internal image registry

The Image Registry Operator installs a single instance of the Red Hat OpenShift Container Platform registry, and manages all registry configuration, including setting up registry storage. Storage is only automatically configured when you install an installer-provisioned infrastructure cluster on Amazon Web Services, GCP, and Microsoft Azure, and thus you can use the internal image registry directly. But for bare metal and vSphere Red Hat OpenShift cluster, you must manually enable the Red Hat OpenShift internal image registry so that it is used by Maximo Application Suite and its applications.

You can enable the Red Hat OpenShift internal image registry by using the existing Red Hat OpenShift Container Storage.

Related information

[Configuring registry](#)

Enabling the Red Hat OpenShift internal image registry by using the existing Red Hat OpenShift Container Storage

If you already have Red Hat OpenShift Container Storage running inside your Red Hat OpenShift Container Platform, you can run the following steps in the bastion server and use OCS storage for the internal image registry.

Procedure

1. Create PVC for internal Registry use.

```
oc project openshift-image-registry
```

```
oc create -f <(echo '{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "name": "image-registry-storage"
  },
  "spec": {
    "storageClassName": "ocs-storagecluster-cephfs",
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": { "storage": "500Gi"
    }
  }
}'
```

2. Update the Registry CR spec with the following command.

```
oc edit configs.imageregistry.operator.openshift.io -n openshift-image-registry
```

3. Change `spec.managementState` from Removed to Managed.
4. Change `spec.storage` from `{}` to:

```
spec:
  managementState: Managed
storage:
  pvc:
    claim: image-registry-storage
```

5. Save and quit.
6. Check that the image registry is available.

```
oc get co image-registry
```

7. Enable the external route.

```
oc -n openshift-image-registry patch configs.imageregistry.operator.openshift.io/cluster
--patch '{"spec":{"defaultRoute":true}}' --type=merge
```

Requirements and capacity planning

Use the IBM Maximo Application Suite sizing calculator to estimate the required sizing for your planned deployment.

The Maximo Application Suite sizing calculator is used to estimate your Red Hat OpenShift worker node configuration requirements, storage requirements, and memory requirements.

1. Download the calculator.
 - [Sizing calculator for Maximo Real Estate and Facilities 9.1](#)
 - [Sizing calculator for 9.0.1](#)
 - [Sizing calculator for 9.0](#)
 - [Sizing calculator for 8.11 and earlier](#)
2. Select or enter values for the yellow fields to match your planned application deployment.

The calculator provides estimated total system requirements in VPCs and Memory (GB) for your configuration in the Resulting Complete Environments Requirements section of the Output table.

Important: The information in this document represents the minimum resources that you need to successfully install Maximo Application Suite. A minimum of 300GB of storage per worker node is recommended for Maximo Application Suite build process. You might need more resources to support your specific workload. If needed, work with your IBM Sales representative to generate more accurate calculations based on your expected workload.

For more information, see [Sizing guidance](#).

Related concepts

[System requirements](#)

Ensure that you understand the system requirements for your Maximo Application Suite installation. System requirements differ depending on the applications that you plan to deploy, and the size requirements of these applications.

Supported software versions

Use the Software Product Compatibility Reports (SPCR) prerequisites tab to understand the architecture, containers, prerequisite versions, supported software, and hardware for Maximo Application Suite and its components.

Running the Software Product Compatibility Reports (SPCR)

1. Go to the [Software Product Compatibility Reports](#) page.
2. From the menu, select the type of report to create, for example **Detailed system requirements**.
3. Search for Maximo Application Suite, then select the product version for which to create a report.
4. Click **Submit** to create the report.

Tip: You can download the created SPCR report as a PDF, or bookmark a permanent link to the page for reference.

Example

For example, if you request the SPCR report for Maximo Application Suite version 8.7, the following link is generated: [IBM Maximo Application Suite 8.7 Software Product Compatibility Report](#).

Verify software entitlement

Before you install Maximo Application Suite, verify that you are entitled to the container software.

Obtaining Red Hat OpenShift Container Platform entitlements

Maximo Application Suite entitlement linking to Red Hat OpenShift Container Platform

Maximo Application Suite includes entitlement to use Red Hat OpenShift Container Platform, Red Hat Enterprise Linux CoreOS (RHCOS), and Red Hat Enterprise Linux. Red Hat Enterprise Linux is an operating system. RHCOS is a version of Red Hat Enterprise Linux that is included in Red Hat OpenShift Container Platform.

To access these entitlements, you must link your Maximo Application Suite to your Red Hat OpenShift account.

You can link your Maximo Application Suite to its Red Hat OpenShift entitlement through IBM Passport Advantage.

For more information, see [“Accessing Red Hat OpenShift entitlements”](#) on page 181.

Maximo Application Suite reserved entitlement linking to BYOL Red Hat OpenShift Container Platform or Hyperscaler Red Hat OpenShift services

If you are using the Maximo Application Suite Reserved entitlements and have bring-your-own-license (BYOL) Red Hat OpenShift Container Platform or hyperscaler Red Hat OpenShift services such as Amazon Web Services or Microsoft Azure, you must use your own agreements for the Red Hat OpenShift versions that the Maximo Application Suite supports.

For more information about supported Red Hat OpenShift versions, see [Software Product Compatibility Reports \(SPCR\)](#).

Note: Maximo Application Suite reserved entitlements are a special license type that requires a BYOL Red Hat OpenShift entitlement or an equivalent service. You must obtain suitable entitlements using other routes and Maximo Application Suite must not be linked to the Red Hat OpenShift account through IBM Passport Advantage.

IBM Storage Fusion Advanced and Data Foundations

If you are using Maximo Application Suite 8.11, the use of IBM Storage Fusion Advanced is permitted for Data Foundations component under the limitations stated in the license document.

Note: IBM Storage Fusion Advanced was previously known as IBM Spectrum Fusion. Data Foundations was previously known as Red Hat OpenShift Data Foundation Essentials or Red Hat OpenShift Container Storage.

If you are using Maximo Application Suite 8.9 or earlier, the use of IBM Spectrum Fusion and Red Hat OpenShift Data Foundation Essentials, previously known as Red Hat OpenShift Container Storage, in

Maximo Application Suite is limited only with Bring-your-own-license (BYOL) hosting on Amazon Web Services (AWS).

Obtaining no charge storage entitlements



Attention: In Maximo Application Suite 8.10, the use of IBM Spectrum Fusion and Red Hat OpenShift Data Foundation Essentials in Maximo Application Suite requires obtaining a no charge storage entitlement.

Maximo Application Suite customers can obtain IBM Spectrum Fusion and Red Hat OpenShift Data Foundation Essentials storage entitlements at no charge.

Storage	Entitlement terms
Red Hat OpenShift Data Foundation Essentials	<p>Red Hat OpenShift Data Foundation Essentials entitlement applies only to self-managed Red Hat OpenShift.</p> <p>You are entitled to use Red Hat OpenShift Data Foundation Essentials with the following limitations:</p> <ul style="list-style-type: none">• You can use up to 6 TB of Red Hat OpenShift Data Foundation Essentials storage.• You can use Red Hat OpenShift Data Foundation Essentials for up to 12 months. <p>If you exceed these terms, a separate license is required.</p> <p>Contact your IBM Sales representative for access to this storage.</p>
IBM Spectrum Fusion	<p>IBM Spectrum Fusion entitlement applies only to self-managed Red Hat OpenShift.</p> <p>You are entitled to use IBM Spectrum Fusion with the following limitations:</p> <ul style="list-style-type: none">• You can use up to 6 TB of IBM Spectrum Fusion storage.• You can use IBM Spectrum Fusion for up to 12 months. <p>If you exceed these terms, a separate license is required.</p> <p>Contact your IBM Sales representative for access to this storage.</p>

Verifying entitlement from the Container software library

1. Go to the [Container software library](#).
2. From the side menu, click **View library**.
3. Ensure that you have access to the IBM Maximo Application Suite container software.

Verifying entitlement by using the command line

Ensure that you can log in to the IBM Entitled Registry by using your entitlement key.

Podman:

```
podman login cp.icr.io/cp/cpd -u cp
Password:<your_entitlement_key>
```

Docker:

```
docker login cp.icr.io/cp/cpd -u cp
Password:<your_entitlement_key>
```

Where `your_entitlement_key` is the value of your entitlement key from the Container software library. For more information, see [“Obtaining your IBM Entitlement key from the IBM Entitled Registry” on page 203](#).

Accessing Red Hat OpenShift entitlements

IBM Maximo Application Suite can include entitlement to use Red Hat OpenShift Container Platform, Red Hat Enterprise Linux CoreOS, and Red Hat Enterprise Linux. To access these entitlements, use IBM Passport Advantage to link your Maximo Application Suite account to your Red Hat OpenShift account.

Procedure

1. On the [IBM Passport Advantage Online for Customers](#) page, click **Sign in to your PAO Site** and log in with your IBMid.
 - The Customer primary contact name on your IBM order is passed to Red Hat OpenShift to accept the Red Hat OpenShift terms and conditions and link entitlements, which are then related to that Red Hat OpenShift account number.
 - If your company has multiple sites, an extra **Sign in** page is shown, so you can select a specific company site number.
2. On the **Passport Advantage Online** page, click **Download Software** in the **Software download** tab.
3. On the **Software downloads** page, confirm your site name and site number and then search for your Maximo Application Suite part number.

You can find your Maximo Application Suite part number in your Proof of Entitlement (POE) document. If you do not have your specific part number available, or if you have more than one Maximo Application Suite that you want to link, search for Maximo Application Suite.
4. Click the **View products** link.
5. Verify the product description and click **Continue**.

Important: Do not change or enter data in the **Version**, **Operating System**, or **Language** fields.
6. In the **Specification** section, locate the **Link with Red Hat OpenShift** heading, and click the order number that corresponds to the order number that you want to link your entitlement to.
 - a) Click **Okay** to open the Red Hat OpenShift login page, so you can map your IBM entitlement to your Red Hat account.

In the **Specifications** section, previously linked orders are shown for the Red Hat OpenShift row.

Important: Do not click **Container Install** or **Software downloads**. These actions are not part of the entitlement linking process.
7. On the Red Hat OpenShift login page, log in with your existing Red Hat OpenShift account or create a new account.

You must have a Red Hat OpenShift account to access the Red Hat OpenShift Cluster Manager. You do not need a paid Red Hat OpenShift subscription entitlement to access any IBM® offering.
8. On the Red Hat OpenShift **Review order summary** page, verify that the information is correct and click **Next**.
9. On the Red Hat OpenShift **Link your Red Hat Account** page, review the account details, select the **Assign the Red Hat subscriptions to this Red Hat account and link my IBM order** check box, accept the enterprise agreement terms, and click **Confirm**.

A message is shown that confirms your Red Hat OpenShift account is linked with your IBM order. Your entitlement is now accessible. You can link more orders on the **Software downloads** page of Passport Advantage Online by clicking the **SDMA** tab in your browser.

What to do next

For more information about viewing and managing your Red Hat OpenShift subscriptions, see [the Red Hat OpenShift Customer Portal page](#).

If these steps do not resolve your Red Hat OpenShift product entitlement issues, contact [IBM eCustomer Care](#).

Networking considerations

IBM Maximo Application Suite uses the networking setup by Red Hat OpenShift Container Platform for its internal communications (that is connections between pods that are running in the Red Hat OpenShift Container Platform).

But you also need:

- Connectivity from the cluster to external endpoints, unless you are running an air-gapped deployment.
- Connectivity into the cluster for Web Browsers to access the Maximo Application Suite control plane and applications or for IoT devices to connect to Maximo Application Suite.
- Connectivity from the Web Browsers to external Internet endpoints

IBM Maximo Application Suite default network policy

IBM Maximo Application Suite uses network isolation to isolate the Red Hat OpenShift project (IBM Cloud Kubernetes Service namespace) where Maximo Application Suite is deployed. The network policies in Maximo Application Suite cannot be customized, and are configured to enable least-privilege access throughout the product with a default deny-all policy. Specific rules are created to grant ingress and egress only where necessary.

To list all network policies in your Maximo Application Suite installation, run the following command:

```
oc -n {namespace} get networkpolicies
```

To review an individual network policy, run the following command:

```
oc -n {namespace} get networkpolicies {policyname} -o yaml
```

For more information, see [Understanding networking](#)

Connectivity from Red Hat OpenShift Container Platform cluster to external endpoints ("Allow listing")

If you are not running the deployment in an air gap environment, your Maximo Application Suite installation and its prerequisites need to access certain external endpoints at run time. These endpoints are typically accessed from port 443 (HTTPS) by using the domain names. The process of granting access to these endpoints is referred to as *allow listing*.

You need to ensure that:

- The Domain Names that are listed are present in the DNS that is being used by Red Hat OpenShift Container Platform.
- Any firewalls that you might have between the Red Hat OpenShift Container Platform cluster and the public Internet allow TCP/IP connections to be made from Red Hat OpenShift Container Platform nodes to port 443 of these external sites.

Firewalls sometimes must be configured by using IP addresses rather than DNS domain names. To help you with this configuration the current values of these IP addresses are included here.

**Attention:**

- These IP addresses might change over time so it is worth checking that the values shown here are still correct. You can check the values by using `nslookup` or `dig` commands.
- For availability reasons a single domain name might map to multiple IP addresses, so if you must configure at the IP address level, you need to ensure that they are all accessible.

Container Registries

Container registries are used to distribute the code for Maximo Application Suite and some of its dependencies. When you initially install Maximo Application Suite, it will try to pull the code from a container registry using the entitlement key that you are granted when you purchased Maximo Application Suite. If a Pod later needs to be restarted, it might again pull its code from an external registry. If you do not have access to these registry endpoints, you see that Pods fail to start and give Image Pull errors. If you get these errors, the Pod's specification will show you where it is trying to pull its images from.

You require access to the following registries:

- `cp.icr.io`
- `quay.io`
- `registry.redhat.io`
- `docker.io`
- `gcr.io`
- `ghcr.io`
- `nvcr.io`
- `registry.connect.redhat.com`
- `registry.redhat.io`
- The IBM Cloud Container Registry - `cp.icr.io`

Maximo Application Suite and SLS use the IBM Cloud Container Registry. This has domain name `cp.icr.io` which then maps to `icr.io` so you need to make sure that both names are available in your DNS.

Here is the output from `dig cp.icr.io`. You can see that the endpoint has 3 IP addresses for the high availability. Your firewall needs to allow access to all three of them.

```
cp.icr.io.      68  IN  CNAME  icr.io.
icr.io.         6   IN  A      169.60.98.86
icr.io.         6   IN  A      169.62.37.246
icr.io.         6   IN  A      169.63.104.236
```

- The `quay.io` registry.

This is used by some of the Maximo Application Suite dependencies and by Red Hat OpenShift Container Platform itself.

It has the following IP addresses:

```
quay.io.       17  IN  A      50.16.140.223
quay.io.       17  IN  A      54.156.10.58
quay.io.       17  IN  A      3.216.152.103
quay.io.       17  IN  A      3.221.13.191
quay.io.       17  IN  A      3.233.133.41
quay.io.       17  IN  A      34.224.196.162
quay.io.       17  IN  A      35.172.159.14
quay.io.       17  IN  A      44.197.21.192
```

As with the IBM Cloud Container Registry you should make sure the firewall (if it requires IP addresses) gives access to all of these.

You might find that you also need to allow access to *.quay.io or specific subdomain like cdn.quay.io or cdn01.quay.io

- The registry.redhat.io registry.

Some Red Hat OpenShift Container Platform images and operator catalog entries are held in this registry.

This uses Akamai content distribution network, and the IP address can vary depending on your geographical location. Use this command to find the address that is relevant to you.

```
dig registry.redhat.io

; <<> DiG 9.10.6 <<> registry.redhat.io
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 42254
;; flags: qr rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;registry.redhat.io.          IN A

;; ANSWER SECTION:
registry.redhat.io. 29 IN CNAME registry.redhat.io.edgekey.net.
registry.redhat.io.edgekey.net. 2032 IN CNAME e14353.g.akamaiedge.net.
e14353.g.akamaiedge.net. 17 IN A 2.18.104.196
```

In this example, the IP address used is 2.18.104.196.

- gcr.io registry pulls a manifest file.

Public network access

If you use a public network to access the IBM Cloud Container Registry and use the domains cp.icr.io and icr.io, add the following hostnames to your firewall rules:

- dd0.icr.io
- dd2.icr.io
- dd4.icr.io
- dd6.icr.io

If you are located in China, also add the following hostnames to your firewall rules:

- dd1-icr.ibm-zh.com
- dd3-icr.ibm-zh.com
- dd5-icr.ibm-zh.com
- dd7-icr.ibm-zh.com

Note: You can also add wildcard characters to hostnames in your allowlist, such as *.icr.io and *.ibm-zh.com.

Other Red Hat OpenShift endpoints

If you are installing your own Red Hat OpenShift cluster, you need to make sure it has access to Red Hat OpenShift cluster Manager. This is used by the Red Hat OpenShift Container Platform installation program and it also provides subscription management and collects telemetry data. This should already have been done if someone set up your Red Hat OpenShift cluster for you.

There are some other domains that you might need access to. They are mentioned in this [OCP Documentation](#) and include:

- *.openshiftapps.com
- sso.redhat.com

- `cert-api.access.redhat.com`
- `api.access.redhat.com`
- `registry.access.redhat.com`
- `infogw.api.openshift.com`
- `console.redhat.com`

DRO Endpoint `iaps.ibm.cloud`

The IBM Data Reporter Operator (DRO) running in your Red Hat OpenShift Container Platform instance connects out to an endpoint running in IBM Cloud . The endpoint DNS name is `iaps.ibm.cloud` and its IP address is `198.41.0.4`.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator ” on page 7](#).

System requirements

Ensure that you understand the system requirements for your Maximo Application Suite installation. System requirements differ depending on the applications that you plan to deploy, and the size requirements of these applications.

Related concepts

[Requirements and capacity planning](#)

Maximo Application Suite requirements

Maximo Application Suite requirements include considerations for your cluster and application components.

For more information, see [“Prerequisite software” on page 5](#).

Maximo Application Suite instance requirements

Review the requirements for planning your instance of Maximo Application Suite. You must provide the instance name, domain name, and workspace ID when you install and set up Maximo Application Suite. These values are used to create the user interface URLs for Maximo Application Suite and its applications.

Instance name

The instance name identifies the Maximo Application Suite installation on your Red Hat OpenShift cluster. Follow these guidelines to create your instance name:

- Must be 3 to 12 characters long
- Must use lowercase letters, numbers, and hyphen (-) symbol
- Must start with a lowercase letter
- Must end with a lowercase letter or a number

The instance name is not included with your Maximo Application Suite URL.

Domain name and DNS server

Determine a domain name and then work with your DNS administrators to ensure that your domain name is connected to your Red Hat OpenShift cluster.

The domain name is included with your Maximo Application Suite URL:`https://<workspace_id>.home.<mas_domain>`

Important: The domain name for your Maximo Application Suite instance must be resolvable within the Red Hat OpenShift cluster that you are deploying to.

You need the following DNS records. <mas_domain> is a domain name.

- admin.<mas_domain>
- api.<mas_domain>
- auth.<mas_domain>
- home.<mas_domain>
- *.home.<mas_domain>
- messaging.iot.<mas_domain>
- *.messaging.iot.<mas_domain>
- assist.<mas_domain>
- *.assist.<mas_domain>
- health.<mas_domain>
- *.health.<mas_domain>
- iot.<mas_domain>
- *.iot.<mas_domain>
- manage.<mas_domain>
- *.manage.<mas_domain>
- monitor.<mas_domain>
- *.monitor.<mas_domain>
- facilities.<mas_domain>
- *.facilities.<mas_domain>
- *.optimizer.<mas_domain>
- *.api.optimizer.<mas_domain>
- predict.<mas_domain>
- *.predict.<mas_domain>
- visualinspection.<mas_domain>
- *.visualinspection.<mas_domain>

Workspace ID

The Maximo Application Suite workspace is a unique collection of configuration settings for your Maximo Application Suite instance. Follow these guidelines to create your workspace ID:

- Must be 3 to 12 characters long
- Must only use lowercase letters and numbers
- Must start with a lowercase letter

The workspace ID is included in your Maximo Application Suite URL. In the following example URL, <workspace_id> is the workspace ID:

```
https://<workspace_id>.home.<mas_domain>
```

Workstation requirements

To accept the license during installation, the workstation on which you run the Maximo Application Suite installer must have a Java Runtime Environment configured.

For more information, see the [“Prerequisite software”](#) on page 5.

Application-specific requirements

Maximo Application Suite application-specific requirements include considerations such as workload calculations and persistent storage requirements.

Application-specific requirements for Maximo Manage

Review the Manage Compatibility Matrix to capacity plan for the Maximo Manage application.

To view compatibility with other Maximo Application Suite applications, add-ons, and industry solutions, download the Maximo Manage Compatibility Matrix file.

- Maximo Manage 9.1 users can download the [Maximo Manage Compatibility Matrix 9.1](#) file.
- Maximo Manage 9.0 users can download the [Maximo Manage Compatibility Matrix 9.0](#) file.
- Maximo Manage 8.11 users can download the [Maximo Manage Compatibility Matrix 8.11](#) file.
- Maximo Manage 8.10 users can download the [Maximo Manage Compatibility Matrix 8.10](#) file.

For more information about generating a compatibility report, see [Software Product Compatibility Reports](#).

Application-specific requirements for Maximo Real Estate and Facilities

In addition to the IBM Maximo Application Suite Software Product Compatibility Report, review the compatibility matrix for Maximo Real Estate and Facilities to see any application-specific requirements.

For more information about generating a compatibility report for IBM Maximo Application Suite, see [Software Product Compatibility Reports](#).

Compatibility matrix for Maximo Real Estate and Facilities

The following software, third-party components, and operating systems, collectively known as environments, are certified with Maximo Real Estate and Facilities.

For a list of updates, see the [Revision history](#) table.

New implementations are advised to install the most current versions of supported third-party software. If you are on an IBM Maximo Real Estate and Facilities version where a required third-party product is not currently supported, you must upgrade to a more current IBM Maximo Real Estate and Facilities release to be in compliance with the supported environments.

Continuous delivery of software updates, such as the maintenance with evergreen browsers such as Chrome and Firefox, might impact existing IBM Maximo Real Estate and Facilities functions. IBM reviews changes, and some situations might require a Maximo Real Estate and Facilities fix pack to restore IBM Maximo Real Estate and Facilities functions.

This document is not a contract. This document, all information contained herein, and IBM's products and services are subject to change at any time with or without notice in IBM's sole and absolute discretion. Without limiting the foregoing, IBM reserves the right to change its support policies and to certify its Maximo Real Estate and Facilities on new or different Environments and to de-certify its application on any Environments. The information contained in this document does not affect or change IBM product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document or from any changes to the information contained in this document.

Supported versions for IBM Maximo Real Estate and Facilities

Name	Supported Version	Notes
IBM Maximo Real Estate and Facilities Platform	9.1	Language packs for the following languages are available for Maximo Real Estate and Facilities: Arabic, Brazilian Portuguese, British English, Czech, Danish, Dutch, Finnish, French, German, Hebrew, Hungarian, Italian, Japanese, Korean, Norwegian, Polish, Simplified Chinese, Spanish, Swedish, Traditional Chinese, and US English.
IBM Maximo Real Estate and Facilities Application	9.1	(N/A)
IBM TRIRIGA Application	11.6, 11.5	(N/A)

Red Hat OpenShift Container Platform

Container Option	Supported Version	Notes
Red Hat OpenShift Container Platform	The Maximo Application Suite 9.1 supported versions are 4.16 to 4.18. Extended Update Support (EUS) releases 4.16 or 4.18 are recommended (Long-term-support versions while they remain in support).	Maximo Real Estate and Facilities is supported on Red Hat OpenShift Container Platform only, either self-managed, or on the following managed options: Red Hat OpenShift Service on Amazon Web Services (ROSA), Microsoft Azure Red Hat OpenShift (ARO), Red Hat OpenShift Dedicated (OSD) on Google Cloud Platform, and Red Hat OpenShift Kubernetes Service (ROKS) on IBM Cloud. You can generate a compatibility report for IBM Maximo Application Suite, see Software Product Compatibility Reports . The following Java and Liberty versions are included: <ul style="list-style-type: none"> • IBM Semeru Runtime Open Edition 17.0.16.0 (build 17.0.16+8) or later. • IBM WebSphere Application Server Version 25.0.0.8 Liberty - (25.0.0.8-cl250820250727-1902) or later.

Database

Maximo Real Estate and Facilities requires an instance of a database management system (DBMS), such as IBM Db2, Oracle Database or Microsoft SQL Server.

Db2 can be installed in the cluster or, if preferred, you can use one of the following external database management systems (DBMS). The database must be accessible from the Red Hat OpenShift cluster with Transport Layer Security (TLS) 1.2.

Maximo Real Estate and Facilities supports future versions of the vendor's database if the vendor does not remove, or explicitly or inadvertently disable, functionality that Maximo Real Estate and Facilities products rely on. Although future versions are supported, any issue that is introduced as a result of a database upgrade might require a Maximo Real Estate and Facilities fix pack or upgrade to be fully supported.

Database Options	Supported Version	Notes
IBM Db2 Advanced Edition	11.5	(N/A)

Database Options	Supported Version	Notes
Oracle Database (Enterprise & Standard)	19c	(N/A)
Microsoft SQL Server	2022	Microsoft SQL Server requires JDBC Driver JTDS 1.3.1. Database performance is central to the Maximo Real Estate and Facilities application. Certain Microsoft SQL Server functionality can reduce the level of performance compared to other database platforms for the Maximo Real Estate and Facilities. For more information, see the Microsoft SQL Server database .

Maximo Application Suite component

Component	Supported Version	Notes
IBM Maximo Monitor	9.1 or later	IBM® Maximo® Monitor for Workplace Analytics expands the Maximo Real Estate and Facilities IoT capabilities.

For more information about suite component compatibility, see the Maximo Application Suite “Compatibility matrix” on page 39.

Desktop client

Third-Party Components	Supported Version	Notes
Microsoft Excel Formats	.xls, .xlsx, .xlsm	Spreadsheets that are saved in Excel formatting are required if you use Maximo Real Estate and Facilities Offline forms.
Microsoft Project	2024, 2021	XML is the only supported method for interchanging Maximo Real Estate and Facilities project data with other project management applications.
Microsoft Windows	11	Minimum hardware requirements: 2 GHz CPU, 8 GB RAM, 100 Mbps NIC, 1280x1024 resolution.
macOS	15, 14, 13	Minimum hardware requirements: 2 GHz CPU, 8 GB RAM, 100 Mbps NIC, 1280x1024 resolution. CAD Integrator is not supported on macOS.
BIRT Report Designer	4.16	Required for BIRT Report development. Requires Java 17.
Microsoft Edge	137	These versions are the minimum versions that are supported for each browser. IBM Maximo Real Estate and Facilities Perceptive Applications require the support of modern web specifications to function optimally. Only Evergreen browsers (Chrome, Edge, Firefox, and Safari) are supported when using Perceptive Applications or creating custom views in the Maximo Real Estate and Facilities UX Framework.
Firefox	139, ESR 115.12.0	
Apple Safari	18	
GoogleChrome	137	

Maximo Real Estate and Facilities UX framework and Perceptive apps

IBM Maximo Real Estate and Facilities supports future versions of a vendor's operating system or browser if the vendor does not remove, or explicitly or inadvertently disable, functionality that Maximo Real Estate and Facilities relies on. Although future versions are supported, any issue that is introduced as a result of an operating system or browser upgrade might require a Maximo Real Estate and Facilities fix pack or upgrade to be fully supported.

Operating System (Browsers)	Minimum Supported Version	Notes
Apple iOS/iPadOS (Safari, Chrome)	17, 16	Applications that are built on the IBM Maximo Real Estate and Facilities UX Framework, including IBM Maximo Real Estate and Facilities Perceptive Applications, support the minimum versions of the operating systems shown. For desktop support requirements, see Desktop Client Compatibility . The UX Framework components that IBM Maximo Real Estate and Facilities delivers and documents on the UX component documentation page are supported. Any custom code that is added around IBM Maximo Real Estate and Facilities components, such as HTML, JavaScript, or third-party code, is the responsibility of the developer. For information about offline support, see Enabling offline mode in the Perceptive apps .
Google Android (Chrome)	13, 12, 11, 10	
macOS (Safari, Chrome)	14, 13, 12	
Microsoft Windows tablet (Edge, Chrome)	11	
Microsoft Windows laptop/desktop (Edge, Chrome, Firefox)	11	<p>The UX Framework supports ReactJS and Google Polymer Web Applications.</p> <p>IBM Maximo Real Estate and Facilities ReactJS components are compatible with the IBM Carbon Design System 11. The IBM Maximo Real Estate and Facilities Polymer UX Applications are rendered with IBM Carbon Design System 11.</p> <p>IBM Maximo Real Estate and Facilities Polymer UX components are available in Polymer 3.1.0 and Polymer 1.6.1. The IBM Maximo Real Estate and Facilities Polymer UX Applications are rendered with Polymer 3.1.0 and Polymer 1.6.1.</p> <ul style="list-style-type: none"> • For Polymer 3.1.0, elements and components from the Polymer Catalog, such as <iron> and <paper>, are the most current versions as of 20 February 2019. • For Polymer 1.6.1, elements and components from the Polymer Catalog, such as <iron>, <paper>, <gold>, and <platinum>, are the most current versions as of 20 October 2016. <p>Applications that are written on the IBM Maximo Real Estate and Facilities UX Framework, including Perceptive Applications, might need to be updated to support the version of Carbon or Polymer that is delivered with IBM Maximo Real Estate and Facilities.</p> <p>For information about Microsoft support for Windows 11, see Lifecycle FAQ - Windows.</p>

Maximo Real Estate and Facilities Advanced Room Search

IBM Maximo Real Estate and Facilities Advanced Room Search is an add-in for Microsoft Outlook. The Advanced Room Search add-in uses Microsoft HTML5 add-in architecture and is delivered by using the Maximo Real Estate and Facilities UX Framework.

Note: To ensure that the Advanced Room Search add-in works correctly on desktop Outlook clients on Windows, you must use Microsoft WebView2 embedded browsers. For more information, see [Browsers and webview controls used by Office Add-ins](#).

For more information, see [Installing the Room Search add-in in Microsoft Outlook](#).

IBM Maximo Real Estate and Facilities Advanced Room Search supports user authentication through standard IBM Maximo Real Estate and Facilities user name and password and IBM Maximo Real Estate and Facilities identity token technology for Single Sign-On (SSO) solutions. It also supports Service Provider (SP) initiated SAML, OAuth, or OpenID Connect (OIDC) SSO. It does not support Identity Provider (IdP) initiated SSO.

Supported Outlook Version	Supported Version	Notes
Windows: Outlook 2019 1902 or later Mac: Outlook 2019 16.66 or later Outlook for Windows 365 (continuous delivery)	9.1, 11.6	IBM Maximo Real Estate and Facilities Advanced Room Search functions interface with Microsoft libraries and APIs which might be subject to continuous revisions by Microsoft . As a result, Advanced Room Search features might be vulnerable to unforeseen impact introduced by such changes. There is a limitation where the Microsoft JavaScript Library does not support all functions (including those required by Advanced Room Search) when using Outlook for Microsoft 365 with Microsoft Exchange servers on premise. Microsoft has not yet published a roadmap that indicates when or if the library will be fully supported on Outlook for Microsoft 365 on premise.

Maximo Real Estate and Facilities CAD Integrator/Publisher

CAD Integrator/Publisher for Autodesk AutoCAD and for Bentley MicroStation supports native, Basic, NTLM, SAML, OAuth, or other Single Sign-On (SSO) technologies for providing user authentication.

The following Autodesk AutoCAD information reflects Autodesk support policies at the time of publication. If Autodesk withdraws support for an AutoCAD version, IBM Maximo Real Estate and Facilities might also cease to support that version.

Product	Supported Version	Notes
CAD Integrator/Publisher	9.1	Installation requires full administrator access. A 64-bit operating system is required (such as Microsoft Windows 11, 64-bit). IBM SDK Java 17 or Oracle Java SE 17 or later is required.
Autodesk AutoCAD Autodesk AutoCAD Architecture	2024, 2023, 2022	Supported by CAD Integrator/Publisher 9.1. 16 GB RAM recommended. For large drawings, 32 GB RAM recommended. AutoCAD requires a 64-bit JRE. AutoCAD LT is not supported.

Product	Supported Version	Notes
Bentley MicroStation	MicroStation 2023 (MS2023)	Supported by CAD Integrator/Publisher 9.1 . 16 GB RAM recommended. For large drawings, 32 GB RAM recommended. All operating systems require a 64-bit JRE, and Microsoft .Net 4.8.

IBM Maximo Real Estate and Facilities Connector for BIM

The following Autodesk Revit list reflects Autodesk support policies at the time of publication. If Autodesk withdraws support for a Revit version, Maximo Real Estate and Facilities might also cease to support that version. The Connector for BIM can be installed for older Revit versions, but it might not support any versions that Autodesk no longer supports.

Product	Supported Version	Notes
IBM Maximo Real Estate and Facilities Connector for BIM	9.1	IBM Maximo Real Estate and Facilities Connector for BIM requires Java 17 64-bit JRE for installation.
Autodesk Revit	2026, 2025, 2024, 2023, 2022	Supported by IBM Maximo Real Estate and Facilities Connector for BIM 9.1.

IBM Maximo Real Estate and Facilities connectors

IBM Maximo Real Estate and Facilities supports connections to these third-party systems. For support of the third-party system itself, contact the vendor directly.

Product	Supported Version	Notes
Esri ArcGIS Enterprise for Kubernetes	11.4, 11.3	(N/A)
ArcGIS Enterprise	11.5, 11.4, 11.3, 11.2, 11.1, 10.9.1, 10.8.1	(N/A)
ENERGY STAR Portfolio Manager	14	The ENERGY STAR Portfolio Manager Web Services API is used for the IBM TRIRIGA Connector for Energy Star Benchmarking. For more information, see ENERGY STAR Portfolio Manager Release Notes .
IBM Maximo Real Estate and Facilities Custom Class Loader	9.1	Integrations written for IBM Maximo Real Estate and Facilities Custom Class Loader must be compatible with Java 17.

Product	Supported Version	Notes
Connector for Business Applications	9.1	Integrations built for IBM Maximo Real Estate and Facilities Connector for Business Applications (CBA) must be compatible with Apache CXF 3.6.4 web service framework. For more information about Apache CXF, see Apache CXF . For more information about IBM Maximo Real Estate and Facilities Connector for Business Applications, see Integrating data with the TRIRIGA Connector for Business Applications .
Microsoft Exchange	Exchange 365 (Online) and hybrid Exchange implementations	(N/A)

IBM Maximo Real Estate and Facilities CMIS

Product	Supported Version	Notes
CMIS	1.1	IBM Maximo Real Estate and Facilities can be configured to store documents in Enterprise Content Management (ECM) systems that support the Content Management Interoperability Services (CMIS) ECM gateway Version 1.1 of the CMIS standard, as established by OASIS. These gateways are specific to the vendor of your particular ECM and should be installed accordingly if one does not already exist in your organization. For more information, see IBM TRIRIGA Platform Configuration - 3.5.2 CMIS Enablement .

Email notifications

Maximo Real Estate and Facilities supports incoming and outgoing email. For support of the third-party system itself, contact the vendor directly.

Category	Product or Component	Notes
Third-Party Components	Mail (Outgoing and Incoming)	Outbound email uses Simple Mail Transfer Protocol (SMTP). The Incoming Mail Agent can be configured for use with IMAPS, POP3.

Cloud login

Maximo Real Estate and Facilities supports authentication with these IBM and third-party identity providers (IdPs). For support of the third-party system itself, contact the vendor directly.

Category	Product or Component	Notes
Third-Party Components	IdP authentication	Cloud login supports Microsoft , Autodesk, Okta, Google, and IBM Security Verify as IdPs. Cloud login has custom options to configure other IdPs, however, the generic cloud-login processing is not guaranteed to work for all IdPs.

Revision history

Date	Document Version	Notes
24 June 2025	V1	This is a new document.

Application-specific requirements for Maximo Monitor and IoT tool

Review the Maximo Monitor and IoT tool typical workload sizes and persistent storage requirements.

Workload sizes

The following table lists typical workload sizes:

	Developer	Small	Medium
IoPoints*	200	5,000	50,000
Max number of simultaneously connected devices	200	5,000	50,000
Max data rate (totaled over all connected devices)	0.4 kB/s	10 kB/s	100 kB/s
Max msg rate (totaled over all connected devices)	4 msg/s	100 msg/s	1,000 msg/s
Max Db2® insert rate	4 inserts/s	100 inserts/s	1,000 inserts/s

*An IoPoint is a data value that is written to the Maximo® Monitor database by a distinct device with a frequency of no more than one write per minute. When you are calculating with IoPoints, you need to consider three dimensions:

- Number of devices
- Number of data points sent in each message
- Number of messages per second

For example, a total of 1,000 IoPoints might be allocated in the following ways:

- 1,000 devices that each send 1 message per minute, where each message contains 1 datapoint.
- 1 device that sends 1,000 messages per minute, where each message contains 1 datapoint.
- 500 devices that each send 2 messages per minute, where each message contains 1 datapoint.
- 1 device that sends 1 message per minute, where each message contains 1,000 data points.

In the preceding table, the data rate and message rate rows show the rates for data that flows from devices to the IoT tool and to the Maximo® Monitor application.

Persistent storage requirements

The IoT Tool requires a PVC to store the MQTT broker state. You set the storage class and size of the PVC when you deploy the IoT tool.

Requirements

To view a list of available storage classes in your cluster, run the following OpenShift console command:

```
oc get storageclasses
```

Choose the appropriate storage class and size for your workload size.

The storage class can be used to dynamically provision a persistent volume with access mode RWO (ReadWriteOnce).

No specific filesystem permissions are required.

Supported and tested storage

The following storage providers have been tested with the IoT Tool:

- OpenShift Container Storage
- IBM Cloud block storage class: `ibmc-block-gold`

The following table lists typical storage requirements:

Product	Storage requirement
Maximo® Application Suite core	RLKS: 5 Gb
Maximo Monitor	None
IoT - Developer	Message Gateway: 64 Gb
IoT - Small	Message Gateway: 64 Gb
IoT - Medium	Message Gateway: 128 Gb
MongoDB	20 Gb
Kafka	ZooKeeper: 20 Gb Kafka: 50 Gb
Cloud Pak for Data (Db2)	Cloud Pak for Data: 20 Gb Db2 Warehouse: (100 Gb * 2) + Db2 user storage based on retention policy

Db2® user data storage

User storage requirements for Db2® are primarily determined by two factors:

- Incoming event data rate
- Retention policy

A typical allocation might be 1 Mb storage per day per 300 bytes/minute (3 IoPoints).

The following table is based on the workload sizes estimates:

	Data rate	IoPoints	Db2® Storage Requirements
Developer	400 bytes/sec	200	80 Mb storage per day
Small	10,000 bytes/sec	5,000	2 Gb storage per day
Medium	100,000 bytes/sec	50,000	20 Gb storage per day

For example, if you wanted to store 1 month's worth (30 days) of data and also match the benchmark figures, the user storage requirements are as follows:

- Developer: 2.4 Gb (30 * 0.08 Gb)
- Small: 60 Gb (30 * 2 Gb)
- Medium: 0.6 Tb (30 * 20 Gb)

Messaging Load Balancer

For messaging workloads with more than 100K device connections, the OCP Ingress load balancer is not an option. Alternatives, such as, the [NLB 2.0](#) load balancer in IBM Cloud®, which is based on IPVS/Keepalived, can scale beyond 100K device connections.

Application-specific requirements for Maximo Health and Maximo Predict

Review the Maximo Health and Maximo Predict typical workload sizes and persistent storage requirements.

Workload sizes

The following table lists typical workload sizes for Maximo Health.

	Small	Medium
Assets	3,000	20,000
Users	1	5

The following table lists typical workload sizes for Maximo Predict:

	Small	Medium
Assets (subset of assets in Health)	600	6,000
Asset groups*	5	10
Users	1	5
Number of sensors per asset	10	30
Sensor readings**	328,000	3,288,000
Failure history for active assets	10,500	75,000

*In each asset group, for a small workload, there were 60 active and 60 decommissioned assets. For medium workload, there were 300 active and 300 decommissioned assets in each group.

**Sensor readings are at a daily level for each active asset over a period of 3 years.

Additionally, Maximo Predict includes a notebook that includes six models under four templates:

	Number of models
Degradation curve	1
Time to failure	2 (1 for each failure mode)
Failure prediction	2 (1 for each failure mode)
Anomaly detection	1

Persistent storage requirements

The following table lists typical storage requirements.

Product	Description
Maximo Health	IBM Db2 Warehouse (dedicated instance) 120 GB
Maximo Predict*	20 GB

*Because Maximo Health is deployed as part of deploying Maximo Predict, storage requirements for Maximo Health are also applicable to a Maximo Predict deployment.

Application-specific requirements for Maximo Visual Inspection

Review the Maximo Visual Inspection recommended hardware configuration and I/O performance.

Recommended hardware configuration and I/O performance

Worker nodes

Maximo Visual Inspection requires the following configuration for each worker node:

- At least one 64-bit x86-compatible processor.
- 60GB of memory.

GPUs

Maximo Visual Inspection needs GPUs to perform deep learning model training and requires GPU configuration. For more information, see [Supported GPU devices](#) in the Maximo Visual Inspection documentation.

Storage

Maximo Visual Inspection uses a persistent volume claim (PVC) to manage storage and requires the following storage configuration and I/O performance:

- A PVC storage class that supports [ReadWriteMany](#) access mode.
 - NFS based, file-storage based, or IBM Cloud File Storage implementations that offer ReadWriteMany access modes are supported.
 - Storage provider implementations that offer ReadWriteOnce access mode, such as most block storage implementations, are NOT supported.
- A PVC storage size of at least 40GB for data sets and models.
 - Increased numbers for data sets and models lead to increased storage requirements.
 - Most storage providers allow storage allocations to be increased.
 - If you cannot increase your allocation, use the following rule to determine a rough estimate of your storage utilization:
$$(100 \text{ KB per image}) \times (\text{number of images}) \times (\text{number of data sets}) +$$
$$1\text{GB} \times (\text{number of models}) +$$
$$10\text{GB}$$

Note: Images and video clips vary greatly in resolution and compression ratios. This rule might not be appropriate for your workload.
- At least 10,000 input/output operations per second (IOPS) of random read/write performance for data set uploads and downloads.
 - Less storage bandwidth means that certain actions are visibly slower.
 - As more storage IOPS are allocated, application performance scales up to approximately 50,000 random read/write IOPS.
 - Configuration of IOPS and storage quotas varies from provider to provider. See how to manage [IBM Cloud File Storage IOPS](#) or refer to your storage provider's documentation.

Application-specific requirements for Maximo Collaborate

Review the Maximo Collaborate typical workload sizes.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Workload sizes

The following table lists typical workload sizes:

	Tiny	Small	Medium	Large
Workers Monitored*	50	250	2,500	10,000

* If worker or asset location tracking is used, the number of workers (or assets) that can be monitored by the solution sizes is reduced to half as many.

Requirements to install in an existing Red Hat OpenShift cluster

Review the following considerations if you plan on installing Maximo Application Suite in an existing Red Hat OpenShift cluster.

Cluster administrator rights

Maximo Application Suite requires global/cluster administrator rights to install. Maximo Application Suite is a cross namespace (cluster wide) application and requires cluster-admin access to create and update custom resource definitions/namespaces, get cluster information for automation and Operator Lifecycle Management related tasks.

The cluster-admin role is required by Maximo Application Suite and used to create and manage resources needed by Maximo Application Suite. However, no global changes will be applied by Maximo Application Suite. Cluster administrators are required to install Operators by default. It is not recommended to allow non-cluster administrators to install Operators.

Cluster-wide settings

Maximo Application Suite installations or upgrades will not change any cluster-wide settings other than the resources created by Maximo Application Suite itself. Maximo Application Suite requires IBM Certificate Manager 1.15 or later to install Maximo Application Suite in the cluster. If IBM Certificate Manager is a shared dependency between the cluster, upgrading it might affect other applications in the cluster.

Sharing common services across Maximo Application Suite instances

In order to most efficiently utilize available resources it may be desirable to share some services across multiple Maximo Application Suite (Maximo Application Suite) instances. The following Maximo Application Suite dependencies provide services which can be shared by multiple Maximo Application Suite instances, either in the same Red Hat OpenShift cluster or across various Red Hat OpenShift cluster.

- [MongoDB](#)
- [“Suite License Service” on page 7](#)
- [“Data Reporter Operator ” on page 7\(DRO\)](#)

Ensure that services are exposed via an Red Hat OpenShift route if they are to be shared outside the Red Hat OpenShift cluster, for example with Maximo Application Suite instances in different Red Hat OpenShift cluster.

A common scenario in which it is cost effective to share common services between multiple Maximo Application Suite instances is development and test. For multiple development and test environments it can be cost effective to share a single instance of MongoDB, DRO, and SLS services. In some cases, for example a load test environment, where there is significant load being placed on the shared services it might be necessary to dedicate a separate instance of MongoDB, DRO, and SLS.

Production environments should not share instances of MongoDB, DRO, and SLS; it is recommended to use dedicated instances of these services for production environments.

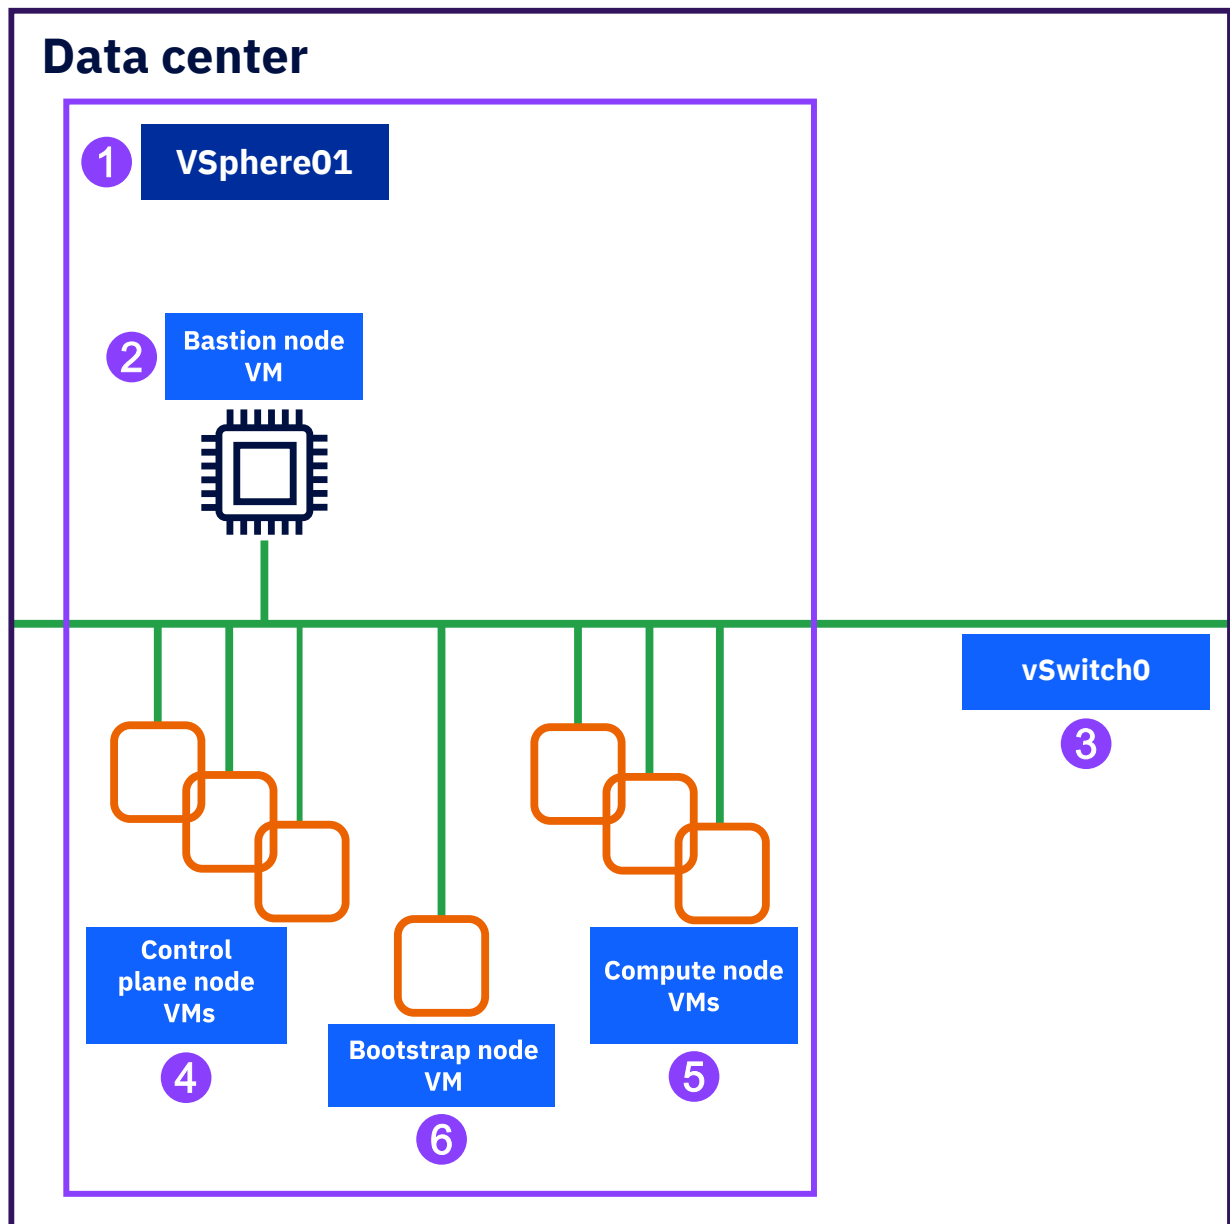
When sharing the SLS service across multiple Maximo Application Suite instances it is important to note that the pool of app points is shared. For example, if there are 300 app points in the pool and one Maximo Application Suite instance consumes 200 app points, then there will only be 100 app points for the remaining Maximo Application Suite instances which are sharing the SLS service.

However, if you have a user defined in one of the Maximo Application Suite instances instances that has the same name and entitlement as a user on the other Maximo Application Suite instances instance, that user can log in to both instances and the app points are only deducted once.

Maximo Application Suite on-premises installation topology

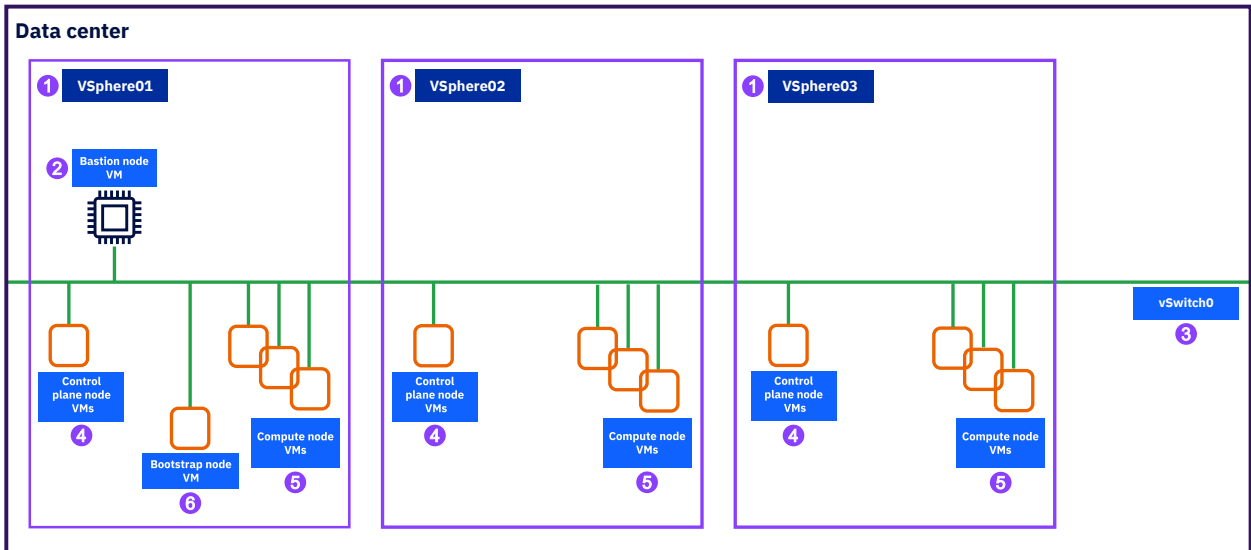
The Maximo Application Suite on-premises installation topology contains sample scenarios for Maximo Application Suite on Red Hat OpenShift Container Platform on vSphere.

The following diagram illustrates a sample installation topology for Maximo Application Suite on Red Hat OpenShift Container Platform on vSphere. In this example, a single bare metal machine with vSphere is used to install Red Hat OpenShift Container Platform.



- 1** The bare metal machine with vSphere hypervisor installed.
- 2** The bastion host VM. Used to provision the OCP bootstrap node. Hosts the PXE, DHCP, DNS, HTTP, TFTP servers for provisioning control plane and compute nodes. For development environments, the bastion host can also be used to host a load balancer. For production environments, it is recommended to use a dedicated VM for load balancing or use a commercial grade L4 load balancer, such as NGINX, F5, or other load balancer.
- 3** vSwitch0 is the virtual network that bridges the VMs with the physical vmnic0 on vSphere. In this topology, a single port group is created on vSwitch0. vmnic0 is attached to a network with internet access.
- 4** Three control plane node VMs used to manage the compute node and the Kubernetes pods in the Red Hat Red Hat OpenShift cluster.
- 5** Compute node VMs used to run Maximo Application Suite dependencies and Maximo Application Suite workloads.
- 6** Bootstrap node VM used to deploy the Red Hat OpenShift Container Platform. After the bootstrap process is complete, the bootstrap node VM is no and can be deprovisioned.

For high availability, use multiple vSphere servers to host control plane and compute node VMs.



- 1 Multiple bare metal machines with vSphere hypervisor installed. Control no longer needed, and compute node VMs are striped across vSphere machines for high availability.
- 2 The bastion host VM. Used to provision the OCP bootstrap node. Hosts the PXE, DHCP, DNS, HTTP, TFTP servers for provisioning control plane and compute nodes. For development environments, the bastion host can also be used to host a load balancer. For production environments, use a dedicated VM for load balancing or use a commercial grade L4 load balancer, such as NGINX, F5, or other load balancer.
- 3 vSwitch0 is the virtual network that bridges the VMs with the physical vmnic0 on vSphere. In this topology, a single port group is created on vSwitch0. vmnic0 is attached to a network with internet access. The servers that are hosted on the bastion host VM are reachable from control plane and compute node VMs on other vSphere servers.
- 4 Three control plane node VMs used to manage the compute node and the Kubernetes pods in the Red Hat Red Hat OpenShift cluster. The control plane node VMs are striped across multiple vSphere servers.
- 5 Compute node VMs used to run MAS dependencies and MAS workloads.
- 6 Bootstrap node VM used to deploy the Red Hat OpenShift Container Platform. After the bootstrap process is complete, the bootstrap node VM is no longer needed, and it can be deprovisioned.

Planning for IBM Cloud

Before you can install Maximo Application Suite on IBM Cloud, you must configure the installation requirements. There are dependencies that you must install before installing the Maximo Application Suite itself and additional dependencies that you can install depending on the Suite application you are planning to deploy. As an alternative, you can make use of automation to provision the Red Hat OpenShift cluster on IBM Cloud and other dependencies by using the Ansible collection.

Creating your IBM Cloud account and configuring permissions

You need an IBM Cloud account with some specific permissions to be able to provision the Red Hat OpenShift cluster and other services.

If you do not have an IBM Cloud account yet you must create one. You need an IBM Cloud account that is created as a Pay-As-You-Go or Subscription accounts. For more information, see the [IBM Cloud Account Types](#) documentation.

To provide and manage the Red Hat OpenShift cluster on IBM Cloud, you need to set the necessary IBM Cloud classic infrastructure permissions:

Note: If you have problems in setting the permissions, contact your IBM Cloud Account administrator.

Procedure

1. Login in your IBM Cloud account through the [IBM Cloud website](#).
2. Go to **Manage** menu and select **Access (IAM)**.
3. Go to **Users** menu and select your user from the list.
4. Go to the **Classic Infrastructure** tab.
5. In the **Permissions** tab, grant the various permissions as follows:
 - a) Expand **Account** and grant the following permissions:
 - Required: Add Server
 - Required: Cancel Server
 - Suggested: Add/Upgrade Storage (StorageLayer)
 - Required: Add/Upgrade Services
 - Required: Cancel Services
 - b) Expand **Devices** and grant the following permissions:
 - Required: View Hardware Details
 - Required: IPMI Remote Management
 - Required: OS Reloads and Rescue Kernel
 - Suggested: Manage Port Control
 - Required: View Virtual Server Details
 - Suggested: Edit Hostname/Domain
 - c) Expand **Network** and grant the following permissions:
 - Suggested: Add IP Addresses
 - Suggested: Manage Network Subnet Routes
 - Suggested: Add Compute with Public Network Port
 - d) Expand **Services** and grant the following permissions:
 - Suggested: Manage DNS
 - Suggested: Storage Manage
 - Suggested: View Certificates (SSL)
 - Suggested: Manage Certificates (SSL)
6. In the **Devices** tab:
 - Suggested: Auto Bare Metal Server Access
 - Suggested: Auto Dedicated Host Access
 - Suggested: Auto Virtual Server Access
7. Click **Apply**.
8. Go to the **Access policies** tab and click **Assign access**.
9. Under the **IAM services** tile, enter or select the following services and assign the corresponding least privileges:
 - a. **IBM Cloud Activity Tracking** – This is required to enable IBM Cloud Activity Tracking integration, which comes as default add-on when you provision your IBM Cloud Red Hat OpenShift cluster.
 - i) Under **Platform access**, add the **Administrator** role.
 - ii) Click **Add**.
 - b. **IBM Cloud Monitoring service** – This is required to enable IBM Cloud Monitoring integration, which comes as default add-on when you provision your IBM Cloud Red Hat OpenShift cluster.
 - i) Under **Platform access**, add the **Administrator** role.

- ii) Click **Add**.
 - c. **IBM Log Analysis** – This is required to enable IBM Log Analysis integration, which comes as default add-on when you provision your IBM Cloud Red Hat OpenShift cluster.
 - i) Under **Platform access**, add the **Administrator** role.
 - ii) Click **Add**.
 - d. **Container Registry** – This is required to enable the clusters that are created by your user to pull the required images from icr.io, which is the main IBM image registry repository.
 - i) Under **Platform access**, add the **Administrator** role.
 - ii) Click **Add**.
 - e. **Databases for MongoDB** - This is an alternative for the MongoDB service available on IBM Cloud . MongoDB is a prerequisite for installing Maximo Application Suite.
 - i) Under **Platform** access, add the **Editor** role.
 - ii) Click **Add**.
 - f. **Event Streams** – this is an alternative for the Kafka service available on IBM Cloud . Kafka service is required if you plan to install IoT tool in your Maximo Application Suite instance. It can also be used by Manage and other Suite applications. Refer to each application documentation for more details.
 - i) Under **Platform** access, add the **Editor** role.
 - ii) Under **Service** access, add the **Writer** role.
 - iii) Click **Add**.
 - g. **Internet Services** - This is an alternative to provide Domain Name Service (DNS) management that is required if you are planning to use custom cluster issuers signed by well-known certificate authorities for your Maximo Application Suite instance.
 - i) Under **Platform** access, add the **Administrator** role.
 - ii) Under **Service** access, add the **Manager** role.
 - iii) Click **Add**.
 - h. **Kubernetes Service** – This is needed to provision and manage a Red Hat OpenShift cluster on IBM Cloud .
 - i) Under **Platform** access, add the **Administrator** role.
 - ii) Under **Service** access, add the **Writer** role.
 - iii) Click **Add**.
10. Click **Assign** to assign all the added permissions to your user.

Results

The permissions are granted to the user.

What to do next

You can get or obtain your IBM Entitlement key from the IBM Entitled Registry.

Obtaining your IBM Entitlement key from the IBM Entitled Registry

The IBM Entitled Registry key is used during the installation to download the container images for the Suite and its applications from the IBM Entitled Registry. It is also required to install some prerequisites. Download this key from the [IBM Container Library](#).

Complete this task to log in to the IBM Container Library. To verify that your entitlement key is valid for Maximo Application Suite.

Procedure

1. Log in to the IBM Container Library with a user ID that has software download rights for your company's entitlement.
 - a) Your entitlement key is displayed in the Container software library. If you do not see the key in the menu bar, select **Get entitlement key**.
 - b) Copy the key to a safe location. Under the Access your container software heading, click **Copy key**.
2. To verify that your entitlement key is valid for Maximo Application Suite, in the Container software library, select **View library** to see the list of products that you are entitled to. If IBM Maximo Application Suite is not listed, or if the **View library** link is not selectable, the identity with which you are logged in to the container library does not have entitlement for Maximo Application Suite. In this case, the entitlement key is not valid for installing the software.

You can verify that you can log in to the IBM Entitled Registry by using your entitlement key by running the following command.

Podman:

```
podman login cp.icr.io/cp/cpd -u cp
Password:<your_entitlement_key>
```

Docker:

```
docker login cp.icr.io/cp/cpd -u cp
Password:<your_entitlement_key>
```

For more information, see [“Verify software entitlement” on page 179](#).

Results

Your IBM Entitlement Key is saved to be used in next installation steps.

What to do next

You can install IBM Cloud CLI in your system and login through command line interface by using an IBM Cloud API key.

Installing IBM Cloud CLI

IBM Cloud CLI provides full management of your IBM Cloud account via command line. Some installation steps described along this guide may need the IBM Cloud Command Line Interface (CLI) available to be performed.

Refer to the [Getting started with the IBM Cloud CLI](#) documentation to install IBM Cloud Command Line Interface in your system.

Note: We recommend you install at minimum these [plugins for IBM Cloud CLI](#).

Creating your IBM Cloud API key

After you installed the IBM Cloud CLI in your system, you need to login in your IBM Cloud account to use the IBM CLI. To do that you need to have a valid IBM Cloud API Key.

Procedure

1. Login in your IBM Cloud account.
2. In the **Manage** menu, select **Access (IAM)**.
3. In the **API keys** menu, click **Create** button.
4. In the **Create IBM Cloud API key** page, enter a name and description for your API Key.

5. In the **Leaked key** section, select either to disable, delete, or not take any action if a key is discovered.
6. In the **Select creation** section, choose whether the API key should create a session in the CLI or not.

Results

The IBM Cloud API key is created, write it down to use anytime you need to login to your IBM Cloud account by using command line.

What to do next

To connect to your IBM Cloud account by using IBM Cloud CLI run:

```
ibmcloud login --apikey $your_ibmcloud_api_key -q --no-region
```

Where `$your_ibmcloud_api_key` is the API key that you have created in this task.

Installing Red Hat OpenShift Container Platform on IBM Cloud

You can install Red Hat OpenShift Container Platform by using the IBM Cloud CLI or through the IBM Cloud Catalog. You will need to specify parameters regarding your cluster details such as cluster name, region, worker nodes capacity and version.

Tip: This task maps to the following Ansible role: `ocp_provision`. For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

Installing Red Hat OpenShift Container Platform by using IBM Catalog

To install OpenShift using IBM Catalog, you must login with your IBM Cloud account in the IBM Cloud website and configure the settings in order to provision your OpenShift cluster. To log in to the IBM Cloud website, go to <https://cloud.ibm.com>.

Procedure

1. Go to the **Catalog** menu.
2. Search for Red Hat OpenShift on IBM Cloud.
3. Click the **Red Hat OpenShift on IBM Cloud** tile.
4. In the **Select your setup** section, ensure that **Manual setup** is selected.
5. In the **Infrastructure** section, ensure that **Classic** is selected.
6. In the **Location** section:
 - a) Select your preferred **Resource Group**. For more information about Resource Groups on IBM Cloud, see [Managing resource groups](#).
 - b) Select your preferred **Geography**. This is the data center region that hosts your Red Hat OpenShift cluster.
 - c) Select your preferred **Availability**. For lesser cost, you can choose Single zone. If you want to have higher availability cluster choose Multi zone.
 - d) Select your preferred **Metro**. This is the data center exact location that hosts your Red Hat OpenShift cluster. When **Availability** is Multizone when you select the **Metro** more than one data center will be automatically selected. When **Availability** is Single zone when you select the **Metro** just one data center will be automatically selected.
7. In the **Worker pool** section you choose the worker flavor and the number of worker nodes per zone. Here you can select some different options available for the configuration of CPU, memory, disk, and operating system of the cluster worker nodes.
 - a) Click in the **Change flavor** button and choose an option that can support the Maximo Application Suite and its prerequisites, according to your needs.

Change worker pool flavor

×

Select a flavor of CPU, memory, and operating system characteristics for the worker pool. At any time later, you can add more worker pools with different flavors to fit the resource needs of your workload.

Machine

Bare metal

Virtual - shared

Virtual - dedicated

Use case

Balanced

Additional storage

GPU

RAM intensive

Data intensive

Size

Small

Medium

Large

	vCPU	Memory	Price	Name	Machine type	Operating system	Primary disk	Secondary disk
<input type="radio"/>	4	16GB	\$0.44/hr	b3c.4x16	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	4	32GB	\$0.51/hr	m3c.4x32	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	8	32GB	\$0.84/hr	b3c.8x32	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	8	64GB	\$0.88/hr	m3c.8x64	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	16	16GB	\$1.35/hr	c3c.16x16	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	16	32GB	\$1.46/hr	c3c.16x32	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	16	64GB	\$1.63/hr	b3c.16x64	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input checked="" type="radio"/>	16	64GB	\$1.63/hr	b3c.16x64.300gb	Virtual - shared	RHEL	25GB SSD	300GB SSD
<input type="radio"/>	16	128GB	\$1.93/hr	m3c.16x128	Virtual - shared	RHEL	25GB SSD	100GB SSD
<input type="radio"/>	32	32GB	\$2.60/hr	c3c.32x32	Virtual - shared	RHEL	25GB SSD	100GB SSD

Items per page: 10 1-10 of 17 items 1 of 2 pages

Cancel
Save worker pool flavor

- b) After chosen the flavor, fill the number of worker nodes per zone. For example, this option in the following example should work for general Maximo Application Suite deployments including all applications and its prerequisites:

Worker pool

Set up a worker pool with the flavor and number of worker nodes that you want to run your first workload. At any time later, you can add more worker pools with different flavors, or resize your worker pools to fit the resource needs of your workloads.

Virtual - shared, RHEL

16

vCPUs

64 GB

Memory

\$1.63/hr

Cost

[Change flavor](#)

Worker nodes per zone

8

±

x 1 zone

= 8 workers total

You always can increase and decrease the number of worker nodes per zone after the OpenShift cluster is provisioned.

8. In the **Worker pool name**, type your preferred worker pool name or leave the default.
9. Ensure that the **Encrypt local disk** option is set to **On**.
10. In the **Master service endpoint** field, select **Public endpoint** only.
11. Under the **Orchestration service** section, select the appropriate version supported by the Maximo Application Suite version you plan to install. For example, for Maximo Application Suite version

8.8.0 select the latest OpenShift 4.8.x version available in the list. For more information about the OpenShift versions that are supported by the Maximo Application Suite version you plan to install, you can [create a System requirements report](#), searching for Maximo Application Suite and selecting the Version you plan to install.

12. Under the **OCF entitlement** section select **Apply my Cloud Pak OCP entitlement to this worker pool**.
13. Based on your customized selections you may need to add additional Infrastructure permissions. Review the results of the **Infrastructure permissions checker** and add the permissions that are missing. For more information about assigning infrastructure permissions, see [the docs](#).
14. In the **Resources details** section, type your preferred **Cluster name** or leave the default. You can also create **Tags** for it. If your user tags are billing related, consider writing tags as key:value pairs, such as `costctr:124`. User tags are visible account-wide. Avoid including sensitive data in the tag name. [Learn more](#).
15. You do not need to configure anything in the Integrations section at this time. If you want to, you can configure them later after the Red Hat OpenShift cluster is provisioned.
 - For more information about **Activity tracking** integration, see [this doc](#).
 - For more information about **Logging** integration, see [this doc](#).
 - For more information about **Monitoring** integration, see [this doc](#).
16. Review the **Summary** and click **Create**.

Results

The Red Hat OpenShift cluster starts to be provisioned on IBM Cloud.

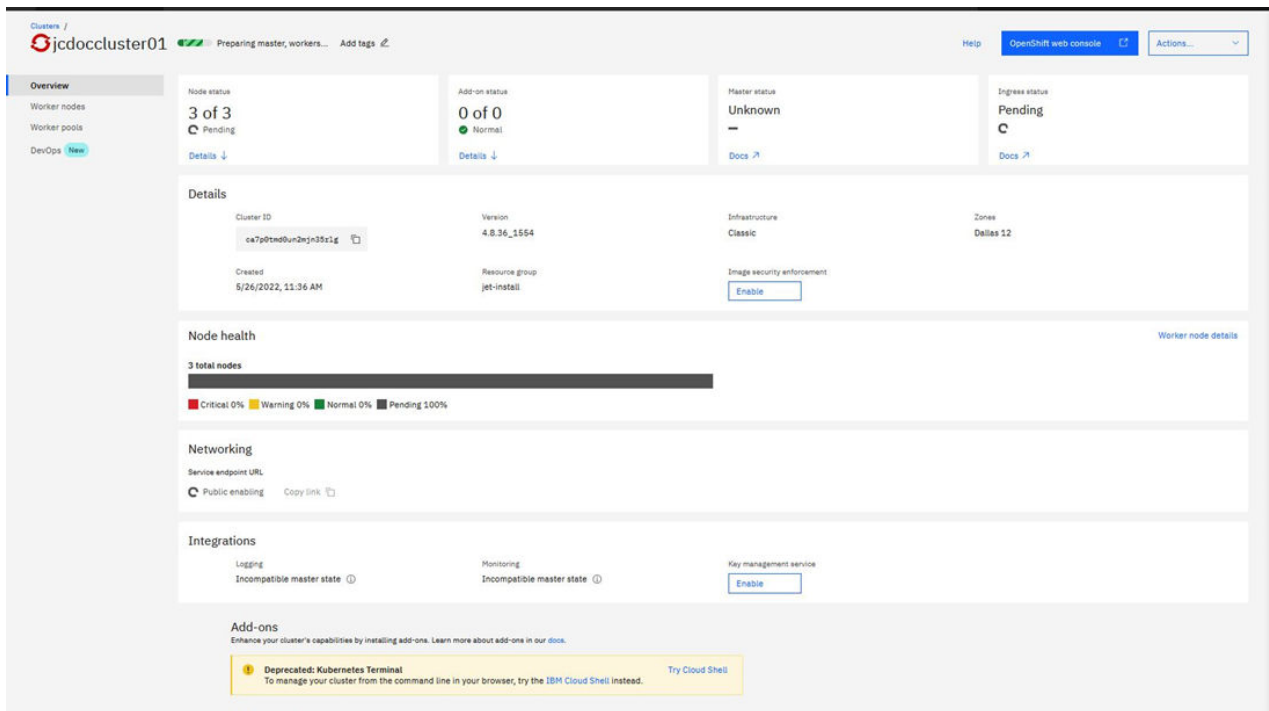
What to do next

You can now monitor the Red Hat OpenShift cluster provisioning.

Monitoring the Red Hat OpenShift cluster provisioning

You can monitor the progress of the cluster while its provisioning after you hit the **Create** button while installing OpenShift Container Platform by using IBM Catalog.

A message that says **Creating** will be shown for some time in the top of the **Create** button. Then, the browser will redirect to the Cluster dashboard, which contain the cluster name in the top left of the page with a message: `Preparing master, workers...` and a progress bar been filled by the green color. You should also see a **Node status** tile showing Pending, a **Master status** tile showing Deploying, and an **Ingress status** tile showing Pending.



After some minutes the message `Finalizing workers...` will be displayed in the side of the progress bar.

Note: During this process you may see some messages indicating the that nodes are in a critical state. You can ignore those messages and wait for the process to complete.

You will see a green icon replacing the progress bar with a `Normal` label in the side of the cluster name. The same green icon should be displayed in the **Node Status** and **Master status** tiles. The **Ingress status** tile might be showing `Unknown` status. It might take several minutes to finish cluster's ingress configuration, but you might be able to access the cluster after some time even while its status still shows the `Unknown` status.

Wait some minutes and try to access the Red Hat OpenShift Container Platform Web Console. Later, if you want to go back to the page from where you can access the Red Hat OpenShift Container Platform Web Console again, refer to [“Accessing Red Hat OpenShift web console”](#) on page 211 for instructions.

Results

The Red Hat OpenShift cluster provisioning is completed.

What to do next

You can now access the Red Hat OpenShift Container Platform Web Console.

Installing Red Hat OpenShift Container Platform by using the IBM Cloud CLI

By using the IBM Cloud CLI you can execute a cluster create command to create the Red Hat OpenShift cluster on IBM Cloud.

Note: For full details on all parameters you can specify while provisioning your Red Hat OpenShift cluster on IBM Cloud using IBM Cloud CLI, run:

```
ibmcloud oc cluster create --help
```


Procedure

1. If not done yet, login in your IBM Cloud Account using your IBM Cloud API key.

```
ibmcloud login --apikey $your_ibmcloud_api_key -q --no-region
```

Where `$your_ibmcloud_api_key` is the API key that you have created in [“Creating your IBM Cloud API key”](#) on page 204.

2. Run the following command:

```
ibmcloud oc cluster create classic --hardware shared --entitlement cloud_pak --name $cluster_name --version $ocp_version --zone $roks_zone --flavor $roks_flavor --workers $roks_workers --private-vlan $private-vlan --public-vlan $public-vlan
```

Where:

Note: --public-vlan and **--private-vlan** parameters: If you do not have a public/private VLAN yet, do not specify this option because one will be automatically created for you.

\$cluster_name

The name of your cluster. For example: `mycluster`.

\$ocp_version

The OpenShift base version to be installed. For example: `4.8_openshift`.

\$roks_zone

The data center region that will host your Red Hat OpenShift cluster. For example: `wdc04`.

\$roks_flavor

The CPU, memory and operating system characteristics of the cluster worker nodes. You can modify it later after your cluster is provisioned at any time. For example: `b3c.16x64.300gb` creates cluster nodes with 16 vCPUs, 64 GB of memory (RAM) and 300 GB of additional storage.

\$roks_workers

The number of worker nodes to be provisioned for your cluster. For example: `8`.

\$private-vlan

Conditional: Specify the ID of the private VLAN. To see available VLANs, run the following command:

```
ibmcloud ks vlan ls --zone $roks_zone
```

Sample output

```
{
  "id": "2981380", # This is the private-vlan ID
  "type": "private",
  "properties": {
    "name": "",
    "note": "",
    "primary_router": "bcr03a.wdc04",
    "vlan_number": "884",
    "vlan_type": "standard",
    "location": "24",
    "local_disk_storage_capability": "true",
    "san_storage_capability": "true"
  }
}
```

\$public-vlan

Conditional: Specify the ID of the private VLAN. To see available VLANs, run the following command:

```
ibmcloud ks vlan ls --zone $roks_zone
```

Sample output

```
{
  "id": "2981378", # This is the private-vlan ID
  "type": "public",
  "properties": {
    "name": "",
    "note": "",
    "primary_router": "fcr03a.wdc04",
    "vlan_number": "837",
    "vlan_type": "standard",
    "location": "24",
    "local_disk_storage_capability": "true",
    "san_storage_capability": "true"
  }
}
```

Results

The Red Hat OpenShift cluster starts to be provisioned on IBM Cloud.

Sample output:

```
root@jcawsb1:~# ibmcloud oc cluster create classic --hardware shared --entitlement cloud_pak --name jcdoccluster03 --version 4.8_openshift --zone dal10 --flavor b3c.4x16 --workers 3 --private-vlan 3218588 --public-vlan 3218586
Creating cluster...
OK
Cluster created with ID cacv29d0f31iv3egkcg
root@jcawsb1:~#
```

What to do next

You can now monitor the Red Hat OpenShift cluster provisioning.

Monitoring the Red Hat OpenShift cluster provisioning by using the IBM Cloud CLI

You can monitor the Red Hat OpenShift cluster provisioning running a command using IBM Cloud CLI.

Procedure

1. If not done yet, login in your IBM Cloud Account using your IBM Cloud API key.

```
ibmcloud login --apikey $your_ibmcloud_api_key -q --no-region
```

Where `$your_ibmcloud_api_key` is the API key that you have created in [“Creating your IBM Cloud API key”](#) on page 204.

2. Run the following command to monitor the progress of the provisioning of the OpenShift cluster:

```
ibmcloud oc cluster get --cluster $cluster_name --output json
```

Where `$cluster_name` is the name of your cluster. For example: `mycluster`

Sample output

```
{
  "location": "Washington D.C.",
  "dataCenter": "wdc04",
  "multiAzCapable": true,
  "vlans": [],
  "worker_vlans": [],
  "workerZones": [
    "wdc04"
  ],
  "id": "ca2k1qgw0ssoqqbpmh0g",
  "name": "mycluster",
  "region": "us-east",
  "state": "normal",
  "status": "All Workers Normal"
}
```

Results

Your Red Hat OpenShift cluster will be ready when the state element value in the sample output is normal.

What to do next

You can now access the Red Hat OpenShift Container Platform Web Console.

Accessing Red Hat OpenShift web console

After the Red Hat OpenShift cluster provisioning is completed you can access the Red Hat OpenShift Administrator Console through the IBM Cloud website. The Red Hat OpenShift web console is a user interface accessible from a web browser. Administrators can use the web console to visualize, browse, and manage the contents of Red Hat OpenShift projects. For more information see [the docs](#).

Procedure

1. Login with your IBM Cloud account in the [IBM Cloud website](#).
2. Click the **Navigation** menu icon, then select **Clusters** in the Red Hat OpenShift console.
3. In the Red Hat OpenShift cluster page, search for the name of the cluster and click its line in the table of clusters.
4. Now in the cluster dashboard, click the **Data**.

Results

The Red Hat OpenShift web console is accessed.

What to do next

Now you can deploy the additional dependencies of Maximo Application Suite.

Installing the Red Hat OpenShift Container Platform Command line Interface

With the OpenShift command-line interface (CLI), the oc command, you can create applications and manage Red Hat OpenShift Container Platform projects from a terminal. For more information about OpenShift CLI see [this doc](#)

Procedure

1. [Access the Red Hat OpenShift Container Platform Web Console](#).
2. Click **Help** and select **Command Line Tools**.
3. Click the appropriate link to download the oc binary that matches with your Operational System and follow its [installation instructions](#) specific to the Operational System of your client machine to be able to run oc commands from a Terminal.

Results

You can now execute oc commands directly from a Terminal on your machine.

What to do next

You can access your Red Hat OpenShift cluster by using the oc command directly from a Terminal in the client machine oc was installed to. In order to do that, first you go in the Red Hat OpenShift web console, then click your Login name and select **Copy login command**. Then, click **View token**. Copy the entire command line under the **Log in with this token** section, paste in the Terminal of the client machine

where oc was installed to and run the command. You should see a message saying that you can access a number of projects and a default project selected. From now on, next time you run oc commands in this opened Terminal in the client machine, it will be running the oc commands in your Red Hat OpenShift cluster.

Adding compute nodes to an existing Red Hat OpenShift Container Platform installation

As additional Maximo Application Suite applications are deployed or as the workload increases, it may be necessary to provision additional compute nodes.

The following are good indicators that it is time to add additional one or more compute node to the cluster:

- Pods cannot be scheduled due to insufficient CPU.
- Pods cannot be scheduled due to insufficient memory.
- Pods are evicted due to disk pressure on compute nodes.

The following instructions show how to add an additional compute node to the existing Red Hat OpenShift Container Platform (OCP) cluster on IBM Cloud.

Procedure

1. Login with your IBM Cloud account in the [IBM Cloud website](#).
2. Click the **Navigation** menu and go to **Resource List**. Filter and select your Red Hat OpenShift cluster by its cluster name, or search for it under the **Clusters** section. When found, click the cluster name in the table of resources.
3. Go to **Worker pools**. Then, you have two options to scale the cluster's capacity:
 - a) Click the three-dot menu of the worker pool that is listed for your Red Hat OpenShift cluster and choose **Resize**. You can scale up or down the current number of worker nodes. Scaling up the number of worker nodes provision new workers and increase cluster's overall resource capacity. In the other hand, scaling down the worker nodes reduce the cluster's overall capacity and might cause performance issues or even downtime to current services and applications deployed in the cluster. To confirm your changes, click **Resize**.
 - b) Or still under **Worker pools**, another option to increase cluster's capacity is to click **Add +** and configure a new worker pool. You can define new worker zones if you want to expand cluster's availability and new worker nodes if you plan to upgrade your existing cluster's capacity. To confirm your changes click **Create**.

Operational mode for installation

From Maximo Application Suite 8.9 or later, you can consider installing and deploying IBM Maximo Application Suite in Production or Non-production mode based on your development and testing requirements and to optimize your AppPoint usage.

You can specify deployments for production or non-production environments by using the **Operational mode** parameter option during installation.

Non-production installations can be used for internal development and testing. All applications, add-ons, and solutions have 0 (zero) install AppPoints in the non-production installations. These specifications are also visible in the metrics that are shared with IBM and on the product UI. Users in non-production environments require AppPoints.

Note:

- The **Operational Mode** parameter is available only for installation. It is not supported for upgrading the IBM Maximo Application Suite.
- By default, production mode is selected as the operational mode for installation.

Related concepts

[IBM Maximo Application Suite installation with Ansible collection](#)

[Installing the Maximo Application Suite on Amazon Web Services](#)

Related tasks

[Installing Maximo Application Suite](#)

Related reference

[Maximo Application Suite Ansible collection examples](#)

IBM Suite License Service

The IBM Suite License Service stores and manages the license of your IBM solution. After you install Suite License Service as part of your IBM solution, you upload the license file to the Suite License Service server.

- [Overview of Suite License Service](#)
- [Installing Suite License Service](#)
- [Configuring Suite License Service](#)
- [Supported pods for workload customization](#)
- [Upgrading Suite License Service](#)
- [Backup and restore](#)

High availability

IBM Maximo Application Suite is an enterprise-grade set of services that are highly scalable for managing large numbers of assets, work orders, incoming sensor data, visual images, and more. The scale and advanced capability bring some degree of service complexity along with multiple persistence stores, each optimized for specific operational performance.

High availability implies redundancy that supports failover. Prolonged application or system outages can have significant business consequences. You can increase the resiliency of the systems to environmental, hardware, or software faults or failure. Techniques and architectural patterns can minimize the impact on the overall system if a planned or unplanned system outage occurs. High availability techniques cannot be implemented at the detriment of other architectural considerations, such as capital cost, ongoing total cost of ownership, or manageability.

Logical architecture

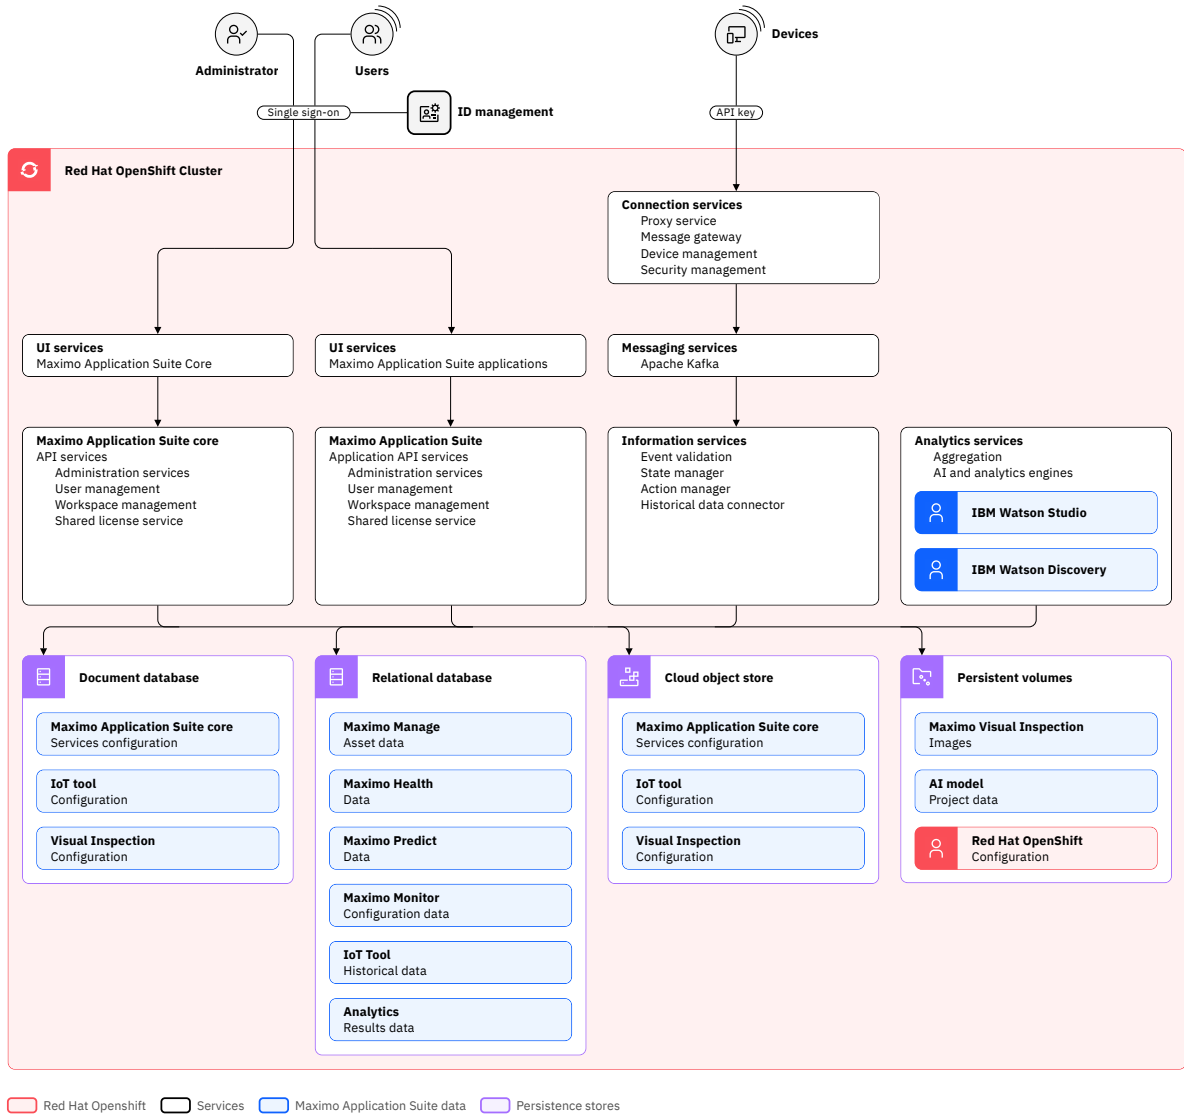
High availability is provided by the main components and the way that they interact in the logical architecture of Maximo Application Suite. You have a choice of multiple application services that provide the primary features for managing, monitoring, inspecting assets, and determining their health and predicting when they require maintenance.

The Maximo Application Suite core services are deployed automatically with any instance. These services handle the basic administration and configuration of the suite, and they store metadata in a deployed MongoDB in a three-node instance.

 [Open image in new tab.](#)

Application services are containerized, assigned to pods, and configured through Kubernetes operators to start and stop based on their configuration.

[Open image in new tab.](#)



Each of these application services uses their own persistence stores that have some flexibility for the instance location. In particular, the application state and user data are spread across the following types of persistence stores:

Document database

A MongoDB variant, which can be enterprise, community, or AWS Document DB, where most of the metadata, preferences, configuration settings, user, and security management are kept. This database is configured with multinode replication and without sharding, which improves resiliency. It also provides backup and restore utilities.

Note: Starting in 8.11, for Amazon Web Services US GovCloud regions, Document DB is not supported.

Relational database

Typically an IBM Db2 database or Oracle Database that holds most of the user data that is necessary for application function, with complete atomicity, consistency, isolation, and durability (ACID) transactional control. Backup services for full and incremental backups and restore services can also apply transaction log changes from the last backup point. This choice enables a wide range of usage patterns and integration with BI Reporting and Replication tools.

Cloud object storage

IBM Cloud Object Storage, AWS S3, or similar object storage that stores big files, which can range from attachments to video images to the actual backup images. The goal is to provide inexpensive storage for large content with high reliability. Because of its high availability through automatic replication, it doesn't require backups, but you might prefer to enable versioning to protect from certain scenarios.

Red Hat OpenShift persistence storage

etcd or other persistent volumes that are attached to the worker nodes that are used by Red Hat OpenShift to operate. This option requires special backup and restore operations that range from Kubernetes operator logic to storage-level utilities.

Given this breadth of data, the processes of backup and recovery are complicated and must be coordinated carefully.

Resilient architecture components

Maximo Application Suite provides resiliency through suite service instances, availability zones, and storage.

Suite service instance resilience

Building Maximo Application Suite on the the Red Hat OpenShift technology stack by using containers and Kubernetes has several advantages. A key advantage is the ability to configure services to automatically restart after a failure and keep multiple instances of the service in operation. This ability is used across the suite, based on the workload size, for resiliency of the system with complete automation. You don't need to define an alert and wait for it to notify operations personnel to restart a service while users wait.

Using replicas, Kubernetes ensures that the configured number of pods are available. This configuration and current operations status are available by using standard Red Hat OpenShift APIs and user interfaces. This configuration creates redundancy of the services within a single physical data center.

Resilience by using availability zones

The purpose of using availability zones on the cloud is to create redundancy of the services across physical data centers, but within low latency network connectivity.

Note: An on-premises deployment requires a similar physical configuration.

Red Hat OpenShift worker nodes are configured across the zones, and Kubernetes automatically schedules redundancy for the different pods.

In addition to the application services, the data services need to be spread across these zones with an adequate level of replication configured. Each data service in Maximo Application Suite has different replication features.

High availability strategy for each type of data

High availability strategies are not available for custom code, running state, and runtime data. The following high availability strategies are available for other types of data:

Application code

Product images can use pods to create multiple redundant copies of the critical microservices that are spread across availability zones.

Red Hat OpenShift Container Platform restarts less critical pods.

Configuration data

Kubernetes configuration secrets and configuration maps are held in etcd, which uses mirroring that is set up by Red Hat OpenShift Container Platform.

Other configuration data is held in MongoDB, which uses mirroring that you set up.

Prerequisites for high availability in Maximo Application Suite

<i>Table 18. High availability prerequisites</i>		
Prerequisite	High availability	References
Red Hat OpenShift	<ul style="list-style-type: none"> • Use a minimum of three availability zones. • Label nodes by using <code>topology.kubernetes.io/zone</code>. 	<ul style="list-style-type: none"> • Red Hat OpenShift Container Platform: Controlling pod placement by using pod topology spread constraints • Kubernetes: Pod topology spread constraints • Kubernetes: Well-known labels, annotations and taints
File system	<ul style="list-style-type: none"> • Red Hat OpenShift Data Foundation (ODF) (OCS) • Portworx 	<ul style="list-style-type: none"> • Understanding OpenShift Data Foundation • Storing data on classic IBM Cloud File Storage • Portworx instructional videos
Db2 Warehouse	<ul style="list-style-type: none"> • IBM Db2 Warehouse SMP high availability and disaster recovery • IBM Data Replication for Db2 Continuous Availability • Built-in high availability feature for IBM Db2 Warehouse massive parallel processing (MPP) deployments with highly available cluster file system across availability zones (ODF or Portworx) 	<ul style="list-style-type: none"> • Db2 Warehouse high availability disaster recovery • Setting up replication • Built-in high availability feature for IBM Db2 Warehouse MPP deployments
MongoDB	<ul style="list-style-type: none"> • Use a replica set with one primary member and two secondaries in each availability zone. 	
Kafka	<ul style="list-style-type: none"> • Strict-based rack aware • Kafka rack: <code>topologyKey: topology.kubernetes.io/zone</code> 	<ul style="list-style-type: none"> • Configuring Strimzi • Using Strimzi

Document database resilience

MongoDB technology is used for its flexibility in schema and resilience, which provides higher availability for the core services and basic operations. A single node failure does not stop activity in Maximo Application Suite

Schema flexibility enhances the product without requiring complex update processes to this database. This flexibility makes the core services more stable and able to direct the updates of the application services during the update process.

Resilience is achieved by using multiple nodes that handle connections and replicate data between the nodes. Maximo Application Suite uses a write-to-primary-only approach with a single data shard to simplify processing because this database has a relatively light transaction load.

Single-node failure

When the primary node fails, one of the remaining secondary nodes is selected as the new primary node, and operations continue. The failed node is eventually restarted, and replication updates it as part of normal operation. Because transactions require a simple majority for commit, this loss of a single node does not require an operational restore action. When a secondary node fails, the operation continues as normal with the same primary node. Again, the failed node is eventually restarted, and replication updates it.

The database node instance failure and the node's disk failure are different. If the failed node is using reliable storage, replication has less data to update when the instance is restarted.

Multinode failure

This situation is rare and can be even rarer if nodes are spread across availability zones. If more than one node fails, including the primary node, a situation where a backup is used to restore the database and restart might occur. In this rare case, loss of data might occur. While a document database activity log exists to help, it does not automate a forward recovery of transactions to the point of failure.

Relational database resilience

A relational database handles most of the data and transaction volume, usually IBM Db2, Oracle Database, or SQL Server.

The workloads that are involved vary depending on which application services are used and how they are used.

- Maximo Manage is the most widely used application and has traditional forms that are processed for work orders, service requests, and updates to asset information and their maintenance.
- Maximo Monitor uses high volume, time-stamped device metrics that are loaded into the database, with aggregation and analytics queries across that data.
- Maximo Visual Inspection processes data at the edge, with more involved model training and advanced GPU processing, and Maximo Health and Predict - Utilities has similar AI model components.

It is possible that all of these workloads share a single relational database management system (RDBMS) instance or are spread across multiple RDBMS instances. Either choice requires a backup scheduling strategy because the operation can be expensive in large volume instances.

Cloud object storage resilience

Cloud object storage is a persistence store, not a database, so its behavior and processes are different. To select a strategy for cloud object storage, you must consider the importance of data content versus storage space expense.

Data categories that use this persistence store type:

- Attachments for Maximo Manage application services, such as receipts, certifications, and invoices
- Backup files for restore
- Data that is related to historical scores in Maximo Health and Maximo Health and Predict - Utilities

Note: On-premises deployment requires a similar strategy.

Red Hat OpenShift persistent resilience

You can use Amazon Elastic Block Store (EBS) or network-attached storage for the ability to access files from worker nodes in different availability zones. Storage choices can provide built-in redundancy for greater hardware protection.

Installing Maximo Application Suite

To install IBM Maximo Application Suite, plan your installation, select a supported installation path, and set up the initial configuration.

Supported installation paths

Use the Maximo Application Suite command line interface (CLI) for a standard installation on multiple platforms such as x86/amd64, IBM Z, and IBM Power. Use the Ansible DevOps collection with CLI for advanced users to build installation topologies beyond what is possible with the Maximo Application Suite CLI.

Alternatively, customize the Maximo Application Suite installation with Amazon Web Services CloudFormation templates for Amazon Web Services user accounts and Microsoft Azure Resource Manager templates for Microsoft Azure user accounts.

The standard installation with CLI works on any Red Hat OpenShift Container Platform instance, such as Installer Provisioned Infrastructure (IPI), User Provisioned Infrastructure (UPI), Single Node Red Hat OpenShift, or IBM Cloud Kubernetes Service. The customized installations by using Amazon Web Services CloudFormation and Microsoft Azure Resource Manager templates work with Red Hat OpenShift Container Platform instances for existing, IPI, or UPI Red Hat OpenShift cluster.

For more information, see [Detailed system requirements](#).

Standard installation with IBM Maximo Application Suite CLI

You can install the IBM Maximo Application Suite by using a command-line interface (CLI) utility.

Before you begin

You can use IBM's container image to install Maximo Application Suite. The container image provides an out of the box environment for managing Maximo Application Suite on Red Hat OpenShift, with numerous dependencies preinstalled.

In a scenario where you need to install the CLI utility, which is an open source tool, on your local system, you must configure the following software.

- Bash (v4)
- Red Hat OpenShift client
- IBM Cloud client with container plug-in enabled

Note: IBM Cloud is not required if you are deploying Maximo Application Suite on an organization's Red Hat OpenShift cluster.

- Ansible
- Python
- Network access to the Red Hat OpenShift cluster

Download the Maximo Application Suite CLI utility onto the bastion host. This utility provides commands to manage the local docker registry, configure policies in the Red Hat OpenShift cluster, and deploy Maximo Application Suite.

For more information about installing the utility, see [Maximo Application Suite CLI Utility](#).

To use the Maximo Application Suite CLI utility, ensure that the bastion host has support for running docker containers.

For more information, see [“Planning for IBM Maximo Application Suite standard installation with CLI” on page 132](#).

Procedure

1. Run the **docker run** command to use the container image that is published.

```
docker run -ti --rm -v ~/.mnt/home --pull always quay.io/ibmmas/cli
```

If you want a specific release of the image, use a specific version tag in the **docker pull** command :

```
docker run -ti -v ~/.mnt/home quay.io/ibmmas/cli:x.y.z
```

For more information on the software included in the container image, see [MAS CLI Base Image](#).

2. Install Maximo Application Suite.

Run the following command from a running CLI docker container.

```
docker run -ti --pull always quay.io/ibmmas/cli mas install
```

For more information, see the [Install command](#).

Tip: To customize Maximo Application Suite instances with extra features that might be unavailable in Maximo Application Suite CLI, you can use ansible-devops for advanced configuration. For more information, see [Maximo Application Suite Devops Ansible Collection](#).

For a list of supported CLI commands, see [CLI commands](#).

You can also run the installation in interactive mode. For more information, see [Interactive install](#).

What to do next

Complete the Maximo Application Suite setup. For more information, see the following topics.

[“Authentication methods” on page 605](#)

Maximo Application Suite supports MongoDB, Lightweight Directory Access Protocol (LDAP) authentication, and Security Assertion Markup Language (SAML) authentication methods for local user authentication.

[“LDAP user registry synchronization” on page 619](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Administering users and user access

The initial superuser account is used to complete the Maximo Application Suite setup. You can add application administrator users or system administrator users for day-to-day administrative tasks.

- [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#)
- [“Administering users and user access in Maximo Application Suite in 9.1” on page 781](#)

[“Getting started” on page 130](#)

With the setup completed, your users can log in and start using Maximo Application Suite.

Related concepts

[Planning for IBM Maximo Application Suite standard installation with CLI](#)

Before you install IBM Maximo Application Suite using the Maximo Application Suite CLI utility, you must consider the dependencies that must be met for installation.

[Foundation service](#)

Related information

[Migration from Maximo Asset Management 7 to Maximo Application Suite 9](#)

[Software Product Compatibility Report](#)

IBM Maximo Application Suite installation in disconnected environments

You can install IBM Maximo Application Suite in an air gap environment, which is also known as disconnected, offline, or restricted network.

Before you begin

Before you install Maximo Application Suite, ensure you download the necessary software and set up the environment.

You must have access to the following components:

- A private image registry setup and running in the restricted network, and secured with certificates. Configure one of the following options:
 - – A bastion host with access to product images on the internet and the restricted network and has support for running docker containers. The docker image that contains the IBM Maximo Application Suite command line utility on the bastion host.
 - – A host outside the restricted network with access to product images on the internet and with support for running docker containers. The docker image that contains the Maximo Application Suite command line utility on the host. Portable disk space sufficient to store the required images. A host inside the restricted network with support for running docker containers and can access the images downloaded to the portable disk space.
- Red Hat OpenShift cluster setup as an air gap cluster for disconnected installation.

Ensure you use the Maximo Application Suite sizing calculator to estimate your Red Hat OpenShift Worker Node configuration, storage, and memory requirements.

For more information, see [“Requirements and capacity planning” on page 178](#).

- IBM entitlement key.
- IBM Maximo Application Suite license file.

For more information, see [“Prerequisites for installing” on page 133](#).

Note: It is recommended that you follow the Red Hat OpenShift instructions for setting up a docker registry.

An alternative is deploying the docker registry into a separate Red Hat OpenShift cluster. This scenario requires two Red Hat OpenShift clusters. One Red Hat OpenShift cluster for the docker registry and a second air gap Red Hat OpenShift cluster for Maximo Application Suite.

You can run the following command to deploy a docker registry into an Red Hat OpenShift cluster.

```
docker run -ti --rm --pull always quay.io/ibmmas/cli mas setup-registry
```

Remember: You must not use the same Red Hat OpenShift cluster for both, the docker registry and Maximo Application Suite.

About this task

Run the following docker commands on a bastion host to install Maximo Application Suite.

Procedure

1. Select one of the static catalogs and automatic approval strategy for your installation.
For more information, see [Catalog selection](#).

2. Mirror the container images.

You can use three modes to mirror the container images.

- **direct** mirrors images directly from the source registry to your private registry.
- **to-filesystem** mirrors images from the source to a local directory.
- **from-filesystem** mirrors images from a local directory to your private registry.

Run the following command to mirror the product images that is necessary to install and run Maximo Application Suite.

```
docker run -ti --pull always quay.io/ibmmas/cli mas mirror-images
```

Tip: You can also use this command to mirror the images for Red Hat OpenShift.

The **mirror-images** command accepts the name of a static catalog to control what is mirrored to your registry. It mirrors the necessary images of the latest package version in that catalog in your private registry.

You are prompted to set the target registry for mirroring the image to choose a catalog and the subset of content that you want to mirror. You can either mirror everything from the catalog or control exactly what is mirrored to your private registry. Controlling what is mirrored reduces the registry storage requirements and the time and bandwidth that is used to mirror the images. For more information, see [Catalog selection](#).

For example, run the **mirror-images** command non interactively to directly mirror images from the internet to your private registry using a bastion host.

```
mas mirror-images \  
-m direct \  
-H myprivateregistry.com -P 5000 -u $REGISTRY_USERNAME -p $REGISTRY_PASSWORD \  
-c v8-221025-amd64 --mirror-core --mirror-iot --mirror-optimizer --mirror-manage \  
--ibm-entitlement $IBM_ENTITLEMENT_KEY \  
--redhat-username $REDHAT_USERNAME --redhat-password $REDHAT_PASSWORD \  
--no-confirm
```

For more information, see [Mirror images command](#).

3. Configure Red Hat OpenShift to use your Private Registry for Maximo Application Suite.

Your cluster must be configured to use the private registry as a mirror for the Maximo Application Suite container images. An `ImageContentSourcePolicy` named `mas-and-dependencies` is created in the cluster, which is also the resource that the Maximo Application Suite install uses to detect whether the installation is a disconnected installation and tailors the options that are presented when you run the **mas install** command.

Run the following command from a running CLI docker container to configure the Red Hat OpenShift instance to use your Maximo Application Suite private registry.

```
docker run -ti --pull always quay.io/ibmmas/cli mas configure-airgap
```

Provide information about the private registry, including the CA certificate necessary to configure your cluster to trust the private registry.

You can run the command in noninteractive mode.

```
mas configure-airgap \  
-H myprivateregistry.com -P 5000 -u $REGISTRY_USERNAME -p $REGISTRY_PASSWORD \  
--ca-file /mnt/local-mirror/registry-ca.crt \  
--no-confirm
```

For more information, see [Configure air gap command](#).

4. Install Maximo Application Suite.

Run the following command from a running CLI docker container to install Maximo Application Suite.

```
docker run -ti --pull always quay.io/ibmmas/cli mas install
```

For more information, see the [Install command](#).

What to do next

Perform the initial configuration tasks for setting up the Maximo Application Suite. For more information, see [“Setting up IBM Maximo Application Suite”](#) on page 281.

Starting in Maximo Application Suite 9.0, 8.11.7, 8.10.10 or later versions, you can also upload your usage data to IBM by using the IBM Data Reporter Operator . For more information, see [“Data Reporter Operator”](#) on page 7.

Related concepts

[Prerequisites for installing](#)

Before you begin, ensure that your environment meets the prerequisites by downloading, and installing the software and interfaces that you use to install IBM Maximo Application Suite.

IBM Maximo Application Suite installation with Amazon Web Services CloudFormation templates

You can install IBM Maximo Application Suite in the Amazon Web Services (AWS) cloud by using the Amazon Web Services CloudFormation templates. Maximo Application Suite is available as a bring-your-own-license (BYOL) and contract pricing product in AWS Marketplace. After you configure the installation requirements and consider your installation preferences, you subscribe to the product, specify the installation parameters, and start the installation.

In your AWS account, the installation process creates the virtual network infrastructure and the Red Hat OpenShift cluster, and then installs the Maximo Application Suite prerequisites and Maximo Application Suite. If you configured a verified email address in the Amazon simple email service (SES), you receive emails that contain the information that you need to access Maximo Application Suite.

Installing the Maximo Application Suite on Amazon Web Services

To install IBM Maximo Application Suite on Amazon Web Services (AWS), you configure the prerequisite components, consider your installation preferences, specify the installation criteria in an AWS CloudFormation stack template, and create the stack.

You must configure prerequisites and gather the information that you need to specify the installation parameters. For more information, see [“Planning to install on Amazon Web Services”](#) on page 138.

In addition, you must consider your installation preferences, such as the type of Maximo Application Suite offering that you want and whether you want to create a Red Hat OpenShift cluster or reuse an existing cluster.

Note: The existing cluster must be created by using the automated deployment option.

For more information, see [Installation considerations](#).

Maximo Application Suite on Amazon is available as a bring your own license (BYOL) and as a paid offering.

Related concepts

[Operational mode for installation](#)

From Maximo Application Suite 8.9 or later, you can consider installing and deploying IBM Maximo Application Suite in Production or Non-production mode based on your development and testing requirements and to optimize your AppPoint usage.

Installing BYOL IBM Maximo Application Suite

The IBM Maximo Application Suite (BYOL) can be installed from the Amazon Web Services Marketplace based on your infrastructure needs. You can install a new Red Hat OpenShift cluster before you install Maximo Application Suite or you can use your existing Red Hat OpenShift cluster to install Maximo Application Suite.

Before you begin

Before you can install Maximo Application Suite on Amazon Web Services, you must configure prerequisites and gather information that you need to complete the installation.

For more information, see [“Prerequisites for installing Maximo Application Suite on Amazon Web Services”](#) on page 143.

Consider other criteria, such as the type of Maximo Application Suite offering that you want and whether you want to create a Red Hat OpenShift cluster or reuse an existing one.

Note: The existing cluster must be created by using the automated deployment option only.

For more information, see [“Preparing to install Maximo Application Suite on Amazon Web Services” on page 146](#).

This product requires an internet connection to deploy properly. The following code is downloaded on deployment for setting up the Maximo Application Suite Red Hat OpenShift Container Platform cluster:

- <https://github.com/ibm-mas/ansible-devops.git>
- <https://github.com/ibm-mas/multicloud-bootstrap.git>

About this task

The following three fulfillment options (CloudFormation templates) are available for installing Maximo Application Suite:

1. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
2. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)
3. Existing Red Hat OpenShift cluster

Note: Starting in 8.11, for US GovCloud regions, you can install Maximo Application Suite in private hosted zones for existing Red Hat OpenShift cluster and New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI).

Procedure

1. In the AWS Marketplace service console, click **Discover product** and search for IBM Maximo Application Suite (BYOL). The product is sold by IBM Maximo.
2. Open Maximo Application Suite.
3. Review the product information to select a fulfillment option and click **Continue to Subscribe**.
4. In the subscription page, to create the subscription, review the terms and conditions and click **Accept Terms**.
5. After the subscription is created, click **Continue to Configuration**.
6. In the configuration page, in the **Region** field, select a supported geographical region where you want to install the Maximo Application Suite.
For the list of supported regions, see [“Amazon Web Services region for installing Maximo Application Suite” on page 147](#).
7. Accept the default values in the other fields and click **Continue to Launch**.
8. In the **Launch** page, select **Choose Action > Launch CloudFormation**.
9. To open the CloudFormation stack wizard, click **Launch**.
10. In the **Create stack** wizard step, accept the default values, and click **Next**.
11. In the **Specify stack details** step, in the **Stack name** section, enter a unique name.
12. In the **Parameters** section, enter the installation parameters by using the information that you gathered when you configured the [“Prerequisites for installing Maximo Application Suite on Amazon Web Services” on page 143](#) and [“Preparing to install Maximo Application Suite on Amazon Web Services” on page 146](#).

a) Enter the mandatory parameters.

- The Maximo Application Suite offering type, such as Maximo Application Suite Core and Cloud Pak for Data or Maximo Application Suite and Maximo Manage.

Starting in 8.11, for US GovCloud regions, you can configure the offering type Maximo Application Suite Core or Maximo Application Suite and Maximo Manage. Cloud Pak for Data is not available for configuration.

- All parameters in the **Cluster and bootnode access** and **Keys and licenses** sections.

b) To reuse an existing Red Hat OpenShift cluster, complete the [Existing Red Hat OpenShift cluster connection details](#) section.

Alternatively, to create a new cluster, complete the steps that are given in [“Connection details for an existing Red Hat OpenShift cluster”](#) on page 149 section.

- c) To reuse an existing network infrastructure, complete the steps that are given in [Connection details for using an existing network infrastructure](#) section.
- d) In the optional parameter groups, such as the group of IBM Maximo Manage database configuration parameters, ensure that you either specify all parameter values or leave all empty.

Note: Follow the instructions commented in the field to know how to fill them. For the field `MASManageDBJdbcUrl`, you can specify it by using one of the following JDBC URL formats. Ensure that the Port that is used contains the SSL enabled port of the database.

IBM Db2

For Db2 SSL database connections in Maximo Manage, you must specify `sslConnection=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB:sslConnection=true;
```

Oracle Database

Starting in 8.11, for US GovCloud regions, you can configure an Oracle Database. You must ensure that the Oracle Database follows the Federal Information Processing Standard (FIPS).

For more information, see [“Configuring Oracle Database ”](#) on page 312

Note: If you choose to install Maximo Application Suite with Manage, you can use the default IBM Db2 instance that is provisioned by IBM instead of configuring your own external Db2 instance. To configure the default IBM Db2, do not add information in the username, password, JDBC URL, certificate URL, and demo data fields.

The internal Db2 configuration is available from Maximo Application Suite 8.10 or later.

13. To configure Amazon Managed Streaming for Kafka, select **Yes**.

The Amazon Managed Streaming for Kafka is configured to process streaming data of applications such as IoT and IBM Maximo Monitor from Maximo Application Suite.

You can configure the default Amazon Managed Streaming from Maximo Application Suite 8.10.

Starting in 8.11, for US GovCloud regions, Amazon Managed Streaming for Kafka configuration is not required.

14. To configure a DocumentDB instance, you can select any of the following options.

- Configure a new MongoDB community edition.
- Use an existing MongoDB

Add existing MongoDB connection details such as the username, password, MongoDB hostname, and CA certificate.

Starting in 8.11, for US GovCloud regions, you must ensure your existing MongoDB connection is FIPS compliant.

- Configure a new Amazon DocumentDB.
- Use an existing Amazon DocumentDB.

Add existing DocumentDB connection details such as the username, password, DocumentDB hostname, and CA certificate.

Add the VPC ID of the region in which the Amazon DocumentDB is deployed. Ensure that the DocumentDB does not have a matching or overlapping IPv4 CIDR block `10.0.0.0/16`.

You can configure the default DocumentDB from Maximo Application Suite 8.10.

Starting in 8.11, for US GovCloud regions, DocumentDB is not available for configuration.

15. In the **Other parameters**, select the **OperationalMode** as Production or Non-production.

You can specify deployments for Production or Nonproduction environments. Nonproduction installations can be used for internal development and testing. The installation AppPoints are unused in the Nonproduction installations. These specifications are also visible in the metrics that are shared with IBM and on the product UI.

16. Click **Next**.
17. In the **Configure stack options** step, configure any additional options that you require. To know more about the stack options, click the "Learn more" links available in each option. Click **Next** when done.
18. In the **Review** step, review the values that you entered and acknowledge the message that relates to identity and access management (IAM) resources.
19. To begin the installation, click **Create stack**.

What to do next

During the installation process, the AWS CloudFormation stack template that you configured is used to create a Bootnode. The Bootnode contains all required resources to complete the installation. To verify that the Bootnode is created successfully, in the **CloudFormation > Stacks** page, confirm that the stack status is updated to `CREATE_IN_PROGRESS`

For more information, see [“Monitoring IBM Maximo Application Suite installation on Amazon Web Services”](#) on page 238.

Installing client managed IBM Maximo Application Suite for public paid offer

The paid AWS product can be installed from the AWS Marketplace based on your infrastructure needs. A new Red Hat OpenShift instance is created as part of the Maximo Application Suite installation.

Before you begin

Before you can install Maximo Application Suite on Amazon Web Services, you must configure prerequisites and gather information that you need to complete the installation.

For more information, see [“Prerequisites for installing Maximo Application Suite on Amazon Web Services”](#) on page 143.

Consider other criteria, such as the type of Maximo Application Suite offering that you want and whether you want to create a Red Hat OpenShift cluster or reuse an existing one.

Note: The existing cluster must be created by using the automated deployment option only.

For more information, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

This product requires an internet connection to deploy properly. The following code is downloaded on deployment for setting up the Maximo Application Suite Red Hat OpenShift Container Platform cluster:

- <https://github.com/ibm-mas/ansible-devops.git>
- <https://github.com/ibm-mas/multicloud-bootstrap.git>

About this task

The Maximo Application Suite client-managed solution is available with two different Marketplace products.

Maximo Application Suite client-managed with Red Hat OpenShift entitlement

This product includes the subscription for the Red Hat OpenShift cluster that is deployed during the installation process.

The following CloudFormation templates are available as fulfillment options:

- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

Maximo Application Suite client-managed without Red Hat OpenShift entitlement

This product does not include the subscription for the Red Hat OpenShift cluster that is deployed during the installation process or the existing one provided by the user. You must have your own Red Hat OpenShift subscription to deploy Maximo Application Suite by using this product.

The following CloudFormation templates are available as fulfillment options:

- An existing Red Hat OpenShift cluster
- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

After you select the Paid product from Amazon Web Services Marketplace, you can complete the following steps for a public paid offering. The public paid offer is the default option that you see after you view the product page.

Procedure

1. In the Amazon Web Services Marketplace service console, depending on if you have the Red Hat OpenShift subscription or not, click **Discover product** and search for Maximo Application Suite client-managed without Red Hat OpenShift entitlement or Maximo Application Suite client-managed with Red Hat OpenShift entitlement. The products are sold by IBM Maximo.
2. Open Maximo Application Suite.
3. Review the product information and click **Continue to Subscribe**.
4. On the **Create an agreement for this software** page, select any one of the contract duration options: 12 months / 24 months / 36 months.

5. Select a minimum of **500 AppPoints** and click **Create Contract**.

6. After 24 hours of subscription, to register the product with your Amazon Web Services account, complete the registration by using an IBMid from <https://www.ibm.com/marketplace/connector/landing/aws/services/amicontractsetup>.

You receive an email notification in few minutes to obtain the artifacts required for the product deployment.

7. After you retrieve the entitlement key, link your Maximo Application Suite subscription with Red Hat OpenShift account by following the steps at <https://www.ibm.com/docs/en/cloud-paks/1.0?topic=iocpc-accessing-red-hat-entitlements-from-your-cloud-paks>, obtain the Red Hat OpenShift pull secret from the Red Hat OpenShift account, return to product page, and click **Continue to Configuration**.

Note: Linking the Maximo Application Suite subscription with Red Hat OpenShift account is required only if you are deploying Maximo Application Suite with new Red Hat OpenShift cluster. If you decide to use existing Red Hat OpenShift cluster, it is not required.

8. In **Fulfillment** option, select the appropriate deployment mode. For Maximo Application Suite deployment with new Red Hat OpenShift cluster, you can select either new network infrastructure (IPI), existing network infrastructure (UPI), or existing Red Hat OpenShift cluster. For more information, refer to the fulfillment options explained earlier in this section.
9. In **Software version**, select the latest version.
10. In **Region**, select a supported geographical region where you want to install the Maximo Application Suite. For the list of supported regions, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.
11. Accept the default values in the other fields and click **Continue to Launch**.
12. In the **Launch** page, select **Choose Action > Launch CloudFormation**.
13. To open the CloudFormation stack wizard, click **Launch**.
14. In the **Create stack** wizard step, accept the default values, and click **Next**.
15. In the **Specify stack details** step, in the **Stack name** section, enter a unique name.
16. In the Parameters section, enter the installation parameters by using the information that you gathered when you configured the [“Prerequisites for installing Maximo Application Suite on Amazon](#)

[Web Services](#)” on page 143 and [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

17. Enter all of the mandatory parameters.

- The Maximo Application Suite offering type, such as Maximo Application Suite Core and Cloud Pak for Data or Maximo Application Suite and Maximo Manage.

Note: If you choose to install Maximo Application Suite with Manage, you can use the default IBM Db2 instance that is provisioned by IBM instead of configuring your own external Db2 instance. To configure the default IBM Db2, do not add information in the username, password, JDBC URL, certificate URL, and demo data fields.

The internal Db2 configuration is available from Maximo Application Suite 8.10 or later.

Starting in 8.11, for US GovCloud regions, Db2 is not supported.

- All of the parameters in the **Cluster and bootnode access** and **Keys and licenses** sections.

18. To create a new cluster, complete the [New OpenShift cluster configuration details](#) section.

19. In the optional parameter groups, such as the group of Maximo Manage database configuration parameters, ensure that you either specify all of the parameter values or leave all of them empty.

20. To configure Amazon Managed Streaming for Kafka, select **Yes**.

The Amazon Managed Streaming for Kafka is configured to process data streaming of applications such as IoT and IBM Maximo Monitor from Maximo Application Suite.

You can configure the default Amazon Managed Streaming from Maximo Application Suite 8.10. or later.

21. To configure a DocumentDB instance, you can select any of the following options.

- Configure a new MongoDB community edition
- Use an existing MongoDB

You must input existing MongoDB connection details such as the username, password, MongoDB hostname, and CA certificate.

- Configure a new Amazon DocumentDB
- Use an existing Amazon DocumentDB

Add existing DocumentDB connection details such as the username, password, DocumentDB hostname, and CA certificate.

You must input the VPC ID of the region in which the Amazon DocumentDB is deployed. Ensure that the DocumentDB does not have a matching or overlapping IPv4 CIDR block 10.0.0.0/16.

You can configure the default DocumentDB from Maximo Application Suite 8.10. or later.

22. In the **Other parameters**, select the **OperationalMode** as Production or Non-production.

You can specify deployments for Production or Non-production environments. Non-production installations can be used for internal development and testing. The installation AppPoints are unused in the Non-production installations. These specifications are also visible in the metrics shared with IBM as well as on the product UI.

23. Click **Next**.

24. In the **Configure stack options** step, configure any additional options that you require. To know more about the stack options, click the "Learn more" links available in each option. Click **Next** when done.

25. In the **Review** step, review the values that you entered and acknowledge the message that relates to identity and access management (IAM) resources.

26. To begin the installation, click **Create stack**.

What to do next

During the installation process, the AWS CloudFormation stack template that you configured is used to create a Bootnode. The Bootnode contains all required resources to complete the installation. To verify

that the Bootnode is created successfully, in the **CloudFormation > Stacks** page, confirm that the stack status is updated to `CREATE_IN_PROGRESS`

For more information, see [“Monitoring IBM Maximo Application Suite installation on Amazon Web Services”](#) on page 238.

Installing client managed IBM Maximo Application Suite for private paid offer

The paid AWS product can be installed from the AWS Marketplace based on your infrastructure needs. A new Red Hat OpenShift instance is created as part of the Maximo Application Suite installation.

Before you begin

Before you can install Maximo Application Suite on Amazon Web Services, you must configure prerequisites and gather information that you need to complete the installation.

For more information, see [“Prerequisites for installing Maximo Application Suite on Amazon Web Services”](#) on page 143.

Consider other criteria, such as the type of Maximo Application Suite offering that you want and whether you want to create a Red Hat OpenShift cluster or reuse an existing one.

Note: The existing cluster must be created by using the automated deployment option only.

For more information, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

This product requires an internet connection to deploy properly. The following code is downloaded on deployment for setting up the Maximo Application Suite Red Hat OpenShift Container Platform cluster:

- <https://github.com/ibm-mas/ansible-devops.git>
- <https://github.com/ibm-mas/multicloud-bootstrap.git>

About this task

The Maximo Application Suite client-managed solution is available with two different Marketplace products.

Maximo Application Suite client-managed with Red Hat OpenShift entitlement

This product includes the subscription for the Red Hat OpenShift cluster that is deployed during the installation process.

The following CloudFormation templates are available as fulfillment options:

- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

Maximo Application Suite client-managed without Red Hat OpenShift entitlement

This product does not include the subscription for the Red Hat OpenShift cluster that is deployed during the installation process or the existing one provided by the user. You must have your own Red Hat OpenShift subscription to deploy Maximo Application Suite by using this product.

The following CloudFormation templates are available as fulfillment options:

- An existing Red Hat OpenShift cluster
- A new Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
- A new Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

After you select the Paid product from Amazon Web Services Marketplace, you can complete the following steps for a private paid offering.

You can avail a private offer for IBM Maximo Application Suite subscription contract with custom configuration and pricing by contacting a IBM sales representative or viewing <https://www.ibm.com/products/maximo/pricing> page. After discussions and agreements, the IBM sales representatives share the offer URL page link.

Procedure

1. Click the private URL shared by the IBM sales representatives, and click **Continue to Subscribe**.
2. Review the number of AppPoints, contract duration, and contract pricing information for the private offer, and click **Create Contract**.
3. Confirm the contract terms on the confirmation dialog.
You see a message that AWS is processing the request. It takes couple of minutes to complete the processing. Until then, **Continue to Configuration** is unavailable and for **Create Contract** displays the label as Pending.
4. After you confirm the contract terms of purchase to register the product with your AWS account, complete the registration by using an IBM ID from <https://www.ibm.com/marketplace/connector/landing/aws/services/amicontractsetup>. You can also get this URL from the Overview section of the product page.

Select the Maximo Application Suite product for which you subscribed, provide the necessary details like the AWS account number, and confirm the pending order by verifying the quotation number. You see the product summary page for your purchase. The browser tab can be closed. You receive an email notification within few minutes to obtain the IBM entitlement key from My IBM portal required for the product deployment.

5. After you retrieve the entitlement key, link your IBM Maximo Application Suite subscription with Red Hat OpenShift account by following the steps at <https://www.ibm.com/docs/en/cloud-paks/1.0?topic=iocpc-accessing-red-hat-entitlements-from-your-cloud-paks>, obtain the Red Hat OpenShift pull secret from the Red Hat OpenShift account, return to product page, and click **Continue to Configuration**.

Note: Linking the IBM Maximo Application Suite subscription with Red Hat OpenShift account is required only if you are deploying IBM Maximo Application Suite with new Red Hat OpenShift cluster. It is not required if you decide to use existing Red Hat OpenShift cluster. (Support for existing Red Hat OpenShift cluster is not available yet.)

6. In **Fulfillment** option, select the appropriate deployment mode.
For Maximo Application Suite deployment with new Red Hat OpenShift cluster, you can select either new network infrastructure (IPI) or existing network infrastructure (UPI), or existing Red Hat OpenShift cluster. For more information, refer to the fulfillment options explained earlier in this section.
7. In **Software version**, select the latest version.
8. In **Region**, select a supported geographical region where you want to install the Maximo Application Suite.
For the list of supported regions, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.
9. Accept the default values in the other fields and click **Continue to Launch**.
10. In the **Launch** page, select **Choose Action > Launch CloudFormation**.
11. To open the CloudFormation stack wizard, click **Launch**.
12. In the **Create stack** wizard step, accept the default values, and click **Next**.
13. In the **Specify stack details** step, in the **Stack name** section, enter a unique name.
14. In the Parameters section, enter the installation parameters by using the information that you gathered when you configured the [“Prerequisites for installing Maximo Application Suite on Amazon Web Services”](#) on page 143 and [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

15. Enter all of the mandatory parameters.

- The Maximo Application Suite offering type, such as Maximo Application Suite Core and Cloud Pak for Data or Maximo Application Suite and Maximo Manage.

Note: If you choose to install Maximo Application Suite with Manage, you can use the default IBM Db2 instance that is provisioned by IBM instead of configuring your own external Db2 instance.

To configure the default IBM Db2, do not add information in the username, password, JDBC URL, certificate URL, and demo data fields.

The internal Db2 configuration is available from Maximo Application Suite 8.10 or later.

- All of the parameters in the **Cluster and bootnode access** and **Keys and licenses** sections.

16. To create a new cluster, complete the [New OpenShift cluster configuration details](#) section.

17. In the optional parameter groups, such as the group of Maximo Manage database configuration parameters, ensure that you either specify all of the parameter values or leave all of them empty.

18. To configure Amazon Managed Streaming for Kafka, select **Yes**.

The Amazon Managed Streaming for Kafka is configured to process data streaming of applications such as IoT and IBM Maximo Monitor from Maximo Application Suite.

You can configure the default Amazon Managed Streaming from Maximo Application Suite 8.10. or later.

19. To configure a DocumentDB instance, you can select any of the following options.

- Configure a new MongoDB community edition
- Use an existing MongoDB

You must input existing MongoDB connection details such as the username, password, MongoDB hostname, and CA certificate.

- Configure a new Amazon DocumentDB
- Use an existing Amazon DocumentDB

Add existing DocumentDB connection details such as the username, password, DocumentDB hostname, and CA certificate.

You must input the VPC ID of the region in which the Amazon DocumentDB is deployed. Ensure that the DocumentDB does not have a matching or overlapping IPv4 CIDR block 10.0.0.0/16.

You can configure the default DocumentDB from Maximo Application Suite 8.10. or later.

20. In the **Other parameters**, select the **OperationalMode** as Production or Non-production.

You can specify deployments for Production or Nonproduction environments. Nonproduction installations can be used for internal development and testing. The installation AppPoints are unused in the Non-production installations. These specifications are also visible in the metrics shared with IBM as well as on the product UI.

21. Click **Next**.

22. In the **Configure stack options** step, configure any additional options that you require . To know more about the stack options, click the "Learn more" links available in each option. Click **Next** when done.

23. In the **Review** step, review the values that you entered and acknowledge the message that relates to identity and access management (IAM) resources.

24. To begin the installation, click **Create stack**.

What to do next

During the installation process, the AWS CloudFormation stack template that you configured is used to create a Bootnode. The Bootnode contains all required resources to complete the installation. To verify that the Bootnode is created successfully, in the **CloudFormation > Stacks** page, confirm that the stack status is updated to CREATE_IN_PROGRESS

For more information, see [“Monitoring IBM Maximo Application Suite installation on Amazon Web Services” on page 238](#).

Installing Maximo Application Suite with Red Hat OpenShift on Amazon Web Services

Starting in IBM Maximo Application Suite 8.11, install the application with Red Hat OpenShift on Amazon Web Services (AWS). You can deploy Maximo Application Suite from the AWS marketplace by using the BYOL or paid offerings so that you can use your existing Red Hat OpenShift cluster.

Before you begin

Create the Virtual Private Cloud (VPC) and subnets for ROSA cluster by using the `pre-req-vpc-subnets.sh` file. Download the compressed file from [IBM GitHub UPI resources](#) page. Refer to the `readme.txt` file in the compressed file folder for instructions to run the `pre-req-vpc-subnets.sh` file.

Procedure

1. Install the ROSA cluster by following the steps in [Getting started](#).

The steps include running the command to create a ROSA cluster, giving the cluster a name, selection the option **Install** on the existing VPC and subnet, and selecting VPC and subnet that was created.

Tip: Make sure that you create a public cluster with Federal Information Processing Standards (FIPS) that is not enabled.

2. Create an Amazon Elastic Files System (EFS) storage by running the `ocp_efs` ansible role. Before you run this ansible role, append the ROSA cluster name to the VPC name used in ROSA cluster creation.

For example, if the VPC name for ROSA cluster creation is `mas-vpc` and the ROSA cluster name is `samplerosacluster`, then the VPC name is `mas-vpc-samplerosacluster`.

3. Run the following commands to set required environment parameters that are necessary to run `ocp_efs` ansible role.

```
export AWS_ACCESS_KEY_ID=<your aws access key id>
export AWS_SECRET_ACCESS_KEY=<your aws access key>
export AWS_DEFAULT_REGION=<aws region where ROSA cluster is created>
export CLUSTER_NAME=<ROSA cluster name>
```

4. Connect to your ROSA cluster by running the `oc login` command.

The following is an example of the `oc login` command:

```
oc login -\-token=<your_server_token> -\-server=https://<server_host>:<port_number>
```

Tip: You can obtain this login command information in your ROSA cluster web console by clicking **cluster-admin > Copy login command**.

5. Run the following commands to install the ansible collection and `ocp_efs` ansible role.

```
ansible-galaxy collection install ibm.mas_devops
export ROLE_NAME=ocp_efs && ansible-playbook ibm.mas_devops.run_role
```

Make sure that the EFS storage is created on the Amazon Web Services console. Verify that the storage class is created on the ROSA cluster console. After the verification, proceed with the stack deployment by completing the **CloudFormation** template page fields.

6. Follow the **Existing Red Hat OpenShift cluster** option available under the IBM Maximo Application Suite (BYOL) AWS marketplace listing to deploy Maximo Application Suite core and IBM Cloud Pak for Data stack or Maximo Application Suite core and IBM Maximo Manage stack.

Note: Only **Existing Red Hat OpenShift cluster** option is supported for ROSA.

7. Enter the required input parameters on **CloudFormation** page to create the stack.

For more information, see [“Installing BYOL IBM Maximo Application Suite” on page 222](#).

Installing Cloud Pak for Data on an Amazon Web Services instance of Maximo Application Suite

To install Cloud Pak for Data on an existing Red Hat OpenShift cluster instance that has been configured to run the automation on Amazon Web Services, you clone a GitHub repository and run a script.

For example, install Cloud Pak for Data on an Amazon Web Services instance on a cluster that has been installed with Maximo Application Suite and Maximo Manage offerings, not on an existing cluster with Cloud Pak for Data already installed.

The following components will be installed:

- [Cloud Pak for Data](#)
- [Db2 Warehouse services](#)
- [Data Management Console](#)

If you installed Maximo Application Suite and Maximo Manage offerings, your Red Hat OpenShift cluster does not include Cloud Pak for Data. If you want to deploy certain applications and add-ons in the Maximo Application Suite, such as Maximo Predict, you must first install Cloud Pak for Data.

To install Cloud Pak for Data, you clone a Git repository, locate the installation script, and run it. When you run the script, you must provide the AWS CloudFormation stack values that identify the Red Hat OpenShift cluster, such as its region, its stack name, and your entitled registry key. The installation script retrieves the Maximo Application Suite <unique-string> and the Red Hat OpenShift cluster details from the stack. If the cluster credentials, that is the username and password, are changed since you installed the Maximo Application Suite, you must provide the updated credentials when you run the script.

You run the script on your local machine or on the Bootnode in the Maximo Application Suite Red Hat OpenShift cluster. For more information, see [“Accessing the Bootnode and Red Hat OpenShift cluster”](#) on page 245.

Note: Starting in 8.11, for US GovCloud regions, the Cloud Pak for Data configuration is not supported.

Before you begin

1. Retrieve the entitled registry key that you provided when you installed the Maximo Application Suite.
If you do not have this key, use the steps in the [MAS download document](#) to download this key from the IBM Container Library.
2. In the AWS CloudFormation console, in the **CloudFormation->Stacks** page, locate the stack that you created when you installed your Red Hat OpenShift cluster.
3. From the **Outputs** tab, record the following values:
 - The instance's stack name
 - OpenShiftConsoleUrl (ocp url)
 - OpenShiftPassword (ocp password)
 - OpenShiftUser (ocp username)
4. In the menu bar, record the region value.
5. On the machine where you want to run the script, ensure that the following CLI packages are installed:
 - Version 4.0 or a later version of [GNU bash](#)
 - [jq](#)
 - [Git](#)
 - [AWS CLI](#). Ensure that this package is configured for authentication with your AWS account. For more information, see [Configuring the AWS CLI](#) in the AWS documentation.

Procedure

Complete the following steps to install Cloud Pak for Data in your Red Hat OpenShift cluster:

1. On the machine where you want to run the script, in a command shell, log in to the AWS service by running the following command:

```
aws configure
```

You are prompted for your identity and access management (IAM) user credentials. Enter the credentials for an IAM user that has the permissions to run the script, such as the IAM user that installed the Maximo Application Suite. For more information, see [Configuring the installation permissions](#).

2. Clone the Git repository that contains the script by running the following command:

```
git clone https://github.com/ibm-mas/mas-on-aws.git
```

3. Make the script executable by running the following commands:

```
cd mas-on-aws
chmod +x deploy-cp4d.sh
```

4. View the script's usage information by running the following command:

```
./deploy-cp4d.sh -h
```

5. Specify the required options and run the script.

- Use the `r` option to specify the region code of the region where the Maximo Application Suite was installed, for example: `-r ap-northeast-3`
- Use the `s` option to specify the Maximo Application Suite CloudFormation stack name, for example: `-s sp-manage-12`
- Use the `e` option to specify the entitled registry key that you provided when you installed the Maximo Application Suite, for example: `-e <entitlement-key>`
- Use the `u` and `p` options to specify the Maximo Application Suite Red Hat OpenShift cluster credentials, for example: `-u <ocp-user> -p <ocp-password>`
- The following sample command installs Cloud Pak for Data by using all of these example options:

```
./deploy-cp4d.sh -r ap-northeast-3 -s sp-manage-12 -e <entitlement-key> -u <ocp-user> -p <ocp-password>
```

The script takes 60 minutes to install Cloud Pak for Data into the Red Hat OpenShift cluster.

For reference, the following screen shots show what the script output looks like during this process:

When the script is running:

```
ec2-user@ip-172-31-2-201:~/mas-on-aws ㉿ 31

ok: [localhost]

TASK [suite_config : Debug information] *****
ok: [localhost] => {
  "msg": [
    "Instance ID ..... mas-5ffkle",
    "Workspace ID ..... wsmasocp",
    "MAS namespace ..... mas-mas-5ffkle-core",
    "MAS config directory ..... /tmp/masconfigdir"
  ]
}

TASK [suite_config : Find *.yaml and *.yml files in the MAS config directory] *****
ok: [localhost]

TASK [suite_config : Debug the list of config files located] *****
ok: [localhost] => (item={"path": "/tmp/masconfigdir/jdbc-db2wh-db01-cpd-services-5ffkle.yml", "mode": "0664", "isdir": False, "ischr": False, "isblk": False, "isreg": True, "isfifo": False, "islnk": False, "issock": False, "uid": 1000, "gid": 1000, "size": 1972, "inode": 13737245, "dev": 66306, "nlink": 1, "atime": 1652195331.4400334, "mtime": 1652195331.1280332, "ctime": 1652195331.4430335, "gr_name": "ec2-user", "pw_name": "ec2-user", "wusr": True, "rusr": True, "xusr": False, "wgrp": True, "rgrp": True, "xgrp": False, "woth": False, "roth": True, "xoth": False, "isuid": False, "isgid": False}) => {
  "msg": "/tmp/masconfigdir/jdbc-db2wh-db01-cpd-services-5ffkle.yml"
}

TASK [suite_config : Apply configs] *****
changed: [localhost] => (item={"path": "/tmp/masconfigdir/jdbc-db2wh-db01-cpd-services-5ffkle.yml", "mode": "0664", "isdir": False, "ischr": False, "isblk": False, "isreg": True, "isfifo": False, "islnk": False, "issock": False, "uid": 1000, "gid": 1000, "size": 1972, "inode": 13737245, "dev": 66306, "nlink": 1, "atime": 1652195331.4400334, "mtime": 1652195331.1280332, "ctime": 1652195331.4430335, "gr_name": "ec2-user", "pw_name": "ec2-user", "wusr": True, "rusr": True, "xusr": False, "wgrp": True, "rgrp": True, "xgrp": False, "woth": False, "roth": True, "xoth": False, "isuid": False, "isgid": False})

TASK [suite_config : Configure MAS workspace] *****
ok: [localhost]

TASK [suite_verify : Fail if mas_instance_id is not provided] *****
skipping: [localhost]

TASK [suite_verify : Configure namespace] *****
ok: [localhost]

TASK [suite_verify : Wait for Suite to be ready (60s delay)] *****
ok: [localhost]

TASK [suite_verify : Lookup MAS superuser credentials] *****
ok: [localhost]

TASK [suite_verify : Lookup Route for admin] *****
ok: [localhost]

TASK [suite_verify : Lookup cluster subdomain] *****
ok: [localhost]

PLAY RECAP *****
localhost      : ok=13   changed=1   unreachable=0   failed=0   skipped=4   rescued=0   ignored=0

===== MAS configuration completed =====
===== Execution completed at Tue May 10 15:08:59 UTC 2022 =====

[ec2-user@ip-172-31-2-201 mas-on-aws]$
```

You may see some **FAILED** messages and **fatal** messages. These messages are normal and are automatically solved during the script execution:

```
ec2-user@ip-172-31-2-201:~/mas-on-aws ㉿ 341
ok: [localhost] => (item=db2wh)
ok: [localhost] => (item=dmc)

TASK [cp4d_install_services : wait : Check if CPD Service Subscriptions operators are running] ***
ok: [localhost] => (item=Subscription is being installed; operator available replicas = 0)
ok: [localhost] => (item=Subscription is being installed; operator available replicas = 0)

TASK [cp4d_install_services : wait : Fail if one or more CPD Service Subscriptions operators are not running] ***
fatal: [localhost]: FAILED! => {"changed": false, "msg": "[0/30] 0 of 2 services are ready"}

TASK [cp4d_install_services : wait : Give up after 30 attempts (approx 30 minutes)] *****
skipping: [localhost]

TASK [cp4d_install_services : wait : Wait for 60 seconds before checking again] *****
Pausing for 60 seconds
(ctrl+C then 'C' = continue early, ctrl+C then 'A' = abort)
ok: [localhost]

TASK [cp4d_install_services : include_tasks] *****
included: /home/ec2-user/mas-on-aws/ansible/playbooks/roles/cp4d_install_services/tasks/cp440/wait_for_subscriptions.yml for localhost

TASK [cp4d_install_services : wait : Set the retry count] *****
ok: [localhost]

TASK [cp4d_install_services : wait : Lookup CPD Service Subscriptions operators] *****
ok: [localhost] => (item=db2wh)
ok: [localhost] => (item=dmc)

TASK [cp4d_install_services : wait : Check if CPD Service Subscriptions operators are running] ***
ok: [localhost] => (item=ibm-db2wh-cp4d-operator-controller-manager operator available replicas = 1)
ok: [localhost] => (item=ibm-dmc-controller-manager operator available replicas = 1)

TASK [cp4d_install_services : wait : Fail if one or more CPD Service Subscriptions operators are not running] ***
skipping: [localhost]

TASK [cp4d_install_services : Install Services CR] *****
changed: [localhost] => (item=db2wh)
changed: [localhost] => (item=dmc)

TASK [cp4d_install_services : include_tasks] *****
included: /home/ec2-user/mas-on-aws/ansible/playbooks/roles/cp4d_install_services/tasks/cp440/wait_for_services.yml for localhost

TASK [cp4d_install_services : wait : Set the retry count] *****
ok: [localhost]

TASK [cp4d_install_services : wait : Lookup CPD Service CRs] *****
ok: [localhost] => (item=db2wh)
ok: [localhost] => (item=dmc)

TASK [cp4d_install_services : wait : Check if CPD Services installation have been completed] ***
ok: [localhost] => (item=cloudpak-db2whservice CR status = In Progress)
ok: [localhost] => (item=dmc-addon CR status = In Progress)

TASK [cp4d_install_services : wait : Fail if one or more CPD Services installation still in progress...] ***
fatal: [localhost]: FAILED! => {"changed": false, "msg": "[0/30] 0 of 2 services are ready"}

TASK [cp4d_install_services : wait : Give up after 30 attempts (approx 3 hours)] *****
skipping: [localhost]

TASK [cp4d_install_services : wait : Wait for 5 minutes before checking again] *****
```

```
ec2-user@ip-172-31-2-201:~/mas-on-aws
TASK [cp4d_install : Install CPD 4.0 Subscription] *****
changed: [localhost]

TASK [cp4d_install : Wait for cpd-platform-operator-manager to be ready (60s delay)] *****
FAILED - RETRYING: Wait for cpd-platform-operator-manager to be ready (60s delay) (10 retries left).
ok: [localhost]

TASK [cp4d_install : Wait for operand-deployment-lifecycle-manager to be ready (60s delay)] ***
ok: [localhost]

TASK [cp4d_install : Apply CloudPak for Data 4.0 - Operand Request] *****
changed: [localhost]

TASK [cp4d_install : Apply CloudPak for Data 4.0 CR] *****
changed: [localhost]

TASK [cp4d_install : Wait for ibm-zen-operator to be ready (60s delay)] *****
FAILED - RETRYING: Wait for ibm-zen-operator to be ready (60s delay) (90 retries left).
ok: [localhost]

TASK [cp4d_install : Wait for ibmcpd CPD 4.0 to be Completed] *****
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (90 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (89 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (88 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (87 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (86 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (85 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (84 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (83 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (82 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (81 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (80 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (79 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (78 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (77 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (76 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (75 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (74 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (73 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (72 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (71 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (70 retries left).
FAILED - RETRYING: Wait for ibmcpd CPD 4.0 to be Completed (69 retries left).
ok: [localhost]

TASK [cp4d_install : Wait for ZenService CPD 4.0 to be Completed] *****
ok: [localhost]

TASK [cp4d_install : Retrieve admin credentials] *****
ok: [localhost]

TASK [cp4d_install : Obtain CP4D dashboard URL] *****
ok: [localhost] => {
  "msg": [
    "CP4D Dashboard ..... https://cpd-cpd-services-5ffkle.apps.masocp-5ffkle.docmanageaws.com",
    "CP4D Username ..... admin",
    "CP4D Password ..... Found in 'admin-user-details' secret under 'cpd-services-5ffkle' namespace"
  ]
}
```

Finally, when the script is completed successfully:

```

ec2-user@ip-172-31-2-201:~/mas-on-aws
Trying to log into OpenShift
The server uses a certificate signed by an unknown authority.
You can bypass the certificate check, but any data you send to the server could be intercepted by others.
Use insecure connections? (y/n): y
Login successful.
You have access to 70 projects, the list has been suppressed. You can list all projects with 'oc projects'
Using project "default".
Welcome! See 'oc help' to get started.
OpenShift Login is successful.
==== Execution started at Tue May 10 14:22:53 UTC 2022 ====
==== CP4D deployment started ====
[DEPRECATION WARNING]: Ansible will require Python 3.8 or newer on the controller starting
with Ansible 2.12. Current version: 3.6.8 (default, Sep  9 2021, 07:49:02) [GCC 8.5.0
20210514 (Red Hat 8.5.0-3)]. This feature will be removed from ansible-core in version 2.12.
Deprecation warnings can be disabled by setting deprecation_warnings=False in ansible.cfg.
[WARNING]: No inventory was parsed, only implicit localhost is available
[WARNING]: provided hosts list is empty, only localhost is available. Note that the implicit
localhost does not match 'all'
PLAY [localhost] *****
TASK [Gathering Facts] *****
ok: [localhost]
TASK [cp4d_install : Debug parameters] *****
ok: [localhost] => {
  "msg": [
    "CPD Version ..... cpd40",
    "MAS Channel ..... 8.7.x"
  ]
}
TASK [cp4d_install : Assert that either cpd_version or mas_channel are defined] *****
ok: [localhost] => {
  "changed": false,
  "msg": "All assertions passed"
}
TASK [cp4d_install : Assert that cpd_version is supported] *****
ok: [localhost] => {
  "changed": false,
  "msg": "All assertions passed"
}
TASK [cp4d_install : Assert that cpd_version and mas_channel are compatible, if both are set] ***
skipping: [localhost]
TASK [cp4d_install : Check CP4D version to be installed] *****
skipping: [localhost]
TASK [cp4d_install : debug] *****

```

6. After the installation, obtain the Cloud Pak for Data instance URL and credentials from the OpenShift console.

One of the many ways to obtain these are:

- Go to **Projects**, and click the Cloud Pak for Data project name which follows the format cpd-services-**<unique-string>**.
- In the **Overview** tab, under **Inventory**, click **Routes**.
- The URL for the Cloud Pak for Data instance is under the "Location" column for the route named `cpd`.
- In the same **Overview** and **Inventory** section click **Secrets**.
- Search for `admin-user-details`, click it, and scroll down the page to see the value for `initial_admin_password`. This is the password to be able to login the Cloud Pak for Data URL. Use `admin` as the username.

What to do next

You can now install and configure dependencies inside Cloud Pak for Data that are required for some Maximo Application Suite applications. For more information about the complete list of prerequisites for each application, see [“Prerequisite software”](#) on page 5.

You also can enable the internet access to a Db2 Warehouse database instance in the Amazon Web Services Red Hat OpenShift cluster. For more information, see [Enabling internet access to a Db2 Warehouse database on an AWS Red Hat OpenShift cluster](#).

Monitoring IBM Maximo Application Suite installation on Amazon Web Services

During the installation process, the AWS CloudFormation stack template that you configured is used to create a Bootnode. The Bootnode contains all required resources to complete the installation. To verify that the Bootnode is created successfully, in the **CloudFormation > Stacks** page, confirm that the stack status is updated to `CREATE_IN_PROGRESS`.

Before you begin

To ensure that the Red Hat OpenShift cluster is created, the Bootnode starts a bootstrap process. This process creates a bootstrap node that uses the Red Hat OpenShift installer to create master and worker nodes. To verify that the bootstrap process is started, in the **CloudFormation > Stacks** page, click the **Events** tab. When the `DeployWaitCondition` event is displayed, the bootstrap process is started.

About this task

You can monitor the progress of the remaining installation tasks in the installation logs. These logs can be viewed in the AWS CloudWatch service or the Bootnode.

Procedure

1. Monitoring the installation logs in AWS CloudWatch.

To monitor the installation progress in AWS CloudWatch, complete the following steps:

- a) In the AWS CloudWatch management console, click **Logs**.
- b) In the log group list, click the group that is named `/ibm/mas/masocp-<unique-string>`.
- c) In the log stream list, open the installation log by clicking the stream that is named `mas-provisioning-logs`.
- d) Optional: In the installation log stream, if the message `Auto retry paused` is displayed, click the `Resume` link to display the latest log updates.

2. Monitoring the installation logs in the Bootnode.

To monitor the installation progress in the Bootnode, complete the following steps:

- a) Connect to the Bootnode by using Secure Shell (SSH) access. For instructions, see [Accessing the Bootnode and Red Hat OpenShift cluster](#).
- b) Run the following command to switch to the root user:

```
sudo su -
```

- c) Monitor the installation log updates by running the following command:

```
tail -f /root/ansible-devops/multicloud-bootstrap/mas-provisioning.log
```

3. After the installation is completed, in the **CloudFormation > Stacks** page, the following indicators confirm that the installation is successful:

- The status of the stack is **CREATE_COMPLETE**.
- In the **Outputs** page, the `DeploymentStatus` parameter displays a message that indicates that the installation succeeded, for example `ID-aws-small-masocp-fjh2sx:SUCCESS#MAS deployment completed successfully`.

4. Proactively check the status of the deployment.

Depending on the parameters that you specified, the installation time might vary.

If the installation is unsuccessful, use the information in the [Troubleshooting installation problems](#) topic to identify and resolve the problem.

Accessing IBM Maximo Application Suite

After you install IBM Maximo Application Suite, to access it, you need the administrator URL, user credentials, and public certificate.

About this task

You must have the following information to access the Maximo Application Suite:

- The Maximo Application Suite administrator URL, which you use to connect to the Maximo Application Suite through a browser.
- Your username and password.
- The public certificate for the Maximo Application Suite.

You import this certificate into your browser's trusted store to ensure secure communication between your browser and Maximo Application Suite.

How you retrieve these items of information depends on whether you configured a verified Amazon SES email address.

If you have a verified Amazon SES email address, you can retrieve the administrator URL, username, and password from the emails that you received. The public certificate is attached to these emails.

If you do not have an Amazon SES email address, you can retrieve the administrator URL from the CloudFormation console for the stack that you created during the installation. However, to retrieve your username, password, and the public certificate, you must connect to the Red Hat OpenShift cluster.

Product versions

Following products are installed as part of the Maximo Application Suite installation. If you provide the existing Red Hat OpenShift cluster, the installation process checks if any of these products are already installed. For details about the installation process behavior if existing products are found, see [“Preparing to install Maximo Application Suite on Amazon Web Services”](#) on page 146.

- Red Hat OpenShift 4.15.x
Supported on Installer Provisioned Infrastructure (IPI), User Provisioned Infrastructure (UPI), and existing Red Hat OpenShift cluster deployment.
- Red Hat OpenShift Service on AWS (ROSA) 4.15.x
Supported on existing Red Hat OpenShift cluster deployment on Amazon Web Services public cloud.
- IBM Cloud Pak foundational services 4.7.0
- IBM Cloud Certificate Manager 3.25.13
- MongoDB (CE) 7.0.12
- IBM Suite License Service 3.10.1
- IBM Data Reporter Operator 2.18.0
- IBM Cloud Pak for Data 6.0.0
- IBM Maximo Application Suite 9.0.5
- IBM Maximo Manage 9.0.5

Procedure

1. In the **Outputs** section of the CloudFormation stack that was created during the deployment, record the following values:
 - The value of the `masAdminUrl` key. This value contains the administrator URL.

- The values of the `openShiftConsoleUrl`. These values contain the URL for the Red Hat OpenShift console.
- Note:** The `openShiftConsoleUrl` is not displayed for an existing Red Hat OpenShift Container Platform.
- The value of the `clusterUniqueString` key. This value contains the cluster unique string to look for the correct resources in Red Hat OpenShift.
2. Retrieve the Red Hat OpenShift Container Platform credentials from the `maximo-ocp-secret`. The Amazon Web Services Secrets Manager secret contains the credentials for Red Hat OpenShift cluster. It consists of a secret named `maximo-ocp-secret` containing Red Hat OpenShift credentials.
 3. Retrieve the Red Hat OpenShift Container Platform `kubeadmin` credentials from `maximo-kubeadmin-secret`.
 4. Connect to the Red Hat OpenShift cluster.
 5. Select **Workloads > Secrets** from the navigation page.
 6. Select the project named `mas-<unique-string>-core`.
 7. Click the secret that is named `<unique-string>-credentials-superuser`.
 8. Click the 'Reveal values' link to get the username and password for Maximo Application Suite.
 9. Click the secret that is named `<unique-string>-cert-public` from the `mas-<unique-string>-core` project.
 10. Click the 'Reveal values' link to get the contents of the certificates.
 11. Retrieve the contents of `ca.crt`, which is the public certificate for Maximo Application Suite.
 12. After you import the public certificate into your browser's trusted store, paste the Maximo Application Suite administrator URL into your browser, and enter the authentication credentials to access the application.
 13. Log in to the Maximo Application Suite to deploy applications, create users, and specify configuration. To resolve loading issues with the initial setup page, you can provision self-signed certificates that are needed to add in the used browser's truststore. For more information, see [IBM Maximo Application Suite - Initial setup page does not load](#).

What to do next

If you installed the Maximo Application Suite core and Maximo Manage, and you want to deploy certain applications and add-ons in the Maximo Application Suite, which requires Cloud Pak for Data, see [“Installing Cloud Pak for Data on an Amazon Web Services instance of Maximo Application Suite” on page 232](#).

If you installed Maximo Application Suite core and Cloud Pak for Data, Cloud Pak for Data is already installed for you, however you might still need to install some Cloud Pak for Data services based on the applications you want to deploy.

You also can enable internet access to a Db2 Warehouse database instance in the AWS Red Hat OpenShift cluster. For more information, see [Enabling internet access to a Db2 Warehouse database on an AWS Red Hat OpenShift cluster](#).

If you want to use well-known certificates that are signed by certificate authority as **Let's Encrypt**, it is recommended that you [uninstall this default Maximo Application Suite instance](#) and install another after you follow the next steps to [configure Let's Encrypt and Route53 on AWS](#).

Note: You cannot install Maximo Application Suite core and Cloud Pak for Data over a cluster that is deployed through Maximo Application Suite core or Maximo Manage.

For more information, see [“Prerequisite software” on page 5](#) and [“Deploying applications, add-ons and industry solutions” on page 291](#).

Configuring Let's Encrypt for Maximo Application Suite on Amazon Web Services

When you install Maximo Application Suite along with a Stack on Amazon Web Services using the automated deployment offerings, Maximo Application Suite uses self-signed certificates.

If you want to use well-known certificates that are signed by Certificate Authority such as Let's Encrypt, install and configure Let's Encrypt and Route53 on Amazon Web Services.

Before you begin

Complete the following tasks:

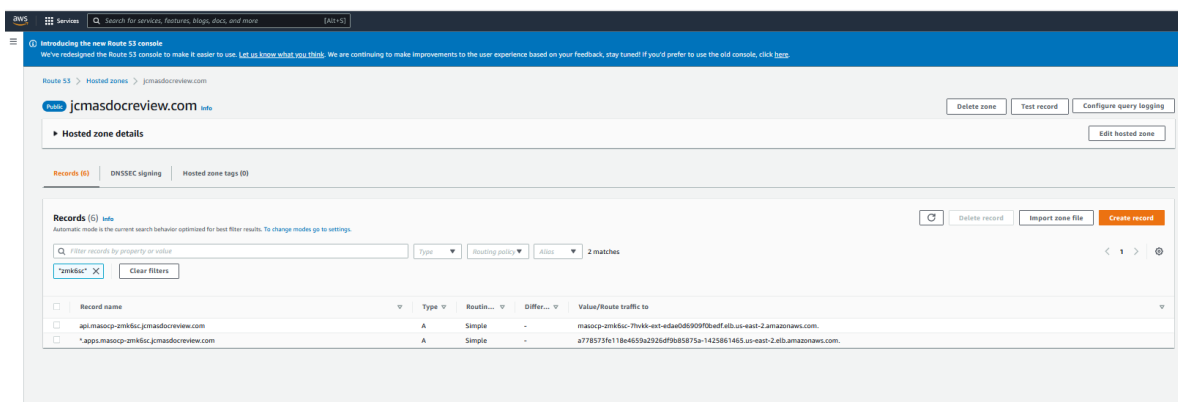
1. Create an [Access Key](#) in the Amazon Web Services console.
2. Create an IAM policy so that certificate manager is able to add records to Route53 in order to solve the DNS01 challenges. To create an IAM policy, complete the following steps:
 - a. Login to the Amazon Web Services console, then search for IAM and click the first option that is displayed.
 - b. Click **Policies** and then click **Create Policy**.
 - c. Go to the **JSON** tab and paste the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "route53:GetChange",
      "Resource": "arn:aws:route53::change/*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "route53:ChangeResourceRecordSets",
        "route53:ListResourceRecordSets"
      ],
      "Resource": "arn:aws:route53::hostedzone/*"
    },
    {
      "Effect": "Allow",
      "Action": "route53:ListHostedZonesByName",
      "Resource": "*"
    }
  ]
}
```

- d. Click **Next: Tags**, then **Next: Review**. Provide Name and Description values.
- e. Click **Create policy**.

Procedure

- Amazon Web Services Route 53 configuration
 - a) In [Route53](#), go to Hosted Zones, click your domain, and then click **Create record**.
For example:



Tip: For all examples, replace the parameters given in the example with your own parameters.

b) Add a CNAME record for your Maximo Application Suite instance ID:

a. Record Name: <mas-instance-id>

b. Record Value: Load Balancer endpoint, which is located under your Hosted Zones. Filter by your cluster unique ID, and then copy the corresponding value for your cluster ingress.

Use the second option. For example, the record name beginning with **.apps.masocp-* and the value beginning with *a77* as shown in the screen shot of the previous step.

Note: Save the instance ID name so you can later use it during the Maximo Application Suite installation.

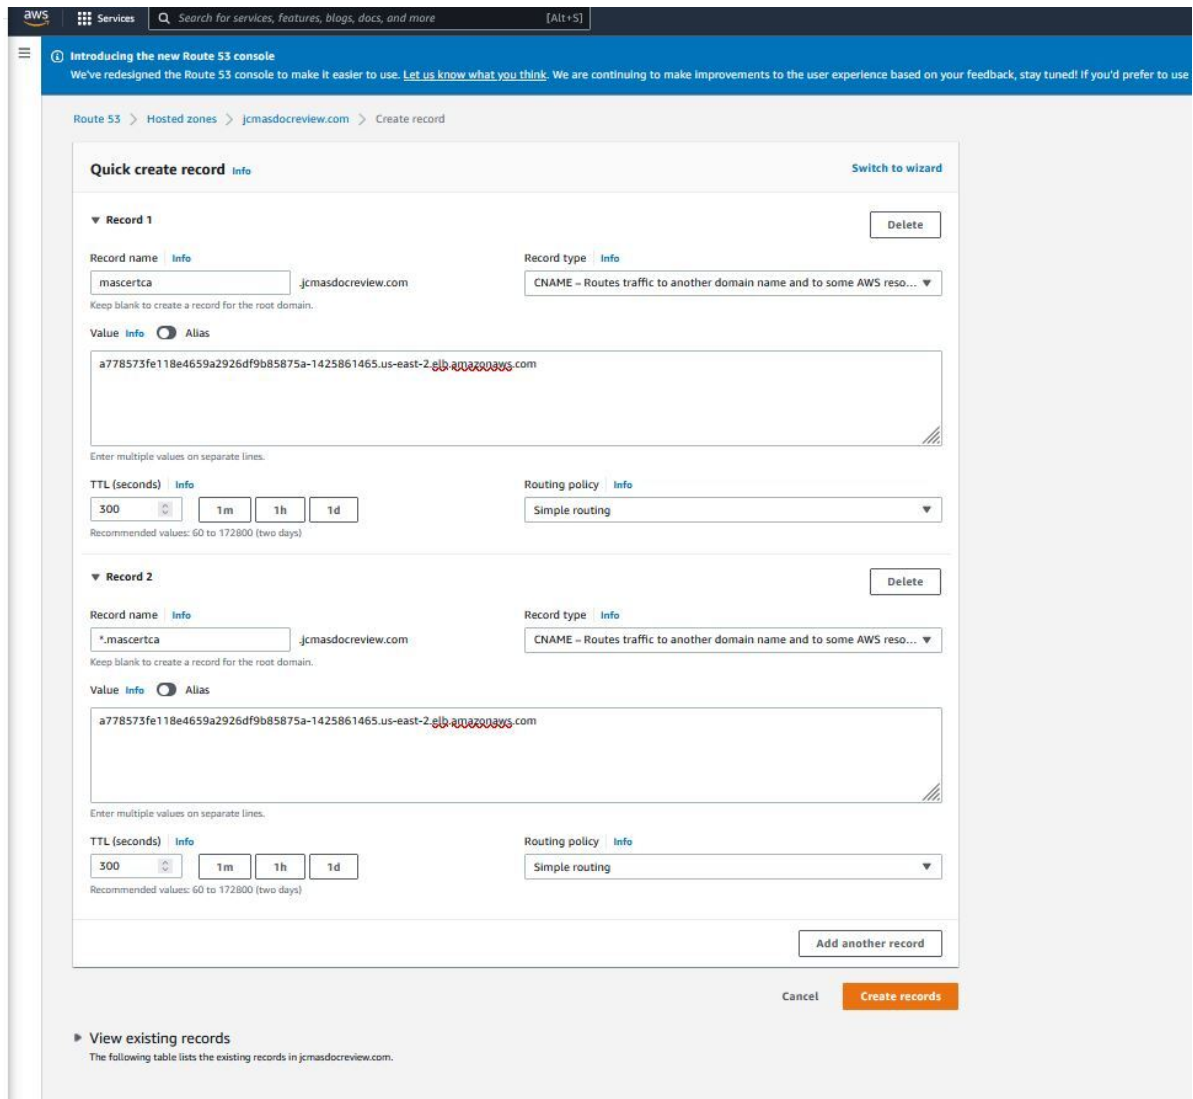
c) Add a new wildcard CNAME record as the following:

a. Record Name: **.<mas-instance-id>*

b. Record Value: Load Balancer endpoint, which is found under your Hosted Zones. Filter by your cluster unique ID, and then copy the corresponding value for your cluster ingress.

Use the second option. For example, the record name beginning with **.apps.masocp-* and the value beginning with *a77* as shown in the screen shot of the previous step.

For example:



- d) Click **Create records**.
- Configure a Let's Encrypt cluster issuer for Maximo Application Suite.
 - a) Run the following script in your terminal. You need to be logged into the cluster via `oc login` command. The script will create a custom cluster issuer named `prod-route53-issuer` in your cluster.

```
# Export the namespace/project where IBM Certificate Manager is installed in your
cluster. Example: ibm-common-services.
CERT_MANAGER_NAMESPACE=ibm-common-services

# Export your AWS secret access key.
SECRET_ACCESS_KEY=<your aws access key>

# Export your AWS secret access ID.
SECRET_ACCESS_ID=<your aws access id>

# Export your email address where you'll get alerts and notifications from Let's Encrypt
certificates
EMAIL_ADDRESS=test@test.com

# Export your route53 hosted zone id.
# Find it under AWS console > Route53 > Hosted Zones > search for your Route53 instance,
the hosted zone id will show at the right hand side.
HOSTED_ZONE_ID=<your route53 hosted zone id>

# Create a secret for secret-access-key
oc create secret generic prod-route53-credentials-secret --from-literal=secret-access-
key=${SECRET_ACCESS_KEY} -n ${CERT_MANAGER_NAMESPACE}
```

```
# create a cluster issuer
cat <<EOF > cluster-issuer.yaml
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: prod-route53-issuer
spec:
  acme:
    email: ${EMAIL_ADDRESS}
    preferredChain: ''
    privateKeySecretRef:
      name: letsencrypt-prod
    server: 'https://acme-v02.api.letsencrypt.org/directory'
    solvers:
      - dns01:
          route53:
            accessKeyID: ${SECRET_ACCESS_ID}
            hostedZoneID: ${HOSTED_ZONE_ID}
            region: us-east-1
            secretAccessKeySecretRef:
              key: secret-access-key
              name: prod-route53-credentials-secret
EOF
oc apply -f cluster-issuer.yaml -n ${CERT_MANAGER_NAMESPACE}
```

To check if this was properly created, log in to your Red Hat OpenShift cluster, go to **Administration > Custom Resource Definitions**, search for **ClusterIssuer > Instances**, search for prod-route53-issuer and click it. The cluster issuer shows the following message:

```
The ACME account was registered with the ACME server.
```

- b) Select the **Suite > YAML** tab and in the **spec** section of the Suite YAML, add **cluster issue** and **domain** parameters.

```
---
spec:
  certificateIssuer:
    duration: 8760h0m0s
    name: prod-route53-issuer
    renewBefore: 720h0m0s
  domain: <<masinstance_id>>.<<domain>>
```

- c) Delete the **finalizer** section from the same Suite YAML to force a reconciliation, and then save the YAML file.

```
finalizers:
  - core.mas.ibm.com/finalizer
```

- d) In **Networking** under **Routes** of project *mas-<mas_instance_id>-core*, wait for the **Routes** to regenerate for the namespace.

- e) Login to the Maximo Application Suite administrator screen and verify the certificate signer.

Note: If IBM Maximo Manage is deployed, the changes might take some time to take effect in Maximo Manage.

- Optional: Configure recursive nameservers

- a) On the Red Hat OpenShift web console, select the *ibm-common-services* project.

- b) In the **Details** tab of **workloads > -> deployment > ibm-cert-manager-operator**, scale down the pod from 1 to 0.

- c) In the **Deployment** tab, select the *cert-manager* controller.

- d) Add the following lines in the yaml file.

```
- '--dns01-recursive-nameservers-only'
- '--dns01-recursive-nameservers=8.8.8.8:53'
```

The following is a sample yaml file.

```

image: >-
  icr.io/cpopen/cpfs/icp-cert-manager-
controller@sha256:1927c16a4dd369c56fa6d2d1897d3ea3d333a3217b8c05ea32b6617c94833a0e
  args:
  - >-
  - --acme-http01-solver-image=icr.io/cpopen/cpfs/icp-cert-manager-
acmesolver@sha256:e8f50ee7b08dc96627e138e9b0d98ed5848c7b4ad92491962c13ef32b2866591
  - '--cluster-resource-namespace=ibm-common-services'
  - '--leader-election-namespace=ibm-common-services'
  - '--dns01-recursive-nameservers-only'
  - '--dns01-recursive-nameservers=8.8.8.8:53'
  serviceAccount: ibm-cert-manager-controller
  dnsPolicy: ClusterFirst

```

Configuring Maximo Application Suite on Amazon Web Services

Use the **Administration** page in Maximo Application Suite to configure the post installation tasks and manage the Maximo Application Suite features.

Setting up Maximo Application Suite environment

The first time that you log in to Maximo Application Suite as the administrator user, no applications are deployed, and no users are logged in to the system. Complete the following tasks to prepare the Maximo Application Suite environment.

About this task

Procedure

In Maximo Application Suite 9.1, select the **Administration** page on side navigation menu from the **Suite** application.

In Maximo Application Suite 9.0 and earlier, select the **Administration** page by clicking the **Administration** icon in the Maximo Application Suite menu bar.

1. On the **Catalog** page, deploy and activate one or more applications.
2. On the **Users** page, add and manage Maximo Application Suite administrators and application users.
3. On the **Configurations** page, administer Maximo Application Suite configurations.
For example, if you deployed the Maximo Manage application, you can add or update the database configuration information.

For more information, see [“Configuring Maximo Application Suite” on page 580](#).

4. If you installed Maximo Application Suite core and Maximo Manage, and you want to deploy certain applications and add-ons in Maximo Application Suite, such as Maximo Monitor, you must first install Cloud Pak for Data.

For more information, see [Installing Cloud Pak for Data on an Amazon Web Services instance of Maximo Application Suite](#).

5. Optional: You can change the authentication to use LDAP or any other supported authentication method.

The default authentication uses Maximo Application Suite authentication.

For more information, see [“Authentication methods” on page 605](#).

Accessing the Bootnode and Red Hat OpenShift cluster

By using Secure Shell (SSH) public key authentication, you can access the Bootnode, the bastion host, and the Red Hat OpenShift cluster nodes.

For operational reasons, you might need command-line access to the Bootnode, the bastion host, or the cluster nodes that are located in the virtual private cloud (VPC) of Maximo Application Suite.

About this task

In the AWS cloud, when you start the Maximo Application Suite installation, a Bootnode is created. By using the required tools and the installation parameters, the Bootnode completes the installation.

In the Red Hat OpenShift cluster that is created during the installation, in a public subnet, a bastion host is created. By using this host, you can connect to the cluster nodes in the private subnets.

The Bootnode, bastion host, and private cluster nodes are all [Amazon EC2 instances](#). To maintain or troubleshooting an EC2 instance, connect to it by using [Secure Shell \(SSH\) public key authentication](#).

Before you installed Maximo Application Suite, you generated a key pair, which consists of a public key and a private key, and uploaded this pair to the Amazon EC2 service. You stored the private key locally. When you specified the installation parameters, you selected the public key in the SSHKey parameter.

During the installation, a copy of the public key is stored in the Bootnode, the bastion host, and the private cluster nodes. Because you have the corresponding private key, you can access these instances by using SSH. In addition, you can use the SSH authentication agent to connect to these instances by using single sign-on authentication.

To use SSH access to connect to the Bootnode, the bastion host, and the private cluster nodes, complete the following steps.

Procedure

1. In your AWS account, connect to the EC2 service console.
2. In the EC2 console, click **Instances**.
3. Retrieve the location of the instance that you want to connect to.
 - a) To retrieve the Bootnode details, search for bootnode.
 - b) To retrieve the bastion host details, search for bastion-host. If you want to connect to the private cluster nodes, you must first connect to the bastion host.
 - c) Click the instance and copy its location from either the **Public IPV4 address** or **Public IPV4 DNS** fields.
 - d) In the **Instance state** column, if the instance is in a shutdown state, click **Start instance**.
4. In your local machine, change the permissions of the private key that you generated before you installed Maximo Application Suite.
For example, for Linux servers, if you stored the private key in the `/tmp/mas-aws-ssh-key.pem` file, run the following command:

```
chmod 0400 /tmp/mas-aws-ssh-key.pem
```

5. If the SSH authentication agent program is not started, run the following command.

```
eval `ssh-agent -s`
```

6. Add your private key file into the SSH authentication agent by running the following command:

```
ssh-add -k /tmp/mas-aws-ssh-key.pem
```

You can now connect to the instance by using single sign-on authentication.

7. By using the instance location that you retrieved in [“3.c” on page 246](#), connect to the instance by running the following command:

```
ssh -A ec2-user@<instance-location>
```

For example, to connect to the instance that is at the IP address 35.161.112.157, run the following command:

```
ssh -A ec2-user@35.161.112.157
```

8. Optional: If you accessed the bastion host, connect to a private cluster node.

- a) Use the Red Hat OpenShift web console to connect to the cluster as an administrator.
- b) In the OCP console, click **Home > Overview**.
- c) In the **Cluster Inventory** card, click the link to the node information.
- d) In the **Nodes** page, click the cluster node that you want to connect to.
- e) Click **Node details** and record the name of the node.
- f) In the bastion host command shell, access the node.

```
ssh core@<node_name>
```

For example, to access the `ip-10-0-132-250.ec2.internal` node, run the following command:

```
ssh core@ip-10-0-132-250.ec2.internal
```

The Bootnode and the bastion host

When you start an IBM Maximo Application Suite installation, a Bootnode is created that controls and completes the installation. In the Red Hat OpenShift cluster, a bastion host is created to allow Secure Shell (SSH) access to cluster nodes.

Note: The bastion host is not created for an existing infrastructure.

During an Maximo Application Suite installation on Amazon Web Services, virtual private clouds (VPC) are created that contain Amazon EC2 instances. For example, a VPC is created to contain the Red Hat OpenShift cluster, and EC2 instances are created in the cluster to represent its master and worker nodes.

The two most important EC2 instances that are created during an installation are the Bootnode and the bastion host.

The Bootnode

In the AWS cloud, after you specify the installation parameters and start the installation, a Bootnode is created. The installation parameters are passed to the Bootnode. In addition, all of the required tools to complete the installation, such as Terraform and Docker, are installed on the Bootnode. By using these tools and parameters, the Bootnode performs the following tasks to complete the installation:

- Creates the virtual network infrastructure, such as the VPC that contains the Red Hat OpenShift cluster.
- Runs the bootstrap process that creates the Red Hat OpenShift cluster.
- Installs the Maximo Application Suite prerequisites.
- Installs the Maximo Application Suite.
- Performs any required postinstallation validation.
- Stores the installation context and Terraform state files both locally and in the Amazon S3 storage bucket that is associated with your AWS account.

Because it is located in its own VPC, the Bootnode is not part of the Maximo Application Suite Red Hat OpenShift cluster. After the installation is complete, you do not need to use the Bootnode to access the cluster or interact with Maximo Application Suite. For this reason, the Bootnode is kept in a shutdown state. However, if required, you can restart it and use it to troubleshoot installation issues.

Note: If you use an existing infrastructure, the Bootnode is in the same VPC of the Red Hat OpenShift Container Platform cluster.

The bastion host

The VPC that the Bootnode creates contains several public and private subnets. In one of the public subnets, a bastion host is created. By using this host, you can connect to the cluster nodes in the private subnets.

After the installation is complete, the bastion host is kept in a shutdown state. However, you can restart it if you want to access cluster nodes by using SSH. For more information, see the [Accessing the Bootnode and Red Hat OpenShift cluster](#) topic.

In addition, if required, you can delete the bastion host and create your own.

Note: No charges apply to Amazon EC2 instances that are in a shutdown state, such as the Bootnode and the bastion host. However, charges apply for their attached EBS GP2 volumes of 10 GB. For more information, see [Amazon EBS pricing](#) in the AWS documentation.

Managing Red Hat OpenShift cluster on Amazon Web Services

You can shut down nodes in your Red Hat OpenShift cluster that is deployed on Amazon Web Services and restart the EC2 instances.

About this task

Note: To avoid certificate issues, do not restart or shut down your cluster within the first 24 hours after you install IBM Maximo Application Suite.

The official Red Hat OpenShift documentation discusses how to gracefully [shut down](#) and [restart](#) clusters. Follow the procedure in the Red Hat OpenShift documentation.

The following procedure describes how to shut down nodes in clusters deployed on Amazon Web Services.

Procedure

1. In the Amazon EC2 Dashboard for the Amazon Web Services region where your Red Hat OpenShift cluster is deployed, click **Instances (running)** or **Instances**.

If you click Instances (running), the resulting list of instances are filtered by Instance state = running. If you click Instances, all instances that are stopped or running, appear in the list.

Instances that are part of the Red Hat OpenShift cluster must have names that have a format `masocp-<unique-string>-<some string>-<nodename>`, where **unique-string** is the installation identifier. For more information, see [“Maximo Application Suite unique identifiers for Amazon Web Services”](#) on page 143.

Tip: If the list includes instances that are not related to your Red Hat OpenShift cluster, you can filter by using the **unique-string**.

2. Select the nodes that you want to stop or restart by checking the box next to the instance names.
3. Click **Instance state** to stop, restart, or terminate the instance.

It is recommended that you stop the worker nodes first before you shut down the primary nodes. When you restart a cluster, start the primary nodes before you start the worker nodes.

Tip: If the cluster is shut down within 24 hours of creation, the following error is shown when you connect to the cluster from a browser or use the **oc login** command:

```
error: Unable to connect to the server: EOF
```

For more information, see [How to renew Master node Certificate in Red Hat OpenShift 4.x](#).

Managing IBM Cloud Pak for Data and IBM Db2 Warehouse

Enable internet access to IBM Db2 Warehouse on a Amazon Web Services Red Hat OpenShift cluster and delete the default Db2 Warehouse database instance that is installed with IBM Cloud Pak for Data.

Procedure

Enable internet access to a Db2 Warehouse on an Amazon Web Services Red Hat OpenShift cluster.

Sometimes it can be helpful to setup internet access to the Db2 Warehouse pods that run in Red Hat OpenShift on Amazon Web Services. Using your database client, you can connect to the database, browse

tables, and run queries. The following task shows how to setup an external route and connect to Db2 Warehouse from a database client on your computer.

1. Get the **NodePort** of the Db2 Warehouse **engn service**.
For example, 32589.

```
oc get svc -n <cp4d-namespace> | grep engn-svc
```

2. Retrieve the Db2 CA certificate from Db2 CA secret, and save it to a temporary pem file.

```
oc get secret zen-ca-cert-secret -n <cp4d-namespace> -o jsonpath="{.data.ca\.crt}" | base64 -d > /tmp/db2-ca.pem
```

3. Transfer the pem file to jks file and set the password.

```
keytool -import -file /tmp/db2-ca.pem -keystore /tmp/db2-ca-truststore.jks
```

4. Retrieve the IBM Maximo Manage Db2 username and password.

```
oc get secret mas-s58gpv-wsmasocp-jdbc-binding-450f5de9 -n <manage-namespace> -o jsonpath="{.data.username}" | base64 -d ; echo
```

```
oc get secret mas-s58gpv-wsmasocp-jdbc-binding-450f5de9 -n <manage-namespace> -o jsonpath="{.data.password}" | base64 -d ; echo
```

5. In Amazon Web Services **Route 53** service, click **Hosted zones** > **<domain name for your cluster>** > **Hosted zones details**, and find the associated VPCs.
6. In **Records**, note the DNS name of the record.
For example, *.apps.masocp-s58gpv.masawsdoc.com
7. In the Amazon EC2 service, click **Load Balancers**, search the VPC name, and click the load balancer that matches the DNS name.
8. In the **Listeners** tab, add the port to the listener.
9. In the **Description** tab, click the load balancer security group to jump to **Security Groups** page.
10. Select the security group with no name, in the **Inbound rules**, and grant the access to the port.
11. Open your database client tool and configure the connection.
For example, by using DBeaver v6.2.1, you can configure the following information.

```
Host: <any text>.apps.masocp-s58gpv.masawsdoc.com
Port: 32589
Database: BLUDB:sslConnection=true;sslTrustStoreLocation=/tmp/db2-ca-truststore.jks;sslTrustStorePassword=123456;
Username: s58gpvmanagedb
Password: s58gpvmanagedbs58gpvmanagedb
JDBC Driver Class Name: com.ibm.db2.jcc.DB2Driver
```

You can now connect to the Db2 instance from outside the Red Hat OpenShift cluster.

Deleting the default Db2 Warehouse database instance installed with Cloud Pak for Data.

When you deploy the Maximo Application Suite and Cloud Pak for Data deployment option or install Cloud Pak for Data by using the script in your Red Hat OpenShift cluster that is provisioned through the Maximo Application Suite and Manage option of the automated deployment, Cloud Pak for Data is installed with the Db2 Warehouse service and creates a database instance. However, that database instance is not displayed in the Cloud Pak for Data user interface. If you want to delete that instance, going through the Cloud Pak for Data and Databases view, and select the option to delete your database instance, you must delete it by following these steps:

12. In Red Hat OpenShift web console **Administration** > **Custom Resource Definitions**, search for db2ucluster.
13. Click the cluster name and access **Instances**.

14. Click three dots icon for the database instance and delete the instance.

The `db2uc1uster cr` is deleted with resources that are related to the database instance.

Note: These steps do not uninstall the Db2 Warehouse service on Cloud Pak for Data. For more information about uninstalling the service, see [Cloud Pak for Data : Uninstalling Db2 Warehouse](#).

Configuring the connection to Amazon Web Services DocumentDb

You can configure an Amazon Web Services DocumentDB instance manually from the default MongoDB instance after you install and deploy IBM Maximo Application Suite on Amazon Web Services.

Note: By default, Maximo Application Suite is configured with MongoDB. Therefore, your only option is to configure a new Maximo Application Suite instance to connect with DocumentDB. You cannot migrate your existing data from MongoDB.

Mongocfg and LicenseService CRD

To view the MongoDB configuration details, on the Red Hat OpenShift consoles, select **Administration > CustomResourceDefinition > MongoCfg CRD** and search for `MongoCfg.config.mas.ibm.com`. Select the **Instances** tab and click the CRD name, for example, **MongoCfg** in `bneluv-mongo-system`.

To view the MongoDB connection details, select the **YAML** tab.

Similarly, IBM Suite License Service is connected with MongoDB by using the `LicenseService.sls.ibm.com` CRD.

To view the MongoDB configuration details in Suite License Service, on the Red Hat OpenShift console, select **Administration > CustomResourceDefinition > LicenseService** and search for `LicenseService.sls.ibm.com`. Select the **Instances** tab and click the CRD name, for example, **LicenseService** in `masocp-bneluv`.

To view the specific variables for DocumentDB that are updated after the DocumentDB instance is provisioned, click the **YAML** tab.

Provisioning Amazon Web Services DocumentDB

You can provision an Amazon Web Services DocumentDB instance after you install and deploy Maximo Application Suite from Amazon Web Services.

Procedure

1. Open the Amazon Web Services console where the Maximo Application Suite is deployed and ensure that a DocumentDB instance is not provisioned.
2. Create a subnet group.
 - a) Select a Virtual Private Cloud (VPC) that contains a Maximo Application Suite unique string. For example, select **masocp-bneluv-vpc** from the list.
 - b) Add all the subnets that are related to this VPC.
3. Create a cluster parameter group.
 - a) Type a new cluster parameter group name.
 - b) Select a document database family. For example, select `docdb4.0` from the list.
4. Create a security group.
 - a) Select the unique string that was created for the VPC.
 - b) For inbound rules, select **Custom TCP** for the rule type and **Custom** for the source.
 - c) For outbound rules, select **All traffic** for the rule type and **Anywhere-PIV4** as the destination.
5. On the Amazon Web Services Key Management Service (KMS) console, create a key.
 - a) Configure the key type, usage, and advanced options, such as key material origin and region.

- b) Add labels, such as the alias, description, and tags.
 - c) Configure the key usage details, such as key administrators and permissions.
 - d) Review your key policy and click **finish**.
6. Create a DocumentDB cluster.
- a) Select **5.0.0** for the engine version to create the cluster.
 - Note:** Maximo Application Suite supports only DocumentDB 5.0.0.
 - b) Set **Show Advanced settings** to on.
 - c) In Network Setting section, select the VPC that contains the unique string for Maximo Application Suite.
 - The subnet group is selected automatically when you select the VPC.
 - d) Select the security group that was created earlier. Remove any security group that is selected by default.
 - e) Under cluster options, select the parameter group that was created earlier.
 - f) Under encryption options, enable encryption and choose the KMS key that was created earlier.
 - g) Create the DocumentDB cluster with other default options, such as maintenance, tags, and deletion protection.
- The cluster is deployed in some time, for example, in 15 minutes.
7. Optional: Verify the DocumentDB connection by using a DB client.
 For example you can use Studio 3T to verify the connection. For more information, see [Connecting to an Amazon DocumentDB cluster from Studio 3T](#).

What to do next

Configure a Maximo Application Suite instance to connect with the Amazon Web Services DocumentDB instance, update the `MongoCfg` CRD and `LicenseService` CRD with the Amazon Web Services DocumentDB connection details.

For more information, see [“Configuring Maximo Application Suite with Amazon Web Services DocumentDB” on page 251](#)

Configuring Maximo Application Suite with Amazon Web Services DocumentDB

To configure a new Maximo Application Suite instance to connect with the Amazon Web Services DocumentDB instance, on the Red Hat OpenShift console, update the `MongoCfg` CRD and `LicenseService` CRD with the Amazon Web Services DocumentDB connection details.

About this task

Note: You can configure a new Maximo Application Suite instance to connect with DocumentDB by using the procedure that is described in this document. However, you cannot migrate your existing data from MongoDB by using this procedure.

Procedure

1. Log in to the Red Hat OpenShift console as an admin user.
2. From the side navigation menu, select **Administration > CustomResourceDefinition** and in the list, click `MongoCfg` CRD.
3. Click **Instances** tab.
4. Update the `MongoCfg` CRD with the Amazon Web Services DocumentDB connection details.
 - By default Maximo Application Suite uses the `MongoCfg` CRD to connect to MongoDB. To connect to the DocumentDB, you must update the connection details on the **YAML** tab on the Red Hat OpenShift console.

a) Create a backup file of the MongoCfg CRD YAML file.

To create a backup file, either copy and paste the MongoCfg CRD YAML file contents to a file or use the following **oc** command:

```
oc get MongoCfg <instance name> -o yaml > instance_name.yaml
```

b) Create a secret *docdb-dbadmin* YAML file in the *mas-[<ClusterUniqueString>-core](#)* namespace.

The following example is a sample secret YAML file in the *mas-vbpoyr-core* namespace.

```
---
# DocumentDB credentials for Core
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: documentdb-admin
  namespace: mas-vbpoyr-core
stringData:
  username: docdbadmin
  password: docdbadmin
```

c) Update the MongoCfg CRD YAML file with the DocumentDB connection details.

- **spec.certificates** – The certificate for the specific DocumentDB. Amazon Web Services DocumentDB provides a PEM file.

If the certificate authority is *rds-ca-2019*, then download the PEM files from [Certificate bundles for specific AWS Regions](#). For example, for *ca-central-1* region, you can use the PEM file from [ca-central-1-bundle.pem](#).

However, if the certificate authority is *rds-ca-rsa2048-g1*, then to download the Root CA RSA2048 PEM files follow steps provided in the [Extracting Root CA RSA2048 PEM files](#) section.

- **spec.config.credentials.secretName** – The secret that contains the DocumentDB admin username and password.
- **spec.config.hosts.host** – DocumentDB hostname.
- **spec.config.hosts.port** – DocumentDB port.
- **spec.config.retryWrites** – Set to false.

The following example is a sample MongoCfg CRD YAML file that is updated with DocumentDB connection details:

```
---
apiVersion: config.mas.ibm.com/v1
kind: MongoCfg
metadata:
  resourceVersion: '369530'
  name: vbpoyr-mongo-system
  uid: 48318d43-b0f6-4848-822b-4443a104b8f0
  generation: 1
  namespace: mas-vbpoyr-core
  ownerReferences:
    - apiVersion: core.mas.ibm.com/v1
      kind: Suite
      name: vbpoyr
      uid: 0df07123-1ad0-492a-aec0-76f88583dd98
  labels:
    mas.ibm.com/configScope: system
    mas.ibm.com/instanceId: vbpoyr
spec:
  certificates:
    - alias: ca
      crt: |
        -----BEGIN CERTIFICATE-----
        MIIEBjCCAu6gAwIBAgIJAMc0ZzaSUK51MA0GCSqGSIb3DQEBCwUAMIGPMQswCQYD
        VQQGEwJVUzEQMA4GA1UEBwwHU2VhdHRsZTETMBEGA1UECAwKV2FzaGluz3RvbjEi
        MCAGA1UECgwZQW1hem9uIFd1YiBTZXJ2aWNlcywgSW5jLjEjETMBEGA1UECwwKQW1h
        em9uIFJEUzEgMB4GA1UEAwwXQW1hem9uIFJEUyBSb290IDlwMTkgQ0EwHhcNMjkw
        ODYMTcwODUwWhcNMjkwODYMTcwODUwWjCBjzELMAkGA1UEBhMCVVMwEDAOBgNV
        BACMB1NlYXR0bGUxZzARBgNVBAGMClhc2hpbmd0b24xIjAgBgNVBAoMGUFtYXpv
        biBZXWlGU2Vydm1jZXMsIEluYy4xZzARBgNVBAsMCKFtYXpvbiBSRFMxIDAEBgNV
```



```

MIIGBTCCA+2gAwIBAgIRAJfKe4Zh4aWnt3bv6ZjQwogwDQYJKoZIhvcNAQEMBAw
gZoxCzAJBGNVBAYTA1VTMSIwIAYDVQQKDB1BbWV6b24gV2ViIFN1cnZpY2VzLCBJ
bmMuMRMwEQYDVQQLDAPBbWV6b24gUkRTMQswCQYDVQQIDAJXQTEzMDUyZjZl
QW1hem9uIFJEUyBjYjY1SjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2IEcxMRAw
DgYD
VQQHDAAdTZWf0dGx1MCAxDTIwMDUyZjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2
IEcxMRAwDgYD
mJlELMAKGA1UEBHMVVMVXVjAgBgNVBAoMGUFTYXpvcjB1BjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2IEcxMRAwDgYD
Yy4xZmZlZmV6b24gUkRTMQswCQYDVQQIDAJXQTEzMDUyZjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2IEcxMRAwDgYD
bWV6b24gUkRTMQswCQYDVQQIDAJXQTEzMDUyZjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2IEcxMRAwDgYD
BACMB1N1YXR0bGUwggIiMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQCPgUH6
bWV6b24gUkRTMQswCQYDVQQIDAJXQTEzMDUyZjZlZW50cmFSLTEgUm9vdCB0SBSU0E0MDk2IEcxMRAwDgYD
Crzd8c0w9prAh2rkQqA0x2vtuI7xX4tmBG4I/um28eBjyVmgwQ1fpq0Zg2nCKS54
Nn0pCmT7f3h6Bvopxn0J45AZXETAjFqXf92NQ3iPth95GVfAJSD7gk2LWmHpmID9
JGQyoGuDPg+hYyr292X6d0madzEkTVVG04mKTF989qEg+ty8+oN0U2fRTTraqa2tZp
iYsmg350ynNopvntsJAfpCO/srwpsqHHLNFZ9jvhTU8uW90wgaK09i31j/mHggCE
+CA0aJCM3g+L8DP1/2QKsb6UkBgaaIwKyRgKSj1IlgRk+0dCBC0gM9jjId4Tqo2j
ZIRRPG16fbn1+etZX+2/tf6tegz+yV0HHQRACkCpaH8AXF44bny9ands1BoNjGx
H6R/3ib4FhPrnBME1zZ5i4+eM/cuPC2huZMBXb/jKGRc/QN1Wm3/nah5FWq+yn+N
tiAF10Ga0BYzVhHDEwZzN7gn38bcY5yi/CjDUNPY00zEe2+dpaBKP1XTaFfn9Nba
CBmXPRF01LGGtPeTagjCju+NEcVa82Ht1pqxyu2sDtbu3J5bXP4RKtj+ShwN8nut
Tkf5Ea9rSmHEY13fzqibZ1QhXaifSKA2ASUwgJP19Putm0XK1BCNSGCoEcemewL
+7Y8FszS4Uu4eaIwwXVqUEE2yf+4ex0hqQ1acQIDAQABo0IwQDAPBgNVHRMBAf8E
BTADAQH/MB0GA1UdDgQWBBSU0E0MDk2IEcxMRAwDgYD
BAMCAYYwDQYJKoZIhvcNAQEMBAQggIBAIpRvxVS0dzoosBh/qw65ghPUGSbP2D4
dm6oYCV5g/zJr4fR7NzEbHOX5a0QnHbQL4M/7veuOCLNPOW1uXwywMg6yY+dbKe
YtPVA1as8G9sUyadeXyGh2uXGszimFXyaESwiAXZyiYyKChS3+g26/7jwECFo5vC
XGhWpIO7H35Yglp8AnwnEAo/PnuXgyt2nvyTSrx1EYa0jus6GZEZd77pa82U1JH
qfHIGmKPWdVLA3+ra1nKnvpWM/xX0pnMznMej5B3RT3Y+k61+kWghJE81IX78T
+tg4jSotgbal53BhtQWBD1yzbbi1qsGE1/DXPXzHVf9yD73fwh2tGWSaVInKYinr
a4tcrB3KDN/PFq0/w5/21lpZjVFyu/eiPj6DmWduHW73XnRwZpHo/20Fkei5R7cT
rn/YdDD6c1dYtSw5YnN56hdCQ3s0iB/xbPRN9VWJa6se79uZ9NLz6RM0r73DNnb2
bhIR9Gf7XAA51YKqQk+A+stokBIT0F65RnkxrXi/6vSiXfCh/bV6B41cf7MY/6YW
ehserSdjhQamv35rTFdM+foJwUKz1QN9n9KZHPxeRmwqPitAV79P1oks0nX25E1N
SlyxdndIoA1wia1HRd26EFm2pqfZ2vtD2EjU3wD42CXX4H8fKVdNa30nNFSYF0yn
jGKc3k6UNxpg
-----END CERTIFICATE-----

```

```

config:
  authMechanism: DEFAULT
  configDb: admin
  retryWrites: false
  credentials:
    secretName: documentdb-admin
  hosts:
    - host: docdb-vbpojr.ctnrsnscupeqf.ca-central-1.docdb.amazonaws.com
      port: 27017
displayName: Document Db in 'mongoce-vbpojr' namespace
type: external

```

- d) Save the MongoCfgr CRD YAML file.
After you apply the changes, reconciliation runs to connect to the new DocumentDB instance. On successful connection, the last reconciliation in the condition section is shown as successful.
- 5. Update the LicenseService CRD with the Amazon Web Services DocumentDB connection details.
 - a) Create a backup file of the LicenseService Suite License Service service instance YAML file.
 - b) Create a secret docdb-dbadmin YAML file in the ibm-sls-<ClusterUniqueString> namespace.
The following example is a sample secret YAML file in the ibm-sls-vbpojr namespace.

```

---
# DocumentDB credentials for SLS
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: documentdb
  namespace: ibm-sls-vbpojr
stringData:
  username: docdbadmin
  password: docdbadmin

```

- c) On the Red Hat OpenShift console, from the side navigation menu, select **CustomResourceDefinition > LicenseService**.
- d) Select the **Instances** tab and click the CRD name. Update the LicenseService CRD YAML file with the following DocumentDB connection details.
 - i) **spec.mongo.certificates** – The certificate for the specific DocumentDB.


```

nodes:
  - host: docdb-vbpoyr.ctnrnscepqf.ca-central-1.docdb.amazonaws.com
    port: 27017
  retryWrites: false
  secretName: documentdb-admin
  reloadCrd:
    fileUpload: -336090572
  settings:
    auth:
      enforce: true
    compliance:
      enforce: true
    registration:
      open: true
    registry: cp.icr.io/cp

```

- e) Save the LicenseService YAML file.
 - f) Verify whether the same status conditions are shown for DocumentDB.
6. In the api-licensing pod that is located in **ibm-sls project > namespace**, verify that the new connection with DocumentDB is established.
 7. Confirm whether the reconciliation is successful.

If a message Licensing system has not been initialized, upload a valid entitlement file to enable the Token pool and License Mgmt APIs is shown, upload the entitlement file again.

 - a. Confirm that the api-licensing pod is running with a ready status.
 - b. To initialize the licensing system, download the entitlement file. On the Red Hat OpenShift console, from the side navigation menu, select **Workloads > Secrets**, select IBM-ls-project, search for `ibm-sls-masocp-<<unique string>>-entitlement`, save, and download the entitlement file.
 - c. In the Maximo Application Suite **Setup** page, in the Settings section, the license key checkmark is disabled.
 - d. Click **Replace** license file and upload the file that was downloaded in the previous step.

After the upload is successful, the license key checkmark is enabled.
 - e. Confirm the LicenseService details are reinitialized.
 - f. After the MongoCfg and LicenseServiceCfg are configured successfully, verify that the following two collections in the new DocumentDB instance are created.
 - `ibm-sls-<<unique_cluster_string>>_masocp-<<unique_cluster_string>>_licensing` specific to Maximo Application Suite core.

For example, `ibm-sls-vbpoyr_masocp-vbpoyr_licensing`.
 - `mas-<<unique_cluster_string>>_core` specific to Suite License Service (SLS).

For example, `mas_vbpoyr_core`.

The SLS-specific token pools and products collection are also created in the new DocumentDB instance that confirms that the DocumentDB was successfully connected to SLS.
 - g. Create an admin user and confirm that the user is shown in the Maximo Application Suite core and Suite License Service collections.

Note: On the Maximo Application Suite **Configurations** page, in Storage section, select MongoDB to view the configured database details. You cannot update the MongoDB settings on the details page.

To create an administrator user, add details, such as the display name, user ID, and password. Select **Authorized** for the access type and add entitlements.

The new user is visible in the Users section of `mas-<<cluster_unique_string>>_core` DocumentDB, which confirms that Maximo Application Suite core is connected with DocumentDB.

The new user, who is an admin user, is added to the `mas-<<cluster_unique_string>>_core` DocumentDB collection.
 8. Delete the existing MongoDB instance by using CLI.

The following example is a sample script to delete the MongoDB instance.

```
# Login to cluster
oc login --token=sha256~xxx --server=https://api.masocp-xupgew.domain.com:6443
# Switch to mongo namespace
oc project ${MONGO_NAMESPACE}

# Note down Certificate resource name
oc get Certificate -n ${MONGO_NAMESPACE}
# Note down StatefulSet name
oc get StatefulSet -n ${MONGO_NAMESPACE}
# Note down mongodbccommunity crd detail
oc get CustomResourceDefinition mongodbccommunity.mongodbccommunity.mongodb.com -n ${MONGO_NAMESPACE}
# Note down pvc name
oc get pvc -n ${MONGO_NAMESPACE} | grep mongo

# Delete Certificate
oc delete Certificate mongo-ca-crt -n ${MONGO_NAMESPACE}
oc delete Certificate mongo-server -n ${MONGO_NAMESPACE}

# Delete MongoDBCommunity
oc delete MongoDBCommunity mas-mongo-ce -n ${MONGO_NAMESPACE}
# Delete MongoDB operator
oc delete deployment mongodb-kubernetes-operator -n ${MONGO_NAMESPACE}
# Delete MongoDB StatefulSet
oc delete StatefulSet mas-mongo-ce
# Delete MongoDB StatefulSet
oc delete CustomResourceDefinition mongodbccommunity.mongodbccommunity.mongodb.com -n ${MONGO_NAMESPACE}

# Delete Mongo secrets
oc delete secrets --all -n ${MONGO_NAMESPACE}
# Delete Mongo configmaps
oc delete configmaps --all -n ${MONGO_NAMESPACE}
# Delete Mongo pvc
oc delete pvc data-volume-mas-mongo-ce-0 -n ${MONGO_NAMESPACE}
oc delete pvc data-volume-mas-mongo-ce-1 -n ${MONGO_NAMESPACE}
oc delete pvc data-volume-mas-mongo-ce-2 -n ${MONGO_NAMESPACE}
oc delete pvc logs-volume-mas-mongo-ce-0 -n ${MONGO_NAMESPACE}
oc delete pvc logs-volume-mas-mongo-ce-1 -n ${MONGO_NAMESPACE}
oc delete pvc logs-volume-mas-mongo-ce-2 -n ${MONGO_NAMESPACE}
# Delete Mongo project
oc delete project ${MONGO_NAMESPACE}
```

Extracting Root CA RSA2048 PEM files

If you see a connection error, you can run the following commands to manually add new CA certificates to your trust stores, and update your existing Amazon DocumentDB instances to use the new CA certificates.

1. Run the **wget** command to download the .pem file for your region.

For example run the command to download the .pem file for Canada (Central) region.

```
wget https://truststore.pki.ids.amazonaws.com/ca-central-1/ca-central-1-bundle.pem
```

2. Split the certificate from the certificate names into separate files.

For example run the following command.

```
cat ca-central-1-bundle.pem | awk 'split_after==1{n++;split_after=0}
/-----END CERTIFICATE-----/
{split_after=1} {print > "temp-ca-central-1" n ".pem"}'
```

3. Search the files that contain the term Root CA RSA2048.

```
for cert in $(ls temp-*.pem); do echo $cert; cat $cert | openssl x509 -noout -text |
grep "Root CA RSA2048 G1"; done
```

Files that contain the term are listed in the output.

4. Copy the file that is required to `root-ca-rsa2048` for your region.

```
cp temp-ca-central-13.pem ../root-ca-rsa2048-ca-central-1.pem
```

5. Verify that the correct `.pem` files are copied.

```
for cert in $(ls root-ca-rsa2048-*.pem); do echo $cert; cat $cert | openssl x509 -noout -text | grep "Root CA RSA2048 G1"; done
```

6. Replace the `spec.certificates.crt` and `spec.mongo.certificates.crt` parameters with the certificate content from the new PEM file.

Configuring IBM Maximo Application Suite to use Amazon MSK

Configure Amazon Managed Streaming for Apache Kafka (Amazon MSK) so that IBM Maximo Application Suite can process streaming data for applications.

Provisioning Amazon MSK

You can provision Amazon Managed Streaming for Apache Kafka (Amazon MSK) so that IBM Maximo Application Suite processes the streaming data for applications.

Procedure

1. On the Amazon Web Services console, search for `msk`.
The Amazon MSK service is listed.
2. To open the Amazon MSK console, select **MSK**.
3. In the **Get started** section, click **Create cluster**.
The **Cluster settings** page is opened.
4. Configure cluster settings, such as creation method, cluster name, cluster type, Apache Kafka version, brokers, and cluster configurations.
 - a) To customize settings, select **Custom create as the Creation method**.
 - b) Assign a name for your cluster, for example, `aws_msk`.
 - c) To specify the number of brokers and the amount of storage that is required for each broker, select **Provisioned as the Cluster type**.
 - d) Select the recommended 2.8.1 Apache Kafka version.
 - e) In the Brokers section, select three zones and one broker for each zone so that your cluster has three total brokers that are distributed evenly across your three availability zones..
 - f) In the Configurations section, select **Custom configuration**.
 - g) On the **Create configuration** page, enter the configuration name and specify the configuration properties.
You must append the `allow.everyone.if.no.acl.found` property and set it to `true` in the Apache Kafka code.
 - h) Click **create**.
The cluster configuration is created and listed on the Amazon MSK console.
 - i) To finish the configuration, on the MSK console, click your new cluster configuration.
You return to the **Configuration** page.
 - j) Click **Next** to configure the network settings.
5. To deploy your brokers, configure network settings, such as Virtual Private Network (VPC), zones, subnet, and security groups.
 - a) Select the virtual networking environment for your cluster.
 - b) Select the subnet for each zone that was configured.

- c) To create the security groups, which act as a virtual firewall for your instances to control inbound and outbound traffic, click **create**.
Before you click create, ensure that you remove the default security group from the list of Chosen security groups.
 - d) On the **Create security group** page, add the security group name and description and select the VPC unique string that was created earlier.
 - e) For inbound and outbound rules, select **Custom TCP** for the rule type and **Custom** for the source.
The security group is created.
 - f) Return to the **Create security group** page to select your custom security group, and click Choose.
 - g) Click **next** to add the security settings.
 - h) In the Security settings, select **SASL/SCRAM authentication** as your access control method and create a custom managed key to encrypt your data.
 - i) Select details, such as Symmetric for the key type, encrypt and decrypt as key usage, KMS as key material origin, and single-region key as region.
 - j) Customize your key configurations for alias, tags, administrators, and usage permissions.
 - k) Assign the key policy and click **finish** to complete the custom key configuration.
6. On the **Monitoring and tags** page, select basic monitoring and click **next**.
 7. On the **Review and create** page, view and verify the configured parameters and click **Create cluster**.
The custom cluster is provisioned and shown on the MSK console.
 8. To associate the Amazon Web Services Secrets Manager secrets with the cluster to configure SASL/SCRAM authentication, click **Associate secrets**.
 - a) On the **Associate secrets** page, click **Create secret**.
 - b) On the **Choose secret type** page, select the encryption secret key and provide details, such as secret type and key or value pairs.
 - c) On the **Configure secret** page, enter the secret name.
 - d) Optional: Configure rotation details for the secret, such as automatic rotation, rotation schedule, and rotation function.
 - e) Review the secret information, such as secret type and sample code.
 - f) Click **store**.
 - g) On the **Associate secrets to cluster** page, choose the secret that was created and click **Associate secrets**.

For more information, see <https://docs.aws.amazon.com/msk/latest/developerguide/msk-password.html>
 9. On the MSK console, verify that the client information is added.

What to do next

Configure the IBM Maximo Application Suite to use a provisioned Amazon MSK instance.

For more information, see [“Configuring IBM Maximo Application Suite with Amazon Web Services MSK” on page 260](#).

Configuring IBM Maximo Application Suite with Amazon Web Services MSK

You can configure Maximo Application Suite to use a provisioned Amazon MSK instance.

Procedure

1. Log in to Maximo Application Suite as an admin user.
2. On the **Configurations** page, select the Apache Kafka option from others list.
3. To add Apache Kafka under the system scope, click **Configure**.

4. Add parameters, such as the hosts, port, SASL mechanism as scram-aha-512, username, and password.
5. Paste the Amazon ca-certificate details that you can copy from <https://www.amazontrust.com/repository/AmazonRootCA1.pem>.
6. Click **Confirm** and then **Save**.
7. Verify the Kafka configuration status in **Configurations > Apache Kafka > System scope**.
8. Optional: To verify that the Kafka configuration is running, open the instances tab of the KafkaCfg CRD for the selected namespace on the Red Hat OpenShift console.

Deploying and activating IoT and Maximo Monitor

After you provision Amazon MSK and configure it to work with IBM Maximo Application Suite, you can deploy and activate Apache Kafka from the IoT tool and Maximo Monitor.

Procedure

1. Log in to Maximo Application Suite as an admin user.
2. On the **Catalog** page, on the **Tools** tab, select IoT.
3. On the **IoT** page, from the toolbar, click the **Administration** icon.
4. Click **Continue** to open the **Administer application upgrades** page.
5. Select an upgrade strategy and subscribe to a channel.
For example, select **Channel subscription** and subscribe to the 8.x channel.
The IoT tool deployment begins.
6. After the IoT tool is deployed successfully, on the **Deploy IoT** page, configure the database connection.
7. Click **Exit**.
8. On the **Configurations** page, to view database details, from the **Storage** list, select the database connection.
You must configure the database connection by adding the Java database connectivity values at the system scope level. The information for connectivity is available in the IBM Maximo Manage Workspace-application scope.
9. Select the Workspace-application scope to edit the JDBC connection information, such as the connection string, username and password, and certificate content.
Tip: You can find the username and password from the JDBC-DB2® namespace on the **Secrets** page on the Red Hat OpenShift console.
10. Click **Save**.
11. Verify that the connection is configured on the Red Hat OpenShift console **CustomResourceDefinitions** page and the Maximo Application Suite **Configurations > Database connection** page.
12. Deploy and activate the IoT tool.
For more information, see [Deploying the IoT tool](#).
13. Repeat the steps for Maximo Monitor.
For more information, see [“Deploying IBM Maximo Monitor” on page 370](#) and [“Activating IBM Maximo Monitor” on page 467](#).

What to do next

Verify that the Apache Kafka connection is established in the Maximo Monitor application by using the IoT device simulator.

For more information, see [“Verifying Apache Kafka connection with IoT simulator” on page 262](#).

Verifying Apache Kafka connection with IoT simulator

Verify that the Apache Kafka connection is established from the Maximo Monitor application by using the IoT device simulator.

Procedure

1. Log in to Maximo Application Suite as an admin user.
2. Create a user by providing details, such as identity, entitlements, and application access to Maximo Monitor and the IoT tool.
3. On the **Catalog** page, on the toolbar, click the **AppSwitcher** icon.
4. Select **Maximo Monitor** and create device types and add metrics.
5. Return to the **Catalog** page, and on the toolbar, click the **AppSwitcher** icon.
6. Select **IoT** and create a device simulator.
7. Configure the device simulator by selecting a device type, configuring the event and payload, and adding device.
8. To verify that the simulation was successful, view the number of events that are shared in Simulation.

IBM Maximo Application Suite installation with Microsoft Azure Resource Manager templates

You can install Maximo Application Suite in the Microsoft Azure cloud by using the Microsoft Azure Resource Manager templates. In Microsoft Azure Marketplace, you subscribe to Maximo Application Suite, configure the installation parameters, and install the application. The network infrastructure, Red Hat OpenShift cluster, and Maximo Application Suite components are created in your Microsoft Azure cloud account.

When you select the **New OpenShift cluster (IPI)** option, the network infrastructure, Red Hat OpenShift cluster, and Maximo Application Suite components are created in your Microsoft Azure cloud account. You can reuse existing Red Hat OpenShift cluster from IPI or UPI with the **Existing OpenShift cluster** option.

Maximo Application Suite is available as a bring-your-own-license (BYOL) and contract pricing product in Microsoft Azure Marketplace. After you configure the installation requirements and consider your installation preferences, you subscribe to the product, specify the installation parameters, and start the installation.

In your Microsoft Azure account, the installation process creates the virtual network infrastructure, the Red Hat OpenShift cluster, the application prerequisites, and the application itself.

If you provide the SMTP configuration during the deployment, you receive emails that contain the information that you need to access Maximo Application Suite.

Installing IBM Maximo Application Suite on Microsoft Azure

The IBM Maximo Application Suite installation process involves configuring the installation prerequisites, finalizing your installation preferences, installing the application, and completing the initial application setup.

Before you begin

You must configure prerequisites and gather the information that you need to specify the installation parameters. For more information, see [“Planning to install on Microsoft Azure” on page 160](#).

In addition, you must consider your installation preferences, such as the type of offering that you want and whether you want to create an Red Hat OpenShift cluster or reuse an existing one.

Note: The existing cluster must have been created by using the automated deployment option.

For more information, see [Installation considerations](#).

Installing Maximo Application Suite

To install the Maximo Application Suite on Microsoft Azure, you configure the prerequisite components, consider your installation preferences, specify the installation criteria in parameters that are provided during the deployment, and deploy the product.

Before you begin

Before you can install Maximo Application Suite on Microsoft Azure, you must configure prerequisites and gather information that you need to complete the installation. For more information, see [“Prerequisites for installing Maximo Application Suite on Microsoft Azure”](#) on page 163.

You must also consider other criteria, such as the type of Maximo Application Suite offering that you want and whether you want to create an Red Hat OpenShift cluster or reuse an existing one.

Note: The existing cluster must be created by using the automated deployment option only.

If you provide the existing Red Hat OpenShift cluster, the installation process checks if any of the following products are already installed:

- Red Hat OpenShift 4.10.35
- IBM Cloud Pak foundational services 4.6.0
- IBM Cloud Certificate Manager 3.21.1 or higher
- MongoDB (CE) 4.2.6 or higher
- IBM Suite License Service 3.4.0 or higher
- IBM Data Reporter Operator
- IBM Cloud Pak for Data 4.0.9
- IBM Maximo Application Suite 8.10
- IBM Maximo Manage 8.6

For more information, see [“Preparing to installing Maximo Application Suite on Microsoft Azure”](#) on page 167.

Related concepts

[Operational mode for installation](#)

From Maximo Application Suite 8.9 or later, you can consider installing and deploying IBM Maximo Application Suite in Production or Non-production mode based on your development and testing requirements and to optimize your AppPoint usage.

Installing BYOL IBM Maximo Application Suite

You can install Maximo Application Suite in Microsoft Azure.

About this task

To install Maximo Application Suite in Microsoft Azure, the following three fulfillment options are available.

1. Existing Red Hat OpenShift cluster
2. New Red Hat OpenShift cluster by using the Installer Provisioned Infrastructure (IPI)
3. New Red Hat OpenShift cluster by using the User Provisioned Infrastructure (UPI)

After you select the fulfillment options from Microsoft Azure Marketplace, complete the following steps.

Procedure

1. In the Microsoft Azure Marketplace, use the search function and search for Maximo Application Suite (BYOL)
2. Open the Suite product.
3. Review the product information and click **Subscribe**.

4. In the Subscribe to IBM Maximo Application Suite (BYOL) step, enter the installation parameters by using the information that you gathered when you configured the [“Prerequisites for installing Maximo Application Suite on Microsoft Azure”](#) on page 163 and considered your [“Preparing to installing Maximo Application Suite on Microsoft Azure”](#) on page 167.

Basics (Required)

- Resource group

This resource group is a new boot node resource group, as the boot node is created in this resource group. The Red Hat OpenShift cluster is created in its own resource group that is created by the Red Hat OpenShift installer when you use the New OpenShift cluster (IPI) option.

From Maximo Application Suite 8.8 or later, when you use the 'New OpenShift cluster, existing network (UPI) option, the Red Hat OpenShift cluster resources are created in the resource group where the existing VNet is configured.

- Region
- Subscription Id
- Public key value
- Boot node NSG Ingress CIDR range
- Bootnode Subnet CIDR IP range

The Bootnode Subnet Classless Inter-Domain Routing (CIDR) IP range is required for UPI and existing OCP deployments.

Tip: The bootnode subnet CIDR range should be in the Vnet CIDR range.

Starting in 8.10, if you are reusing existing Red Hat OpenShift cluster or starting in 8.11, if you are using existing Microsoft Azure Red Hat OpenShift cluster, the following options are available.

- BootNodeVnetId
- BootNodeVnetResourceGroup
- Microsoft Azure service principal ID
- Microsoft Azure service principal client secret

Application Settings (Optional)

- Public or hosted domain. This is required only if you are provisioning new Red Hat OpenShift cluster in IPI or UPI mode. For an existing Red Hat OpenShift cluster, public domain field is not available.

Public domain is changed to **Hosted domain** from Maximo Application Suite 8.10 or later.

If you provision a new Red Hat OpenShift cluster by using the IPI option, the hosted domain value is fetched from the Public DNS zone. If you provision an existing Red Hat OpenShift cluster by using the UPI option, the hosted domain value is fetched from the Private DNS zone for installing a private cluster and DNS Zone for a public cluster.

- Offering type
- From Maximo Application Suite 8.8 or later, Cluster size. This is required only if you are provisioning new Red Hat OpenShift cluster in IPI or UPI mode.
- Entitled registry key
- Red Hat OpenShift pull secret . This is required only if you are provisioning new Red Hat OpenShift cluster in IPI or UPI mode. If using existing Red Hat OpenShift cluster, keep the value empty.
- Maximo Application Suite license URL
- From Maximo Application Suite 8.9 or later, Operational Mode - specify the parameter as Production or Non-production.

You can use the non-production installations for internal development and testing. The installation AppPoints are unused in the non-production installations. These specifications are also visible in the metrics shared with IBM as well as on the product UI.

Existing network infrastructure

Existing VNet to use, For more information, see [“SSH key pair” on page 164](#).

Existing Infrastructure

The following information is required only if you are reusing existing Red Hat OpenShift cluster.

- Red Hat OpenShift cluster API URL.
- Red Hat OpenShift user
- Red Hat OpenShift password

For more information, see [Existing network infrastructure by using Red Hat OpenShift UPI mode](#).

- SLS endpoint URL
- SLS registration key
- SLS public certificate URL
- DRO endpoint URL
- DRO API key
- DRO public certificate URL

Database Settings

- Maximo Application Suite Manage DB user
- Maximo Application Suite Manage DB password
- Maximo Application Suite Manage DB JDBC URL
- Maximo Application Suite Manage DB certificate URL
- Starting in 8.11, VNetId of the database that is provisioned
- Import demo data

Note: If you choose to install Maximo Application Suite with Manage, you can use the default IBM Db2 instance that is provisioned by IBM instead of configuring your own external Db2 instance. To configure the default IBM Db2, do not add information in the username, password, JDBC URL, certificate URL, and demo data fields.

The internal Db2 configuration is available from Maximo Application Suite 8.10 or later.

Email Settings

- Email notification
- SMTP host
- SMTP port
- SMTP username
- SMTP password
- Notification email addresses

5. To begin the installation, click **Review + > Create**.

Installing customer managed IBM Maximo Application Suite

You must first subscribe to IBM Maximo Application Suite to complete the transactional aspect of the IBM Maximo Application Suite customer managed purchase process.

About this task

You can install either by using the public paid offer or private paid offer.

Procedure

- **Installing with public paid offer**

- a) In the Microsoft Azure, use the search function and search for Maximo Application Suite (customer-managed).
- b) Open the IBM Maximo Application Suite product.
- c) Review the product information and click **Subscribe**.
- d) In the Subscribe to IBM Maximo Application Suite (customer-managed) step, enter the parameters based on the following guidelines.

Basics (Required)

- Resource group: This resource group is used to keep the application instance representing the purchased instance of IBM Maximo Application Suite (customer-managed) product.
 - Region
 - Microsoft Azure location
 - Name: Name of the application instance representing the purchased IBM Maximo Application Suite (customer-managed) product.
 - Recurring billing: Whether you want the renew or terminate the billing after subscription is ended.
- e) To review and initiate the purchase process, click Review and Subscribe.
Note: The public offer is a fixed contract of 12 months with 500 AppPoints.
 - f) To initiate the purchase process, click Subscribe.
 - g) You see the message that subscription is in progress.
 - h) After the subscription process completes, you will get an auto-generated email from Microsoft to activate IBM Maximo Application Suite (customer-managed) subscription.
 - i) You can click the **Activate now** in the email or **Configure account now** available for the offer. Clicking on the button will take you to the IBM registration page where you will need to provide the following details.

Company

Your company name.

Full name

Your full name.

Email address

Your corporate email address.

IBM ID

Email address that you have registered as your IBM ID.

Offer type

Select the public offer type.

IBM Quote number

Leave this empty. This is required only for private offer.

Company address

This is required for the public offer.

State

State where the company is located. This is required for the public offer.

Zip code

Zip code of the address. This is required for the public offer.

Country

Country where the company is located. This is required for the public offer.

- j) Click the COMPLETE REGISTRATION button.

You will get an auto-generated email from Tackle.io about the process initiation for setting up your account.

- k) After the account setup is completed, you will get two emails from IBM. First email describes how to retrieve IBM Entitled Registry key from My IBM. Second email describes how to get the IBM Maximo Application Suite license and Red Hat OpenShift pull secret.
- l) After these artifacts are retrieved, perform the actual product deployment by following the installations steps for IBM Maximo Application Suite (BYOL) product.

- **Installing with private paid offer**

- a) Contact your IBM Sales Representative for a customized subscription contract (Private Offer) with recommended configuration and negotiated pricing.

For more information, see <https://www.ibm.com/products/maximo/pricing>.

- b) After you agree upon the price and number of AppPoints, IBM sends a link so that you can accept the private offer.

- c) Only users with the permissions shown in this table can access and sign the offer contract.

- For Microsoft Customer Agreement accounts, the user role must be billing account owner or billing account contributor.

- For EA accounts, the user role must be EA admins.

For more information about the type of account you have, see [Billing accounts and scopes](#).

Accepting a private offer simply means you've agreed to the terms and prices listed in the offer. No purchase has been made and no money has exchanged yet.

- d) After the private offer is accepted, it will appear in the marketplace under **My Marketplace** > **Private products** section.

- e) Select the appropriate offer in case you have multiple private offers for same or different products.

- f) Review the product information and click Subscribe.

- g) In the Subscribe to IBM Maximo Application Suite (customer-managed) step, enter the parameters based on the following guidelines.

- **Basics (Required)**

- Resource group: This resource group is used to keep the application instance representing the purchased instance of IBM Maximo Application Suite (customer-managed) product.

- Region

- Microsoft Azure location

- Name: Name of the application instance representing the purchased IBM Maximo Application Suite (customer-managed) product.

- Recurring billing: Whether you want the renew or terminate the billing after subscription is ended.

- h) To review and initiate the purchase process, click Review and Subscribe.

- i) To initiate the purchase process, click Subscribe.

- j) You see the message that subscription is in progress.

- k) Once subscription process completes, you will get an auto-generated email from Microsoft to activate IBM Maximo Application Suite (customer-managed) subscription.

- l) You can click **Activate now** in the email or **Configure account now** available for the offer. Clicking on the button will take you to the IBM registration page where you will need to provide the following details.

- **Company**

- Your company name.

- **Full name**

- Your full name.

Email address

Your corporate email address.

IBM ID

Email address that you have registered as your IBM ID.

Offer type

Select the public offer type.

IBM Quote number

Provide the IBM quote number

Company address

Leave this empty. This is required only for the public offer.

State

Leave this empty. This is required only for the public offer.

Zip code

Leave this empty. This is required only for the public offer.

Country

Leave this empty. This is required only for the public offer.

- m) Click the COMPLETE REGISTRATION button.

You will get an auto-generated email from Tackle.io about the process initiation for setting up your account.

- n) After the account setup is completed, you will get two emails from IBM. First email describes how to retrieve IBM Entitled Registry key from My IBM. Second email describes how to get the IBM Maximo Application Suite license and Red Hat OpenShift pull secret.
- o) After these artifacts are retrieved, perform the actual product deployment by following the installations steps for IBM Maximo Application Suite (BYOL) product.

Installing Cloud Pak for Data on an Azure instance of Maximo Application Suite

To install Cloud Pak for Data on an existing Maximo Application Suite instance on Azure, you clone a Git repository and run a script.

About this task

If you installed Maximo Application Suite core and Maximo Manage, your Maximo Application Suite instance does not include Cloud Pak for Data. If you want to deploy certain applications and add-ons in the Suite, such as Maximo Monitor, you must first install Cloud Pak for Data.

To install this application, you clone a Git repository, locate the installation script, and run it. When you run the script, you must provide the Azure deployment values that identify the Maximo Application Suite instance, such as boot node's resource group, and your entitled registry key. The installation script retrieves the Suite's <unique-string> and the Red Hat OpenShift cluster details from the resource group's deployment object. If the cluster credentials, that is the username and password, are changed since you installed the Suite, you must provide the updated credentials when you run the script.

You run the script on your local machine or on the boot node in the Suite's Red Hat OpenShift cluster.

Installing Cloud Pak for Data

Complete the following steps to install Cloud Pak for Data in your Maximo Application Suite instance:

Procedure

1. On the machine where you want to run the script, in a command shell, log in to the Azure service by running the following command.

```
az login
```

It opens a new browser window where you can log in with the Azure credentials. After login, you can close the browser window and continue the next steps from the command shell.

2. Clone the Git repository that contains the script by running the following command:

```
git clone https://github.com/ibm-mas/multicloud-bootstrap.git
```

3. Make the script executable by running the following commands:

```
cd multicloud-bootstrap/azure
chmod +x deploy-cp4d.sh
```

4. View the script's usage information by running the following command:

```
./deploy-cp4d.sh -h
```

5. Specify the required options and run the script.

- Use the `r` option to specify the resource group name of the boot node, for example: `-r mas-ocp-deploy-rg`
- Use the `e` option to specify the entitled registry key that you provided when you installed the Suite, for example: `-e <entitlement-key>`
- Use the `u` and `p` options to specify the Suite's Red Hat OpenShift cluster credentials, for example: `-u <ocp-user> -p <ocp-password>`
- The following sample command installs Cloud Pak for Data by using all of these example options:
- The script takes 60 minutes to

```
./deploy-cp4d.sh -r <bootnode-resource-group> -e <entitlement-key> -u <ocp-user> -p <ocp-password>
```

installation Cloud Pak for Data into the Maximo Application Suite instance.

- Verify that the script completed successfully. If the script is successful, output that is similar to the following text is displayed:

```
:: Script Inputs ::
  Resource group = mas-ocp-deploy-rg
...
:: OpenShift Details ::
...
OpenShift Login is successful.
...
==== MAS configuration started ====
==== MAS configuration completed ====
==== Execution completed at Mon Mar 25 14:02:56 IST 2022 ====
```

- Maximo Application Suite core with Cloud Pak for Data is deployed.
- Db2 Warehouse service of Cloud Pak for Data is enabled.

For information about troubleshooting, see [IBM Support notes for zen-databases failure](#).

6. Log in to the Cloud Pak for Data admin console.
7. Select the service catalog and search for Db2 Warehouse.
8. Click **Provision instance** of Db2 Warehouse service.
9. Configure the database.

What to do next

You can now deploy the following applications and industry solutions that depend on Cloud Pak for Data:

- Maximo Application Suite
- Maximo Monitor /IoT tool
- Maximo Health
- Maximo Predict

- Maximo Collaborate
- Maximo Health and Predict - Utilities

Monitoring installation logs

During the installation process, the ARM template that you configured is used to create a boot node. The boot node contains all of the required resources to complete the installation. To verify that the boot node is created successfully, in the **Deployment** section of the boot node resource group, you will see a link similar to *3 Succeeded*. After clicking that link, you will see a deployment that is related to the operation you submitted from the Marketplace. Click that deployment and make sure it is completed successfully. This is only the boot node creation part. The actual Maximo Application Suite deployment continues running in the background on the boot node.

Procedure

- **Monitoring the installation logs using Azure Log Analytics**
 - In the Azure **Log Analytics workspaces** service, Click the workspace created in the boot node resource group.
 - Click **Logs** from the side navigation in the workspace window.
 - Close the two pop-ups that appear by default. First is for **Welcome to Log Analytics**, and second is for **Queries**.
 - Expand the **New Query 1 > Custom Logs** section.
 - Click the log table name (there would be only one), and click **Run**.
 - The logs from the installation log file are shown in the table format.

The logs do not update automatically, you will need to refresh the page to get the updated logs.

- You can export the logs to the `.csv` file to review in detail.
- Once you use the message of either provisioning completed (for successful deployment) or failed (for failed deployment), the deployment can be treated as complete.

```
===== PROVISIONING COMPLETED =====
```

OR

```
===== PROVISIONING FAILED =====
```

- **Monitoring the installation logs in the boot node**
 - Connect to the boot node by using Secure Shell (SSH) access. For instructions, see [Accessing the boot node and Red Hat OpenShift cluster](#).
 - Run the following command to switch to the root user:

```
sudo su -
```

- Monitor the installation log updates by running the following command:

```
tail -f /root/ansible-devops/multicloud-bootstrap/root/mas-on-aws/mas-provisioning.log
```

Results

Proactively check the status of the deployment. Depending on the parameters that you specified, the installation time might vary.

If the installation is unsuccessful, use the information in the [Troubleshooting installation problems](#) topic to identify and resolve the problem.

Accessing IBM Maximo Application Suite

After you install Maximo Application Suite, you need the following items of information to access it:

- The Maximo Application Suite administrator URL, which you use to connect to Maximo Application Suite through a browser.
- Your username and password.
- The public certificate for Maximo Application Suite. You import this certificate into your browser's trusted store to ensure secure communication between your browser and Maximo Application Suite.

How you retrieve these items of information depends on whether you opted for the email notification during the deployment.

If you have provided the correct SMTP configuration during the deployment, you can retrieve the administrator URL, username, and password from the emails that you received. The public certificate is attached to these emails.

If you have not opted for the email notification, you can retrieve the administrator URL and credentials of the Red Hat OpenShift cluster from the deployment created in the boot node's resource group. However, to retrieve your username, password, and the public certificate of Maximo Application Suite, you must connect to the Red Hat OpenShift cluster.

Procedure

Complete the following steps to retrieve all of the required information:

1. In the **Outputs** section of the deployment that was created in the bootnode resource group, record the following values:
 - The value of the `masAdminUrl` key. This value contains the administrator URL.
 - The values of the `openShiftConsoleUrl`. These values contain the URL for the Red Hat OpenShift console.
 - The value of the `clusterUniqueString` key. This value contains the cluster unique string to look for the correct resources in Red Hat OpenShift.
2. The Microsoft Azure Vault contains the credentials for Red Hat OpenShift cluster. Look for the vault named `maximo-vault-<unique-string>` in the Microsoft Azure Key Vaults service. It consists of two secrets named `maximo-ocp-secret` containing Red Hat OpenShift credentials and `maximo-mas-secret` containing Maximo Application Suite credentials. Retrieve the Red Hat OpenShift Container Platform credentials from this secret.

Note: By default no user has access to view the secrets. You must provide appropriate access from the **Access policies** section of the key vault.
3. Connect to the Red Hat OpenShift cluster.
4. Select **Workloads > Secrets** from the navigation page.
5. Select the `mas-<unique-string>-core` project.
6. Click the `<unique-string>-credentials-superuser` secret.
7. Click the 'Reveal values' link to get the username and password for Maximo Application Suite.
8. Click the `<unique-string>-cert-public` secret from the `mas-<unique-string>-core` project.
9. Click the 'Reveal values' link to get the contents of the certificates.
10. Retrieve the contents of `ca.crt` file, which is the public certificate for Maximo Application Suite.
11. After you import the public certificate into your browser's trusted store, paste the Maximo Application Suite administrator URL into your browser and enter the authentication credentials to access the application.

Results

You can now log in to Maximo Application Suite to deploy applications, create users, and specify configuration.

In addition, if you installed the Maximo Application Suite core, Maximo Application Suite core with Cloud Pak for Data is deployed.

- IBM Db2 Warehouse service of Cloud Pak for Data is enabled.

To configure the Cloud Pak for Data database:

1. Login to the Cloud Pak for Data admin console.
2. Select the service catalog and search for Db2 Warehouse.
3. Click Provision instance of Db2 Warehouse service.
4. Configure the database.

If you installed Maximo Manage offering type, and you want to deploy certain applications and add-ons in Maximo Application Suite, such as Maximo Monitor, you must first install Cloud Pak for Data. For more information, see [Installing Cloud Pak for Data on an Amazon Web Services instance of Maximo Application Suite](#).

Configuring Let's Encrypt for Maximo Application Suite on Microsoft Azure

When you install Maximo Application Suite on Microsoft Azure, Maximo Application Suite uses self-signed certificates. If you want to use well-known certificates signed by Certificate Authority as Let's Encrypt, you can install and configure Let's Encrypt on Microsoft Azure.

Before you begin

To configure Let's Encrypt, a service principal in Microsoft Azure must be created. For more information, see <https://cert-manager.io/docs/configuration/acme/dns01/azuredns/#service-principal>

About this task

The cert-manager can create and then delete **DNS-01** records in Microsoft Azure DNS. However, the DNS needs to authenticate with Microsoft Azure first. The following method uses the Microsoft Azure Service principal authentication to configure Let's Encrypt.

Procedure

1. Create a service principal in Microsoft Azure by using the service principal connection parameters. For example:

```
AZURE_DNS_ZONE_RESOURCE_GROUP=masperf
AZURE_DNS_ZONE=mas4azure.com
AZURE_CERT_MANAGER_SP_APP_ID=3xx721x5-xx3x-4x10-x39x-xxx31335405
AZURE_CERT_MANAGER_SP_PASSWORD=X87xXxX6q6xxXxXxx7nYhZmbXxxxX~Tho
AZURE_TENANT_ID=xxx67057-50x9-4xx4-98x3-xxxx64xxx9x9
AZURE_SUBSCRIPTION_ID=x2xx5467-2502-4b05-x78x-744604x6531x
```

Tip: For all examples, replace the parameters given in the example with your own parameters.

2. Log in to Microsoft Azure by using the service principal connection details.

```
az login --service-principal -u $AZURE_CERT_MANAGER_SP_APP_ID -p
$AZURE_CERT_MANAGER_SP_PASSWORD --tenant $AZURE_TENANT_ID
```

3. Create the **DNS Contributor** role to associate DNS zone with the service principal. For example:


```
DNS_ID=$(az network dns zone show --name $AZURE_DNS_ZONE --resource-group
$AZURE_DNS_ZONE_RESOURCE_GROUP --query "id" --output tsv)

az role assignment create --role "DNS Zone Contributor" --assignee-object-id
$AZURE_CERT_MANAGER_SP_APP_ID --assignee-principal-type ServicePrincipal --scope $DNS_ID
```

- In **DNS Record**, click your domain, create **A** record set
`*.<<Cluster_unique_String>>.<<DNS_NAME>>`.
 For example, `*.i417mh.mas4azure.com`, where **i417mh** is the cluster unique string and **mas4azure** is the DNS name.

Tip: Use the same value that is used in the **A** record.

- Login to the Red Hat OpenShift cluster.
 For example:

```
oc login --token=<<token_number>> --server=https://api.masocp-i417mh.mas4azure.com:6443
```

- Create a secret **azuredns-config**, which contains the service principal password.
 For example:

```
oc create secret generic azuredns-config --from-literal=client-
secret=$AZURE_CERT_MANAGER_SP_PASSWORD -n ibm-common-services
```

- In the Red Hat OpenShift console, create a **ClusterIssuer** from the **Instances** tab of the **Home** > **API Explorer** cert-manager.io group.
 For example:

```
apiversion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: letsencrypt-prod
  namespace: ibm-common-services
spec:
  acme:
    server: https://acme-v02.api.letsencrypt.org/directory
    email: username.ibm.com
    privateKeySecretRef:
      name: letsencrypt-prod
  solvers:
  - dns01:
    azureDNS:
      clientID: 3xx721x5-xx3x-4x10-x39x-xxx31335405
      clientSecretSecretRef:
        name: azuredns-config
        key: client-secret
      subscriptionID: x2xx5467-2502-4b05-x78x-744604x6531x
      tenantID: xxx67057-50x9-4xx4-98x3-xxxx64xxx9x9
      resourceGroupName: masperf
      hostedZoneName: mas4azure.com
      environment: AzurePublicCloud
```

Note: In Maximo Application Suite 8.10, wait for the routes to regenerate and verify the generated routes to check if the certificate is signed by Let's encrypt.

- In the Red Hat OpenShift console, from **Home** > **API Explorer** search for a suite in your namespace.
- In the Red Hat OpenShift console from **Administration** > **CustomResourceDefinition**, select the **Instances** tab for your **Suite CRD**.
- Click your Custom Resource, and in the **Instances** tab, select **YAML** to add **cluster issue** and **domain** parameters in the **spec** section.

```
---
spec:
  certificateIssuer:
    duration: 8760h0m0s
    name: prod-route53-issuer
    renewBefore: 720h0m0s
    domain: <<masinstance_id>>.<<domain>>
```

11. Delete the **finalizer** section from the same Suite YAML to force a reconciliation, and then save the YAML file.

```
finalizers:  
- core.mas.ibm.com/finalizer
```

12. In **Networking** under **Routes** of project *mas-<mas_instance_id>-core*, wait for the **Routes** to regenerate for the namespace.

Note: The **Routes** regeneration takes some time.

The Certificate in the routes is signed by Let's encrypt.

13. Login to the Maximo Application Suite administrator screen and verify the certificate signer.

Administering Maximo Application Suite on Microsoft Azure

Use the **Administration** page to configure and manage the suite level features.

You access the **Administration** page from the **Suite** application in the side navigation menu.

Getting started

The first time that you log in to Maximo Application Suite as the administrator user, no applications are deployed and no users are logged in to the system.

Complete the following tasks to get your Maximo Application Suite environment ready:

- Use the **Catalog** to deploy and activate one or more applications.
- Use the **Users** page to add and manage Maximo Application Suite administrators and application users.
- Use the **Configurations** page to administer suite-level configurations. For example, if you deployed the Maximo Manage application, you can add or [update the database configuration information](#).
- If you installed Maximo Application Suite core and Cloud Pak for Data its service Db2 Warehouse is enabled by default.
- If you install Maximo Manage, and you want to deploy certain applications and add-ons in the Maximo Application Suite, such as Maximo Monitor, you must first install Cloud Pak for Data.

For more information, see [“Installing Cloud Pak for Data on an Azure instance of Maximo Application Suite”](#) on page 268.

Boot node and cluster administration

For operational reasons, you might need command-line access to the boot node, the bastion host, or the cluster nodes that are located in the Maximo Application Suite Microsoft Azure Virtual Network (VNet). For more information, see [Accessing the boot node and Red Hat OpenShift cluster](#) and [boot node and bastion node](#).

Maximo Application Suite authentication

The default authentication uses the Maximo Application Suite authentication. However, you can change the authentication to use LDAP or any other supported authentication method.

For more information, see [“Authentication methods”](#) on page 605.

Accessing the boot node and Red Hat OpenShift cluster

By using Secure Shell (SSH) public key authentication, you can access the boot node and the Red Hat OpenShift cluster nodes.

About this task

The Maximo Application Suite deployment configures the Azure Bastion service to provide the SSH access to the Red Hat OpenShift cluster nodes.

1. Go to resource group where the Red Hat OpenShift Container Platform cluster virtual machines are deployed, for example, masocp-<uniquistr>-rg.

If you use the New OpenShift cluster (IPI) deployment mode, the resource group name format is masocp-<uniquistr>-rg. For the New OpenShift cluster, existing network (UPI) deployment mode, the resource group is the same as the one where existing network infrastructure is available, of which the VNet was provided at the time of deployment.
2. Select the virtual machine to which you want to connect and choose to connect option on azure portal.
3. Choose bastion option from the connect menu and it takes you to the page where you need to fill the details like username and SSH private key.
4. Use the **“core”** username. You can either upload your private key from local machine or you can paste the content of private key on azure portal textbox.
5. Press connect button. You are redirected to a new browser window where you can access the VM command line.

In the Azure cloud, when you start the Maximo Application Suite installation, a boot node is created. By using the required tools and the installation parameters, the boot node completes the installation.

In the Red Hat OpenShift cluster that is created during the installation. The boot node has the public IP address so it can be accessed directly from outside using SSH client. However, the Red Hat OpenShift cluster nodes are not assigned with the public IP address.

The boot node, and private cluster nodes are all [Azure virtual machines](#). If you need to perform maintenance or troubleshooting tasks in a virtual machine, you can connect to it by using [Azure Bastion service](#).

Before you installed Maximo Application Suite, you generated a key pair, which consists of a public key and a private key. You stored the private key locally. When you specified the installation parameters, you selected the public key in the sshKey parameter.

During the installation, a copy of the public key is stored in the boot node, and the private cluster nodes. Because you have the corresponding private key, you can access these instances by using SSH over Bastion service.

Procedure

Accessing the boot node and private cluster nodes

To use SSH access to connect to the boot node and the private cluster nodes, complete the following steps:

1. In your Azure account, go to the Virtual machines service.
2. Click the virtual machine that belongs to the Red Hat OpenShift cluster you have provisioned.
 - You can find the virtual machine based on the resource group name where the virtual machine is available. The resource group name will start with masocp-<unique-identifier>.
 - All the virtual machines are in the same resource group. In case of the IPI deployment mode, the resource group name format is masocp-<uniquistr>-rg. In case of UPI deployment mode, the resource group is the same as the one where existing network infrastructure is available, of which the VNet was provided at the time of deployment.
3. Retrieve the node details.

For bootnode

- a. To retrieve the boot node details, get the public IP address from *Public IP address field* displayed on the virtual machine's *Overview* page.
- b. Perform the SSH from your workstation.

```
ssh azureuser@<bootnode-public-ip-address>
```

For Red Hat OpenShift cluster nodes

- a. Choose *Bastion* option from the *Connect* menu and it will take you to the page where you need to fill the details like Username and SSH private key.
- b. Use username as “**core**” and you can either upload your private key from local machine or you can paste the content of private key on azure portal text-box.
- c. Click the connect button and you will get redirected to a new browser window with a command shell to the Red Hat OpenShift cluster

Boot node and bastion service on Microsoft Azure

When you start a Maximo Application Suite installation, a boot node is created that controls and completes the installation. In the Red Hat OpenShift cluster, a bastion service is configured to allow Secure Shell (SSH) access to cluster nodes.

During a Maximo Application Suite installation on Azure, virtual networks (VNet) are created that contain Azure virtual machines. For example, a VNet is created to contain the Red Hat OpenShift cluster, and virtual machines are created in the cluster to represent its master and worker nodes.

The boot node

In the Azure cloud, after you specify the installation parameters and start the installation, a boot node is created. The installation parameters are passed to the boot node. In addition, all of the required tools to complete the installation, such as Terraform and Docker, are installed on the boot node. By using these tools and parameters, the boot node performs the following tasks to complete the installation:

- Creates the virtual network infrastructure, such as the VNet that contains the Red Hat OpenShift cluster.
- Runs the bootstrap process that creates the Red Hat OpenShift cluster.
- Installs the Maximo Application Suite prerequisites.
- Installs Maximo Application Suite.
- Performs postinstallation validation.
- Stores the installation context and Terraform state files both locally and in the Azure Blob storage container that is associated with your Azure subscription.

Because it is located in its own VNet, the boot node is not part of the Maximo Application Suite Red Hat OpenShift cluster. After the installation is complete, you do not need to use the boot node to access the cluster or interact with Maximo Application Suite. For this reason, the boot node is kept in a shutdown state. However, if required, you can restart it and use it to troubleshoot installation issues.

The bastion service

Maximo Application Suite deployment uses the Azure Bastion service to enable the SSH access to the Red Hat OpenShift cluster nodes. This service is configured during the deployment and can be used to access the Red Hat OpenShift cluster nodes from the Azure portal. See the [“Accessing the boot node and Red Hat OpenShift cluster”](#) on page 274 section for the steps.

Note: No charges apply to Azure virtual machine that are in a shutdown state, such as the boot node. However, charges apply for their attached Premium SSD LRS Azure disk volumes of 30 GB. For more information, see [Azure managed disk pricing](#) in the Azure documentation.

IBM Maximo Application Suite installation with Ansible collection

To automate some of the manual steps that are involved with installing Maximo Application Suite and its components, use the Ansible collection roles that match your installation path or use case.

The IBM Maximo Application Suite development team maintains a public Ansible collection that automates the installation and configuration of Maximo Application Suite and its dependencies. The Maximo Application Suite DevOps Ansible collection includes a number of automated tasks, referred to as roles and playbooks. These automated tasks can be used to streamline processes from having a simple

Red Hat OpenShift cluster, to having Maximo Application Suite including multiple applications and its required dependencies.

It also provides a docker container, which contains all the prerequisites to run the Ansible automation on any local machine.


For example, from Maximo Application Suite 8.9, you can use the role variable `mas_annotations` to install the Maximo Application Suite in production or nonproduction mode. The `mas_annotations` is an optional variable, which accepts a comma-separated list of annotations that need to be added to the Maximo Application Suite CR. To deploy your Maximo Application Suite in nonproduction mode, set this variable to `mas.ibm.com/operationalMode=nonproduction`. For more information, see “[Maximo Application Suite Ansible collection examples](#)” on page 278.

Note: The Maximo Application Suite Ansible collection is developed by the IBM Maximo Application Suite development team. If you need help or have issues, contact [IBM Support](#) or [raise an issue directly in the GitHub repository](#).

For more information about how to help with the development of new roles and collections or improvements to the existing ones, see [Contributing](#).

- Ansible role documentation contains terms and variables, prefixed by W3 or ARTIFACTORY that are intended for internal IBM use only. Ignore these variables if they do not apply to your role, use-case, environment, or scenario.

Mapping documentation tasks to Ansible roles

Each documentation task that contains an Ansible role begins with a  Tip that links to the Ansible role that applies to the task.

For reference, the following documentation tasks map to Ansible roles:

Documentation task	Ansible role
Setting up Red Hat OpenShift Container Storage	ocs
“ IBM operator catalog ” on page 927	ibm_catalogs
Installing IBM Cloud Pak foundational services for IBM Cloud Pak for Data	common_services
Ansible dev-ops role - cert_manager	cert_manager
Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO). For more information, see “ Data Reporter Operator ” on page 7.	uds
IBM Data Reporter Operator	dro
Installing IBM Cloud Pak for Data	cp4d
“ Db2 Warehouse ” on page 21	db2
“ Watson Studio ” on page 20	cp4d_service Specify ws1 service name.
“ Watson Machine Learning ” on page 20	cp4d_service Specify wm1 service name.
Installing MongoDB	mongodb

Documentation task	Ansible role
IBM Cloud Object Storage	cos Specify <code>cos_type</code> as <code>ibm</code> for IBM Cloud Object Storage, or <code>ocs</code> Red Hat OpenShift Container Storage.
Installing Suite License Service	sls
Recommended: Installing Suite Applications from Operator Hub	suite_app_install
“Activating applications” on page 465	suite_app_config
“Installing Red Hat OpenShift Container Platform on IBM Cloud” on page 205	ocp_provision
IBM Cloud Internet Services	suite_dns
“Installing Apache Kafka for IBM Maximo Manage” on page 363 “Installing Apache Kafka for IoT tool” on page 403	kafka
“Installing the NVIDIA operator” on page 394	nvidia_gpu
“Configuring Red Hat OpenShift cluster monitoring” on page 824 “Installing Grafana” on page 826	cluster_monitoring
“Converting IBM Maximo Application Suite from manual deployment to channel subscription” on page 482 This Operator Lifecycle Manager (OLM) conversion script and role is available from Maximo Application Suite 8.10.	convert_to_olm

Related concepts

Operational mode for installation

From Maximo Application Suite 8.9 or later, you can consider installing and deploying IBM Maximo Application Suite in Production or Non-production mode based on your development and testing requirements and to optimize your AppPoint usage.

Maximo Application Suite Ansible collection examples

You can use Ansible collections to install Maximo Application Suite, its applications and prerequisites, and complete other related tasks. Review the examples to learn how to run the ansible roles and playbooks.

For more information, see [Ansible DevOps documentation](#).

You can run Ansible playbooks and roles in two ways:

- Install prerequisite software and ensure it is available on your workstation. For more information, see [MAS DevOps Ansible Collection](#).
- Use a Docker container that contains everything you need. For more information, see [Docker container](#). Ensure that [Docker](#) is installed on your workstation.

The following examples use the Maximo Application Suite DevOps Docker container to run roles and playbooks. You set values for environment variables, then run the corresponding role or playbook.

Running the role to install the IBM Certificate Manager

1. In your local machine with Docker installed, run this command to pull the image and initiate the Docker container:

```
docker run -ti --rm --pull always quay.io/ibmmas/cli
```

Tip: Use the **--pull always** command to pull the newest image.

Note: The **--rm** command ensures that the container is removed after you run the command. However, running the **--rm** command is optional.

For more information, see [ibmmas/cli](#).

1. In the command line inside the container, you can run the following commands:
 - a. Run the **oc login** command. You can access your Red Hat OpenShift cluster by using the **oc** command directly from a terminal in the client machine **oc** was installed to.
 - i) Go to the Red Hat OpenShift web console.
 - ii) Click your login name and select the option: Copy login command.
 - iii) Click View token.
 - iv) Copy the entire command line under the Log in section with this token and paste in the command line, running it from inside the docker container.
 - b. Export the environment variable used by the script. In the example, IBM Certificate Manager is installed.

```
export MAS_CHANNEL=8.8.x
```

- c. Run the **cert_manager role**.

```
ROLE_NAME=cert_manager ansible-playbook ibm.mas_devops.run_role
```

After several minutes, IBM Certificate Manager is installed.

Running a playbook to provision an IBM Cloud Red Hat OpenShift cluster

This example shows how to run a playbook to provision an IBM Cloud Red Hat OpenShift cluster.

Note: This example shows a playbook that just runs one role. The playbook sets the environment variables and the specific role. You might do the same using the steps that are shown in the previous example. But an advantage of playbooks is that you can run more than one role and also set the variables for all them as needed in the same file. In other words, the playbook is usually used to orchestrate an execution of a set of roles.

For example, a playbook can install Maximo Application Suite and perform more configurations after it is installed. For more information, see [suite_install](#).

1. In your local machine with Docker installed, run this command to pull the image and initiate the Docker container:

```
docker run -ti --rm --pull always quay.io/ibmmas/cli
```

2. In the command line inside the Docker container, run the following commands:
 - a. Update **ibmcloud cli**, including plug-ins, to the newest version.

```
curl -sL https://raw.githubusercontent.com/IBM-Cloud/ibm-cloud-developer-tools/master/linux-installer/idt-installer | bash
```

- b. Now, export environment variables necessary to provision your IBM Cloud Red Hat OpenShift cluster.

Note: For more information about variables to customize your IBM Cloud Red Hat OpenShift cluster, see [Role Variables - ROKS](#).

- c. Export your IBM Cloud API Key.

```
export IBM_CLOUD_APIKEY=<your IBM Cloud API Key>
```

Note: To create a key, see [Creating your IBM Cloud API key](#)

- d. Export the name of your cluster.

```
export CLUSTER_NAME=<my own cluster>
```

- e. Run the command to run the playbook:

```
ansible-playbook ibm.mas_devops.ocp_roks_provision.yml
```

3. The IBM Cloud cluster is provisioned after some time.

Customizing a playbook to run existent roles

In this example, customize a playbook to run a set of roles. The playbook installs IBM catalog, IBM Cloud Pak for Data, and IBM User Data Services, which is a Maximo Application Suite prerequisite.

To perform the task, the playbook sets the variables and runs the specific roles to install each one of them.

1. In your local machine with Docker installed, run this command to pull the image and initiate the Docker container:

```
docker run -ti --rm --pull always quay.io/ibmmas/cli
```

2. In the command line inside the container, you can run the following commands:

- a. Run the `oc login` command. You can access your Red Hat OpenShift cluster by using the `oc` command directly from a terminal in the client machine `oc` was installed to.

- i) Go to the Red Hat OpenShift web console.

- ii) Click your login name and select the option: Copy login command.

- iii) Click View token.

- iv) Copy the entire command line under the Log in section with this token and paste in the command line, running it from inside the docker container.

3. Create a `custom_dro_playbook.yml` file with the following content. (You can use an editor such as `vi` and insert the following sample code.)

```
...- hosts: localhost
  any_errors_fatal: true
  vars:
    dro_contact:
      email: "{{ lookup('env', 'DRO_CONTACT_EMAIL') }}"
      first_name: "{{ lookup('env', 'DRO_CONTACT_FIRSTNAME') }}"
      last_name: "{{ lookup('env', 'DRO_CONTACT_LASTNAME') }}"
  roles:
    # 1. Install DRO
    - ibm.mas_devops.ibm_catalogs
    - ibm.mas_devops.common_services
    - ibm.mas_devops.dro
```


4. Export the required environment variables to install Data Reporter Operator .

```
export DRO_CONTACT_EMAIL=john.doe@test.com
```

```
export DRO_CONTACT_FIRSTNAME=John
```

```
export DRO_CONTACT_LASTNAME=Doe
```

5. Run the playbook with the following command:

```
ansible-playbook custom_dro_playbook.yaml
```

6. Data Reporter Operator was installed along with its prerequisites.

Installing on nonproduction environment

This example describes how to install the IBM Maximo Application Suite in production or nonproduction environment by using the *mas_annotations* role. The *mas_annotations* is an optional variable, which accepts a comma separated list of annotations that need to be added to the Maximo Application Suite CR.

1. In your local machine with Docker installed, run a command to pull the image and initiate the Docker container:

```
docker run -ti --rm --pull always quay.io/ibmmas/cli
```

2. In the command line, inside the Docker container, run the following command to deploy the Maximo Application Suite in nonproduction mode for development and testing deployment.

Set environment variable MAS_ANNOTATIONS for nonproduction mode:

```
export MAS_ANNOTATIONS=mas.ibm.com/operationalMode=nonproduction
```

3. Run the following command to install the Maximo Application Suite:

```
...
- hosts: localhost
  any_errors_fatal: true
  vars:
    mas_instance_id: "inst1"
    mas_config_dir: "/home/david/masconfig"
    mas_entitlement_key: "{{ lookup('env', 'IBM_ENTITLEMENT_KEY') }}"
    mas_annotations: "mas.ibm.com/operationalMode=nonproduction"

  roles:
    - ibm.mas_devops.suite_install
    - ibm.mas_devops.suite_config
    - ibm.mas_devops.suite_verify
```

Related concepts

[Operational mode for installation](#)

From Maximo Application Suite 8.9 or later, you can consider installing and deploying IBM Maximo Application Suite in Production or Non-production mode based on your development and testing requirements and to optimize your AppPoint usage.

Setting up IBM Maximo Application Suite

After you install IBM Maximo Application Suite, the setup program guides you through the initial configuration.

Before you begin

1. Complete the installation.

Obtain the link to the Maximo Application Suite setup program and the login credentials that you need to complete the setup process.

For more information about obtaining the login credentials, see [how to locate the default username and password](#).

2. Enable login for Maximo Application Suite self-signed certificates.

If you are using self-signed certificates in a development or test environment, you must manually enable login by using either of the following methods.

- Download the certificates from the cluster and add them to your local certificate manager.
- In your browser, go to the Maximo Application Suite API URL `https://api.<mas_domain>/` and accept the certificate security risks. After you accept the risks, an AIUC01999E error is displayed. This message is expected. You can now continue with the setup process.

If the Maximo Application Suite dashboard does not load after you login for the first time and instead see a spinning wheel, see [how to troubleshoot the issue](#).

About this task

The Maximo Application Suite setup configurations are set at the System scope. For more information about configuration scopes, see [Configure Maximo Application Suite](#).

Procedure

1. Log in to the Maximo Application Suite setup program by using the superuser credentials that were created during installation.

`https://admin.<mas_domain>/initialsetup`

Important: Treat the superuser account the same way that you treat the root account on your servers. Use it only for the initial setup. As part of the setup, you create a default administrator user account that has access to the Maximo Application Suite administrative interface. Use this administrative account to add and manage users, deploy applications, and more.

For more information about obtaining the superuser credentials, see [how to locate the default username and password](#).

2. Configure MongoDB.

MongoDB is used as the data dictionary for Maximo Application Suite and its applications. It is also used as the default user registry.

Specify the following MongoDB information:

Hostname and port

You can configure one or more MongoDB hostname and port combinations.

Authentication mechanism

Specify the mechanism that is used to authenticate Maximo Application Suite when it connects to MongoDB. Select the closest match to the mechanism that is configured for your MongoDB cluster. For example, if your cluster uses the SCRAM-SHA-256 mechanism, select **DEFAULT (SCRAM)**.

To authenticate by using LDAP, specify **PLAIN** as the authentication mechanism.

Auth db

Provide the name of the authentication database. If you are authenticating with LDAP, the value must be `$external`.

MongoDB login credentials

At a minimum, the MongoDB administrator needs table creation privileges.

Note: The MongoDB verification might take up to a minute. The configuration cannot be modified after the MongoDB verification is complete. MongoDB is a prerequisite for Maximo Application Suite. Changing the configuration requires careful coordination and possible data migration to avoid service outages. System administrators can change the configuration in the Red Hat OpenShift console. For assistance with changing the MongoDB configuration, contact your IBM representative.

For more information, see [Installing MongoDB](#).

3. Upload a CA certificate.

If the service uses the transport layer security (TLS) communication protocol and is not secured with a certificate that is issued by a well-known certificate authority (CA), then provide the certificate of the CA that issued the service's certificate. Because the CA might use intermediate CAs, you can provide more than one certificate.

For each certificate that you provide, the following details are displayed:

- The name of the certificate issuer.
- The name of the subject, such as the organization, that the certificate is issued to.
- The start and end dates of the certificate's validity period. If the validity of any certificate that you provide expires soon, a warning message appears.

You can automatically retrieve or manually add certificates.

Important: If your MongoDB cluster uses self-signed CA certificates that you must retrieve or add a certificate.

- Automatically retrieving certificates

In the certificates section, click **Retrieve**. If the connection credentials that you specify are correct, all CA certificates that are configured on the server are automatically retrieved and displayed.

These certificates are not validated. Verify that only the correct certificates are retrieved and remove any unexpected certificates.

After you retrieve the certificates, you can manually add more certificates.

- Manually adding certificates

In the certificates section, click **Add manually** and specify the following values for each certificate that you want to add:

Alias

An alphanumeric identifier that is in the range 3 - 50 characters long.

Certificate content

The content of a certificate file in either the X.509 or PEM formats.

For more information, see [“Configuring certificate authority certificates”](#) on page 590.

4. Configure a Simple Mail Transfer Protocol (SMTP) server connection to enable email notifications for system events, such as new user welcome emails and password reset communication. For more information, see [“Setting up email notifications”](#) on page 634.

5. Configure analytics data.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator ”](#) on page 7.

- If you are using Maximo Application Suite 8.11, 8.10 or earlier versions, you must migrate your User Data Services to Data Reporter Operator . For more information, see [“Migrating Maximo Application Suite from User Data Services to Data Reporter Operator ”](#) on page 8.
- If you are using Maximo Application Suite 9.0, 8.11.7, 8.10.10 or later versions, configure IBM Data Reporter Operator.

The IBM Data Reporter Operator accepts events and transforms them into reports that are submitted to the Data Service of the IBM Metrics Operator.

- a. Enter the following information to configure Data Reporter Operator for Maximo Application Suite:
 - **URL** - This URL is the DRO URL endpoint. To find it, go to your Red Hat OpenShift console, switch to `ibm-common-services` project, then **Networking** > **Routes**. Copy the URL displayed under the Location column for the `dro-endpoint` route.

For example, `https://dro-endpoint-ibm-common-services.<your-cluster-domain>`
 - **API Key** - This API key is the DRO API Key credential. To find it, go to you Red Hat OpenShift console, switch to `ibm-common-services` project, then **Workloads** > **Secrets** > **Search and select the secret named dro-api-key**. Under the **Data** section, copy the `apikey` value.

For example, `k2wnQY...`
 - **Email** - Enter a contact email address to use for DRO communication. The email address does not have to match an existing Maximo Application Suite user.
 - **Given Name** - Enter the given name of the owner of the provided contact email address that is used for DRO communication.
 - **Surname** - Enter the surname of the owner of the provided contact email address that is used for User Data Services communication.
 - **Certificates** - Enter the chain of SSL certificates for your DRO. To retrieve the certificates, you can click the **Retrieve button (under Certificates section)** while configuring DRO into Maximo Application Suite. The DRO certificates to configure in Maximo Application Suite will vary according to the cloud service provider's cluster that is hosting your DRO installation.
- b. Click **Add** to add the **intermediate of the certificate chain**.
- c. Enter an **alias**. **Example:** `drocertpart1`.
- d. Enter the **Certificate content**. Include the **Let's Encrypt R3 intermediate certificate**, issued to **US, Let's Encrypt, R3**. For more information, see [certificate content](#). **Example:**

```
-----BEGIN CERTIFICATE-----
MIIF5jCCBM6gAwIBAgISA0Y...
-----END CERTIFICATE-----
```

- e. Click **Confirm**. The first part of this certificate should include valid dates and look like the following example:

```
Issued to: US, Let's Encrypt, R3
Issued by: US, Internet Security Research Group, ISRG Root X1
Valid from: Thu Aug 01 2024
Valid to: Mon Sep 15 2025
```

This is the intermediate certificate which is required for the SSL connection to DRO endpoint.

- f. Click **Add** to add the **root of the certificate chain**.
- g. Enter an **alias**. **Example:** `drocertpart2`.
- h. Enter the **Certificate content**. Include the **ISRG Root X1 self-signed certificate**. For more information, see [certificate content](#). **Example:**

```
-----BEGIN CERTIFICATE-----
MIIFazCCA10gAw...
-----END CERTIFICATE-----
```

- i. Click **Confirm**. The **second part of this certificate** should have valid dates and look like the following example:

```
Issued to: US, Internet Security Research Group, ISRG Root X1
Issued by: US, Internet Security Research Group, ISRG Root X1
```

Valid from: Thu Jun 04 2015
Valid to: Mon Jun 04 2035

This is the root certificate which is required for the SSL connection to DRO endpoint.

- j. **Save** the DRO configuration.
- k. Now, wait for the DRO configuration to reconcile, this process might take up to 10 minutes. The configuration will be successfully completed when the configuration status is set to Ready.
Example:

Configuration Ready - DRO configuration was successfully verified

6. Configure the Suite License Service.

The Suite License Service (SLS) stores and manages the Maximo Application Suite license.

Each Maximo Application Suite instance can be connected to a unique SLS instance. Two or more Maximo Application Suite instances can also share an SLS and the corresponding license file.

Enter the following SLS information to configure Maximo Application Suite:

- URL - The URL for the SLS server.
- Registration key - Enter the SLS registration key.

Depending on your environment, the SLS configuration might take 10 minutes or more to complete.

7. Optional: Upload your license key file.

If the IBM Suite License Service that you configured for use with Maximo Application Suite includes a valid license file, you do not need to upload a license file. You can continue with the next configuration step.

To activate Maximo Application Suite, you must provide your license key from the [IBM License Key Center](#). The login information is provided in the license Key Center welcome letter. For more help on licensing, see the [IBM Support - Licensing](#) page.

- a) Log in to the license Key Center.
- b) Select your company name.
- c) Select the **IBM AppPoints** product line.
- d) Select the IBM Maximo Application Suite... license key name.
- e) Select the product or sales order for which to create the license key.
- f) Enter the number of keys to generate. These correspond to the AppPoints that are allocated to the license key.
- g) Provide the Maximo Application Suite license server parameters.

Use the parameters that are displayed in the **Advanced settings > license key** section of the Maximo Application Suite setup program, or provide the following parameters:

- For Configuration, specify a Single License Server.
- For Host ID type, specify the Ethernet address.
- For Host ID, specify the host ID that was generated when you installed the Suite License Service (SLS). To display this ID, connect to your Red Hat OpenShift cluster and run the following command:

```
oc -n <sls_project_namespace> get licenseservice sls
```

For example, if the namespace of the SLS project is mas-sls-dev5, run the following command:

```
oc -n mas-sls-dev5 get licenseservice sls
```

In the command output, the host ID is displayed in the LICENSEID column.

- For Hostname, specify a hostname of your choice, for example: sls-mas
- For Port, specify 27000.

- h) Download the key and then upload it to the Maximo Application Suite setup program.

8. Create the workspace.

The Maximo Application Suite workspace is a unique collection of configuration settings for your instance of Maximo Application Suite. Enter the following information to create your Maximo Application Suite workspace:

- **Workspace ID**

The workspace ID forms part of the Maximo Application Suite URL, for example:

```
https://<workspace_id>.home.<mas_domain>
```

Note: The workspace ID must be 3 - 12 characters in length, and can contain only lowercase letters and numbers. The first character must be a letter.

- **Workspace display name**

The display name is shown in your Maximo Application Suite user interface.

9. Review the setup configuration.

Your Maximo Application Suite setup is now complete. Verify that all configuration settings are done and then click **Finish** to complete the setup.

What to do next

After the Maximo Application Suite setup is complete, you can start to use your environment by going to the Maximo Application Suite administration or the Maximo Application Suite navigator page:

```
https://admin.<mas_domain>  
https://<workspace_id>.home.<mas_domain>
```

As the Maximo Application Suite superuser, you can now continue configuring your environment to suite your enterprise needs:

- [Configure authentication](#)

Maximo Application Suite supports local user authentication by MongoDB and authentication by using LDAP or SAML.

- [Configure LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

- [Create administrator user accounts](#)

The initial superuser account is used to complete the Maximo Application Suite setup. You can add application administrator users or system administrator users for day-to-day administrative tasks.

- [Getting started](#)

With the setup completed, your users can log in and start to use Maximo Application Suite.

Related concepts

[Simple Mail Transfer Protocol](#)

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

Related tasks

[Migrating Maximo Application Suite from User Data Services to Data Reporter Operator](#)

As an IBM Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance. New and existing Maximo Application Suite users can install or migrate to DRO by using the IBM Maximo Application Suite command line interface (CLI), ansible role, or manually.

Uninstalling

The method that you use to uninstall IBM Maximo Application Suite depends on your environment.

Customer-managed **Uninstalling Maximo Application Suite**

Starting in 8.10, you can uninstall the IBM Maximo Application Suite with the **mas uninstall** command by using the CLI command utility. By uninstalling Maximo Application Suite, you remove the core application and all deployed applications from your environment.

About this task

The uninstallation process sequentially removes the Maximo Application Suite applications, industry solutions, add-ons, and tools that were installed.

- After the industry solutions, add-ons, and tools are removed, all config maps and secrets are also removed.
- If you use Strimzi as the Kafka operator, the Kafka topics can be removed if needed.
- All entries that are related to the instance can be removed from the MongoDB instance if needed. During this phase, double confirmation is required for each database removal to ensure no unexpected data loss. A confirmation string is created at the start of this phase to provide the confirmation for each removal.
- The default ClusterIssuer resource is also removed if one was not provided on the installation. If you created and provided your own ClusterIssuer resource to the installation script, it is not removed.
- Finally, all projects and namespaces that are related to this instance are removed.

Even though multiple instances of Maximo Application Suite might be installed on a single cluster, the uninstaller removes just one instance at a time. Run the uninstall multiple times to remove each instance.

Note: Supporting components, such as MongoDB and IBM Cloud Pak for Data, are not removed by the uninstaller and must be removed separately if needed.

You can also automate the Maximo Application Suite uninstallation. First run the ansible role `suite_app_uninstall` to remove the applications and then use the ansible `uninstall core` playbook to uninstall Maximo Application Suite core platform and its dependencies from your cluster. For more information, see [suite_app_uninstall role](#) and [uninstall core playbook](#).

Procedure

Run the **mas uninstall** [options] command either in interactive or non interactive mode.

1. Run the command in interactive mode.

```
mas uninstall -i|--id MAS_INSTANCE_ID
```

where,

-i refers to interactive mode

--id MAS_INSTANCE_ID refers to the Maximo Application Suite that is to be uninstalled

Other options include **--no-confirm** to launch the uninstall without prompting for confirmation and **-h|--help** to display a help message

Alternatively, you can run the command in non interactive mode.

```
mas uninstall -i MAS_INSTANCE_ID -\--no-confirm
```

2. If you are not connected to a Red Hat OpenShift cluster, you are prompted to provide the server URL and token, and whether to verify the server certificate or not.

If you are already connected to a cluster you can opt to change to another cluster.

Deleting the Maximo Application Suite stack on Amazon Web Services

Download and run a script that deletes the IBM Maximo Application Suite stack, virtual infrastructure, VPCs, and EC2 instances from your Amazon Web Services (AWS) account.

When you install an instance of the Maximo Application Suite, several virtual infrastructure resources are created in your AWS account, such as virtual private clouds (VPC), Amazon EC2 instances, and a CloudFormation stack. To uninstall the Maximo Application Suite, you download and run a script that deletes these resources. You can use this script to uninstall the Maximo Application Suite regardless of whether the installation succeeded or failed.



Attention: Do not delete the CloudFormation stack before running the script to uninstall Maximo Application Suite. The script deletes the stack. If the stack is deleted manually, the uninstallation might fail.

You run the script on your local machine or on a server that is not located in a VPC that the installation process created, such as the VPC that contains the Red Hat OpenShift cluster.

Note: Do not run the script in any of the EC2 instances that the Maximo Application Suite installation process created, such as the Bootnode, the bastion host, or any of the cluster nodes. The script deletes these EC2 instances. If you run it in any of them, it fails.

Before you begin

- On the machine where you want to run the script, ensure that the following CLI packages are installed:
 - Version 4.0 or a later version of [GNU bash](#)
 - [jq](#)
 - [AWS CLI](#)

Ensure that the AWS CLI package is configured for authentication with your AWS account. For more information, see [Configuring the AWS CLI](#) in the AWS documentation.

- If Amazon DocumentDB or Amazon MSK are configured in the Maximo Application Suite stack, you must delete the Amazon DocumentDB or Amazon MSK instance.

For more information, see [Deleting an Amazon DocumentDB cluster](#) or [Deleting an Amazon MSK cluster](#).

Procedure

1. In a browser window, [open the script](#), right-click the page and save the script to your local machine by using the name `cleanup-mas-deployment.sh`.
 - a. If you do not want to run the script locally, use SCP or any file transfer tool to copy the script to the machine where you want to run it.
 - b. On the machine where you want to run the script, in a command shell, log in to the AWS service by running the following command:

```
aws configure
```

You are prompted for your identity and access management (IAM) user credentials. Enter the credentials for an IAM user that has the permissions to run the script, such as the IAM user that installed the Maximo Application Suite. For more information, see [Configuring the installation permissions](#).

- c. Make the script executable by entering the following command:

```
chmod +x cleanup-mas-deployment.sh
```

- d. View the script's usage information by running the following command:

```
./cleanup-mas-deployment.sh -h
```

2. Run the script.

- You must specify the region code of the region where the Maximo Application Suite was installed by using the `-r` option, for example: `-r ap-northeast-3`
- To delete the virtual resources by using the CloudFormation stack name, use the `-s` option.

For example, if the CloudFormation stack name is `sp-manage-12` and the region code is `ap-northeast-3`, run the following command:

```
./cleanup-mas-deployment.sh -s sp-manage-12 -r ap-northeast-3
```

3. Verify that the script completed successfully.

If the script is successful, output that is similar to the following text is displayed:

```
$ ./cleanup-mas-deployment.sh -s sp-manage-12 -r ap-northeast-3
Stack name: sp-manage-12
Unique string:
Region: ap-northeast-3
Supported region provided
Deleting by stack-name sp-manage-12
Execution started at Mon Mar 7 22:46:48 IST 2022
MAS instance unique string: nove9h
Checking for EC2 instances
...
EC2 instances found for this MAS instance
...
Terminate request submitted
Waiting for instances to be terminated
Deleted EC2 instances
Checking for volumes
...
Found volumes for this MAS instance
...
Checking for VPC
VPC_ID = vpc-0851c8fc0523cac86
Found VPC with Id vpc-0851c8fc0523cac86 for this MAS instance, it will be deleted at the end
Checking for NAT gateways
...
Found NAT gateways for this MAS instance
...
Checking for EIPs
...
Checking for load balancers
...
Checking for v2 load balancers
...
Checking for network interfaces
...
Checking for internet gateways
...
Checking for subnets
...
Checking for routing tables
...
Checking for network ACLs
...
Checking for security groups
...
Checking for S3 buckets
...
Checking for IAM users
...
Checking for IAM instance profiles
...
Checking for IAM policies
...
Checking for IAM roles
...
Checking for private hosted zones
...
Checking for CloudWatch log groups
...
Checking for CloudFormation stack
...
Execution completed at Mon Mar 7 23:02:47 IST 2022
```

4. In the AWS CloudFormation console, verify that the stack that you created when you installed the Maximo Application Suite is deleted.

5. In the AWS VPC console, verify that the virtual infrastructure that was created when the Maximo Application Suite was installed is deleted.
6. Verify that no VPCs exist that contain `<unique-string>` in the VPC name, for example: `masocp-
<unique-string>-vpc`
7. Verify that no EC2 instances exist that contain `<unique-string>` in the EC2 instance name.
For more information on `<unique-string>` and other identifiers that are used in this documentation, see [Unique identifiers](#).

Related information

[Troubleshooting installation problems for Amazon Web Services](#)

An unsuccessful Maximo Application Suite installation has many possible causes, such as missing or invalid installation parameters, Bootnode creation failures, or cluster creation problems.

Uninstalling Maximo Application Suite on Microsoft Azure

To uninstall IBM Maximo Application Suite on Microsoft Azure, all the infrastructure resources that are created during the deployment must be deleted. Additionally, if a paid product is installed, unsubscribe the product from your Microsoft Azure account.

When you install an instance of the Maximo Application Suite, several virtual infrastructure resources are created in your Microsoft Azure account. These resources include virtual network (VNet), Microsoft Azure virtual machines, storage, and other resources. To uninstall the Maximo Application Suite, you download and run a script that deletes these resources. You can use this script to uninstall the Maximo Application Suite regardless of whether the installation succeeded or failed.

You run the script on your local computer or on a server that is not located in a VNet that the installation process created, such as the VNet that contains the boot node or the Red Hat OpenShift cluster.

Note: Do not run the script in any of the virtual machines that the Maximo Application Suite installation process created, such as the boot node, or any of the cluster nodes. The script deletes these virtual machines. If you run it in any of them, it fails.

Before you begin

On the computer or server where you want to run the script, ensure that the following CLI packages are installed:

- [GNU bash](#) version 4.0 or later
- [jq](#)
- [Microsoft Azure CLI](#)

Ensure that this package is configured for authentication with your Microsoft Azure account. For more information, see [Configuring the Azure CLI](#) in the Microsoft Azure documentation.

Procedure

Uninstall Maximo Application Suite depending on your product type.

1. Delete the infrastructure resources.

This step is applicable to BYOL and paid products.

- a) In a browser window, [open the script](#), right-click the page and save the script to your local machine by using the name `cleanup-mas-deployment.sh`.
- b) If you do not want to run the script locally, use SmartCloud Provisioning, or any file transfer tool to copy the script to the machine where you want to run it.
- c) On the machine where you want to run the script, in a command shell, log in to the Microsoft Azure service by running the following command.

```
az login
```

A browser opens where you can log in with the Microsoft Azure credentials. After login, you can close the browser window and continue the next steps from the command shell.

d) Run the following command:

```
chmod +x cleanup-mas-deployment.sh
```

e) View the script's usage information by running the following command:

```
./cleanup-mas-deployment.sh -h
```

f) Specify your preferred options and run the script.

You must specify either the boot node resource group where the Maximo Application Suite was installed by using the `-r` option, for example: `-r mas-ocp-deploy-rg` or the unique string that is associated with the Red Hat OpenShift cluster, for example: `-u nove9h`

To delete the virtual resources by using the resource group name, use the `-r` option. For example, if the boot node resource group name is `mas-ocp-deploy-rg`, run the following command:

```
./cleanup-mas-deployment.sh -r mas-ocp-deploy-rg
```

g) To delete the virtual resources by using the installation identifier that is `<unique-string>`, use the `-u` option.

For example, if `<unique-string>` is `nove9h`, run the following command:

```
./cleanup-mas-deployment.sh -u nove9h
```

The script takes 15 - 20 minutes to delete the cluster and the virtual network infrastructure.

h) Verify that the script completed successfully. If the script is successful, output that is similar to the following text is displayed:

```
$ ./cleanup-mas-deployment.sh -r mas-ocp-deploy-rg
==== Execution started at Tue May 10 13:17:10 EDT 2022 ====
Script Inputs:
  Bootnode resource group = test-mas-1
  Unique string =
SUB_ID: b2ca5467-2502-4b05-b78e-744604c6531d
Trying to delete OCP cluster resource group
Deleting by 'bootnode-resource-group' test-mas-1
UNIQ_STR: vt57ov
...
==== Execution completed at Tue May 10 13:26:53 EDT 2022 ====
```

i) On the Microsoft Azure portal in the Resource group, verify that both the resource groups (boot node resource group and Red Hat OpenShift cluster resource group) are deleted.

2. Unsubscribe the paid product.

This step is applicable to the paid product only.

For more information about canceling the Maximo Application Suite subscription, see <https://learn.microsoft.com/en-us/azure/cost-management-billing/manage/cancel-azure-subscription>.

Customer-managed **Deploying applications, add-ons and industry solutions**

By using the **Applications** tab in the Maximo Application Suite catalog, administrators who have system configuration privileges can add and remove applications.

Related concepts

Maximo Application Suite Industry solutions

On the Industry solutions tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove industry solutions.

Applications that are included with your installed Maximo Application Suite version can be deployed into your environment. Customize your environment by deploying the applications that you need.

About this task

The deployment process consists of two steps:

1. Deploy the application.

Tip: This task can also be done by using the following Ansible role: [suite_app_install](#). For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

2. Activate the application.

Tip: This task can also be done by using the following Ansible role: [suite_app_config](#). For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

As part of the application deployment process, you configure your application by selecting an upgrade method, configuring any prerequisites, and configuring advanced settings.

Depending on the application, prerequisites and advanced settings might be included in the deployment step, the activation step, or both.

Application dependencies are prerequisites of the application, such as databases or other applications that interact with or extend the application.

Important: If the prerequisite is another application, the deployment process for that application opens in a new tab. Complete the deployment and start the activation process for the prerequisite application before you activate the application that you are deploying.

Advanced application settings might include databases, languages, and other customizations. To access the advanced settings, set **Advanced settings** to **On**.

By default, the system manages advanced settings, and the suggested default values are set. If the system-managed default values change in the release of an updated version of the application, values are updated when you update your application.

If your environment requires different values for these settings, set **System managed** to off and update the values.

Each deployed application is available from the Suite navigator and is also available at specific URLs.

For more information, see [“Maximo Application Suite application URLs”](#) on page 77

After you activate the application, you must grant users access to it.

Important: First log in to Maximo Real Estate and Facilities with the mandatory initial FACILITIESADMIN user and set up other users. If you try to log in to Maximo Real Estate and Facilities with any other user before the FACILITIESADMIN user, you get a blank screen and an error message that says the user is invalid, see [“Administering Maximo Real Estate and Facilities users”](#) on page 385.

Deploying IBM Maximo Collaborate

By deploying and activating the Maximo Collaborate application, you make it available for use in Maximo Application Suite.

Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate. If Maximo Assist is deployed in Maximo Application Suite 9.0 or earlier and you upgrade to Maximo Application Suite 9.1, the name is automatically changed in the user interface.

Note: In the Maximo Application Suite documentation, Maximo Assist is now referred to as Maximo Collaborate.

Before you begin



Attention:

Starting in Maximo Application Suite 9.0, Watson Discovery, which is used to support the query and diagnose functions, is no longer available as a dependency in Maximo Assist. If Maximo Assist is already deployed and activated with Watson Discovery, and you are upgrading to Maximo Application Suite 9.0, before you can complete the upgrade, you must contact IBM Support to help with the manual removal of Watson Discovery.

Before you deploy Maximo Collaborate, ensure that the following dependencies are prepared and available for use:

Watson Discovery

This configuration applicable for Maximo Collaborate deployed in Maximo Application Suite 8.11 and earlier.

- [Install Cloud Pak for Data](#)
- [Install Watson Discovery on Cloud Pak for Data](#).

S3-Compatible object storage

Choose a S3-compatible object storage solution. For more information about S3-compatible object storage, see [“Selecting S3-compatible object storage for IBM Maximo Collaborate”](#) on page 295

About this task



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription”](#) on page 482.

Complete the per-deployment and deployment steps before you activate the application and then complete the post-activation steps. During deployment, select your application update method. To later change from channel subscription versioning to manual versioning, you must first delete and then redeploy the application. After you activate the application, you must grant users [access](#) to it.

Starting in 8.2, CouchDB is embedded and automatically deployed.

Configuration parameters

The following parameters are configurable:

- URL, for example, `https://<mycouchdb><mydomain.com>`
- Username
- Password

Required by

- Maximo Collaborate at the Application scope.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Applications** and then click **Collaborate**.

On the Collaborate page, verify the information. Click **Continue**.

Note: If the available AppPoints are insufficient to deploy this application, you can still click **Continue** to complete the application configuration. The application automatically is deployed when the required number of AppPoints are available.

For more information, see [Upgrade methods](#).

2. Select your application update method.

- a) Select an update method.
To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.
To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.
- b) Subscribe to a channel by selecting a version from the list.
For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
- c) Click **Subscribe to channel <version>**
3. On the **Object storage** tile, in the **Dependencies** section, click **Configure** and enter the object storage information for your object storage.
 - a) Enter the object storage S3 route URL.
For example, `https://ceph-s3-rook-ceph.apps.test.os.yourcompany.com`. Or, for AWS S3, `https://s3.{region}.amazonaws.com`
 - b) Enter the username, which is the S3-compatible object storage access key ID.
 - c) Enter the password, which is the S3-compatible object storage secret key.
 - d) Optional: Upload a certificate. For more information, see [Upload a certificate](#).
4. Configure Apache CouchDB and Redis.
 - a) Click **Show advanced settings**.
 - b) Configure CouchDB by specifying the parameters for storage size, replica, and storage class.
Specify the Apache CouchDB Storage Class that can be used from the Red Hat OpenShift cluster. For example, use the existing `ocs-storagecluster-cephfs` or `portworx-couchdb-sc` storage class in the cluster.
 - c) Configure the Redis Server by specifying the parameters for storage size and class. Enter an expected disk size of at least 30 Gi or more. Next, enter an existing storage class from the Red Hat OpenShift cluster.
The storage class can be used to dynamically provision a persistent volume with access mode RWO. For example, set `ocs-storagecluster-cephfs` as the storage class name for the embedded Redis server for Maximo Collaborate.
5. Click **Deploy** and then **Begin deployment**. The estimated deployment time is an estimate of the time that it takes to configure and deploy the application. The time includes both processing and configuration. You can track the deployment process on the details page.
6. On the **Applications** page you can monitor the deployment status for Collaborate. Deployment is complete when the **Application** card displays the Collaborate is ready message and the **Activate** button is displayed.

You can also check the Collaborate deployment status by using the Red Hat OpenShift client command line to log in to your Red Hat OpenShift Container Platform cluster.

```
oc login <OCP_cluster>

oc get collaborateapp -n {{Collaborate_project}}
NAME          VERSION    STATUS    AGE
masdev        9.1.0     Ready    73d
```

The **Collaborate Operator** uses the **Ready** condition reason code of the custom resource to indicate progress. The following reason codes are used:

Ready

The Maximo Collaborate `collaborateapp`, `collaborateworkspace`, `collaboratebackup`, or `collaboraterestore` Operator custom resource is up-to-date and ready to use.

InvalidConfiguration

One or more configurations are not available or incorrect. Refer to the fail message for details.

CollaborateTooOld

Migration installation of the existing Maximo Collaborate application is not supported.

CollaborateorCouchDBNotReady

The Maximo Collaborate application is not ready or the embedded CouchDB is not available for backup or restore.

CertificateNotReady

The secret for the certificate is not yet available.

Note:

Horizontal Pod Autoscale (HPA) is enabled by default during deployment and is based on the resource utilization for Collaborate pods. Maximo Collaborate automatically scales in or out the internal components.

If you want to manually scale in or out your pod numbers for each deployment, you need to disable Horizontal Pod Autoscale by setting **settings.common.podautoscale** to `false` in the CollaborateApp custom resource.

What to do next

Activate Maximo Collaborate. For more information, see [Activating Maximo Collaborate](#).

Related concepts

[Interactive Connectivity Establishment \(ICE\) server](#)

[IBM Maximo Collaborate](#)

Related tasks

[Activating IBM Maximo Collaborate](#)

[Updating Maximo Collaborate](#)

[Watson Discovery](#)

Related information

[Getting started](#)

[Maximo Assist](#)

Selecting S3-compatible object storage for IBM Maximo Collaborate

IBM Maximo Collaborate supports S3-compatible object storage. You can specify your S3-compatible object storage solution to ensure that it is set up for Maximo Collaborate deployment.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

About this task

After you select a S3-compatible object storage solution, it is automatically set up for the Maximo Collaborate deployment.

Note: If you choose to use Red Hat OpenShift Container Storage, ensure that its independent entitlement is separate from Maximo Application Suite. For more information, see [Red Hat OpenShift Container Storage 4.5 installation with Object Storage Gateway](#).

Procedure

1. Select an object storage solution.

You can use any of the following object storage:

- IBM Cloud Object Storage
- AWS object storage
- Ceph® Object Storage on Red Hat OpenShift Container Platform 4 with the Rook-Ceph Operator or Red Hat OpenShift container storage operator.

2. After you install the object storage, get the access key and secret access key.

Storage options	Procedure
AWS object storage	Create an AWS S3 account before deployment: <ol style="list-style-type: none"> Log in to the AWS console, for example at <code>https://us-east-1.console.aws.amazon.com/</code>. Click Service > IAM > User. On the User page, click the Security credentials tab. Click Create access key to generate an S3 access key ID and secret access key.
Ceph Object Storage	<ol style="list-style-type: none"> Get the access key and secret access key of the S3 user account that you previously created. Locate the Ceph Object Storage endpoint. In Red Hat OpenShift, set the <code>tls.termination</code> to <code>edge</code> to create a Red Hat OpenShift route for the <code>rook-ceph-rgw</code> service.
IBM Cloud Object Storage	<ol style="list-style-type: none"> In IBM Cloud, locate the endpoint URL for your Cloud Object Storage instance. <p>Note: If you are Cloud Object Storage in IBM Cloud, locate the Cloud Object Storage instance. From the side navigation menu, click Resource list and select the object storage instance created. Then, click Endpoints, and choose the endpoint for the region that you want to use, for example, <code>s3.us.Cloud-object-storage.appdomain.cloud</code>.</p> Locate the username. Locate the password. Locate the certificates. Enter the chain of SSL certificates for your Cloud Object Storage instance. To retrieve the certificates, in the Certificates section, click Retrieve during the configuration of object storage for Maximo Application Suite.

What to do next

After you select your S3-compatible object storage solution, deploy Maximo Collaborate. For more information, see [“Deploying IBM Maximo Collaborate” on page 292](#).

Customer-managed **Deploying IBM Maximo Health**

Improve the reliability of your assets by proactively monitoring and managing asset health by using Maximo Health.

Before you begin

To prepare for deployment, you must complete several tasks, such as configuring the database for Maximo Health. If you are deploying in multiple languages, review the languages that are available.

For more information, see [Planning and preparing to deploy Maximo Health](#).

About this task

These instructions cover the deployment of Maximo Health as a stand-alone Suite application. If Maximo Manage is deployed, follow the steps in the Maximo Manage documentation to deploy Maximo Health as a Maximo Manage component. For more information about deploying Maximo Health as a Maximo Manage component, see [“Deploying IBM Maximo Manage” on page 297](#).

Note: You cannot install the Maximo Health stand-alone application and the Maximo Health extension for Maximo Manage in the same Maximo Application Suite instance.

Note: Maximo Health stand-alone will not be available in Maximo Application Suite 9.1. If you want to upgrade Maximo Health 9.0 stand-alone to 9.1, you manually upgrade to Maximo Manage with Health. For more information see [Upgrading Maximo Health 9.0 stand-alone to Maximo Manage 9.1 with Health](#).

Procedure

As a system administrator, you deploy and activate Maximo Health in Maximo Application Suite by selecting the version, connecting to the configured database, and then applying the Java database connectivity (JDBC) configurations. You can specify the configuration settings, such as language preferences, customization archive, and external file storage. For more information, see [Deploying and activating Maximo Health](#).

What to do next

As an application administrator, if Maximo Health is not deployed as part of Maximo Manage, you must load data into the application. For more information, see [Loading data](#).

For more information about setting up Maximo Health, see [Getting started as an application administrator](#).

Related concepts

[Application database](#)

To deploy Maximo Health, Maximo Real Estate and Facilities, or Maximo Manage, a database instance must be configured and running. The applications support Db2, Db2 Warehouse, Microsoft® SQL Server, or Oracle Database. If Maximo Health is deployed as part of Maximo Manage, the two applications share a database.

[IBM Maximo Health](#)

IBM Maximo Health is an application in Maximo Application Suite. By using Maximo Health, you can improve your asset's reliability by understanding asset health and taking action. You can review your assets' performance and condition indicators, such as the last failure date and the maintenance-to-replacement ratio (MRR), and take action by creating work orders and service requests. You can use work queues to improve the quality of your asset's details and related data. You can also configure scoring for assets' health, criticality, and risk.

[Maximo Health and Predict - Utilities](#)

Related tasks

[Updating IBM Maximo Health](#)

When new versions of IBM Maximo Health become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

Related information

[Getting started with Maximo Health and Predict - Utilities](#)

Customer-managed **Deploying IBM Maximo Manage**

To prepare for deployment, you must complete several tasks, such as configuring the database for Maximo Manage. Verify the compatibility of the industry solutions and add-ons if you want to deploy them

and complete other required and optional preparation steps as needed. For example, if you are deploying in multiple languages, review the support languages that are available.

Related concepts

Application database

To deploy Maximo Health, Maximo Real Estate and Facilities, or Maximo Manage, a database instance must be configured and running. The applications support Db2, Db2 Warehouse, Microsoft® SQL Server, or Oracle Database. If Maximo Health is deployed as part of Maximo Manage, the two applications share a database.

IBM Maximo IT

IBM Maximo IT provides a single point of user support and enterprise service management of information technology (IT) and operational technology (OT) assets and processes.

IBM Maximo Manage

IBM Maximo Manage is an application in Maximo Application Suite. By using Maximo Manage, you can get a comprehensive view of all of your asset types, their conditions and locations, and the work processes that support them, to support optimal planning, control, audit, and compliance capability.

Deploying Maximo Manage in Maximo Application Suite

By deploying and activating Maximo Manage, you make it available for use in IBM Maximo Application Suite.

A system administrator can do the following tasks:

- Configure your database for Maximo Manage.
- Select and deploy industry solutions and add-ons.
- Complete other configurations as needed.
- After the deployment is complete, grant users access permission to use Maximo Manage.
- Complete other post-deployment and administration tasks.

Maximo Application Suite is a suite of asset management and data analytic offerings.

As a system administrator, you can deploy and activate Maximo Manage and related industry solutions and add-ons in Maximo Application Suite. After you install Maximo Manage, the architecture is as shown in the following figure.

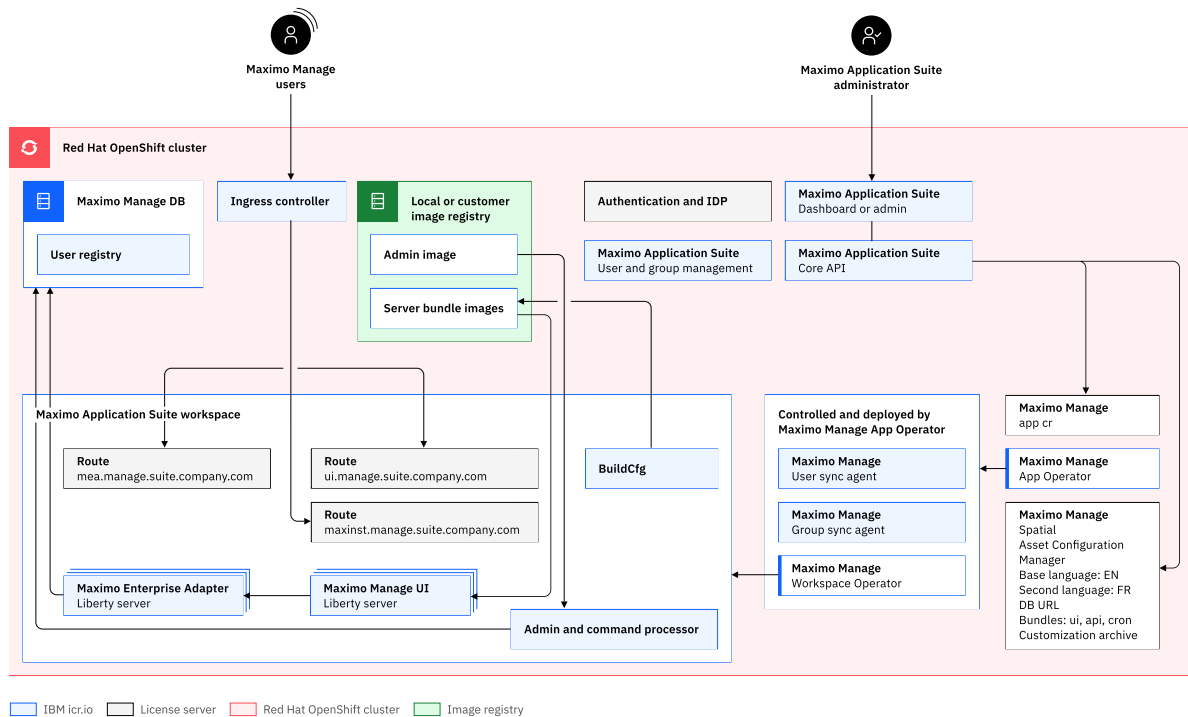


Figure 1. Maximo Manage in Maximo Application Suite

When you want to connect to a particular database or install an add-on, a custom resource is generated to specify exactly what you want to do. The Maximo Manage operator generates the admin image and the server bundle images. For example, if you want to install industry solutions, the server bundle images are the Liberty servers because the server must be deployed there. The operator completes these tasks. Then, those finalized images are placed in the local image registry or a customer-provided image registry. The operator issues commands to the Admin pod and the database to configure the database. When the configuration is done, the operator starts the Liberty servers and makes the routes available, so users can start to connect to it. The operator automates the whole deployment.

Before you deploy, verify that the following tasks are complete:

- Verify that your environment meets the system requirements. For more information, see [system requirements](#).
- Verify that Maximo Application Suite is installed and configured according to your requirements.
- Learn about preparing your database, provisioning persistent storage, configuring your customization archive, and other tasks for the deployment. For more information, see [Plan and prepare Maximo Manage](#).

Then, you can deploy and activate Maximo Manage in Maximo Application Suite. For more information, see [Deploying and activating Maximo Manage](#).

Preparing to deploy

To prepare for deployment, complete several tasks, such as configuring the database for IBM Maximo Manage and determining the compatibility of the industry solutions and add-ons that you want to deploy. If you are deploying in multiple languages, review the support languages that are available.

System requirements

To deploy in Maximo Application Suite, your environment must meet the hardware and software requirements for Maximo Manage and Maximo Application Suite. For more information, see [IBM Maximo Application Suite system requirements](#).

If you are creating a Maximo Manage database, see the following guidelines:

- Use hard disks that have over 200 MB/second throughput.
- If you are using IBM Cloud Storage, choose a custom performance block storage of 100+ IOPS during setup.
- If you are using AWS cloud storage and Elastic File System (EFS) for your database, consider Provisioned mode for the constant throughput. For more disk options, see [Amazon EFS performance](#).

Check the detailed system requirements. For more information, see [Software Product Compatibility Reports](#).

Configuring an internal image repository

Create a repository for storing images that are created by the operator during Maximo Manage deployment.

Procedure

1. Log in to a Bastion host.
2. Create a network file system (NFS) for the image repository.

```
mkdir /disk1/nfs/image-registry && chmod 777 /disk1/nfs/image-registry/
```

3. Edit the `/etc/exports` file and add an export statement.

```
/disk1/nfs/image-registry
10.176.245.0/24(rw,no_subtree_check,sync,no_wdelay,insecure,no_root_squash)
```

4. Restart the server.

```
systemctl restart nfs-server
```

5. Create the `image-storage.yaml` file.

```
---
apiVersion: v1
kind: PersistentVolume
metadata:
  name: image-registry-pv
spec:
  capacity:
    storage: 500Gi
  claimRef:
    apiVersion: v1
    kind: PersistentVolumeClaim
    name: image-registry-storage
    namespace: openshift-image-registry
  accessModes:
  - ReadWriteMany
  nfs:
    path: /disk1/nfs/image-registry
    server: 10.176.245.2
    persistentVolumeReclaimPolicy: Retain
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  finalizers:
  - kubernetes.io/pvc-protection
  name: image-registry-storage
  namespace: openshift-image-registry
spec:
  accessModes:
  - ReadWriteMany
  resources:
    requests:
      storage: 500Gi
```

6. Apply the `image-storage.yaml` file.

```
oc apply -f image-storage.yaml
```

7. Edit the operator.

```
oc edit configs.imageregistry.operator.openshift.io
```

8. Change the `spec.managementState` value from `Removed` to `Managed`.
9. Update the `spec.storage` value.

storage

```
managementState: Managed
```

pvc

```
claim: image-registry-storage
```

10. Save and quit.
11. Verify that the image repository is available.

```
oc get co image-registry
```

12. Enable the route.

```
oc patch configs.imageregistry.operator.openshift.io/cluster --patch '{"spec": {"defaultRoute":true}}' --type=merge
```

Setting up your database

Before you can deploy IBM Maximo Manage, you must configure your database and determine how your database is encrypted.

Preparing your database for deployment

Maximo Manage supports IBM Db2, IBM Db2 Warehouse, Microsoft SQL Server, and Oracle Database. Configure your database and gather the information that is needed when you deploy the application.

Before you begin

Tip: Set up your database at the same site as IBM Maximo Application Suite to reduce the chances of database failures because of connection issues and data latency. If you have issues when you set up the database during Maximo Manage deployment, troubleshoot the deployment. For more information, see [“Troubleshooting database deployment”](#) on page 325.

Procedure

Confirm that one of the following databases is configured.

- [Configuring IBM Db2](#)
- [Configuring IBM Db2 Warehouse](#)
- [Configuring Oracle Database](#)
- [Configuring Microsoft SQL Server](#)

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

Configuring IBM Db2

Before you deploy Maximo Manage in Maximo Application Suite, configure IBM Db2 for use by Maximo Manage.

Before you begin

Before you configure the database, install and deploy it. For more information, see [“Configuring database instances”](#) on page 316.

For information about Db2 version compatibility, see [Maximo Application Suite detailed system requirements](#).

See the following guidelines on how to configure a database instance.

- Configure separate system, data, and backup storage when you create a Db2 instance.
- Increase the maxsequence cache to 50.
- Run **REORG INDEXERS/TABLES** and **RUNSTATS** daily.
- Separate system storage, user storage, backup storage, transaction logs storage, and temporary table space storage on different disks.
- Maximo Manage requires row-organized tables. By default, the IBM Db2 Warehouse database setting uses column-based table organization. Update the setting as needed.
- Maximo Manage does not support Massively Parallel Processing (MPP) or table partitioning. Archive records that are over a year old. InfoSphere® Optim Data Growth Solution can be used for archiving. For more information, see [IBM Maximo Archiving 7.5.1 for IBM Maximo Asset Management](#) .
- An issue can occur when you load a large amount of data by using the Maximo Integration Framework. Increase the concurrently running statements that are allowed for a Db2 application. For more information, see [How many concurrently running statements allowed for a Db2 Java application and how to increase it?](#).
- If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:
 - Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
 - Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
 - Secure operator versions and disable auto-update for all components.
 - Use the IAM and LDAP service instead of the default authentication methods.
 - For more information, see [DB2 Performance Insight](#)®.

Note:

- Starting in 9.0.5 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x architecture, Db2 is not configured.
- Starting in 9.0.12 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM Power (ppc64le) architecture, internalDb2 is not configured. Db2 can be used as an external service.

For information about supported database versions, generate a Software Product Compatibility Report. For more information, see [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. On the **Supported Software** tab of the report, check for the supported database versions.

Configure your database with the following operating systems:

- Linux or UNIX
- Microsoft Windows

About this task

The commands in this task can be used to configure a Db2 database outside of the Red Hat OpenShift cluster, by using different operating systems such as Microsoft Windows, Linux, or UNIX.

Note: The commands in this task are not applicable for the configuration of Db2 Warehouse instance.

The commands in this task are examples of the commands that you must run. For example, maxdb80 is the name of the database. If maxdb80 is not your database name, ensure that you replace all instances with the correct database name.

Procedure

1. Log in to the system as a user that has administrative permissions.
2. If system users do not exist on the system, create the system users.

- Windows
 - db2admin
 - maximo
- Linux or UNIX
 - maximo for the Maximo database user
 - ctgfenc1 for the Db2 fenced user
 - ctginst1 for the Db2 instance owner

The most used administrative user is assigned the primary group of the instance owner to complete some of the following steps.

3. At the Db2 installation directory, set up the command-line environment.

- For Windows , run the following command:**db2cmd**
- For Linux or UNIX , ensure that the /opt/ibm/db2/V11.5/bin, /opt/ibm/db2/v11.5/instance, and /opt/ibm/db2/V11.5/adm directories are added to your PATH.

4. Run the following commands to create the database instance.

- Windows

Where *<administrator_password>* with the Db2 administrator password.

```
db2icrt -s ese -u db2admin,<administrator_password> -r 50005,50005 ctginst1
set db2instance=ctginst1
db2start
db2 update dbm config using SVCENAME 50005 DEFERRED
db2stop
db2set DB2COMM=tcPIP
db2start
```

- Linux or UNIX

```
db2icrt -s ese -u ctgfenc1 -p 50005 ctginst1
./home/ctginst1/sqlib/db2profile
db2start
db2 update dbm config using SVCENAME 50005 DEFERRED
db2stop
db2set DB2COMM=tcPIP
db2start
```

5. Run the following commands to create the database:

```
db2 create db 'maxdb80' ALIAS 'maxdb80' using codeset UTF-8 territory US pagesize 32 K
db2 connect to 'maxdb80'
db2 GRANT DBADM ON DATABASE TO USER db2admin (windows only)
db2 GRANT SECADM ON DATABASE TO USER db2admin (windows only)
db2 connect reset
```

6. Run the following command according to your operating system and bit size:

Operating system	Command
32-bit Microsoft Windows	db2 update db cfg for maxdb80 using MAXFILOP 32768 DEFERRED #32-bit Windows

Operating system	Command
64-bit Windows	db2 update db cfg for maxdb80 using MAXFILOP 65335 DEFERRED #64-bit Windows
32-bit UNIX	db2 update db cfg for maxdb80 using MAXFILOP 30720 DEFERRED #32-bit UNIX
64-bit UNIX	db2 update db cfg for maxdb80 using MAXFILOP 61440 DEFERRED #64-bit UNIX

7. Run the following commands to configure the database:

```

db2 update db cfg for maxdb80 using SELF_TUNING_MEM ON
db2 update db cfg for maxdb80 using APPGROUP_MEM_SZ 16384 DEFERRED
db2 update db cfg for maxdb80 using APPLHEAPSZ 2048 AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using AUTO_MAINT ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_TBL_MAINT ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_RUNSTATS ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_REORG ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_DB_BACKUP ON DEFERRED
db2 update db cfg for maxdb80 using CATALOGCACHE_SZ 800 DEFERRED
db2 update db cfg for maxdb80 using CHNGPGS_THRESH 40 DEFERRED
db2 update db cfg for maxdb80 using DBHEAP AUTOMATIC
db2 update db cfg for maxdb80 using LOCKLIST AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using LOGBUFSZ 1024 DEFERRED
db2 update db cfg for maxdb80 using LOCKTIMEOUT 300 DEFERRED
db2 update db cfg for maxdb80 using LOGPRIMARY 20 DEFERRED
db2 update db cfg for maxdb80 using LOGSECOND 100 DEFERRED
db2 update db cfg for maxdb80 using LOGFILSIZ 8192 DEFERRED
db2 update db cfg for maxdb80 using SOFTMAX 1000 DEFERRED

db2 update db cfg for maxdb80 using PCKCACHESZ AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using STAT_HEAP_SZ AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using STMTHEAP AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using UTIL_HEAP_SZ 10000 DEFERRED
db2 update db cfg for maxdb80 using DATABASE_MEMORY AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using AUTO_STMT_STATS OFF DEFERRED
db2 update db cfg for maxdb80 using STMT_CONC LITERALS DEFERRED
db2 update alert cfg for database on maxdb80 using db.db_backup_req SET THRESHOLDSCHECKED
YES
db2 update alert cfg for database on maxdb80 using db.tb_reorg_req SET THRESHOLDSCHECKED YES
db2 update alert cfg for database on maxdb80 using db.tb_runstats_req SET THRESHOLDSCHECKED
YES
db2 update dbm cfg using PRIV_MEM_THRESH 32767 DEFERRED
db2 update dbm cfg using KEEPFENCED NO DEFERRED
db2 update dbm cfg using NUMDB 2 DEFERRED
db2 update dbm cfg using RQRIOLBK 65535 DEFERRED
db2 update dbm cfg using HEALTH_MON OFF DEFERRED
db2 update dbm cfg using AGENT_STACK_SZ 1000 DEFERRED
db2 update dbm cfg using MON_HEAP_SZ AUTOMATIC DEFERRED
db2set DB2_SKIPINSERTED=ON
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_MINIMIZE_LISTPREFETCH=Y
db2set DB2_EVALUNCOMMITTED=YES
db2set DB2_FMP_COMM_HEAPSZ=65536
db2set DB2_SKIPDELETED=ON
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON

```

8. For Linux or UNIX , log in to the system.

For example, log in as the ctginst1 user and then restart the Db2 command-line environment:

```

su - ctginst1
db2

```

9. Run the following command to stop the database:

```

db2stop force

```

10. Run the following command to start the database:

```

db2start

```


11. Run the following command to reconnect to the database:

```
db2 connect to 'maxdb80'
```

12. Run the following commands to create a buffer pool:

```
db2 CREATE BUFFERPOOL MAXBUFPOOL IMMEDIATE SIZE 4096 AUTOMATIC PAGESIZE 32 K
db2 CREATE REGULAR TABLESPACE MAXDATA PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
db2 CREATE TEMPORARY TABLESPACE MAXTEMP PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
BUFFERPOOL MAXBUFPOOL
db2 CREATE REGULAR TABLESPACE MAXINDEX PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
db2 GRANT USE OF TABLESPACE MAXDATA TO USER MAXIMO
```

13. Run the following command to create the schema:

```
db2 create schema maximo authorization maximo
```

14. Run the following commands to grant authority to the Maximo user:

```
db2 GRANT
DBADM,CREATETAB,BINDADD,CONNECT,CREATE_NOT_FENCED_ROUTINE,IMPLICIT_SCHEMA,LOAD,CREATE_EXTERN
AL_ROUTINE,QUIESCE_CONNECT,SECADM ON DATABASE TO USER MAXIMO
db2 GRANT USE OF TABLESPACE MAXDATA TO USER MAXIMO
db2 GRANT CREATEIN,DROPIN,ALTERIN ON SCHEMA MAXIMO TO USER MAXIMO
```

15. Run the following command to break the database connection:

```
db2 connect reset
```

Example

For example, you can configure the database by using an Amazon Web Services EC2 instance.

Related information

[Software Product Compatibility Report](#)

Configuring IBM Cloud Pak for Data

If you select Db2 as your database, you must configure Cloud Pak for Data.

About this task

If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:

- Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
- Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
- Secure operator versions and disable auto-update for all components.
- Use the IAM and LDAP service instead of the default authentication methods.
- For more information, see [DB2 Performance Insight](#).

Procedure

1. Open the Red Hat OpenShift console to log in to IBM Cloud Pak for Data.

If you do not have the Cloud Pak for Data admin password the first time you log in, you can get it through the Red Hat OpenShift console.

- a) In the Red Hat OpenShift console, from the side navigation menu, click **Workloads > Secrets**.
- b) In the **Project** field, select a Cloud Pak for Data namespace.

If you installed Maximo Application Suite on Amazon Web Services, the Cloud Pak for Data namespace is **cpd-services-uniqueid**.

- c) Filter for `admin-user-details`.
- d) Click the name of the administrator account.
- e) Copy the value of the `initial_admin_password` field.
2. From the side navigation menu, click **Networking > Routes**.
3. On the **Routes** page, from the **Project** field, select a Cloud Pak for Data namespace.
4. Click the **Location** link to open the Cloud Pak for Data login page in a new browser tab.
5. Log in to Cloud Pak for Data as an administrator.
6. Select **Databases** and then click **Create Database**.
7. On the **Select a database** page, click **Next**.
8. On the **Configure** page, select **Single location for all data** and then click **Next**.
9. On the **Advanced** page, click **Next**.
10. On the **Storage** page, select the storage class and then click **Next**.
For Amazon Web Services customer-managed Red Hat OpenShift clusters that are provisioned through an automated deployment offering, select **oc-storagecluster-cephs**.
11. On the **Finalize** page, update the display name if needed, and then click **Create**.

Results

After the instance is created, a green icon appears on the database tile.

Configuring IBM Db2 Warehouse

Create a IBM Db2 Warehouse database for exclusive use by Maximo Manage. You cannot reuse a Db2 Warehouse database on Cloud Pak for Data that is already used by another deployed Maximo Application Suite application.

Before you begin

Create a Db2 Warehouse database on IBM Cloud Pak for Data. For more information, see [Creating Db2 instance by using IBM Cloud Pak for Data console](#).

After Db2 Warehouse databases are provisioned, a parameter is configured to specify user table organization on creation set as a column-organized table. However, to successfully deploy Maximo Manage on Db2 Warehouse, you must change this parameter so that the tables are created as row-organized tables.

Db2 Warehouse can have more than one provisioned database instance. Ensure that you choose the one that is used by Maximo Manage to set this configuration and use when you configure the database. If your instance of Db2 Warehouse was installed through Cloud Pak for Data, you can complete the following steps to find the Db2 administrator pod:

1. Log in to the Cloud Pak for Data interface as an administrator.
2. Select **Databases**.
3. On the tile of the Db2 Warehouse database instance that you provisioned for Maximo Manage, select **Details** from the menu.
4. Search for the deployment ID value, for example, `db2wh-1652220906500619`.
5. Copy the suffix of the deployment ID, for example, `1652220906500619`.
6. Search for the database name. The name of the database is the value that you use in place of the `$DB_NAME` variable in command-line examples.
7. Open the Red Hat OpenShift console and from the side navigation menu, click **Workloads > Pods**.
8. From the **Project** menu, select your Cloud Pak for Data namespace.

If you installed Maximo Application Suite on Amazon Web Services, the Cloud Pak for Data namespace is **cpd-services-uniqueid**.

9. In the **Filter** field, enter the suffix of the deployment ID that you copied and append it to the administrator pod value, for example, 1652220906500619-db2u-0. The pod that is displayed is the db2u administrator pod where you run the **db2inst1** command.

See the following guidelines on how to configure a database instance.

- Configure separate system, data, and backup storage when you create a Db2 instance.
- Increase the maxsequence cache to 50.
- Run **REORG INDEXERS/TABLES** and **RUNSTATS** daily.
- Separate system storage, user storage, backup storage, transaction logs storage, and temporary table space storage on different disks.
- Maximo Manage requires row-organized tables. By default, the IBM Db2 Warehouse database setting uses column-based table organization. Update the setting as needed.
- Maximo Manage does not support Massively Parallel Processing (MPP) or table partitioning. Archive records that are over a year old. InfoSphere Optim Data Growth Solution can be used for archiving. For more information, see [IBM Maximo Archiving 7.5.1 for IBM Maximo Asset Management](#) .
- An issue can occur when you load a large amount of data by using the Maximo Integration Framework. Increase the concurrently running statements that are allowed for a Db2 application. For more information, see [How many concurrently running statements allowed for a Db2 Java application and how to increase it?](#).
- If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:
 - Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
 - Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
 - Secure operator versions and disable auto-update for all components.
 - Use the IAM and LDAP service instead of the default authentication methods.
 - For more information, see [DB2 Performance Insight](#).

Note: The following commands use the variable `$DB_NAME` to indicate the name of the database that you defined. Before you run the commands, ensure that you replace the `$DB_NAME` with the name of your database in the **export** command, which is available in the list of the commands.

```
su - db2inst1
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
db2 update db cfg for $DB_NAME using dft_table_org row
db2 terminate
db2 deactivate db $DB_NAME
db2stop
db2start
db2 activate db $DB_NAME
db2 connect to $DB_NAME
db2 get db cfg | grep DFT_TABLE_ORG
```

About this task

The following information is an example of the steps that you must complete and applies only to Cloud Pak for Data. If you complete different steps, ensure that a database schema and table spaces are configured and that the **ddl_constraint_def** parameter is set to yes.

The following steps and commands also use the variable `$DB_NAME` to indicate the name of the database that you defined. Before you run the commands, ensure that you replace the `$DB_NAME` with the actual name of your database.

Procedure

1. Confirm that the database is created.
 - a) As an administrator, log in to Red Hat OpenShift and in the Red Hat OpenShift cluster, from the navigation menu, click **Workloads** > **Pods** .

- b) Locate and open the `db2wh-string-db2u-0` pod. *string* is a randomly generated set of numbers.
- c) On the **Terminal** tab, run the following commands:

```
su - db2inst1
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
```

If the database is configured, the database connection information is returned. The following text is an example of this information, where `BLUDB` is the name of the database. You need these values later in the configuration.

```
sh-4.2$ su - db2inst1
Last login: Tue May 26 14:29:55 UTC 2020
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ db2 connect to BLUDB

Database Connection Information

Database server = DB2/LINUX8664 11.5.2.0
SQL authorization ID = DB2INST1
Local database alias = BLUDB
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

2. Create an administrative Cloud Pak for Data user.

To set Maximo Manage to connect to the database, you can use the `db2inst1` user that comes by default with Db2 Warehouse, or you can create a different administrative user.

If you installed Db2 Warehouse without Cloud Pak for Data, you can create a different administrative user. For more information, see [Authentication options for Db2U](#).

If you installed Db2 Warehouse through Cloud Pak for Data, you can create a new administrative user through Cloud Pak for Data.

- a) As an administrator, log in to Cloud Pak for Data and from the side navigation menu, click **Administration > User management**.
- b) Create a user by specifying the following information:
 - Specify a user and username.
 - Specify an email address and password.
 - Select the **Administrator** role.
- c) From the side navigation menu, click **Data > Databases**.
- d) Click the three-dot icon for your Maximo Manage database and then click **Details**.
- e) From the drop-down menu, click **Manage access**.
- f) In the Maximo user row, click the edit icon.
- g) In the **Role** field, select **Admin** and then click **Save**.

Note: Take note of this username because it is used as the value of the `$DB_USERNAME` variable in commands in this procedure.

3. Prepare the database for the `maxinst` program. Run the following commands from a database browser or command-line:

- a) Run the following command to connect to the database as the `db2inst1` user:

```
su - db2inst1
```

- b) Run the following command to connect to the database:

```
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
```

- c) Optional: For administrators who connect from outside Red Hat OpenShift, run the following command to connect to the Db2 pod. Replace the variable with the string value for your `db2wh-string-db2u-0` pod that was accessed during step 1.

```
oc rsh -n <db2wh namespace> c-db2wh-*<string>*-db2u-0 /bin/bash
```

Note: The name of the c-db2wh-*<string>*-db2u-0 pod can be different if you installed Db2 Warehouse without Cloud Pak for Data. For example, if you installed the DB2u operator through Db2, the name of the pod is similar to c-*DB2_INSTANCE_NAME*-db2u-0.

4. Configure the database.

a) Run the following commands to configure the database:

Note: Choose the APPHEAPSZ db2 value to use in the following commands and replace the value of the *\$HEAPVALUE* variable by at least 2048 in its specific **export** command line. If you plan to install many Manage extensions at the same time, set the *\$HEAPVALUE* variable to 16384 to prevent failure of the **maxinst** or **updated** processes.

```
su - db2inst1
export DB_NAME=$DB_NAME
export HEAPVALUE=$HEAPVALUE
db2 connect to $DB_NAME
db2 update db cfg for $DB_NAME using "APPHEAPSZ $HEAPVALUE AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_MAINT ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_TBL_MAINT ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_RUNSTATS ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_REORG ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_DB_BACKUP ON DEFERRED"
db2 update db cfg for $DB_NAME using "CATALOGCACHE_SZ 800 DEFERRED"
db2 update db cfg for $DB_NAME using "CHNGPGS_THRESH 40 DEFERRED"
db2 update db cfg for $DB_NAME using "DBHEAP AUTOMATIC"
db2 update db cfg for $DB_NAME using "LOCKLIST AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "LOGBUFSZ 1024 DEFERRED"
db2 update db cfg for $DB_NAME using "LOCKTIMEOUT 300 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGPRIMARY 20 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGSECOND 100 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGFILSIZ 8192 DEFERRED"
db2 update db cfg for $DB_NAME using "SOFTMAX 1000 DEFERRED"
db2 update db cfg for $DB_NAME using "MAXFILOP 61440 DEFERRED"
db2 update db cfg for $DB_NAME using "PCKCACHESZ AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "STAT_HEAP_SZ AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "STMTHEAP AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "UTIL_HEAP_SZ 10000 DEFERRED"
db2 update db cfg for $DB_NAME using "DATABASE_MEMORY AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_STMT_STATS OFF DEFERRED"
db2 update db cfg for $DB_NAME using "STMT_CONC LITERALS DEFERRED"
db2 update alert cfg for database on $DB_NAME using "db.db_backup_req SET THRESHOLDSCHECKED YES"
db2 update alert cfg for database on $DB_NAME using "db.tb_reorg_req SET THRESHOLDSCHECKED YES"
db2 update alert cfg for database on $DB_NAME using "db.tb_runstats_req SET THRESHOLDSCHECKED YES"
db2 update dbm cfg using "PRIV_MEM_THRESH 32767 DEFERRED"
db2 update dbm cfg using "KEEPFENCED NO DEFERRED"
db2 update dbm cfg using "NUMDB 2 DEFERRED"
db2 update dbm cfg using "RQRIOBLK 65535 DEFERRED"
db2 update dbm cfg using "HEALTH_MON OFF DEFERRED"
db2 update dbm cfg using "AGENT_STACK_SZ 1000 DEFERRED"
db2 update dbm cfg using "MON_HEAP_SZ AUTOMATIC DEFERRED"
db2 update db cfg using "DDL_CONSTRAINT_DEF YES"
db2set DB2_SKIPINSERTED=ON
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_MINIMIZE_LISTPREFETCH=Y
db2set DB2_EVALUNCOMMITTED=YES
db2set DB2_FMP_COMM_HEAPSZ=65536
db2set DB2_SKIPDELETED=ON
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
```

b) Run the following command to create the buffer pool:

```
db2 CREATE BUFFERPOOL MAXBUFPOL IMMEDIATE SIZE 4096 AUTOMATIC PAGESIZE 32 K
```

c) Run the following commands to create the table spaces:

```
db2 CREATE REGULAR TABLESPACE MAXDATA PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOL
db2 CREATE TEMPORARY TABLESPACE MAXTEMP PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
BUFFERPOOL MAXBUFPOL
```

```
db2 CREATE REGULAR TABLESPACE MAXINDEX PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFFERPOOL
```

d) Run the following command to create the schema:

Note:

- The `$DB_SCHEMA` variable is the name you give to the schema.
- The `$DB_USERNAME` variable is the user with administrative rights, which was created in step 2.

```
export DB_SCHEMA=$DB_SCHEMA
export DB_NAME=$DB_NAME
export DB_USERNAME=$DB_USERNAME
db2 CREATE SCHEMA $DB_SCHEMA AUTHORIZATION $DB_USERNAME
```

e) Run the following commands to grant authority to the database user:

```
db2 GRANT
DBADM,CREATETAB,BINDADD,CONNECT,CREATE_NOT_FENCED_ROUTINE,IMPLICIT_SCHEMA,LOAD,CREATE_EXTE
RNAL_ROUTINE,QUIESCE_CONNECT,SECADM ON DATABASE TO USER $DB_USERNAME
db2 GRANT USE OF TABLESPACE MAXDATA TO USER $DB_USERNAME
db2 GRANT CREATEIN,DROPIN,ALTERIN ON SCHEMA $DB_SCHEMA TO USER $DB_USERNAME
```

f) Run the following command to break the database connection:

```
db2 connect reset
```

5. Verify that the configuration is successful.

- a) In the Red Hat OpenShift console, from the side navigation menu, click **Workload > Pods**.
- b) In your Db2 Warehouse database instance, click the **Db2u** pod.
- c) Select the **Terminal** tab and run the following commands:

```
su - db2inst1
Export DB_NAME=<yourdbname>
Export DB_USERNAME=<yourdbusername> db2 connect to $DB_NAME user $DB_USERNAME using
<password>
```

Tip:

- Replace `<yourdbname>` with the database name that you obtained in a previous step.
- Replace `<yourdbusername>` with the username that you obtained in a previous step.
- Replace `<password>` with the username password that you set when you created and configured the user in a previous step.

The details for the database connection are displayed after the database is connected. The following text is an example of this information, where `BLUDB` is the database name and `MAXIMO` is the username.

```
sh-4.2$ su - db2inst1
Last login: Tue May 26 14:29:55 UTC 2020
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ export DB_NAME=BLUDB
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ echo $DB_NAME BLUDB
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ export DB_USERNAME=MAXIMO
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ db2 connect to $DB_NAME user
$DB_USERNAME using MAXIMO
```

Database Connection Information

```
Database server      = DB2/LINUX8664 11.5.2.0
SQL authorization ID = MAXIMO
Local database alias = BLUDB
```

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

d) Run the following command to list the table spaces:

```
db2 list tablespaces
```

Confirm the Db2 MAXINDEX, MAXTEMP, and MAXDATA table spaces are listed. The following text is an example of the table spaces details:

```
NAME = MAXDATA
TYPE = Database Managed Space
CONTENTS = All permanent data. Regular table space.
STATE = 0x0000
Detailed explanation = Normal
```

```
TABLESPACE ID = 5
NAME = MAXTEMP
TYPE = System Managed Space
CONTENTS = System Temporary data
STATE = 0x0000
Detailed explanation = Normal
```

```
TABLESPACE ID = 6
NAME = MAXINDEX
TYPE = Database Managed Space
CONTENTS = All permanent data. Regular table space.
STATE = 0x0000
Detailed explanation = Normal
```

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

- e) Run the following command to view the default table organization:

```
db2 get db cfg | grep DFT_TABLE_ORG
```

Confirm that the configuration `db2 get db cfg | grep DFT_TABLE_ORG` is set to ROW. The following text is an example of the configuration:

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$ db2 get db cfg | grep DFT_TABLE_ORG
Default table organization (DFT_TABLE_ORG) = ROW
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

- f) Run the following command to view the default application heap:

```
db2 get db cfg | grep APPLHEAPSZ
```

Confirm that the value for APPLHEAPSZ in `db2 get db cfg | grep APPLHEAPSZ` is the same that you set. The following text is an example of the configuration:

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$ db2 get db cfg | grep APPLHEAPSZ
Default application heap (4KB) (APPLHEAPSZ) = AUTOMATIC(16384)
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

6. Disable the non-SSL Db2 database instance port for security.

- a) Run the following command to edit the Db2 service configuration:

```
oc edit svc -n project c-service_name-db2u-engn-svc
```

Where

project

The name of the Red Hat OpenShift project where Db2 is deployed.

service_name

The identifier for the Db2 service instance.

For example, `c-db2o1tp-1605022957148004-db2u-engn-svc`.

- b) Remove the following text from the `spec.ports` section:

```
- name: legacy-server
  nodePort: 30279
  port: 50000
  protocol: TCP
  targetPort: 50000
```

- c) Save the service.

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

Configuring Oracle Database

To configure Oracle Database for use with Maximo Manage, you create table spaces, create a database user, and configure database settings.

Before you begin

For information about installing and deploying Oracle Database, review the Oracle Database product documentation.

For information about supported database versions, generate a Software Product Compatibility Report. For more information, see [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. Check for the supported database versions on the **Supported Software** tab of the report.

Configure your database with the following operating system:

- Linux or UNIX
- Microsoft Windows

For more information about system performance, see [Best practices for system performance](#).

Procedure

1. Log in as the Oracle software user. Typically, this user is named `oracle`.
2. To manage requests to connect to the database, create the database listener.
3. Create a database for use by Maximo Manage.

For the database initialization parameters, change the values of the following parameters:

nls_length_semantics

Change this value to CHAR.

open_cursors

Change this value to 1000.

cursor_sharing

Set this value to FORCE.

Note: Ensure that the database character setting is set to the AL32UTF8 character set, which is required for Oracle Database.

4. In SQL*Plus, create a table space by running the following command. Replace the directory with the path to the database location.

```
Create tablespace maxdata datafile
'C:\oracle\product\12.1.0.1\db_1\dbs\maxdata.dbf'
size 1000M autoextend on;
```

To create table spaces for indexes, repeat the command and use a similar syntax.

5. Create a temporary table space by running the following command. Replace the directory with the path to the database location.


```
create temporary tablespace maxtemp tempfile
'C:\oracle\product\12.1.0.1\db_1\dfs\maxtemp.dbf'
size 1000M autoextend on maxsize unlimited;
```

6. To create the Maximo user and grant permissions, run the following command:

```
create user maximo identified by maximo default tablespace maxdata temporary
tablespace maxtemp;
grant connect to maximo;
grant create job to maximo;
grant create trigger to maximo;
grant create session to maximo;
grant create sequence to maximo;
grant create synonym to maximo;
grant create table to maximo;
grant create view to maximo;
grant create procedure to maximo;
grant alter session to maximo;
grant execute on ctxsys.ctx_ddl to maximo;
alter user maximo quota unlimited on maxdata;
```

If you created a separate table space for indexing, you must also grant access to that index table space to the Maximo user.

For example, if you created a separate table space for indexing that is called **TSI_MAM_OWN**, then run the following command:

```
alter user maximo quota unlimited on TSI_MAM_OWN
```

7. Create an Oracle preference arbitrary that is called **MAXIMO_STORAGE** and store the Oracle text indexes in dedicated table spaces.

- Store implicit indexes in the MAXINDEX table space.
- Store implicit tables in the MAXDATA table space.
- Store implicit LOB tables in the LOB table space.

For example, run the following preference definition to split the implicit objects in a text index across three table spaces: MAXDATA, MAXINDEX, and MAXLOBS.

```
begin
ctx_ddl.create_preference('MAXIMO_STORAGE', 'BASIC_STORAGE');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'I_TABLE_CLAUSE',
'tablespace MAXDATA LOB(token_info) store as (tablespace MAXLOBS
enable storage in row)');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'I_INDEX_CLAUSE',
'tablespace MAXINDEX compress 2');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'K_TABLE_CLAUSE',
'tablespace MAXINDEX');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'R_TABLE_CLAUSE',
'tablespace MAXDATA LOB(data) store as (tablespace MAXLOBS
cache)');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'N_TABLE_CLAUSE',
'tablespace MAXINDEX');
end;
```

To create an Oracle text index, the preference definition must be specified in the **CREATE INDEX** clause as shown in the following example.

```
create index pm_ndx6 on pm (description) indextype is
ctxsys.context parameters ('lexer global_lexer language column
LANGCODE storage MAXIMO_STORAGE');
```

8. Set up Oracle text preferences and sublexer definitions.

- a) Use an SQL query tool to log on to the database as the maximo user, which is the schema owner, and run the following set of calls.

```
call ctx_ddl.drop_preference('global_lexer');
call ctx_ddl.drop_preference('default_lexer');
call ctx_ddl.drop_preference('english_lexer');
call ctx_ddl.drop_preference('chinese_lexer');
call ctx_ddl.drop_preference('japanese_lexer');
```

```

call ctx_ddl.drop_preference('korean_lexer');
call ctx_ddl.drop_preference('german_lexer');
call ctx_ddl.drop_preference('dutch_lexer');
call ctx_ddl.drop_preference('swedish_lexer');
call ctx_ddl.drop_preference('french_lexer');
call ctx_ddl.drop_preference('italian_lexer');
call ctx_ddl.drop_preference('spanish_lexer');
call ctx_ddl.drop_preference('portu_lexer');
call ctx_ddl.create_preference('default_lexer','basic_lexer');
call ctx_ddl.create_preference('english_lexer','basic_lexer');
call ctx_ddl.create_preference('chinese_lexer','chinese_lexer');
call ctx_ddl.create_preference('japanese_lexer','japanese_lexer');
call ctx_ddl.create_preference('korean_lexer','korean_morph_lexer');
call ctx_ddl.create_preference('german_lexer','basic_lexer');
call ctx_ddl.create_preference('dutch_lexer','basic_lexer');
call ctx_ddl.create_preference('swedish_lexer','basic_lexer');
call ctx_ddl.create_preference('french_lexer','basic_lexer');
call ctx_ddl.create_preference('italian_lexer','basic_lexer');
call ctx_ddl.create_preference('spanish_lexer','basic_lexer');
call ctx_ddl.create_preference('portu_lexer','basic_lexer');
call ctx_ddl.create_preference('global_lexer','multi_lexer');
call ctx_ddl.add_sub_lexer('global_lexer','default','default_lexer');
call ctx_ddl.add_sub_lexer('global_lexer','english','english_lexer','en');
call ctx_ddl.add_sub_lexer('global_lexer','simplified chinese','chinese_lexer','zh');
call ctx_ddl.add_sub_lexer('global_lexer','japanese','japanese_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','korean','korean_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','german','german_lexer','de');
call ctx_ddl.add_sub_lexer('global_lexer','dutch','dutch_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','swedish','swedish_lexer','sv');
call ctx_ddl.add_sub_lexer('global_lexer','french','french_lexer','fr');
call ctx_ddl.add_sub_lexer('global_lexer','italian','italian_lexer','it');
call ctx_ddl.add_sub_lexer('global_lexer','spanish','spanish_lexer','es');
call ctx_ddl.add_sub_lexer('global_lexer','portuguese','portu_lexer',null);

commit;

```

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

If you modified the default role sets assigned to the user ID used to connect to the database, then you must explicitly grant the role sets to the Maximo user. If you restricted the default privileges that are granted to user IDs, you must also explicitly grant the role sets to the Maximo user. For example, if you do not grant a role such as the **select_catalog_role** role, you must explicitly grant that role to the Maximo user. Make the assignment by running the following SQL*Plus command:

```
grant select_catalog_role to maximo
```

Configuring Microsoft SQL Server

To configure Microsoft SQL Server for Maximo Manage, you create table spaces, create a database user, and configure database settings.

Before you begin

For information about configuring Microsoft SQL Server, review the Microsoft SQL Server product documentation.

See the following guidelines for configuring the database:

- If the original database was created in a version earlier than Microsoft SQL Server 2019, set the compatibility level to the older version to maintain the execution plan.
- Set the transactions isolation level by using the following commands:

```
ALTER DATABASE MyDatabase
    SET ALLOW_SNAPSHOT_ISOLATION ON

ALTER DATABASE MyDatabase
    SET READ_COMMITTED_SNAPSHOT ON
```

For information about supported database versions, you can generate a [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. Check for the supported database versions in the **Supported Software** tab of the report.

Procedure

1. Configure the listener port.

The default instance of the Microsoft SQL Server Database Engine listens on TCP port 1433. Named instances of the Microsoft SQL Server Database Engine and Microsoft SQL Server Compact Edition are configured for dynamic ports, which means they select any available port when the service starts. When you connect to a named instance across a firewall, configure the Database Engine to listen on a specific port to open this port in the firewall.

2. Verify that you enabled the Full-text Search setting during the installation of Microsoft SQL Server.
3. Create a Microsoft SQL Server database.

- a) In Microsoft SQL Server Management Studio, select **New Database** from the databases folder.
- b) Specify a unique database name.
For example, enter maxdb80
- c) For the maxdb80 Logical Name, change the **Initial Size (MB)** field to 500 and also set the value of the **Autogrowth / Maxsize** field to **By 1 MB, Unlimited**.
- d) Optional: Modify the log settings to accommodate your production environment.
- e) To deploy Maximo Manage in a specific language, choose the default collation for the database.
For example, to deploy the application in English, select **Latin1_General_100_CI_AS_KS_SC_UTF8**.

Starting from Maximo Application Suite 9.0, Maximo Manage supports a Microsoft SQL Server database that uses Unicode. You must select a collation that has UTF8 in its name. Microsoft SQL Server supports multiple languages in the same database in instances where the chosen languages support the same Microsoft SQL Server collation. For example, English and French can be installed because both languages support the same Microsoft SQL Server collation. However, English and Japanese cannot be installed together because they have different Microsoft SQL Server collations.

For more information, review the Microsoft SQL Server Collation and Unicode support documentation.

4. Create the Maximo user for Microsoft SQL Server.

- a) In Microsoft SQL Server Management Studio, from the SQL Server Configuration Manager navigation, click **Databases**.
- b) Right-click the **maxdb80** database and select **New Query**
- c) Enter the following command to create the Maximo database user MAXIMO with a password that adheres to the password policy of the system.

```
sp_addlogin MAXIMO,password
go
```

This value is case-sensitive.

- d) Enter the following command to change the database owner to MAXIMO.

```
sp_changedbowner MAXIMO
go
```

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

Configuring database instances

Maximo Manage supports multiple databases. You must configure your database and gather the information when you deploy the application.

Before you begin

If you intend to use the database instance for production, you can select **Deploy database on dedicated nodes** and adjust the number of nodes as needed on the **Configure** page.

When you are considering the size the database, remember the following guidelines:

- The storage class that you select can affect the size of your database. For example, if you plan to use Red Hat OpenShift Container Storage, it increases the storage space that you need for your database.
- Be aware of space requirements for services, such as IBM Cloud Pak for Data. If you deploy in a clustered environment, IBM Cloud Pak for Data space requirements change.
- Managed versus on-premises services, such as MongoDB and Kafka, impact your disk space needs.
- All-in-one pod bundles have different disk space requirements when compared to configuring separate resources for UI, report, cron, Maximo integration framework, and Maximo Mobile.
- The number of concurrent users affect resource needs. For example, for the UI resource, configure a Java virtual machine (JVM) that has two cores to support 50–75 concurrent users.

About this task

Depending on the type of database, you must configure one of the following databases:

- [Configuring IBM Db2](#)
- [Configuring IBM Db2 Warehouse](#)
- [Configuring Oracle Database](#)
- [Configuring Microsoft SQL Server](#)

For information about database version compatibility, see [Maximo Application Suite detailed system requirements](#).

Database encryption

When you deploy Maximo Manage, fields that require security, such as passwords and API keys, are encrypted or reencrypted to provide security.

Database encryption overview

When you configure Maximo Manage, specify encryption keys and encryption algorithms to determine how the fields that require security are encrypted.

Important: Save the encryption secret, which contains the encryption keys, after the Maximo Manage deployment is completed. The same keys are used for configuration when you reinstall Maximo Manage with the same database.

The following table describes the Crypto and CryptoX encryption keys for Maximo Manage:

<i>Table 19. Encryption keys</i>	
Key	Description
MXE_SECURITY_CRYPTO_KEY	Use it to encrypt Crypto fields, such as passwords. For Crypto encryption, if you specify a MXE_SECURITY_CRYPTO_KEY value that matches the MXE_SECURITY_OLD_CRYPTO_KEY value that was used in the previous deployment, no reencryption occurs. If you specify a key value during deployment that does not match the MXE_SECURITY_OLD_CRYPTO_KEY value, the database is reencrypted.
MXE_SECURITY_OLD_CRYPTO_KEY	Specifies the value for the previous Crypto encryption key that was used for the database.
MXE_SECURITY_CRYPTOX_KEY	Used to encrypt CryptoX fields, including API keys, such as the electronic signature key. For CryptoX encryption, if you specify a MXE_SECURITY_CRYPTOX_KEY value that matches the MXE_SECURITY_OLD_CRYPTOX_KEY value that was used in the previous deployment, no encryption changes occur. CryptoX values cannot be decrypted, and the original value cannot be determined. If you specify a key value in a deployment that does not match the MXE_SECURITY_OLD_CRYPTOX_KEY value, CryptoX values are set to null when encryption is run.
MXE_SECURITY_OLD_CRYPTOX_KEY	Specifies the value for the previous CryptoX encryption key that was used for the database.

The following encryption properties are also supported:

<i>Table 20. Encryption properties</i>	
Encryption property	Description
MXE_SECURITY_CRYPTO_ALGORITHM	The default value is AES.
MXE_SECURITY_CRYPTO_MODE	The default value is CBC.
MXE_SECURITY_CRYPTO_MODULUS	
MXE_SECURITY_CRYPTO_PADDING	The default value is PKCS5Padding.
MXE_SECURITY_CRYPTO_SPEC	The length must be a multiple of 8.
MXE_SECURITY_CRYPTOX_ALGORITHM	The default value is AES.
MXE_SECURITY_CRYPTOX_MODE	The default value is CBC.
MXE_SECURITY_CRYPTOX_MODULUS	
MXE_SECURITY_CRYPTOX_PADDING	The default value is PKCS5Padding.

Table 20. Encryption properties (continued)	
Encryption property	Description
MXE_SECURITY_CRYPTOX_SPEC	The length must be a multiple of 8.

Note: After the database is installed, only the Crypto and CryptoX encryption keys can be changed.

When you configure the database settings for deployment before you activate the application, you can add a value for the **MXE_SECURITY_CRYPTO_KEY** or **MXE_SECURITY_CRYPTOX_KEY** encryption keys. If you do not specify an encryption key secret in the Maximo Manage configuration when you activate, the system automatically generates keys. The system names the secret in the keys by using the following naming convention:

```
<workspaceId>-<appId>-encryptionsecret
```

For more information about how to specify the encryption secret in the Maximo Manage configuration, see [“Adding encryption key secrets” on page 322](#).

Because your database functions only with valid encryption keys, implement the following practices:

- Maintain your encryption keys in a vault or other secure management system for secrets.
- Specify your own values for encryption keys instead of using system-generated values. If you use system-generated values and do not create a backup, you cannot retrieve the keys. Without the keys, you cannot use your database.

Database encryption scenarios

Your deployment scenario determines your encryption and reencryption options. Scenarios include deploying a new database, deploying a previously encrypted database, or changing the encryption keys for deployment.

Deploying a new database

If you deploy a new database, you have two options for database encryption:

Table 21. Encryption for deploying a new database		
Option	Action	Result
Provide your own values for the encryption keys.	Specify the values that you want to use for the Crypto and CryptoX encryption keys when you configure your database during deployment.	The database is then encrypted by using the key values that you provide.
Use system-generated keys.	Do not specify key values when you configure your database.	If you do not specify the keys, a secret is automatically generated that contains the new MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. Later, if you need the keys, you can view the keys in the secret.

Deploying a previously encrypted database

The following scenarios can occur if your deployment includes a database that was previously encrypted. For example, you might be upgrading from IBM Maximo Asset Management 7.6.0.10 or 7.6.1.2, which are not installed on Red Hat OpenShift but can be upgraded to Maximo Manage. You might also be upgrading from a previous Maximo Manage version. The scenarios apply both to a database that previously implemented the default encryption for Maximo Manage or a database that was encrypted by using a different algorithm or keys.

An existing database that used the default encryption for Maximo Manage and no values were provided for the Crypto or CryptoX key

The database for Maximo Manage no longer uses a default set of keys for encryption. If you have an existing database that used the default encryption, provide the **MXE_SECURITY_CRYPTOX_KEY** and **MXE_SECURITY_CRYPTO_KEY** values so that the database can be reencrypted. If your database was previously encrypted by using keys other than the default encryption for Maximo Manage, provide the old **MXE_SECURITY_OLD_CRYPTOX_KEY** and **MXE_SECURITY_OLD_CRYPTO_KEY** encryption keys. As a result, the database can be decrypted and then reencrypted.

When you configure the database, select one of the following options for reencryption:

Note: Reencryption always occurs in this scenario.

<i>Table 22. Encryption for existing databases where no values were provided for the encryption keys.</i>		
Option	Action	Result
Provide your own values for the encryption keys.	<ol style="list-style-type: none"> 1. Enter the MXE_SECURITY_CRYPTOX_KEY and MXE_SECURITY_CRYPTO_KEY values. 2. Do not specify any key values when you configure your database. 	The database is reencrypted by using the values that you specified.
Use system-generated keys.	<ol style="list-style-type: none"> 1. Do not specify any key values when you configure your database. 	The system generates new keys, and the database is reencrypted with the system-generated keys.

An existing database that used values from your own Crypto or CryptoX keys

When you configure the database, select one of the following options for encryption:

Note: The first option is less likely to result in errors. You can change the keys later.

<i>Table 23. Encryption for existing databases where you provided your own values for the encryption keys.</i>		
Option	Action	Result
Do not reencrypt the database.	<ol style="list-style-type: none"> 1. Specify the Maximo Manage security properties, including the MXE_SECURITY_OLD_CRYPTOX_KEY and MXE_SECURITY_OLD_CRYPTO_KEY encryption keys. 2. Specify the same values for the MXE_SECURITY_CRYPTOX_KEY and MXE_SECURITY_CRYPTO_KEY encryption keys. 	Because you specified the same values for the old and new keys, the database is not reencrypted.

Table 23. Encryption for existing databases where you provided your own values for the encryption keys.
(continued)

Option	Action	Result
Reencrypt the database.	<ol style="list-style-type: none"> 1. Specify the Maximo Manage security properties, including the MXE_SECURITY_OLD_CRYPTO_KEY and MXE_SECURITY_OLD_CRYPTOX_KEY encryption keys. 2. Select one of the following options: <ul style="list-style-type: none"> • Specify values for the new MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. • To use system-generated keys, do not specify encryption key values. 	The database is reencrypted.

Changing the encryption keys for deployment

The following table describes the tasks to complete reencrypting the database when you want to change the **MXE_SECURITY_CRYPTO_KEY** and **MXE_SECURITY_CRYPTOX_KEY** encryption keys.

Table 24. Encryption by using new encryption keys

Option	Action	Result
Reencrypt the database by using new encryption keys.	<ol style="list-style-type: none"> 1. Set the MXE_SECURITY_OLD_CRYPTO_KEY and MXE_SECURITY_OLD_CRYPTOX_KEY encryption keys to the values that the database currently uses for the MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. 2. Select one of the following options: <ul style="list-style-type: none"> • Specify values for the new MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. • To use system-generated keys, do not specify values for the MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. 	The database is reencrypted by using the new encryption keys.

Viewing database encryption history

You can view the history for database encryption. If you encrypt or reencrypt the database, a new entry is logged in the syschangetracker table in the Maximo database.

The syschangetracker table contains the following columns:

- PROCESSNAME
- MESSAGE
- CHANGEBY
- CHANGEDATE
- WORKSPACEID
- APPID
- INSTANCENAME
- SYSCHANGETRACKERID
- ROWSTAMP

If you have issues with encryption keys, the record in the syschangetracker table helps to identify whether reencryption occurred and when it occurred. The MESSAGE column includes the first 2 bytes and last 2 bytes of the encryption key. The MESSAGE column can help you identify the encryption keys that are being used.

Disabling automatic generation of encryption keys

By default, if encryption keys are not specified when you activate Maximo Manage with a fresh database, new encryption keys are automatically generated. If you do not want to automatically generate encryption keys, set the **autoGenerateEncryptionKeys** property to false.

About this task

The **autoGenerateEncryptionKeys** property controls whether encryption keys are automatically generated when you activate Maximo Manage with a fresh database. By default, this property is set to `true`, and encryption keys are automatically generated if no value is specified for them. Set the property to `false` if you do not want to generate keys automatically. If you set the **autoGenerateEncryptionKeys** property to `false` and you do not provide encryption keys, deployment fails. You also receive an error message that the property is set to `false` and encryption keys are missing.

If encryption keys are automatically generated, you can easily lose track of them, especially in development and test environments where databases are reused. If you set the **autoGenerateEncryptionKeys** property to `false`, users are forced to enter the key. You are less likely to lose keys that you generate and maintain. Securely store the encrypted keys, for example, by using a password keeper.

This property takes effect when you activate the application by using an API call. Updating the property for an already activated instance does not produce an effect until you make a change that is related to the encryption keys, such as when you delete the keys to trigger reencryption.

Procedure

1. In the Red Hat OpenShift console, from the side navigation menu, click **Administration > Custom Resource Definitions**.
2. On the **CustomResourcesDefinitions** page, select the ManageWorkspace custom resource definition record.
3. On the **CustomResourceDefinition details** page, on the **Instances** tab, select the instance for which you want to disable automatic generation of encryption keys.
4. On the **YAML** tab for the instance, set `spec.settings.deployment.autoGenerateEncryptionKeys` to `false`.
5. Save the custom resource.

Adding encryption key secrets

For a new installation, the encryption secret is used as the key to encrypt the database. For already encrypted databases, the encryption secret is used to restore the encryption keys in the future.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. In the side navigation menu, from the **Suite** application, select **Administration > Workspace**.
3. On the **Manage workspace details** page, click **Actions**, and select **Update configuration**.
4. In the Database connection row, click the edit icon.
5. Click **Show advanced settings** and set the **System managed** toggle to off.
6. Click **Add property** to add a row for an encryption key property.
7. In the Key column, enter the encryption key property.
For example, `MXE_SECURITY_CRYPT0_KEY`.
8. In the Value column, enter the value for the added encryption key property.
9. Click **Apply changes** to save and apply all the configuration changes.

Resetting the Crypto and CryptoX fields in the database

If you lose your Crypto and CryptoX encryption keys, you can run a reset script that clears the Crypto and CryptoX fields in the Maximo Manage database. Then, you can restart the server and reset your data.

About this task

When you run the `resetcryptocryptox.sh` script, the Crypto and CryptoX fields are cleared. After you restart the server, you can review the data that was cleared and reset the values that were cleared, such as properties or API keys.

Procedure

1. In the Red Hat OpenShift Container Platform, from the side navigation menu, click **Workloads > Pods**.
2. On the **Pods** page, open the `maxinst` pod.
3. On the **Terminal** tab, enter `resetcryptocryptox.sh`.
4. Restart the Maximo Manage server.
5. Add any variables, such as API keys or properties that the `resetcryptocryptox.sh` script cleared.
6. Redeploy or reactivate the application.

Reducing system downtime

Configure a database state to reduce Maximo Manage system downtime when you upgrade to Maximo Manage 9.0 and later.

Before you begin

`onlineUpgrade` is supported on Db2 only from Maximo Manage 8.6.2. Support for other databases is provided only after Maximo Manage 9.0.0. The operator displays an error if you attempt to run an online or offline upgrade, and checks if `maxupg` indicates a version older than 8.6.2.

About this task

The `ManageWorkspace` custom resource (CR) has the `spec.settings.db.upgrade.upgradeType` object, which has `enum` property values that correspond to the following states for reduced system downtime:

enum value	Description
<code>regularUpgrade</code>	The regular type of upgrade where the server is down during the entire database upgrade.
<code>onlineUpgrade</code>	Two-phased upgrade with offline phase triggered. The server is down only during the offline phase of the database upgrade.

`onlineUpgrade` reduces the system downtime period and provides better control when the system is down.

Note:

- The upgrade setting affects how the database is updated during version update, as well as when a new add-on is added to Maximo Manage.
- Before you add an industry solution, set the `upgradeType` value to `regularUpgrade`. After the industry solution is installed, set the `upgradeType` to `onlineUpgrade` for future updates.

Procedure

1. Log in to Red Hat OpenShift Container Platform.
2. From the side navigation menu, click **Administration > CustomResourceDefinitions**.
3. On the **CustomResourceDefinitions** page, search for `ManageWorkspace`.
4. Click **ManageWorkspace** and then click the **Instances** tab.
5. Select an instance and click the **YAML** tab.
6. In the `settings.upgrade.failureControl` section, set a value to control the upgrade failure:

retry

The default value for all reduce downtime states.

Retries the upgrade if the target status is not reached for a particular stage of the upgrade.

rollbackForOnlineUpgrade

Once it is set, and the upgradeType is not regular upgrade, the operator attempts to roll back the change.

Once rolled back, the operator does not consider it as a reconcile failure.

It reflects the result in the status, but not in retry.

Use the `settings.upgrade.failureControl` section to control the operator behavior if there are upgrade failures. Set it to `rollbackForOnlineUpgrade` to roll back the online portion of the upgrade or to stop the retrying. It gives you the time to fix the database.

Note: The roll back can be done only for the online portion of the database change. If a failure happens during the offline portion of the upgrade, the process can continue only when the database is fixed, or manually recovered by using a backup of the original state, or at the state when the online upgrade is completed.

7. If you want to run the version of Maximo Manage that you had before the upgrade, here are the following options:

- Starting in Maximo Application Suite 9.1, you can remain in the current state but it has two limitations:
 - If there is a serious issue with the cluster that results in the loss of the **ManageServerBundle** custom resources, the system does not revert as-is.
 - You cannot change the system apart from the `podTemplates` and the replica size of the bundle servers.
- Roll back the Maximo Manage operator. Set the **ManageApp** custom resource version field to the previous Maximo Manage version.
- Retain the already upgraded operator, and set the **ManageWorkspace** custom resource's base and other components to the earlier version.
- Retain the already upgraded operator and set the **ManageWorkspace** custom resource's `spec.settings.deployment.buildTag` to the previous build tag to match the version of Maximo Manage before the upgrade. You can find the build tag from `ImageStream`. Make a note of the build tag and make sure that it is not purged before you start the upgrade. When you are ready to do the upgrade after you fix the database, set the build tag to `latest`.

Starting in Maximo Application Suite 9.1, on failure of the online upgrade stage, if roll back is set for `failureControl`:

- The operator shuts down the servers, rolls back the database, and restarts the servers.
 - The server bundles run with the older version image after the roll back.
8. Optional: To trigger the offline portion of the upgrade, use `toolsAPI` or directly use the Maximo Manage admin console to issue the **start-offline-upgrade.sh** command. Alternatively, directly update the **ManageOfflineUpgrade** custom resource's value from `waiting` to `requested`.

When `onlineUpgrade` is used, after the online portion of the upgrade is completed, the operator creates a `waiting` value in the **ManageOfflineUpgrade** custom resource in the Maximo Manage namespace.

```
spec
  stage: waiting
  status:
```

Starting in Maximo Application Suite 9.1, after the online upgrade is started, and before the offline portion is finished, the changes to the deployment of the server bundles is limited. Scaling the server bundles by changing the replica size is possible only through `podTemplates`.

Results

The stage of the upgrade and any upgrade failures are reported in the DBReady condition of the **ManageWorkspace** custom resource.

In the MAXVAR table, the SEAMLESSUPGRADE value indicates the status of the database during, before, and after the upgrade:

Condition	Value
When the system is started from 862	The value is left to be 0.
When the online portion of the upgrade starts	The value is set to 1.
When the online portion of the upgrade is completed	The value is set to 2. If there is an error in the online upgrade, the value remains at 1.
When the offline portion of the upgrade starts	The value is set to 3.
When the offline portion of the upgrade is completed	The value is set to 0.
When the roll back starts	The value is set to 4. It can be rolled back only if the MAXVAR is set to 1. If the value is 2, you can only go forward or restore the database.
When the roll back is completed	The value is set to 0.

Note:

0

The database is consistent, there is no partial upgrade.

1

The database started the online phase of the upgrade, but it did not complete or failed.

2

The database online upgrade phase finished completed and stayed at this stage.

3

The offline portion of the upgrade started but it did not finish or failed.

4

The roll back or the online portion of the upgrade is in progress.

The upgrade stops and operator reports a failure if the database is of SEAMLESSUPGRADE=1 or SEAMLESSUPGRADE=2, indicating a failed or partially complete upgrade, while the MASIMAGEIDSTART recorded by the database does not match the current image ID. It is an indication that the image might mismatch. Therefore, the operator cannot ensure that it can roll back or upgrade the database, and reports an error.

Troubleshooting database deployment

When you deploy the Maximo Manage application, issues that are related to the database deployment might cause the deployment process to fail.

To prevent or solve some database configuration issues that might result in issues during the Maximo Manage deployment, follow these practices:

- Review the database settings that you used to configure Maximo Manage. If something is wrong with the values, edit your database configuration on the **Manage Status** page and activate it again.
- Ensure that the JDBC URL format and details are correct and are using the correct JDBC protocol, database URL, name, and port. Test your connection to the database by using a database tool.
- If you are using the SSL-enabled database and port, verify that you selected the SSL option.
- To ensure that your database is ready, verify that you are using the correct username, password, schema name, tablespace name, and index tablespace name.

- Do not skip preparation steps. For example, the deployment process for Maximo Manage databases does not work if you skip the step to set its table organization to be row-based instead of column-based.
- If you plan to deploy languages, ensure that your database is properly configured. Set the configuration to support Maximo Manage or SQL for the language that is used for the databases.
- Confirm that you used the correct database certificate for the database in the Maximo Manage configuration and that the certificate is not expired.
- If you are deploying Maximo Manage and reusing a database that was previously used by Maximo Manage, pass the encryption keys that are saved from the previous deployment. Otherwise, the process generates new keys, which do not match the ones that are used to encrypt the database, and the deployment process does not proceed.
 - For example, you deactivated Maximo Manage and deleted the application. Now, you are redeploying the Maximo Manage application and activating it again, but you are still using the previously deployed Maximo Manage database.
- To deploy multiple Maximo Manage components at the same time, if you are using IBM Db2 Warehouse databases or other Db2 versions, increase the **APPHEAPSZ** value to at least 16384 to avoid failures during the database deployment. Alternatively, you can deploy fewer components simultaneously by deploying Maximo Manage components one at a time.
- Save your database in the same location as Maximo Manage. If the location of your database is different, the chances of database failures during the Maximo Manage deployment due to connection issues and latency are higher.

During the Maximo Manage deployment, database update interruptions might occur due to a situation, such as lost connection. The Maximo Manage operator reconciliation process usually resolves the problem.

If the failure persists, make sure that your database connection is working correctly. Before you remove, re-create the database, and retry the Maximo Manage deployment, you can try the following procedures:

1. Select the **Bypass upgrade version** checkbox and then reactivate the Maximo Manage application. If you select this option, the deployment can run, regardless of the current version of the installed database. It skips version validations that might be preventing the deployment to progress.
2. If the failure persists, review the log files to assess the situation and note the script number for the failed deployment. To review the log files, complete the following steps:
 - a. In the Red Hat OpenShift console, from the side navigation menu, click **Workloads > Pods**.
 - b. From the **Project** list, select the Manage project.
 - c. On the **Pods** page, select the administrative pod. The administrative pod includes manage-maxinst in the pod name, as shown in the following example:


```
env-managedev-manage-maxinst-7fd3c77492-kmj6z
```
 - d. To review the logs, on the **Pod details** page, click the **Logs** tab. If you want to view all the logs, click the **Terminal** tab and view the logs in the `/opt/IBM/SMP/maximo/tools/maximo/log` directory.
 - e. After you review the log files, manually correct the problems and run the database installation again.

Deployment of industry solutions and add-ons

When you deploy Maximo Manage, you can also include industry solutions and add-ons for deployment.

Before you select any Maximo Manage industry solutions or add-ons for deployment, it is important that you check the system requirements. For more information, see [system requirements](#). Select the IBM Maximo Application Suite version from the list. Then, see the **IBM Maximo Application Suite - Manage Application Compatibility Report** to validate that the selected components can be deployed together.

You can deploy the following industry solutions and add-ons with Maximo Manage:

Industry solutions

- Maximo Aviation
- Maximo Civil Infrastructure
- Maximo Nuclear
- Maximo Oil & Gas
- Maximo Transportation
- Maximo Utilities

Add-ons

- Maximo Asset Configuration Manager
- Maximo Asset Investment Planning
- Maximo Connector for Oracle Applications
- Maximo Connector for SAP Applications
- Maximo Health, Safety and Environment
- Maximo Service Provider
- Maximo Spatial
- **Customer-managed** Maximo Connector for Workday Applications
- Maximo Anywhere
Available only in Maximo Application Suite 8.8 and earlier. For more information, see [Maximo Anywhere documentation](#).
- **Customer-managed** Maximo Connector for Envizi
- **Customer-managed** Maximo Connector for TRIRIGA
- Maximo IT
- Reliability Strategies

Maximo Health

You can deploy Maximo Health as part of Maximo Manage to automatically connect and share data between both applications. When deployed as part of Maximo Manage, all deployment settings are shared with Maximo Health.

Note: If you deploy Maximo Health as part of Maximo Manage, you cannot deploy the Maximo Health stand-alone application on the same Maximo Application Suite instance or the other way around.

Maximo IT

You can deploy Maximo IT as part of Maximo Manage to automatically connect and share data between both applications. When deployed as part of Maximo Manage, all deployment settings are shared with Maximo IT.

Language support

Maximo Manage is available in a number of languages. You can specify languages when you activate the application.

The following settings control the languages that are displayed in the user interface:

Browser language setting

If content is available for the language that users select for their browser, the application shows that language. If the language that is selected in the browser is not available, the application displays the base language.

Preferred language setting

In IBM Maximo Application Suite, users can select a preferred locale, language, and time zone when they set up their profile. If you change the user's preferred language setting in Maximo Manage, future synchronizations do not overwrite the setting from Maximo Application Suite.

Application language settings

When Maximo Manage is deployed, you can select a base language, which is the default language for the application. By default, the base language is English. If you select a different language, the language is permanently set as the base language, and you cannot change it in subsequent deployments. If the base language is not specified when the application is deployed, the default value is used instead. If the default value is used, you can change the base language when you redeploy.

The base language can be changed only once. After it is changed, it cannot be changed back to English or any other language.

You can select other languages when you activate the application. You cannot remove languages after the application is activated.

If you use IBM Db2 and plan to install a language other than English, select the checkbox for **Db2 Vargraphic** when you configure your database.

Although Maximo Application Suite does not support bidirectional languages, Maximo Manage provides language support for bidirectional languages.

To add a language, on the **Activation** page, in the **Base** field and the **Additional** field, select the language code.

Note: If you want to use the database to deploy Maximo Manage only in English, you can choose Latin1_General_CI_AI collate for SQL Server. If you plan, for example to deploy Maximo Manage in Japanese, the choice is Japanese_CI_AI. For more information about which collate to use for each language, see [SQL Server documentation](#). Although SQL Server supports Unicode, Maximo Manage does not support SQL Server, so you cannot set it. Therefore, you cannot deploy Maximo Manage with multiple languages in case those languages require different collates, for example: simplified Chinese and Russian.

The following languages are supported:

- Arabic (AR)
- Brazilian Portuguese (PT-BR)
- Croatian (HR)
- Czech (CS)
- Danish (DA)
- Dutch (NL)
- English (EN)
- Finish (FI)
- French (FR)
- German (DE)
- Hebrew (HE)
- Hungarian (HU)
- Italian (IT)
- Japanese (JA)
- Korean (KO)
- Norwegian (NO)
- Polish (PL)
- Simplified Chinese (ZH_CN)
- Slovak (SK)
- Slovenian (SL)

- Spanish (ES)
- Swedish (SV)
- Traditional Chinese (ZH_TW)
- Turkish (TR)

Note: Maximo Manage supports Db2, Oracle Database, and Microsoft SQL Server. However, it cannot support Microsoft SQL Server when the database collation is set to Turkish.

Because Maximo Manage does not handle Turkish *i* characters that are stored in Db2 and Oracle Database, those characters are not supported.

Server bundle overview

The application can be deployed in one or more workloads, which are called server bundles. A server bundle isolates the workload processes so that they can be independently managed.

Server bundles can be independently scaled and managed based on your needs. For example, the ui server bundle allows users to access the user interface by using a web browser. A crontask server bundle allows users to run background jobs. When a ui workload and crontask workload are deployed as two separate server bundles, the CPU and memory that are used by these workloads do not affect each other.

A server bundle is composed of one or more pods. Each server bundle can be individually configured for its additional server configuration, jvm options, and bundle level properties. Different server bundles can associate to the same or different persistent volumes.

All the server bundles connect to the same database.

One server bundle that you deploy is designated as the default server. The server bundle that you designate as the default is used to establish the default URL, or route that links to the Manage application in the workspace.

When you create your server bundles, remember the following requirements and behaviors:

- You can configure multiple server bundles based on your workload need. The bundle type determines the capabilities of the server. For example, a cron bundle cannot serve UI requests but only to run crontasks.
- You can have multiple server bundles of the same bundle type if grouping of workloads for different configurations is needed. For example, you can define two **all** server bundles where one server bundle is used for the first group of users and one for the second group.
- If only one server bundle is specified, you must specify **all** as the bundle type. If you do not specify **all** as the type for the default, the deployment might appear successful. However, the application does not function properly.
- The default server bundle must have a bundle type of either **all** or **ui**. If you do not specify one of these types for the default, the deployment might appear successful, but the application does not function properly. If your default server bundle is a **ui** bundle, you must also have a **mea** server bundle because the **mea** server bundle is needed for user synchronization.
- If you create only one server bundle and do not specify a default, that bundle is selected as the default.
- If you create multiple server bundles and do not specify a default, a server bundle that has the **all** type is randomly selected as the default server.

When you deploy the application, you specify parameters that control deployment of the server bundle pods and their associated services and routes.

You can specify the following server bundle parameters to configure the deployment:

Name

The user-defined name for the server bundle. You can also specify whether the server bundle is the default server bundle.

Pod count

The number of pods to deploy for the server bundle.

Type

The type of the server bundle. The following six bundle types are available:

all

Includes ui, cron, mea, and report bundle types.

ui

The user interface components.

cron

The components that are needed for cron tasks.

mea

The enterprise web services API. This server bundle type is needed for user synchronization.

report

The components for reports.

standalonejms

Specify `standalonejms` for the server bundle type to add a Java Message Service (JMS) server bundle during deployment of Maximo Manage. For more information, see [“Configuring JMS servers for Maximo Manage”](#) on page 335.

Additional server properties

As an option, you can add the following properties:

Route subdomain

If you specify a subdomain, it is added as a prefix to the main domain when the route is created for the server bundle. For example, if you specify `maximoui` as a subdomain, your full domain is similar to the following example: `maximoui.workspaceid.manage.domain.com`.

Additional server configurations

To configure the WebSphere® Application Server Liberty application server, you can provide custom parameters, as shown in the following example:

```
server-custom.xml: |-
  <!-- Enable features -->
  <featureManager>
    <feature>jsp-2.3</feature>
  </featureManager>

  <!-- To access this server from a remote client add a host attribute to the
  following element, e.g. host="*" -->
  <httpEndpoint id="defaultHttpEndpoint"
    httpPort="9080"
    httpsPort="9443" />
  <!-- Automatically expand WAR files and EAR files -->
  <applicationManager autoExpand="true"/>
```

Bundle-level properties

To add Manage system properties to specific server bundles, you can specify property names and values when you deploy the application.

Provisioning storage

Maximo Manage supports both ephemeral storage and persistent storage.

Ephemeral storage is transient and does not persist when you redeploy or reactivate applications.

Persistent storage is used for data, such as file attachments, that you want to persist throughout multiple deployments.

Ephemeral storage

By default, ephemeral storage is allocated for the administrative and server bundle builds when Maximo Manage is deployed. Administrators can modify the requests and also limit the amount of ephemeral storage that is provisioned.

By default, the administrative and server bundle builds request 30 GB of storage, and the builds assign a limit of 100 GB of storage, as shown in the following example:

```

ephemeralStorage:
  requests:
    adminBuild: 30Gi
    serverBundleBuild: 30Gi
  limits:
    adminBuild: 100Gi
    serverBundleBuild: 100Gi

```

Administrators can configure different requests and limits by specifying different values in the custom resource for the application.

For more information, see [Understanding ephemeral storage](#) in the Red Hat OpenShift documentation.

Configuring ephemeral storage

To configure ephemeral storage, you can add parameters to the ManageWorkspace custom resource definition by using the Red Hat OpenShift administrative console.

About this task

By default, the administrative and server bundle builds request 30 GB of storage and assigns a limit of 100 GB of storage. You can configure different requests and limits by specifying different values in the custom resource for the application.

Procedure

1. In the Red Hat OpenShift console, from the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
2. On the **CustomResourcesDefinitions** page, select the ManageWorkspace custom resource definition record.
3. On the **CustomResourceDefinition details** page, on the **Instances** tab, select the instance for which you want to specify ephemeral storage.
4. On the **YAML** tab for the instance, in the **spec.settings.deployment.ephemeralStorage** section, specify values for the storage request and storage limits, as shown in the following example:

```

"deployment": {
  "buildTag": "latest",
  "mode": "up",
  "persistentVolumes": [],
  "serverBundles": [
    {
      "name": "default",
      "bundleType": "all",
      "isDefault": true,
      "isUserSyncTarget": true,
      "isMobileTarget": true,
      "routeSubDomain": "all",
      "replica": 1
    }
  ],
  "ephemeralStorage": {
    "limits": {
      "adminBuild": "100Gi",
      "serverBundleBuild": "100Gi"
    },
    "requests": {
      "adminBuild": "30Gi",
      "serverBundleBuild": "30Gi"
    }
  }
},

```

5. Save the custom resource.

Persistent storage

You use persistent volume storage for data, such as file attachments, that you want to persist across multiple deployments. Administrators must provision the persistent storage systems that are needed for your applications.

Creating persistent volumes gives you greater control over your storage if you want to increase the amount of storage or to add existing files. Creation of external file storage systems for data and objects is outside the scope of the Maximo Manage deployment. To provision persistent storage for Maximo Manage, you can use the Kubernetes persistent volume (PV) framework.

For more information about persistent storage, see [Understanding persistent storage](#) in the Red Hat OpenShift documentation.

Set up persistent volume storage before you deploy the application.

The following guidelines describe how to set up Network File System (NFS) storage. A similar process can be applied to other types of persistent storage.

1. Deploy an NFS server. Then, ensure that it is enabled and started.
2. Export an NFS share and ensure that it is accessible.
3. Configure the storage.
 - a. Deploy a persistent volume provisioner.
 - b. Define a storage class. Record the storage class name. When you deploy Maximo Manage, you must supply the class name for the persistent volume storage that you created.
 - c. Create a volume for the storage. Record the name of the volume so that you can use that name for the persistent volume claim when you deploy Maximo Manage.
 - d. Set the **persistentVolumeReclaimPolicy** for your storage volume to `Retain` to ensure that the data is not deleted when a persistent volume claim is released. Then, you can reclaim the content manually. For example, for test environments, you can use the `Retain` option to manually reclaim and reuse storage assets. For more information, see [Persistent Volumes](#) in the Kubernetes documentation.

You can configure persistent volume claims for Maximo Manage deployments on different platforms, such as IBM Cloud . For more information about Persistent Volume Claims configuration, see [“Configuring persistent volume claims”](#) on page 332.

You can configure a persistent volume claim through the deployment process or by manually updating the Maximo Manage workspace custom resource (CR).

A persistent volume claim is dynamically provisioned in Maximo Manage when you declare a persistent volume and do not specify a persistent volume claim. The claim is dynamically provisioned by using the storage class that you configured for the volume.

If necessary, you can configure a persistent volume claim for a specific server bundle. The mount path that you specify for a server bundle overrides any default mount path that was specified for the deployment.

After deployment, if you want to adjust the size of a persistent volume claim, you can specify a new value for the size and redeploy the application.

If the amount of storage that you need exceeds the available amount for the storage volume and you want to increase storage capacity after deployment, create a new volume for the storage. Then, you can move your data to the new volume. And then, you can redeploy the application and provide the persistent volume claim parameters for the new storage.

Configuring persistent volume claims

You can configure a persistent volume claim (PVC) through the deployment process or by manually updating the `ManageWorkspace` custom resource.

Before you configure a PVC, ensure that you prepared a persistent storage volume and can provide the volume name, class name, mount path, and size for your PVC.

If you declare a persistent volume, you can choose not to specify a PVC when you deploy Maximo Manage. If you do not specify a persistent volume claim, it is dynamically provisioned by using the storage class that you configured for the volume.

If you configure a PVC when you deploy the application, the storage in the volume that you provisioned is available to all server bundles in the workspace. You can also configure persistent volume claims for specific server bundles in a deployment. If you configure a PVC for a server bundle, the mount path that you specify for the server bundle PVC overrides the path that you specify for the deployment.

Configuring persistent volume claims for server bundles

To configure a persistent volume claim (PVC) for a specific server bundle, you can add parameters to the ManageWorkspace custom resource by using the Red Hat OpenShift administrative console. The mount path that you specify for the server bundle PVC overrides any default path that you configured for a persistent volume during deployment.

Procedure

1. In the Red Hat OpenShift console, from the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
2. On the **CustomResourcesDefinitions** page, select the ManageWorkspace custom resource definition record.
3. On the **CustomResourceDefinition details** page, on the **Instances** tab, select the instance for which you want to create persistent storage.
4. On the **YAML** tab for the instance, in the **spec.settings.deployment.persistentVolumes** section, specify values for the persistent volume that is used for Maximo Manage. You can also specify the following properties:

<i>Table 25. Persistent volume properties</i>	
Property name	Description
accessModes	The corresponding storage access mode for your persistent volume. Typically, it is either <code>ReadWriteMany</code> or <code>ReadWriteOnce</code> .
pvcName	The persistent volume claim name.
size	The persistent volume claim size.
storageClassName	The persistent volume claim's storage class.
mountPath	The local mounted path for your persistent volume.
volumeName	If you have a preconfigured persistent volume, you can specify it. If you do not specify this property, then Maximo Manage tries to dynamically provision a persistent volume based on the chosen storage class that is provisioned.

Note: You can provide the parameters from Table 1 in the next step.

The following example shows how to add persistent volume parameters for deployment.

```
spec:
  settings:
    deployment:
      persistentVolumes:
      - accessModes:
        - ReadWriteMany
        pvcName: my_pvc_name
        mountPath: /my_pv_path
        size: 20Gi
        storageClassName: my-storage-class-name
```

5. In the **spec.settings.deployment.serverBundles** section, specify the PVC name and mount path for a server bundle, as shown in the following example.

```
spec:
  settings:
  ...
  serverBundles:
    - name: server_bundle_name
      bundleType: bundle_type
      pvcMount:
        - pvcName: my_server_bundle_pvc_name
          mountPath: /mount path for the server bundle storage
```

6. Save the custom resource.

Backing up and restoring Maximo Manage

To avoid disruptions to service and unplanned downtime, create regular Maximo Manage backups that you can restore as needed.

For more information, see [Back up and restore the Manage application in Maximo Application Suite](#).

Adding images to an external registry

To add build images for Maximo Manage to an external registry, configure the custom resource for the Maximo Manage workspace.

Before you begin

Before you can configure Maximo Manage to store build images in an external registry, you must create a secret that is based on Docker credentials for the external registry repository.

1. Log in to the Red Hat OpenShift console.
2. In the workspace where the Maximo Manage custom resource is located, use the command line interface or the Red Hat OpenShift dashboard to create an image pull secret.

For example, you can use the command line interface to create the secret:

```
kubectl create secret docker-registry regcred
--docker-server=your_docker_registry_server
--docker-username=your_your_docker_registry_account
--docker-password=your_your_docker_registry_api_key
```

About this task

When Maximo Manage is deployed or updated, the build process generates an administrative image and a server bundle image for each bundle type. Bundle types are specified in the custom resource for the workspace. Images are stored in an internal registry.

Sometimes, you might want to store the build images in an external repository. For example, you might want to run security scans on images that have customizations, or you might want to test builds independently before you deploy the application in Maximo Application Suite. Administrators can configure Maximo Manage to store images in an external repository.

To configure Maximo Manage to add build images to an external registry, you must add the **registryForGeneratedImages** properties to the **spec.settings** section of the custom resource and specify values for the following fields:

registryPath

The path to the repository where you want to store the image.

secretName

The secret that contains the credentials for accessing the repository.

Procedure

1. In the Red Hat OpenShift console, from the side navigation menu, click **Administration > Custom Resource Definitions**.
2. Open the Maximo Manage workspace and then on the **Instances** tab, click the instance for your external registry.
3. On the **YAML** tab in the **spec.settings** section, add the following lines:

```
registryForGeneratedImages:  
  registryPath:  
  secretName:
```

4. Specify values for the **registryPath** and **secretName** fields, as shown in the following example:

```
spec:  
  settings:  
    ...  
    languages:  
      baseLang:  
      secondaryLangs:  
    registryForGeneratedImages:  
      registryPath: your_docker_registry_server  
      secretName: your-image-pull-secret-for-the-repository
```

5. Save the custom resource.

Results

The Maximo Manage workspace operator automatically picks up the change to the custom resource. A build is started, and the new images are pushed to the external registry. You can use the images to deploy Maximo Manage on a local computer.

If you no longer want to push build images to an external repository, remove the lines of code that you added for the external registry from the custom resource for the workspace.

Configuring JMS servers for Maximo Manage

To add a JMS server bundle, specify `standalonejms` for the server bundle type when Maximo Manage is deployed.

About this task

This configuration creates a single-pod Liberty JMS server.

By default, ephemeral storage for the JMS server uses the `/jmsstore` directory in the container. Generally, users require persistent storage for messaging. To provide persistent storage, provision a persistent storage volume for Maximo Application Suite. Then, configure a persistent volume claim for the JMS bundle when you deploy the application. You can specify a different directory for the JMS storage when you set the server configuration by using the additional server configuration parameters for the bundle.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. On the **Suite administration** page, select **Workspaces** from the side navigation menu and then select the **Manage** tile.
3. On the **Manage workspace details** page, click **Actions**, and then select **Update configuration**.
4. In the Server bundles row, click the edit icon.
5. In the Server bundles section, set **System managed** to off. Then, click **Add bundle**.
6. In the **Name** column, enter a name for the bundle, such as `jmsserver`.
7. In the **Type** column, select **standalonejms**.

8. Configure the queues.

- a) In the **Additional properties** column for your JMS server bundle, click **View**.
- b) Optional: Set a different route subdomain.
- c) In the **Additional server config** field, use XML to specify default and custom queues.

You can specify a queue as:

- Outbound sequential
- Outbound continuous
- Inbound sequential
- Inbound continuous

Note: You can specify a queue as continuous or sequential in , by selecting **Add/Modify Queues** from the External Systems application. Specify a queue as continuous by leaving the **Sequential** checkbox deselected. Specify a queue as outbound by leaving the **Inbound** checkbox deselected.

```
<?xml version="1.0" encoding="UTF-8"?>
<server description="new server">

  <!-- Enable features -->
  <featureManager>
    <feature>wasJmsSecurity-1.0</feature>
    <feature>wasJmsServer-1.0</feature>
  </featureManager>
  <applicationManager autoExpand="true"/>
  <wasJmsEndpoint host="*" wasJmsSSLPort="7286" wasJmsPort="7276" />
  <messagingEngine>
    <fileStore path="/jmsstore"/>
    <queue id="sqoutbd" maintainStrictOrder="true" maxMessageDepth="100000"
failedDeliveryPolicy="KEEP_TRYING" maxRedeliveryCount="-1"/>
    <queue id="sqinbd" maintainStrictOrder="true" maxMessageDepth="200000"
failedDeliveryPolicy="KEEP_TRYING" maxRedeliveryCount="-1"/>
    <queue id="cqinerrbd" maxMessageDepth="100000" exceptionDestination="cqinerrbd"/>
    <queue id="cqinbd" maxMessageDepth="100000" exceptionDestination="cqinerrbd"/>
    <queue id="cqoutererrbd" maxMessageDepth="100000"
failedDeliveryPolicy="KEEP_TRYING"/>
    <queue id="cqoutbd" maxMessageDepth="100000" exceptionDestination="cqoutererrbd"/>
    <queue id="notferrbd" maxMessageDepth="100000"
failedDeliveryPolicy="KEEP_TRYING"/>
    <queue id="notfbd" maxMessageDepth="100000" exceptionDestination="notferrbd"/>
  </messagingEngine>
</server>
```

Note: The `fileStore` path can be modified to use a different directory to store JMS messages. If you are using persistent volumes, make sure that the `fileStore` path is configured correctly in the persistent volume directory. For example, if your persistent volume is mounted on the `/nfs` directory, then the `fileStore` path might be `/nfs/jmsstore`.

- d) Click **Save**.

The queues are established in the JMS server.

9. Update the Liberty server that runs the Maximo server bundles. On the Maximo server, update the `remoteServerAddress` with the following information before you use the XML that is outlined in the following steps.

<InstanceId>

The name of your Maximo Application Suite instance.

<workspaceId>

The name of your workspace.

<serverbundlename>

The name of your JMS server bundle, for example, `jmsserver`.

Note: Your `remoteServerAddress` must be in the following format:

```
remoteServerAddress="<InstanceId>-<workspaceId>-<serverbundlename>.mas-<InstanceId>-
manage.svc:7276:BootstrapBasicMessaging"
```


For example, if your Maximo Application Suite server URL is `https://main.home.ivt810x-01.ibmماسivt.com/`, your environment is configured with specific values.

- Your instance ID is `ivt810x-01`.
- Your workspace ID is `main`.
- Your JMS server bundle name is `jmsserver`.

The resulting `remoteServerAddress` address is:

```
remoteServerAddress="ivt810x-01-main-jmsserver.mas-ivt810x-01-manage.svc:7276:BootstrapBasicMessaging"
```

10. If you have a single server (all) bundle, prepare the other Maximo Manage server bundles for the JMS queues.

- a) In the **Additional properties** column for your JMS server bundle, click **View**.
- b) In the **Additional server config** field, specify features for the queues that you want to add and the `jmsActivationSpec` details, as shown in the following example.

```
<?xml version='1.0' encoding='UTF-8'?>
<server description="new server">
<featureManager>
<feature>jndi-1.0</feature>
<feature>wasJmsClient-2.0</feature>
<feature>jmsMdb-3.2</feature>
<feature>mdb-3.2</feature>
</featureManager>

  <jmsQueueConnectionFactory jndiName="jms/maximo/int/cf/
intcf" connectionManagerRef="mifjmsconfact"><properties.wasJms
remoteServerAddress="InstanceId-<workspaceId>-<serverbundlename>.mas-<InstanceId>-
manage.svc:7276:BootstrapBasicMessaging"/></jmsQueueConnectionFactory>
  <connectionManager id="mifjmsconfact" maxPoolSize="20"/>

  <jmsQueue jndiName="jms/maximo/int/queues/sqout"><properties.wasJms
queueName="sqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/sqin"><properties.wasJms
queueName="sqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqin"><properties.wasJms
queueName="cqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqinerr"><properties.wasJms
queueName="cqinerrbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqout"><properties.wasJms
queueName="cqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqoutererr"><properties.wasJms
queueName="cqoutererrbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notf"><properties.wasJms
queueName="notfbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notferr"><properties.wasJms
queueName="notferrbd"/></jmsQueue>

  <jmsActivationSpec id="maximo-all/mboejb/JMSContQueueProcessor-1"
maxEndpoints="5"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqin"
maxConcurrency="5" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximo-all/mboejb/JMSContQueueProcessor-2"
maxEndpoints="1"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqinerr"
maxConcurrency="1" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximo-all/mboejb/JMSContOutQueueProcessor-1"
maxEndpoints="5"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqout"
maxConcurrency="5" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximo-all/mboejb/JMSContOutQueueProcessor-2"
maxEndpoints="1"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqoutererr"
maxConcurrency="1" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
</server>
```

Remember: You can specify `jmsActivationSpec` for both outbound and inbound continuous queues and outbound and inbound continuous error queues.

11. If you have a `mea` bundle, prepare the other Maximo Manage server bundles for the JMS queues.

- a) In the **Additional properties** column for your JMS server bundle, click **View**.
- b) In the **Additional server config** field, specify features for the queues that you want to add and the `jmsActivationSpec` details, as shown in the following example.

```
<?xml version='1.0' encoding='UTF-8'?>
<server description="new server">
  <featureManager>
    <feature>jndi-1.0</feature>
    <feature>wasJmsClient-2.0</feature>
    <feature>jmsMdb-3.2</feature>
    <feature>mdb-3.2</feature>
  </featureManager>

  <jmsQueueConnectionFactory jndiName="jms/maximo/int/cf/
intcf" connectionManagerRef="mifjmsconfact"><properties.wasJms
remoteServerAddress="<InstanceId>-<workspaceId>-<serverbundlename>.mas-<InstanceId>-
manage.svc:7276:BootstrapBasicMessaging"/></jmsQueueConnectionFactory>
  <connectionManager id="mifjmsconfact" maxPoolSize="20"/>

  <jmsQueue jndiName="jms/maximo/int/queues/sqout"><properties.wasJms
queueName="sqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/sqin"><properties.wasJms
queueName="sqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqin"><properties.wasJms
queueName="cqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqinerr"><properties.wasJms
queueName="cqinerrbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqout"><properties.wasJms
queueName="cqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqoutererr"><properties.wasJms
queueName="cqoutererrbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notf"><properties.wasJms
queueName="notfbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notferr"><properties.wasJms
queueName="notferrbd"/></jmsQueue>

  <jmsActivationSpec id="maximomea/mboejb/JMSContQueueProcessor-1"
maxEndpoints="5"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqin"
maxConcurrency="5" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximomea/mboejb/JMSContQueueProcessor-2"
maxEndpoints="1"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqinerr"
maxConcurrency="1" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximomea/mboejb/JMSContOutQueueProcessor-1"
maxEndpoints="5"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqout"
maxConcurrency="5" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
  <jmsActivationSpec id="maximomea/mboejb/JMSContOutQueueProcessor-2"
maxEndpoints="1"><properties.wasJms destinationLookup="jms/maximo/int/queues/cqoutererr"
maxConcurrency="1" maxBatchSize="20" connectionFactoryLookup="jms/maximo/int/cf/
intcf"/></jmsActivationSpec>
</server>
```

Remember: As you can see in the example, you must specify `jmsActivationSpec` for both outbound and inbound continuous queues and outbound and inbound continuous error queues.

12. If you have report, cron, or ui bundles, configure the server.

- a) In the **Additional properties** column for your JMS server bundle, click **View**.
- b) In the **Additional server config** field, specify features for queues that you want to add, as shown in the following example:

```
<?xml version='1.0' encoding='UTF-8'?>
<server description="new server">
  <featureManager>
    <feature>jndi-1.0</feature>
    <feature>wasJmsClient-2.0</feature>
    <feature>jmsMdb-3.2</feature>
    <feature>mdb-3.2</feature>
  </featureManager>

  <jmsQueueConnectionFactory jndiName="jms/maximo/int/cf/
intcf" connectionManagerRef="mifjmsconfact"><properties.wasJms
remoteServerAddress="<InstanceId>-<workspaceId>-<serverbundlename>.mas-<InstanceId>-
manage.svc:7276:BootstrapBasicMessaging"/></jmsQueueConnectionFactory>
```

```

<connectionManager id="mifjmsconfact" maxPoolSize="20"/>
  <jmsQueue jndiName="jms/maximo/int/queues/sqout"><properties.wasJms
queueName="sqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/sqin"><properties.wasJms
queueName="sqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqin"><properties.wasJms
queueName="cqinbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqinerr"><properties.wasJms
queueName="cqinerrbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqout"><properties.wasJms
queueName="cqoutbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/cqouter"><properties.wasJms
queueName="cqouterbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notf"><properties.wasJms
queueName="notfbd"/></jmsQueue>
  <jmsQueue jndiName="jms/maximo/int/queues/notferr"><properties.wasJms
queueName="notferrbd"/></jmsQueue>
</server>

```

Remember: You can specify the continuous and sequential type for both outbound and inbound queues.

13. Click **Save**.

14. After you complete the workspace updates, activate the changes.

What to do next

After you configure the stand-alone JMS servers, you can log in to Maximo Manage and configure your JMS queues. For more information, see [JMS queue configuration](#).

Scaling JMS servers

By default, the server bundle for JMS servers is a single pod or server bundle in Maximo Manage. You can add server bundles if you need them.

Procedure

Create more instances of the standalonejms server bundle and the mea server bundle and choose one of the following options to add JMS server bundles:

- a) Create another instance of the standalonejms server bundle.
- b) Create the same set of queues that existed in your first server bundle.
- c) Create another instance of the mea server bundle that uses the new standalonejms server bundle.
- d) To distribute the processing load across multiple servers, you can configure an external system to use either of the mea cluster URLs to send the messages.
- e) You can add more mea server instances for the same standalonejms server. This process effectively creates message-driven beans (MDB) that can consume messages in parallel from the JMS server.

Deployment planning worksheet

Use the planning worksheet to record the information that you need during the configuration and deployment of Maximo Manage. Information that can be recorded includes database parameters, server bundle files, and customization archive location.

The following table contains settings for deployment.

Setting	Default	Your value
Db2 hostname		
Db2 port		
Oracle hostname		
Oracle port		

<i>Table 26. Settings for deployment (continued)</i>		
Setting	Default	Your value
SQL Server hostname		
SQL Server port		
Db2 Warehouse hostname		
Db2 Warehouse port		
Maximo Manage database name		
Maximo Manage database username		
Maximo Manage password		
Database schema name	maximo	
Database table space	MAXDATA	
Database index space	MAXINDEX	
Industry solution and add-on components	For more information, see “Deployment of industry solutions and add-ons” on page 326	
Languages	EN For more information, see Language support	
Server bundles	All For more information, see Server bundles	
Customization archive location	For more information, see Customizing the application	
Mode	On	
Persistent volume claim (PVC) name	For more information, see Persistent storage	
PVC volume name		
PVC size		
PVC mount path		
PVC storage class name		
Build tag	Latest	
Import certificates		
Server time zone	GMT	

Deploying and activating Maximo Manage

As a system administrator, you deploy and activate Maximo Manage in Maximo Application Suite. Select the application version, connect to the configured database, and apply the Java database connectivity (JDBC) configurations to deploy and activate the application.

Before you begin

- Confirm that the database is configured and you have its information available, including the correct JDBC URL in the correct format. Review the database encryption and choose the database encryption strategy so that you understand how to avoid losing the encryption keys after the database is first encrypted. For more information, see [Database encryption](#).
- Confirm that your customization archive is properly configured for your Maximo Manage deployment.
- If needed, configure persistent storage. For more information, see “[Provisioning storage](#)” on page 330.

To record the information that you need during the deployment process, you can also use the deployment planning worksheet. For more information, see “[Deployment planning worksheet](#)” on page 339.

Before you select industry solutions or add-ons, review the compatibility matrix to validate that the selected components are compatible to be deployed together. For more information, see [compatibility matrix](#)

About this task

To enable Maximo Manage, first, you deploy Maximo Manage in Maximo Application Suite by selecting the application version. This part of the process usually takes a few minutes to complete.

Then, when you activate the application, you configure Maximo Manage for use in Maximo Application Suite.

After you confirm the Maximo Manage activation, the main deployment process starts. This process takes a couple of hours or more to finish.

Deploying Maximo Manage

You deploy Maximo Manage in the Maximo Application Suite user interface.

Procedure

1. On side navigation menu, in the **Suite > Administration** page, click **Catalog** and on the **Applications** tab, select the **Manage** tile.
2. On the **Setup** tab, review the information.
 - a) If you plan to co-deploy industry solutions or add-ons with Manage, click **Show all**. Next, click the plus icon on the respective tile of each industry solution or add-on that you want to co-deploy to reserve the necessary AppPoints for them. For more information, see [AppPoints](#).
 - b) Click **Continue**.
3. Select your application update method.

Note: If you choose an automatic upgrade strategy, required downtime might occur before you have a chance to take preparatory action, such as reviewing changes or backing up the database. For production Maximo Manage deployments, set the approval mode to **Off**.

With manual updates, you must trigger the updates yourself. Review the changes, run backups of the Maximo Manage configuration and custom resource definitions, schedule the update, and communicate the scheduled downtime to users.

- a) Select an update method.

To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.

To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.

- b) Subscribe to a channel by selecting a version from the list.
For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
- c) Click **Subscribe to channel <version>**

Wait for the deployment to be completed on the **Manage** page.

Activating Maximo Manage

Before you can grant users access and start working with Maximo Manage, you must activate the application. Activating Maximo Manage triggers the second phase of the application's deployment.

What to do next

- Complete the post-deployment configurations, such as giving permission to the Maximo Manage admin user to connect to Maximo Manage. After the Maximo Manage admin user is synchronized, log out and log in again to Maximo Application Suite as the admin user and access Maximo Manage.
- When new versions are available, system administrators can update the deployed application. To update to a new version, in Maximo Application Suite, in **Suite administration**, select **Applications** from the side navigation menu, and click **Update available for Manage**.

When updates are required, system administrators can also reconfigure and update initial implementations on the **Manage workspace details** page. To reconfigure and update changes, select **Workspaces** from the side navigation menu, and click **Manage**. On the **Manage workspace details** page, click **Actions**, and select **Update configuration**.

You can perform a rolling upgrade from a previous version of Maximo Manage to a later version of Maximo Manage. With a rolling upgrade, the interruption to Maximo Manage is minimal when you upgrade both operator and operands. During a rolling upgrade from one version to another version, the user session becomes invalid. Therefore, you must authenticate to continue your work. For more information, see [Upgrade and Rollback](#).

Configuring and activating Maximo Manage

Activate Maximo Manage to configure it and make it available in your Maximo Application Suite environment.

Before you begin

- You can check the actual supported matrix of your current Maximo Manage application that is deployed and the matrix of compatibility between the components, through the Red Hat OpenShift console:
 1. In the Red Hat OpenShift console, from the side navigation menu, click **Workloads > Pods** and from the **Projects** list, select the namespace of your Maximo Manage instance, such as **mas-*<yourmasinstancename>-manage***, and then click the **ibm-mas-manage-operator-*<somestring>*** pod.
 2. On the **Terminal** tab, connect to the webhook container.
 3. Open the supported matrix JSON file that displays the components version that each Maximo Manage application version supports. Then, look for the version that you previously installed. To determine the version of the base, click **Details** and click **Manager** in the **Containers** section. The version is displayed in the **Image version** field. Return to the **Terminal** tab and use this command to open the file:

```
cat /manage-admission/metadata/supported-versions-matrix.json
```

See the following example of how the file format appears:

```
"8.7.4": [
  {
    "name": "aviation",
    "versions": [
      "8.1.12"
    ]
  }
]
```

Note: In this example, the Manage application version is 8.1.12.

4. Open the dependency matrix JSON file that displays the compatibility between each component by using the following command: **cat dependency-matrix.json** The file displays the components with the following attributes:

```
"aviation": {
  "description": "Maximo Aviation",
  "includesForbidden": [
    "hse",
    "serviceprovider"
  ],
  "includesCoexist": [
    "acm",
    "transportation"
  ],
  "conflict": [
    "civil",
    "health",
    "oilandgas",
    "oracleadapter",
    "nuclear",
    "sapadapter",
    "spatial",
    "utilities"
  ]
}
```

Note: The component **includesForbidden** cannot be co-deployed with the listed components. The component **includesCoexist** can be co-deployed with the listed components. The component **conflict** cannot be co-deployed with the listed components.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Applications** and then click **Manage**. On the **Manage** page, click **Activate**.
2. Click **Show advanced settings** to configure Maximo Manage.
3. Configure the database connection information for Maximo Manage.
 - a) In the **Dependencies and Integrations** section, on the **Database connection** tile, select **View**.
 - b) On the **Database connection** page, click **Edit**.
 - c) In the **JDBC connection information** section, specify the following fields:

Note:

If you are upgrading from Maximo Asset Management to Maximo Application Suite, refer to the `maximo.properties` file of your Maximo Asset Management folder to get the values for hostname, port, and database name to use in the jdbc URL as required by Maximo Manage.

- i) If you want to use an SSL-enabled connection, specify this field by using one of the following JDBC URL formats. Ensure that the port that is used contains the SSL-enabled port of the database.

Oracle Database

TCPS is the protocol to use for Oracle SSL connections. For Oracle SSL database connections in Maximo Manage, you must specify `SID=<TNS Service ID>`. You can use the following URL as an example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(Host=mymaximodb.com)
(Port=2484))(CONNECT_DATA=(SID=MAXDB)))
```

Microsoft SQL Server Database

For SQL Server SSL database connections in Maximo Manage, you must specify `encrypt=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:sqlserver://mymaximodb.com:1433;databaseName=MAXDB;encrypt=true;
```

IBM Db2 database

For Db2 SSL database connections in Maximo Manage, you must specify `sslConnection=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB:sslConnection=true;
```

- a) In the **User name** field, specify the database username.
 - b) In the **Password** field, specify the database user password.
 - c) Ensure that you select the **SSL Enabled** option.
- ii) If you want to use non-SSL enabled connection, specify the **Connection string** field by using one of the following JDBC URL formats, depending on the database you are using.

Oracle Database

You can use the following URL as an example:

```
jdbc:oracle:thin:@mymaximodb.com:1521:MAXDB
```

Microsoft SQL Server database

You can use the following URL as an example:

```
jdbc:sqlserver://  
mymaximodb.com:1433;databaseName=MAXDB;integratedSecurity=false;encrypt=false;
```

IBM Db2 database

You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB
```

- a) In the **Username** field, specify the database username.
 - b) In the **Password** field, specify the database user password.
 - c) Ensure that you do not select the **SSL Enabled** option.
- d) In the **Additional driver options** section, in the **Driver options** field, add more driver options, which are separated by a semicolon.
- Typically, you can specify JDBC options as part of the URL for the database. However, in some cases you might want to specify JDBC options in the **Driver options** field.
- For example, your URL might exceed the maximum length that is allowed. You might also want to configure a JDBC option that cannot be included in the connection URL. However, you cannot specify the same JDBC option in both the URL and the **Driver options** field. If you do, JDBC driver errors can cause the connection to fail.
- If you specify an extra JDBC option for your database, the CustomProxyDriver acts as a proxy driver that routes the database requests to the actual driver for your type of database.
- e) If you chose to use an SSL-enabled database connection, in the **Security** section, click **Add +** to display the fields to include in your database certificate.
- i) In the **Alias** field, specify an alias name to identify the certificate, for example, DB2WHcert.
 - ii) In the **Certificate content** field, copy and paste your certificate in the format that is mentioned in the field content. You can retrieve a Privacy Enhanced Mail (PEM) certificate for your database. The file must be a Base-64 encoded X.509 file. You do not need to retrieve a private key. For more information, see the documentation for your database. After you copy and paste the text into the field, including the BEGIN CERTIFICATE, and END CERTIFICATE text, click **Confirm**.
- f) Click **Save**. The **Database connection** page is closed.

Note: Click **Save and Select** if you do not want to wait for the database connection verification after you complete the fields.

- g) Click the **Database connection** tile to verify your database connection. Expand the **Status** section that is loading in the **Configuration Scope: Workspace-application** section to display some tiles. Click **Select** when the **Status** icon displays the **Ready** message ready.
- h) If you plan to deploy Maximo Optimizer, see [Deploying Maximo Optimizer](#).
4. In the **Components** section, select the industry solutions or add-ons and the version that you are also activating with Maximo Manage.

- If you select **latest** from the **New version** list for the selected industry solution or add-on, the version of the component that is supported by the current Maximo Manage application version is co-deployed with Maximo Manage base component.

Note: If you set your configuration settings to the latest version for the components before you upgrade the application, the components will be automatically updated after the upgrade. If you select a specific version instead, then the components are not updated until you manually change their versions to the current, supported version of the new application version. You can do this for both by selecting the exact version or by selecting **latest**.

- If you select one component to deploy by selecting the **latest** option, you must select **latest** for any other component you co-deploy. If you select one component to deploy with an exact version number, you must select the exact version number for any other extra component that you co-deploy. You cannot mix exact version numbers and the latest version in the components you want to co-deploy.
- You also can click in the **New version** column for a component and select **Select version**. In the **Select unlisted version** modal, you can specify a valid version that is supported by the current application version that is deployed in the **Version** field. Click **Save**.

For more information, see [Deployment of industry solutions and add-ons](#). The documentation includes information on access to the compatibility matrix that shows the compatibility between application version and components versions. Some components might not be compatible with each other.

5. Click **Show advanced settings** to view and specify the configuration settings, such as database, server bundle, language settings, and others.

- a) In the **Database** section, set **System managed** to off and manually configure the database.

Schema

Enter the name of the schema that is configured in your database. For more information on database configuration for Maximo Manage, see [Setting up your database](#).

Encryption secret (optional)

This value is optional if you are deploying Maximo Manage in your database for the first time, and your database is not encrypted. Enter your encryption keys for this parameter. For more encryption settings information, see [Database encryption](#).

A table with **Key** and **Value** column titles is displayed. Click **Add property +**. In the **Key** column, enter `MXE_SECURITY_CRYPTOX_KEY`. In the **Value** column, enter your encryption key value.

In the Key/Value table, click **Add property +** again. In the **Key** column, enter `MXE_SECURITY_CRYPTOX_KEY`. In the **Value** column, enter your encryption key value.

Table space

If the default value does not match your database configuration, enter the name of the table space that was configured in your database. For more information about database configuration for Maximo Manage, see [Setting up your database](#).

Index space

If the default value does not match your database configuration, enter the name of the index space that was configured in your database. For more information about database configuration for Maximo Manage, see [Setting up your database](#).

Install demo data

If you are deploying a test or demonstration environment for Maximo Manage, you can install sample data.

The sample data in the demo database is useful for development or test environments.

To set up a test or development environment with demo data, install an instance of Maximo Application Suite specifically for testing or development. Then, when you configure the database settings for your Maximo Manage deployment, select the option to install demo data.

Note: You cannot add sample data after Maximo Manage is deployed because the database is updated without sample data. To add the sample data after deploying Maximo Manage, you must re-create or clean your database, and reconfigure Maximo Manage.

Db2 text search

This property is only applicable if the IBM Maximo Asset Management database is Db2 and it is not containerized Db2. For any other cases, this property is ignored. If containerized Db2 is used, text search cannot be enabled. The property only takes effect when the Maximo Manage database is being installed. In this case, the `-q` parameter passes to **maxinst** to enable text search for all columns that are flagged as searchable. After the Maximo Manage database is installed, any change to this property has no effect and does not alter the database. The default value is `False`.

Db2 Vargraphic

If you use Db2 and you plan to install a language other than English for your base language or as an extra language, select this option. If you intend to add more languages later, select this option during your initial deployment. This option does not affect the Maximo Manage deployment if it is selected and you are using a database other than Db2.

6. If you want to set Maximo Manage with a language other than English, or to include other languages in your Maximo Manage deployment, set the **System managed** checkbox in the **Languages** section to off.

Then, in the **Base** field, select your preferred language to be the base language. In the **Additional** field, order the list of other languages. For more information, see [Language support](#).

Note: If you are selecting other languages, make sure that you do not select a language in the **Additional** field that was selected in the **Base** field. For example, if you set the base as **EN**, do not select **EN** in the **Additional** field.

7. Configure server bundles for your deployment.

- a) To deploy Maximo Manage with more than one server bundle, or with customized configurations for it, in the **Server Bundles** section, set **System managed** to off.

A table with **Name**, **Pod count**, **Type**, and **Additional Properties** column titles is displayed. This table has **Default**, **User synchronization**, and **Mobile** optional fields available.

When you set **System managed** to off, the default settings for the **Name** and **Type** fields are **all**. The default **Pod count** is **1**.

- b) Click **Add bundle** to add more server bundles.
- c) Select the server bundle that you want to set for the **Default**, **User synchronization**, and **Mobile** fields.
- d) Optional: Change the name, pod count, type, route subdomain name and other customized configurations according to your preference.
- e) In the **Additional Properties** column, click **View** to view the **Route subdomain** value.

The **Additional server bundle properties** page for your selected server bundle is displayed with the **Route subdomain** and **Additional server config** fields available. You can specify properties by clicking **Add property +** in the **Bundle level properties** table. For more information, see [Server bundle overview](#).

8. If you want to include specific customizations through a customization archive, in the **Customizations** section set **System managed** to off. Click **Add customization archive +**. In the

File address field, specify the location of the customization archive. If you must enter credentials to access the file, specify them in the **Credentials (optional)** field. To include more customization archive files, click **Add customization archive +**. For more information, see the related sections in [Customizing the application](#).

- a) In the **Customization** section of the configuration window, specify the URL for the customization archive file.

The following URL protocols are supported:

- HTTP
- HTTPS
- FTP
- FTPS

- b) Optional: If you applied password security to the file, in the **Credentials** field, specify the user ID and password in the following format:

```
user=your user name password=your password
```

9. If you do not want the server bundles to start after the database operations of the Maximo Manage deployment are completed, set **System managed** to off in the **Server mode** section. Next, set **Mode** to **Off** to prevent access to the Maximo Manage application after deployment. You can restart the server bundle or bundles when you change the configuration to **On** and activate Maximo Manage again.
10. To connect to Persistent Volume Claims (PVC), in the **Persistent volume claims** section, set **System managed** to off, and click **Add PVC +**. A table with the following columns is displayed.

Option	Description
PVC name	The user-defined name of the persistent volume claim with a maximum of 63 characters.
Volume name	Leave blank as it is provisioned dynamically.
Size	The amount of storage that is required for this persistent claim, for example, 60G.
Mount path	The mount path for the volume within the Maximo Manage pod.

When you configure the PVCs in OpenShift Container Platform cluster on deployments, use the default storage class name `StorageClasses ocs-storagecluster-cephfs` to create a **ReadWriteMany** (rwx) PVC. The storage in the volume that you provisioned are available to all server bundles in the workspace. You can also configure a PVC for specific server bundles in a deployment. If you configure a PVC for a server bundle, the mount path that you specify for the server bundle PVC overrides the path that you specify for the deployment.

Tip: To configure a PVC in OpenShift Container Platform cluster on IBM Cloud platform, use the default `StorageClasses ibmc-file-gold-gid` (instead of `StorageClasses ocs-storagecluster-cephfs`) to create a **ReadWriteMany** PVC.

11. If you select **Asset Configuration Manager** or **Aviation** in the list of components, the **Build data interpreter** section is displayed. If you want to customize the configuration for the build data interpreter (BDI), set **System managed** to off. You can then specify a **BDI version** instead of latest, the default setting. Click **Add instance +**. You can customize each instance by selecting **View** in the **Configuration** column. The **BDI Configuration** page is displayed. You can change and save the configuration. Then, you can return to the **BDI configuration** page, select **Reset to Defaults**, and click **Save** to return to the default settings, if needed.
12. If you want to use an earlier build for deployment, in the **Build tag** section, set **System managed** to off. Then, in the **Build tag** field, specify the build tag. Build images are tagged with a timestamp, for example `buildtag: 202011092887843`.

Note: You can see all available build tag or `imagestreamtags` by going to the Red Hat OpenShift web console. Locate your the build and navigate to image streams. Click on the `maxadmin` image stream

to see all available imagestreamtags. You can use the Red Hat OpenShift command line interface to find the same information.

13. To connect to any external systems that Maximo Manage with is integrated with, set **Imported certificates** to off. Then, click **Add +** and import the certificate for the system and click **Confirm**.
14. Optional: You can specify the time zone that your database server is configured to use. In the **Server time zone** section, set the **System Managed** to off. In the **Time Zone** field, select the time zone of your database server.
15. If you are deploying Maximo Health as part of Maximo Manage , you can enable asset investment optimization. In the **Asset investment optimization** section, set the **System managed** to off. Then, select **Asset investment optimization**. When asset investment optimization is enabled, the **Asset investment optimizer** page is available in Maximo Health. Ensure that you deploy and configure Maximo Scheduler Optimization before you enable asset investment optimization.
16. If you want to configure Maximo Manage to send server bundle logs to Simple Storage Service (S3) object storage, set **Simple storage service** to on.
17. If you want to specify information for IBM Watson Studio, set **IBM Watson Studio** to on.
18. If you want to specify information for IBM App Connect, set **IBM App Connect** to on.
19. Click **Start activation** to start the activation process.

Note: You can reconfigure the activation process as needed. Follow these steps.

- a. Select **Actions > Workspace Details > Update configuration**.
- b. Select the specific configuration that you want to change in the **Update Manage Configuration** page. The configuration option is displayed. You can reconfigure a specific option and other options as you need.
- c. Click **Activate** again.

If you click **Exit**, you are redirected again to the **Application Details** page. You can continue to monitor the status. View the Workspace Activation detailed process, or under Activate in workspace, you can click Go to workspace details page link.

Results

Maximo Manage is activated, the second part of its deployment process starts, and is eventually completed. To access Maximo Manage, first you must configure at least the admin user to have permission to access Maximo Manage so that you can connect to Maximo Manage with it.

Maximo Manage tools

When Maximo Manage is deployed, certain tools become available.

- An admin pod is deployed when Maximo Manage is deployed.
- The command **kubect1 exec** can be run on the admin pod container, or you can directly use the container console to run the tools such as the integrity checker.
- Tools logged to **stdout** are forwarded to the Elasticsearch, Fluentd, and Kibana (EFK) stack, if it is configured.

Elasticsearch is an object store that stores logs. Fluentd gathers logs from nodes and feeds them to Elasticsearch. Kibana is the user interface (UI) for Elasticsearch.

After it is deployed in a cluster, the EFK stack aggregates logs from all nodes and projects into Elasticsearch. You can use EFK to view logs. Cluster administrators can view all logs. Application developers can view logs solely for projects that they have permission to view. The stack components communicate securely.

Monitoring Maximo Manage application deployment

You can monitor the Maximo Manage operator activation and deployment either in Maximo Application Suite or in the Red Hat OpenShift console.

Monitoring Maximo Manage activation by using the Maximo Application Suite interface

Monitor the IBM Maximo Manage activation in the Maximo Application Suite user interface.

About this task

You can monitor the Maximo Manage activation on the **Manage Status** page in the Maximo Application Suite.

Note: If you didn't log out and are following the previous step to deploy Maximo Manage, you are now in the **Workspace details** page. If you leave the page, you can access it again by following the steps in this task.

Procedure

1. In Maximo Application Suite, click the **Administration** icon to go to the **Administration** page, then click the **Applications** tile.
Starting in Maximo Application Suite 9.1, to access the **Applications** tab, on the side navigation menu, select the **Suite > Administration > Suite** page.
2. On the **Applications** tab, view the progress of deployment displayed in the **Status** column of the Maximo Manage application.
3. Select the Maximo Manage application.
4. In the **Active in workspaces** section, click **Go to workspace details**.
5. Review the status of each step of the Maximo Manage activation process in the tiles view.
 - View the detailed information, such as the current activation status. Status processes begin with the build in progress phase of the deployment process, which is when the Maximo Manage image is created. The `ear` or `war` files are also generated to deploy the application in the server bundle later, and the server bundle images are created.
 - When the build is completed, a green icon and a message with the build tag is displayed on the **Workspace** tile.
 - After some minutes, the **Deployment ready** tile starts to change the message that is displayed, and eventually shows the `maxinst/updatedb` is in progress.
 - After the `maxinst/updatedb` process is completed, the server bundles are started, the Maximo Manage `war` or `ear` files are deployed. When this last deployment step is completed, all tiles display green icons, and all the Maximo Manage components are in the Ready status. The Maximo Manage tile also displays a green icon, and the status is active.

Monitoring Maximo Manage activation by using the Red Hat OpenShift console

Monitor the IBM Maximo Manage activation in the Red Hat OpenShift console.

Procedure

After the Maximo Manage activation starts, confirm if the `ManageWorkspace` custom resource of the Maximo Manage instance is created from the Red Hat OpenShift web console.

- a) In Red Hat OpenShift console, from the side navigation menu, click **Administration > CustomResourceDefinitions**.
- b) In the `CustomResourceDefinitions` list search field, specify `ManageWorkspace` and then click **ManageWorkspace**.
- c) On the **CustomResourceDefinitions details** page, on the **Instances** tab, click the name of the custom resource for your instance.
For Amazon Web Services customer-managed automated deployment, the name is **mas-`<uniqueid>`-wsmasocp**.
- d) Check the section conditions.
If error messages are displayed, wait several minutes to see whether the Maximo Manage operator resolves the error. In the Red Hat OpenShift console, from the side navigation menu, click **Workloads**

> **Pods** and check whether the new pod admin-build-config pod was started. Otherwise, review the logs of the entity-mgr-ws pod and look for keywords such as `level:error`, `FAIL`, `FATAL` strings. These messages might clarify the reason that the process is stuck.

- e) If no errors occurred to prevent the deployment process to progress after activation, an admin-build-config pod is created. The admin-build-config pod is where the Maximo Manage image is created based on the settings that you chose.

Tip:

- You can check the progress of the build image by clicking the name of the admin-build-config pod and selecting the **Logs** tab. If it takes too long to view the admin-build-config pod that is being created, select the **Logs** tab, and you can access the entitymgr-ws pod. Review the log and look for error messages, such as `level:error`, `FATAL`, `FAIL`, or error strings.
- f) After the admin-build-config pod changes its status to complete, the server bundle build config pods are created. Their quantity and names might vary depending on how you configure the server bundle options before you activate Maximo Manage. After several minutes, the pod or pods change status to Complete, and the deployment process creates the maxinst pod, which installs Maximo Manage in the previously configured database.

- i) You can monitor the status of these additional build config pods in the Red Hat OpenShift console by clicking **Workloads > Pods** and in the **Project** list, select the **mas-*<yourinstance>*-manage** namespace.

Tip: Click in the **Created** column and sort by descending order so the most recent pod appears first in the list of pods displayed.

- ii) If your existent database is empty, the maxinst pod installs a new database. Otherwise, the maxinst pod validates whether updates to the database are supported in that version.

Tip:

- You can check the progress of the maxinst process, or updatedb process in the maxinst pod, by accessing the pod and selecting the **Logs** tab. After some time, the logs can run quicker showing several commands and outputs, such as scripts that were run. You can select the **Terminal** tab and use the following commands to show you the logs that were updated during the process:

```
cd log
ls -ltr
```

- If completed successfully, you see the message `maximo install or update completed`. The message is displayed as one of the most recent messages on the **Logs** tab of the maxinst pod.

- g) After completing the maxinst process, its pod continues to display the Running status until new server bundle pods are created. If you sorted the pods in descending order, the new pods appear listed first in the maxinst pod. The pod names and number vary depending on how you configured them before you activate Maximo Manage. When all containers for the server bundle pods display **Ready**. Usually, you see 2/2 displayed in the **Ready** column of the pod view in the standard server bundle configuration, which uses only one pod. Then, the Maximo Manage deployment is complete.

Tip:

- If you access Administration/Custom Resource Definitions through the Red Hat OpenShift console, search for ManageWorkspace, click its name and select the **Instances** tab, then select the CR of your Maximo Manage instance. Review the **Conditions** and the following information is displayed:
 - i) Line Ready with the column Status True, Reason Ready and Message: `Manage deployed in workspace <yourworkspace>`
 - ii) Line Failure with the status **False**
 - iii) Line BuildReady with the status **Build** completed.

- iv) DeploymentReady with the status **True** and Reason Ready and Message: UpdateDB/Maxinst Completed, all - X of X pods ready... This message specifies that the updatedb or maxinst was completed and the server bundle are all started and ready.
- You can check the logs of the server bundle pods and the entity-mgr-ws pod to look for error messages in case the statuses do not change to **Ready** after several minutes.

Monitoring Maximo Manage application deployment

You can monitor the Maximo Manage operator deployment on the **Application details** page in Maximo Application Suite.

Procedure

if you logged out after you activated the application, log in to Maximo Manage. Otherwise, you are now in the **Application details** page. If you logged out, you can access it again by following the steps:

- a) Log in to Maximo Manage by using your administration user ID and password.

Note: You can find the link of the Maximo Manage **Administration** page in the Red Hat OpenShift console by clicking **Networking > Routes** from the side navigation menu. Select mas-*<yourinstancename>*-core in the **Project** list. The link is in the **Location** column of the listed routes.

- b) In Maximo Application Suite, from the side navigation menu, click **Applications**.

The Maximo Manage application is listed. If the deployment is still in progress, it is indicated in the **Status** column.

- c) In the **Name** column, click **Manage** to return to the **Application details** page.

The **Application details** page contains tiles that show the status of each step of the Maximo Manage deployment process.

- The **Operator** tile displays the version of the Maximo Manage operator that is being deployed.
- Some messages are displayed inside the tiles during the deployment process. Deployment is complete when the **Application** card displays the Monitor is ready message and the **Activate** button is displayed. Click **Activate**.

Monitoring Maximo Manage application deployment by using the Red Hat OpenShift console

You can monitor the Maximo Manage operator deployment by using the Red Hat OpenShift console.

Procedure

1. In the Red Hat OpenShift console, click **Workloads** and select the mas-*<yourinstancename>*-manage namespace from the **Project** list.
2. In the pods list, click in the **Created** column to change the order of the pods that are displayed to be descending. Sorting the pods by descending order displays the most recent pods that are created at the beginning of the list.

When the Maximo Manage deployment starts, you can see an ibm-mas-manage-operator pod and an ibm-truststore-mgr-controller-manager pod. After some time, more pods are created, such as the usersyncagent, groupsyncagent, and other entitymgr pods.

3. Confirm the completion of the Maximo Manage operator deployment by using the Red Hat OpenShift console.
 - a) From the side navigation menu, click **Administration > Custom Resource Definitions**.
 - b) In the **Search** field of the custom resource definitions list, specify ManageApp and click **ManageApp** in the **Name** column of listed custom resource definitions.
 - c) On the **Instances** tab, click the name of the custom resource of your instance.
 - d) Check the **Conditions** section. The Maximo Manage operator deployment is complete when the message Ready is displayed in the **Reason** and **Message** columns, and the message True is displayed in the **Ready** column. Then, you can configure Maximo Manage.

Note: If some error occurred, the error message is displayed in the **Conditions** section. You can also look for error messages in the Red Hat OpenShift console. From the side navigation menu, click **Workloads > Pods** and select the `mas-<yourinstancename>-manage` project. Click the **ibm-mas-manage-operator** pod and select the **Logs** tab. You can download the logs until the current point and then search for keywords, such as `error`, `FAIL`, or `FATAL`.

- e) When the status of the operator custom resource is ready, configure Maximo Manage and activate it.

Note: You might deploy by using some automated deployment mode, where the Maximo Manage configuration was already passed through automation scripts at an earlier point. In this case, the deployment process progresses automatically from this point, depending on how your automated process was configured. Otherwise, the process waits for the configuration to be completed and for Maximo Manage to be activated.

Deploying Maximo Manage with externally built images

When you configure Maximo Manage, you can instruct the operator not to create the admin and bundle images automatically. Instead, the operator can provide the admin and bundle images directly to be deployed for Maximo Manage.

Note: This feature is available only in the feature channel. In Maximo Application Suite, customer-managed users can use the feature channel to update their nonproduction instances to preview new features. For more information, see [“What's new in the Maximo Application Suite feature channel ” on page 43](#)

Deploying Maximo Manage with externally built images is reserved for advanced users who want to control the image build outside their production Red Hat OpenShift cluster environment. You can also use it in the development and test environments for a more flexible image build process and to speed up the build and deployment process.

Note: If the automatic image build process by the operator is skipped, check that the image is generated correctly and conforms to the Maximo Manage deployment standards. IBM Support might require the image to be generated without the flag and test it to rule out any image problems caused by the external build process.

To instruct the operator to skip the build, set `spec.settings.skipImageBuild` to true in the **manageworkspace** custom resource. Instead of generating the images and updating the ImageStreams with the newly generated tag, the operator generates only a configmap with Dockerfile templates for the admin and bundle image builds. Follow the instructions in the Dockerfile template and understand the instructions to modify the Dockerfile and generate your own images.

The externally generated images can be put into external repositories if Maximo Manage can access the images during deployment with provided pull secrets.

Maximo Manage admin and bundle server deployment is put on hold until the images are provided. During the hold period, the operator reports the status to indicate that the images are not ready.

You can use two ways to provide the externally built images for Maximo Manage deployment.

- Specify the images by using the `spec.settings.usePrebuiltImages` section in the **manageworkspace** custom resource. Explicitly specify the repository image reference and the pull secret to access the images. Using `spec.settings.usePrebuiltImages` is the simpler option.
- Create ImageStream and Image Tag in the Red Hat OpenShift Maximo Manage namespace. Follow the same format and naming convention if `skipImageBuild` were not set to point to externally built images. Use `skipImageBuild` if you want to use `buildTag` for flexibly switching between images, especially in the test and development environments.

With either option, the images for the `maxinst` admin server, and all the bundles that are specified in the **manageworkspace** custom resource must be provided.

The steps in this section show how to set up the complete process by using Red Hat OpenShift web console and manually modifying the **manageworkspace** custom resource. However, to automate the process, you can use the Red Hat OpenShift command line interface and any script of your choice.

Disabling Maximo Manage default image build process

Disable the Maximo Manage default image build process by editing the **manageworkspace** custom resource.

Procedure

1. Log in to Red Hat OpenShift Container Platform.
2. From the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
3. On the **CustomResourceDefinitions** page, search for **manageworkspace**.
4. Click **manageworkspace** and then click the **Instances** tab.
5. Select an instance and click the **YAML** tab.
6. In the `spec.settings` section, add `skipImageBuild: true`.

When you set `skipImageBuild: true`, Maximo Manage skips the default build and generates a Dockerfile template. Use the Dockerfile template to build the image separately or locally.

Once the reconciliation is complete, access the Dockerfile template for the next steps. The following status message is displayed in the **manageworkspace** custom resource after you skip the image build.

```
skipImageBuild is set to True. Skipping the build. Find the docker template under config map:
managebuild-buildinfo
```

Creating external images

Build the external images outside Maximo Application Suite and Maximo Manage.

Procedure

1. Create a Dockerfile for the `maxadmin` image.
The following example shows a configmap for `admin` image.

```
#Installing industry solution. Remove the section or add corresponding rows for additional
industry solutions.
# Section End
FROM docker-na-public.artifactory.swg-devops.com/wiotp-docker-local/manage/
manageadmin@sha256:663a049e0213a651a3507007c6f209868ac2e6fd959e43cbf7448f17473c9ae4 AS ADMIN
WORKDIR /opt/IBM/SMP/maximo

# Copy file for industry solutions, remove the section or add corresponding rows for
additional industry solutions.
# Section End
RUN mkdir -p /opt/IBM/SMP/maximo/additional-server-files

# tmp_build_files is a mounted volume. The OOB dockerfile relies on the content of this
volume for truststore and credential of retrieving customization archive if it exists,
and the additional-server-files come to this folder as well. Do the necessary additional
content process using a similar mount volume, or use the directory in relation to
the dockerfile for copying files such as customization package or additional server files.

COPY --chown=maximoinstall:0 tmp_build_files .

#RUN rm -rf customizationCredentials && rm trust.p12 && rm truststorePassword
WORKDIR /opt/IBM/SMP/maximo
# Remove translation files that will not be used.
RUN rm -f translation_files.zip lang/MaximoLangPkgXliff_Ar.zip lang/
MaximoLangPkgXliff-Cs.zip lang/MaximoLangPkgXliff_Da.zip lang/MaximoLangPkgXliff_De.zip
```

Note: Do not omit any required components in the resulting image. The image must conform to what Maximo Manage expects.

Follow the instructions to incorporate your own customization or change the location of the content. Do not alter the structure.

Create your own build process to build and store all the required images.

For more information, see [“Disabling Maximo Manage default image build process”](#) on page 353.

Note:

- If you use all as the server bundle type, create a Dockerfile from all_dockerfile data in the template. Similarly, for admin build, the Dockerfile must be based on the admin_dockerfile from the template.
 - Anything that is referenced in the Dockerfile template configmap must have the corresponding images that are built externally.
 - You do not build or provide the image for the standaloneJMS bundle type. The image for the JMS server is fixed by default.
2. Create a Dockerfile for each bundle type, except for standaloneJMS, for each of the bundle types based on the corresponding Dockerfile provided in the generated configmap. The following example shows a configmap for all image.

```
# Update the below FROM command to include the image of manage admin.
FROM image-registry.openshift-image-registry.svc:5000/mas-<INSTANCE_ID>-manage/<INSTANCE_ID>-<WORKSPACE_ID>-admin@sha256:<image ID of the manage admin image>
AS ADMIN

FROM docker-na-public.artifactory.swg-devops.com/wiotp-docker-local/manage/ubi-wlp-
manage@sha256:7b7c5f75cb67733df11c40bea092de0c68366583cbcd8b3ff93be14ec837c6c AS LIBERTY
#ARG VERBOSE=true
USER root
RUN mkdir -p /dmroot/maximo/migration/outbound RUN mkdir -p /dmroot/maximo/migration/preview
RUN mkdir -p /dmroot/maximo/migration/inbound
RUN chmod -R 777 /dmroot/maximo/migration/outbound chmod -R 777 /dmroot/maximo/migration/
preview chmod -R 777 /dmroot/maximo/migration/inbound
RUN mkdir -p /airoot
RUN chmod -R 777 /airoot
COPY --chown=1001:0 maximo-all-server/ /config
RUN rm /opt/ibm/wlp/usr/servers/defaultServer/server.env && mkdir /managefiles/additional-
server-files
COPY --chown=1001:0 additional-server-files/ /managefiles/additional-server-files/
ENV MXE_MASDEPLOYED=1
ENV MZE_USESQLSERVERSEQUENCE=1
ENV LC_ALL=en_US.UTF-8
USER 1001
# This is overridden by container cmd when deployed in OpenShift
# CMD /opt/ibm/wlp/bin/server run defaultServer
```

Providing externally built images

You can use usePrebuiltImages or ImageStream and ImageStream tag to provide the externally built images. Create the images beforehand and store them in an image registry that is accessible on the network by the Red Hat OpenShift cluster that Maximo Manage uses.

Procedure

- Use usePrebuiltImages to provide the externally built images.
 - a) Create the image pull secret that contains the credentials to access the external repository in the Maximo Manage namespace.
 - b) Log in to Red Hat OpenShift Container Platform.
 - c) From the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
 - d) On the **CustomResourceDefinitions** page, search for manageworkspace.
 - e) Click **manageworkspace** and then click the **Instances** tab.
 - f) Select an instance and click the **YAML** tab.
 - g) In the spec.settings section, add usePrebuiltImages.
 - h) Specify the adminImage and a list of bundle images.

Note: Both tag and image ID are supported.

The following example shows an adminImage and an all bundle image.

```
skipImageBuild: true
usePrebuiltImages:
  adminImage: 'na.artifactory.swg-devops.com/iot-maximo-openshift-team-health-a-docker-
local/abc/prebuiltadmin:latest'
```

```
imagePullSecrets:
  - name: prebuilt-image
serverBundleImages:
  - bundleType: all
    image: 'na.artifactory.swg-devops.com/iot-maximo-openshift-team-health-a-docker-
      local/abc/prebuiltbundle:latest'
```

- i) Specify the `imagePullSecrets`.
 - j) Save the **manageworkspace** custom resource.
Saving triggers the Maximo Manage operator reconciliation. The maxinst pod is deployed by using the specified images and later bundle servers are deployed.
- Alternatively, reconstruct the ImageStream and image tag to reference the externally built images, similar to how the operator looks up the default image if it is not added to the **manageworkspace** custom resource.

If the images are not found, the deployment of maxinst and bundle servers is put on hold. Compared to the `usePrebuiltImages` option, ImageStream and Image tag are useful if you want to use `buildTag` to flexibly switch between a set of images in test and development environments.

- a) Push the locally built and tagged image to the internal repository of the cluster where Maximo Application Suite is installed, or the external repository in the **manageworkspace** custom resource `registryForGeneratedImages`.

Use the same tag name for all images. The expected format for a tag name is `<alphanumeric string without special characters>-<major.minor version of Manage>`. For example, `externalbuild-9.1`.

- b) In Red Hat OpenShift Maximo Manage's namespace, create an ImageStream for the maxadmin image, and one for each bundle image type.
- c) Name the ImageStreams as `<masinstance>-<workspace>-<admin or bundle type>`.
- d) Use the **oc tag** to tag the image. Alternatively, directly create an Image Stream with the tags by following the Red Hat OpenShift API.

For more information on adding tags to the Image Stream for each externally built image, see [Managing image streams](#).

Note: All admin images and bundle images must have the same tag, for example, `externalbuild-9.1`.

- e) In the `spec.settings.deployment.buildTag` section of the **manageworkspace** custom resource, specify the tag, for example, `externalbuild-9.1`.

The **manageworkspace** reconciles and after a few minutes, you can see the admin build pod that is created in the Maximo Manage namespace. Similarly, bundle servers are created by using the images of the tag that points to the externally built images.

Configuring Maximo Manage after deployment

You must give users access to the Maximo Manage application. Application administrators can configure the application to set up the required data structure for managing assets, including an organization, general ledger account structure, and configuring email notifications. You can also complete other post-deployment configurations.

Administering users

Maximo Manage users are created and authenticated in Maximo Application Suite. User data is synchronized to Maximo Manage at scheduled intervals.

For more information about administering users, see the following topics:

- [Administering user access in Maximo Application Suite](#)
- [Configuring authentication in Maximo Application Suite](#)
- [Understanding application points](#)

Configuring initial data after deployment

After Maximo Manage is deployed, you must complete several data configuration tasks before you use Maximo Manage.

Before you begin

- Create security groups and add users to the security groups.
- If you are using a directory server as part of your deployment, ensure that the user names are created in your Lightweight Directory Access Protocol (LDAP) repository.
- To ensure that you have the accounting information that you need, consult your organization's financial department.

Procedure

1. To send users email notifications of system events, configure the Simple Mail Transfer Protocol (SMTP) server and then enable the **mail.smtp.host** and **mxe.adminEmail** system properties.
 - a) In Maximo Application Suite, configure the Simple Mail Transfer Protocol (SMTP) server. For more information, see [Simple Mail Transfer Protocol server](#).
 - b) In Maximo Manage, from the side navigation menu, click **System Configuration > Platform Configuration > System Properties**.
 - c) In the Global Properties section, click the **Filter** icon to search for and expand the **mail.smtp.host** property.
 - d) In the **Global Value** field, specify the SMTP host name.
 - e) Select the **mail.smtp.host** checkbox.
 - f) From the **Common Actions** menu, select **Live Refresh** and click **OK**.
 - g) In the Global Properties section, click the **Filter** icon to search for and expand the **mxe.adminEmail** property.
 - h) In the **Global Value** field, specify the email address for the user.
 - i) Select the **mxe.adminEmail** checkbox.
 - j) From the **Common Actions** menu, select **Live Refresh** and click **OK**.
2. Define a currency code.
 - a) From the side navigation menu, click **Financial > Currency Codes**
 - b) Add a row and specify a currency code and a description.
For example, enter USD for United States of America Dollar.
3. Define item sets and company sets.
 - a) From the side navigation menu, click **Administration > Sets**.
 - b) Add a row and specify an item set name.
For example, enter IT Items.
 - c) In the **Type** field, specify ITEM.
 - d) Add a row and specify a company set name.
For example, enter IT Comps.
 - e) In the **Type** field, specify COMPANY and save your changes.
4. Create an organization.
 - a) From the side navigation menu, click **Administration > Organizations**.
 - b) Click the **New Organization** icon and specify an organization name and a long description.
For example, enter EAGLENA.
 - c) Specify the base currency that you previously defined.
 - d) Specify the item set and the company set that you defined previously defined.
 - e) In the **Default Item Status** field, set the status to PENDING.

- f) Click the **Sites** tab, and in the **Sites** section, add a row and specify a site name.
For example, enter `Factory01` and enter a long description.
5. Create a general ledger account component.
 - a) From the side navigation menu, click **System Configuration > Platform Configuration > Database Configuration**.
 - b) From the **More Actions** menu, click **GL Account Configuration**.
 - c) Add a row and specify a component name, length for the component, and a type for the component.
For example, enter the name `MYCOMPONENT`, the length `5`, and the type `Alphanumeric`.
6. Apply changes to the database.
 - a) From the **More Actions** menu, select **Manage Admin Mode**.
 - b) Select **Turn Admin Mode ON** and click **OK**.
This task takes several minutes to complete. You can click **Refresh Status** to view the progress.
 - c) From the **More Actions** menu, click **Apply Configuration Changes**.
In the status column of the listed objects, ensure that the status `To Be Changed` is not displayed.
 - d) From the **More Actions** menu, select **Manage Admin Mode**.
 - e) Select **Turn Admin Mode OFF** and click **OK**.
If you do not turn off Admin Mode, cron tasks fail.
7. Create a general ledger account.
 - a) From the **Go To** menu, click **Financial > Chart of Accounts**.
 - b) On the **Organizations** page, select your organization.
 - c) From the **More Actions** menu, select **GL Component Maintenance**.
 - d) In the **Components** section, select the component that you previously entered.
 - e) Add a row and specify a GL component value and a description, and then click **OK**.
 - f) In the **GL Accounts** section, add a row and specify a general ledger account.
 - g) From the **Go To** menu, select **Administration > Organizations**.
 - h) Find your organization and select the record.
 - i) In the **Clearing Account** field, specify the general ledger account that you created.
 - j) Select **Active** and save your changes.
8. Authorize a security group to modify a general ledger component type.
 - a) From the **Go To** menu, select **Security > Security Groups**.
 - b) Select the group that provides authorization.
For example, select **FINANCE**.
 - c) On the **GL Components** tab, for each GL component, select the **Authorized** checkbox.
 - d) For each GL Component that is listed, select the **Authorized** checkbox and save your changes.
You can select **Authorize Group to Change All GL Component Types**.
9. Update the company-related accounts.
 - a) From the **Go To** menu, select **Financials > Chart of Accounts**.
 - b) Select your organization and from the **More Actions** menu. Then, select **Company-Related Accounts**.
 - c) Click **New Row** and specify the company type `Courier`.
 - d) Enter an account number in the **RBNI Account**, **AP Suspense Account**, and **AP Control Account** fields.
You can specify the same account number in each field.
 - e) Click **OK**.

- f) From the **More Actions** menu, select **Update Database** and click **OK**.
10. Create a default insert site.
 - a) From the **Go To** menu, select **Security > Users**.
 - b) Find **maxadmin** and select the record.
 - c) In the **Default Insert Site** field, specify the site name that you previously created.
 - d) In the **Storeroom Site for Self-Service Requisitions** field, specify the same site name.
 - e) Click **Save User**.
 11. Define work types, denote the importance of the work task.
 - a) From the **Go To** menu, select **Administration > Organizations**.
 - b) Find your organization and select the record.
 - c) From the **More Actions** menu, select **Work Order Options > Work Type**.
 - d) Click **New Row**.
 - e) In the **Work Order Class** field, select **WORKORDER** and specify a **Work Type**.
For example, enter **MAJOR**.
 - f) Set the start status to **In Progress**.
 - g) Set the complete status to **Completed**.
 - h) Click **New Row** and repeat steps f-i to create another work order class that has a different work type.
For example, enter **MINOR**.
 - i) Click **New Row** and repeat steps f-i to create a **CHANGE** work order class that has a different work type.
For example, enter **SIG** to represent a significant change.
 - j) Click **OK** and click **Save Organization**.

Changing database credentials

If you change the username and password of your database after you deploy Maximo Manage, you must update the database connection in Maximo Application Suite. Updating ensures synchronization of the change to the new credentials and prevents database connection errors.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. From the side navigation menu, click **Workspace** and then click the **Manage** tile.
3. On the Manage details page, click **Actions** and then select **Update configuration**.
4. Click the **Edit** icon for **Database connection**.
5. In the **Dependencies and integrations** section, on the **Database connection** tile, click **View** and then click **Edit**.
6. Enter the new database username and password and save your changes.
7. To update the deployment with the new database credentials, click **Apply changes**.

Switching between Maximo Spatial and Maximo Spatial with ESRI

You can switch between Maximo Spatial and Maximo Spatial with ESRI without having to deactivate Maximo Manage.

Procedure

1. Log in to Red Hat OpenShift web console by using your administrator credentials.
2. From the side navigation menu, click **Administration > CustomResourceDefinitions**.
3. On the **CustomResourceDefinitions** page, search for **manageworkspace**.
4. Click **manageworkspace** and then click the **Instances** tab.

5. Select an instance and click the **YAML** tab.
6. In the `spec.components` section, change from Maximo Spatial to Maximo Spatial with ESRI by adding `esri`.
For example, for Maximo Spatial it must be:

```
spatial:  
  version: latest
```

For Maximo Spatial with ESRI, it must be:

```
spatialesri:  
  version: latest
```

7. Click **Save**.

Customer-managed **Optional: Installing IBM MQ for IBM Maximo Manage**

Set up Maximo Manage to connect and use IBM MQ as a messaging provider. It is assumed that you installed and configured an existing queue manager and that the user, password, and connection information are known before proceeding.

Before you begin

Create a queue manager by completing the following steps in IBM Cloud.

1. Click **Create** to create Queue manager.
2. Click **Queue manager**.
3. Click **Administration**.
4. Click **Launch IBM MQ Console**.
5. On the **Queue Manager** page, click **Connection information**. Choose any format and download connection information for future configuration.

Next, you can continue to create the queues. For more information about how to configure and manage IBM MQ, see [Getting started with IBM MQ](#).

Procedure

1. Configure queues in your queue manager.

When you create queues on the queue manager, each queue has properties that need to be set.

- a) Create a queue by using the properties and values shown:
 - Queue name: `sqout`
 - Default persistence: `Persistent`
 - Queue type: `Local`
 - Message delivery sequence: `FIFO`
 - b) Repeat the same steps for the `sqin`, `cqin`, `cqinerr` queues.
 - c) For the `cqinerr` queue, set the following properties:
 - Backout queue: `cqinerr`
 - Backout Threshold: `5`
 - d) Add Pass all context authorization for the application user on this queue, which gives authorization to the user to move messages from `cqin` to `cqinerr`. It also provides Browse and Inquire authorization.
 - e) Configure the `CLOUD.APP.SVRCONN` communication channel's `mcauser` value to the value of your IBM MQ user.
2. Download the IBM MQ resource adapter for queuesLiberty.

- a) Download the `wmq.jmsra.rar` IBM MQ resource adapter file that is needed by Liberty to connect to IBM MQ resources. You can add this file to your customization archive file to deploy. For more information, see [Where can I download the WebSphere MQ resource adapter?](#).
 - b) To add the downloaded RAR file to the customization archive, create the following folder structure. These paths are relative to `/opt/IBM/SMP/maximo`.
 - `additional-server-files\`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\mboejb\ejbmodule\META-INF\`
 - c) Place the file `wmq.jmsra.rar` in the `additional-server-files` folder.
3. Configure the client to enable message-driven beans (MDBs).

To modify your existing Maximo Manage application to serve as an IBM MQ client, complete the following steps:

- a) Copy the `ejb-jar.xml` file or its contents from the following location on the `maxinst` pod:
 - `/opt/IBM/SMP/maximo/deployment\was-liberty-default\config-deployment-descriptors\maximo-all\mboejb\ejbmodule\META-INF\ejb-jar.xml`
- b) Place the file in the folder structure that you created:
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\mboejb\ejbmodule\META-INF\`
- c) Edit the file and enable the MDBs by uncommenting them. The resulting `ejb-jar.xml` file contents are similar to the following example:

The enterprise-beans file contents:

```
<message-driven id="MessageDriven_JMSContQueueProcessor_1">
  <ejb-name>JMSContQueueProcessor-1</ejb-name>
  <ejb-class>psdi.iface.jms.JMSContQueueProcessor</ejb-class>
  <transaction-type>Container</transaction-type>
  <message-destination-type>javax.jms.Queue</message-destination-type>
  <env-entry>
    <env-entry-name>MESSAGEPROCESSOR</env-entry-name>
    <env-entry-type>java.lang.String </env-entry-type>
    <env-entry-value>psdi.iface.jms.QueueToMaximoProcessor</env-entry-value>
  </env-entry>
</message-driven>

      <message-driven id="MessageDriven_JMSContQueueProcessor_2">
    <ejb-name>JMSContQueueProcessor-2</ejb-name>
    <ejb-class>psdi.iface.jms.JMSContQueueProcessor</ejb-class>
    <transaction-type>Container</transaction-type>
    <message-destination-type>javax.jms.Queue</message-destination-type>
    <env-entry>
      <env-entry-name>MESSAGEPROCESSOR</env-entry-name>
      <env-entry-type>java.lang.String </env-entry-type>
      <env-entry-value>psdi.iface.jms.QueueToMaximoProcessor</env-entry-value>
    </env-entry>
    <env-entry>
      <env-entry-name>MDBDELAY</env-entry-name>
      <env-entry-type>java.lang.Long </env-entry-type>
      <env-entry-value>30000</env-entry-value>
    </env-entry>
    <env-entry>
      <env-entry-name>ERRORQUEUE</env-entry-name>
      <env-entry-type>java.lang.String </env-entry-type>
      <env-entry-value>1</env-entry-value>
    </env-entry>
  </message-driven>
```

The assembly-descriptor file contents:

```
<container-transaction>
  <method>
    <ejb-name>JMSContQueueProcessor-1</ejb-name>
    <method-name>*</method-name>
```



```

        </method>
        <trans-attribute>Required</trans-attribute>
    </container-transaction>
    <container-transaction>
        <method>
            <ejb-name>JMSContQueueProcessor-2</ejb-name>
            <method-name>*</method-name>
        </method>
        <trans-attribute>Required</trans-attribute>
    </container-transaction>

```

- d) Save the changes.
4. Create and deploy the customization files.
- Using a compressed file tool, create a compressed file from the deployment directory that you created earlier.
 - Place the compressed file on a file transfer protocol (FTP) server or persistent volume that Maximo Application Suite has access to.
 - Log in to the Maximo Application Suite as an administrative user.
 - From the side navigation menu, click **Applications** and then click the **Manage** tile.
 - Click **Actions**, and then click **Update Configuration**.
 - Click the edit icon for **Customization**.
 - If needed, enter the file address where the compressed file is stored and provide any credentials.
 - Click **Apply changes**.

The applications take time to be built and deployed. After the deployment is complete, you see that in the application pods such as mes and ui, the resource adapter is deployed to the following location:

- /managefiles/additional-server-files/wmq.jmsra.rar

5. Configure the Liberty servers to enable IBM MQ resources.

The XML that is shown here adds the extra feature classes, JMS resources, and the IBM MQ resource adapter to the Liberty server. You add the XMLs to the server bundles in the Maximo Application Suite for Maximo Manage.

- Replace the bolded field values with the correct values for your user and queue manager.

```

<?xml version="1.0" encoding="UTF-8"?>
<server description="new server">
  <featureManager>
    <feature>javaMail-1.6</feature>
    <feature>jdbc-4.2</feature>
    <feature>jaxws-2.2</feature>
    <feature>jndi-1.0</feature>
    <feature>wasJmsClient-2.0</feature>
    <feature>ssl-1.0</feature>
    <feature>webProfile-8.0</feature>
    <feature>openidConnectClient-1.0</feature>
    <feature>transportSecurity-1.0</feature>
    <feature>monitor-1.0</feature>
    <feature>wmqJmsClient-2.0</feature>
    <feature>jmsMdb-3.2</feature>
    <feature>jsonp-1.1</feature>
    <feature>ejbRemote-3.2</feature>
    <feature>ejbHome-3.2</feature>
  </featureManager>
  <logging traceSpecification="JMSApi=all:WAS.j2c=all"/>
  <variable name="wmqJmsClient.rar.location" value="/managefiles/additional-server-files/wmq.jmsra.rar"/>
  <jmsConnectionFactory jndiName="jms/maximo/int/cf/intcf" connectionManagerRef="MIFJMS">
    <properties.wmqJms

```

```

        transportType="CLIENT"
        hostName="<hostname>"
        port="<portnumber>"
        channel="<channelname>"
        applicationName="<username>"
        queueManager="<qmgrname>"/>
    </jmsConnectionFactory>
    <connectionManager id="MIFJMS" maxPoolSize="20"/>
    <jmsQueue id="sqout" jndiName="jms/maximo/int/queues/sqout">
    <properties.wmqJms baseQueueName="sqout" baseQueueManagerName="<qmgrname>"/>
</jmsQueue>
    <jmsQueue id="sqin" jndiName="jms/maximo/int/queues/sqin">
    <properties.wmqJms baseQueueName="sqin" baseQueueManagerName="<qmgrname>"/>
</jmsQueue>
    <jmsQueue id="jms/maximo/int/queues/cqin" jndiName="jms/maximo/int/queues/cqin">
    <properties.wmqJms baseQueueName="cqin" baseQueueManagerName="<qmgrname>"/>
</jmsQueue>
    <jmsQueue id="jms/maximo/int/queues/cqinerr" jndiName="jms/maximo/int/queues/cqinerr">
    <properties.wmqJms baseQueueName="cqinerr" baseQueueManagerName="<qmgrname>"/>
</jmsQueue>
    <jmsActivationSpec id="maximomea/mboejb/JMSContQueueProcessor-1">
    <properties.wmqJms
        transportType="CLIENT"
        destinationRef="jms/maximo/int/queues/cqin"
        destinationType="javax.jms.Queue"
        hostName="<hostname>"
        port="<portnumber>"
        maxSequentialDeliveryFailures="-1"
        channel="<channelname>"
        queueManager="<qmgrname>"/>
    <authData id="auth1" user="<user>" password="<password>"/>
    </jmsActivationSpec>
    <jmsActivationSpec id="maximomea/mboejb/JMSContQueueProcessor-2">
    <properties.wmqJms
        transportType="CLIENT"
        destinationRef="jms/maximo/int/queues/cqinerr"
        destinationType="javax.jms.Queue"
        hostName="<hostname>"
        port="<portnumber>"
        maxPoolDepth="1"
        maxSequentialDeliveryFailures="-1"
        channel="<channelname>"
        queueManager="<qmgrname>"/>
    <authData id="auth1" user="<user>" password="<password>"/>
    </jmsActivationSpec>
</server>

```

- b) After you modify the XML for your IBM MQ server and user, log in to the Maximo Application Suite.
 - i) From the **Catalog**, click **Manage**.
 - ii) From the **Actions** menu, click **Update Configuration**.
 - iii) Click edit for **Server Bundles**.
 - iv) Click the view link for the UI bundle under **Additional properties**.
 - v) In the **Additional server config** field, enter the XML code and click **Save**.
 - vi) Complete the same action for the Maximo Enterprise Adapter (MEA) bundle.
 - vii) Click **Apply changes**.

The preceding XML takes effect only in ui and mea server bundles.
 - c) The deployment takes time to complete. After the servers are restarted, proceed to the next section.
6. Configure Maximo Manage.
- a) Log in to Maximo, open the **System Properties** application and ensure that `mxe.int.disablejmsessionctx` is set to 1.
 - b) Apply the changes.
 - c) The JMS queues Java Naming and Directory Interface (JNDI) names in the XML used in previous steps are predefined in Manage. You can create new queues if you have used different queue names in the External Systems application, using the **Add/Modify Queues** action.

Installing Apache Kafka for IBM Maximo Manage

Apache Kafka provides a buffer for messages that are sent to and received from external interfaces. Apache Kafka is not required if the IBM Maximo Manage software does not interface with external systems.

Red Hat AMQ Streams operator, which is based on the Strict operator, is the preferred way to install Kafka for on-premises installations. It can also be used to install Kafka in cloud-based Maximo Application Suite installations when a managed Kafka service by the cloud provider is not desirable. For more information, see [Red Hat AMQ Streams operator](#) and [Strimzi operator](#).


Tip: This task can also be done by using the following Ansible role: `Kafka`. For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276 and [kafka](#).

What to do next

Configure Apache Kafka Suite parameters. For more information, see [“Apache Kafka”](#) on page 23.

Installing by using the Red Hat OpenShift Container Platform web console

Procedure

1. In Red Hat OpenShift Container Platform, from the side navigation menu, click **Home > Projects** and then click **Create Project**. Enter the name `kafka`, and click **Create** to provision the new namespace for Kafka.
2. In the global navigation bar, click the **Import YAML** icon . Enter the following YAML.

```
---
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: "kafka"
  namespace: "kafka"
spec:
  targetNamespaces:
    - "kafka"
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-streams
  namespace: "kafka"
spec:
  channel: amq-streams-1.8.x
  installPlanApproval: Automatic
  name: amq-streams
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

Tip: For Maximo Application Suite on AWS (BYOL) version 8.7, change `amq-streams-1.8.x` to `amq-streams-1.7.x` to match the version of AMQ streams that is installed in the BAS namespace.

3. Click **Create** to create the operator group and subscription resources in the `kafka` namespace.
4. From the side navigation menu, click **Operators > Installed Operators**. Search for AMQ Streams and verify that the operator status is set to **Succeeded**.
5. In the global navigation bar, click the **Import YAML** icon. Enter the following YAML code.

```
---
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: "maskafka"
  namespace: "kafka"
spec:
  # -----
```

```

kafka:
  version: 2.7.0
  replicas: 3
  resources:
    requests:
      memory: 4Gi
      cpu: "1"
    limits:
      memory: 4Gi
      cpu: "2"
  jvmOptions:
    -Xms: 3072m
    -Xmx: 3072m
  config:
    offsets.topic.replication.factor: 3
    transaction.state.log.replication.factor: 3
    transaction.state.log.min.isr: 2
    log.message.format.version: "2.7"
    log.retention.hours: 24
    log.retention.bytes: 1073741824
    log.segment.bytes: 268435456
    log.cleaner.enable: true
    log.cleanup.policy: delete
    auto.create.topics.enable: false
  storage:
    type: jbod
    volumes:
      - id: 0
        type: persistent-claim
        class: "ocs-storagecluster-ceph-rbd"
        size: 100Gi
        deleteClaim: true
  authorization:
    type: simple
  listeners:
    - name: tls
      port: 9094
      type: route
      tls: true
      authentication:
        type: scram-sha-512
# -----
zookeeper:
  replicas: 3
  resources:
    requests:
      memory: 1Gi
      cpu: "0.5"
    limits:
      memory: 1Gi
      cpu: "1"
  jvmOptions:
    -Xms: 768m
    -Xmx: 768m
  storage:
    type: persistent-claim
    class: "ocs-storagecluster-ceph-rbd"
    size: 10Gi
    deleteClaim: true
# -----
entityOperator:
  userOperator: {}
  topicOperator: {}

```

Modify the specified storage class `ocs-storagecluster-ceph-rbd` to use a supported storage class for your cluster.

6. Click **Create** to create the Kafka cluster.
7. From the side navigation menu, click **Workloads > StatefulSets** and switch to the **kafka** project. You see two statefulsets: **maskafka-kafka**, which is the Kafka brokers, and **maskafka-zookeeper**, which the Kafka ZooKeepers. Select each statefulset and verify that each has three pods that are in Ready state.
8. In the global navigation bar, click the **Import YAML** icon. Enter the following YAML.

```

---
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaUser

```

```

metadata:
  name: "maskafkauser"
  labels:
    strimzi.io/cluster: "maskafka"
  namespace: "kafka"
spec:
  authentication:
    type: scram-sha-512
  authorization:
    type: simple
  acls:
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: topic
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: group
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: cluster
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: transactionalId

```

9. Click **Create** to create a Kafka user, which is used by Maximo Application Suite to authenticate connections to Kafka.
10. From the side navigation menu, click **Workloads > Secrets** and switch to the **kafka** project. Verify that the user entity operator created the **maskafkauser** secret.
11. From the side navigation menu, click **Networking > Routes** and switch to the **kafka** project. Verify that the **maskafka-kafka-tls-bootstrap** route was created.
12. Get the Kafka information.
 - a) To get the Kafka host and port, input the following code:

```
oc get Kafka.kafka.strimzi.io maskafka -o jsonpath="{.status.listeners[0].addresses[0]}"
```

Sample output

```
{"host": "maskafka-kafka-tls-bootstrap-kafka.apps.cluster1.example-cluster.com", "port": 443}
```

- b) Get the Kafka CA certificate.

```
oc get Kafka.kafka.strimzi.io maskafka -o jsonpath="{.status.listeners[0].certificates[0]}"
```

Sample output

```
-----BEGIN CERTIFICATE-----
MIIFLTCCAxWgAwIBAgIUtEXU12XrdIPy6vZAtk9toGh2jbEwDQYJKoZIhvcNAQEN
BQAwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEw
MDAeFw0yMjA1MTEyMTAyMzFaFw0yMzA1MTEyMTAyMzFaMCAwEzARBgNVBAoMcm1v
LnN0cm1temkxXjAUBgNVBAMMDWnsdXN0ZXItY2EgdjAwGSIiMA0GCsqGSIb3DQEB
AQUAA4ICDwAwggIKAoICAQDh6bYIudhZQ1/rR9I9Sb7pZqTvtRiN0vzmnZPdtVtT
q71NLyTPqR6uuCIrhpuR0CPb++Rvjp2QrWgXr5VWBktT1MLk8WzDfX3+qxd5xkC8
B00EKneBZkhohxBdb0co8ipxDpQAFTy+SeXhuR0d5vwlEuh30JeZMEUfTcNfUbvo
J/IHUIGeDmhK//DumQE79z3vflC2EcQgenMo0VoBy4ooQ2o4B7Y3p1XHuStvt6h
lam30rSA+p3nKskrMDDPnkAdHtmCzWl/rZZBFYb7DTdUpi69NeW3TEMRXGG3dMdk
YYTDKN0zkB5BTvRx5FC6GX+cz/Uq3Snx1SmWB1DT+2n1nlwzVAgbNdsW4HiDUIcI
FBjYQDqWTH9e7aUv3Rz1rT4c995YBTfh1Jdvq5mzneMf61ab7iZow1hGYQLRRC5y
v8iTycwHd7EEGf/tjGrJ/s5nWPgGv/DE0g95/UvTRz9dZUURwHCFAnd0LaFw/HdF
qkhuivZOKNXqfr7zxnCw/F+0408+vcR43HKUTwId7vql+F+EgjT69U5pDF4sh6ep
SgLTHoCGd/bekq5HHkrylC0ty+ZU9EEWp4fQD+wN3RzGxJ080AA3RjkqsXmHbd5e
aX1nhDB68mWpohFuJ6YcInBBX1C/2HhDeR7PiMD9Zj0/7A3UHZj4hHXcSQoCnSW7
mwIDAQABo0UwQzAdBgNVHQ4EFgQU6yQKlZ+FEJYmKjsPxmHERps1vgwEgYDVR0T
AQH/BAGwBgEB/wIBADAQBGNVHQ8BAf8EBAMCAQYwDQYJKoZIhvcNAQENBQADggIB
AEfcrS4I2xsbTuULMth10GLgv7Mo+aJ80s+vCE+MvSMVrsVslVnigzE6aSvi7Ys
TTpstmAhIf0cEEqldRa5GcG6Az6NwLbskZXfftojWtjnZevkuRnn/xICdizX+mj4
A3Wl/GOVpTAWVUa5+1Uh1AzFWhBw5kDvMxHyQhmpegt98ptxNpj5n9cHSWwJpjXl
boNil+Y5kA4rawGa6gE0E0lwmLyS5pj0WCTCTD2MvldNakYPMq0bVPE4DNia4qa1
huxOyxdr51KNBc7yVgQ1Fa7ZD+rF1a6aa6GwvWAKYNoxd7VW7fmZBSckpuWer9+R
YCVvgE2a4vLnc5zLfw0fhjqazSiIX0PMEmkHx1ZTriVg0GVZ8beU+I9BxUQsJyJU
S4z9UaHexmYu/YRAQXK0Dw1xhqqR6oW2+CXyrtUvzN6kamFh8jN3AKf4PKA+TmjL
maWOM7FVp+0Erne59hBcZhKG0QYx4AkjCwKclRwDBXxcBTcmXduDFeGzLub0napJ
Ucz0z2URQ7L6qPew9Guh001dnGp+kgi8T8kt/DniMvQBWDK3GvFi0A5mVjLQqMHQ
HvAPzshx7Si1045hepGK4fxQMcvAHw6c1V3j10R8RHh7bckld5mJ5Nh/BjZhk/LK
N5K1fwoek0QSVAXQfnX1YtJfrHfz5+TYx0NNTcgX6fE
-----END CERTIFICATE-----
```

- c) Get the Kafka username and password.

```
oc extract secret/maskafkauser -n kafka --keys=sasl.jaas.config --to=-
```

- d) Get the Simple Authentication and Security Layer (SASL) mechanism.

```
oc get Kafka.kafka.strimzi.io maskafka -n kafka -o
jsonpath='{.spec.kafka.listeners[0].authentication}' | jq -r .
```

Sample output:

```
{
  "type": "scram-sha-512"
}
```

Note: The SASL mechanism is SCRAM-SHA-512.

Installing by using the Red Hat OpenShift command-line interface (CLI)

Procedure

1. From the bastion host, create the YAML file `kafka-sub.yaml`, which contains the namespace, `OperatorGroup`, and `Subscription` resources that are used to install Kafka:

```
---
apiVersion: v1
kind: Namespace
metadata:
  name: "kafka"
---
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: "kafka"
  namespace: "kafka"
```

```

spec:
  targetNamespaces:
    - "kafka"
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: amq-streams
  namespace: "kafka"
spec:
  channel: amq-streams-1.8.x
  installPlanApproval: Automatic
  name: amq-streams
  source: redhat-operators
  sourceNamespace: openshift-marketplace

```

Tip: For Maximo Application Suite on AWS (BYOL) version 8.7, change `amq-streams-1.8.x` to `amq-streams-1.7.x` to match the version of AMQ streams that is installed in the BAS namespace.

2. Apply the `kafka-sub.yaml` file to the Red Hat OpenShift Container Platform cluster:

```
oc apply -f kafka-sub.yaml
```

3. Verify that the AMQ Streams operator was successfully deployed:

```
oc get csv -n kafka -l operators.coreos.com/amq-streams.kafka
```

Sample output

NAME	DISPLAY	VERSION	REPLACES	PHASE
amqstreams.v1.8.4	Red Hat Integration - AMQ Streams	1.8.4	amqstreams.v1.8.3	Succeeded

4. From the bastion host, create the YAML file `kafka-cluster.yaml`, which contains the Kafka resource that describes the configuration of the Kafka cluster:

```

---
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: "maskafka"
  namespace: "kafka"
spec:
  # -----
  kafka:
    version: 2.7.0
    replicas: 3
    resources:
      requests:
        memory: 4Gi
        cpu: "1"
      limits:
        memory: 4Gi
        cpu: "2"
    jvmOptions:
      -Xms: 3072m
      -Xmx: 3072m
    config:
      offsets.topic.replication.factor: 3
      transaction.state.log.replication.factor: 3
      transaction.state.log.min.isr: 2
      log.message.format.version: "2.7"
      log.retention.hours: 24
      log.retention.bytes: 1073741824
      log.segment.bytes: 268435456
      log.cleaner.enable: true
      log.cleanup.policy: delete
      auto.create.topics.enable: false
    storage:
      type: jbod
      volumes:
        - id: 0

```

```

        type: persistent-claim
        class: "ocs-storagecluster-ceph-rbd"
        size: 100Gi
        deleteClaim: true
    authorization:
        type: simple
    listeners:
    - name: tls
      port: 9094
      type: route
      tls: true
      authentication:
        type: scram-sha-512
# -----
zookeeper:
  replicas: 3
  resources:
    requests:
      memory: 1Gi
      cpu: "0.5"
    limits:
      memory: 1Gi
      cpu: "1"
  jvmOptions:
    -Xms: 768m
    -Xmx: 768m
  storage:
    type: persistent-claim
    class: "ocs-storagecluster-ceph-rbd"
    size: 10Gi
    deleteClaim: true
# -----
entityOperator:
  userOperator: {}
  topicOperator: {}

```

Ensure that you modify the specified storage class `ocs-storagecluster-ceph-rbd` to use a supported storage class for your cluster.

5. Apply the `kafka-cluster.yaml` file to the Red Hat OpenShift cluster:

```
oc apply -f kafka-cluster.yaml
```

6. Verify that the Kafka cluster was successfully deployed. The Kafka CR is in the Ready state.

The Kafka CR specified in the following command is fully qualified with its API group name `kafkas.kafka.strimzi.io` to avoid ambiguity with the Kafka CR that is provided by `kafkas.ibmevents.ibm.com`.

```
oc get kafkas.kafka.strimzi.io -n kafka
```

Sample output

NAME	DESIRED KAFKA REPLICAS	DESIRED ZK REPLICAS	READY	WARNINGS
maskafka	3	3	True	

7. From the bastion host, create the YAML file `kafka-user.yaml`. The file contains the `KafkaUser` resource that describes the configuration of the Kafka user that is used by Maximo Application Suite to authenticate connections to Kafka:

```

---
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaUser
metadata:
  name: "maskafkauser"
  labels:
    strimzi.io/cluster: "maskafka"
  namespace: "kafka"
spec:
  authentication:
    type: scram-sha-512

```



```

authorization:
  type: simple
  acs:
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: topic
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: group
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: cluster
    - host: '*'
      operation: All
      resource:
        name: '*'
        patternType: literal
        type: transactionalId

```

8. Apply the `kafka-user.yaml` file to the Red Hat OpenShift cluster:

```
oc apply -f kafka-user.yaml
```

9. Verify that the user entity operator is created the `maskafkauser` secret.

```
oc get secret maskafkauser -n kafka
```

Sample output

NAME	TYPE	DATA	AGE
maskafkauser	Opaque	2	2m14s

10. Get the Kafka information.

a) Get the Kafka host and port.

```
oc get Kafka.kafka.strimzi.io maskafka -o jsonpath="{.status.listeners[0].addresses[0]}"
```

Sample output:

```
{"host": "maskafka-kafka-tls-bootstrap-kafka.apps.cluster1.example-cluster.com", "port": 443}
```

b) Get the Kafka CA certificate.

```
oc get Kafka.kafka.strimzi.io maskafka -o jsonpath="{.status.listeners[0].certificates[0]}"
```

Sample output:

```

-----BEGIN CERTIFICATE-----
MIIFLTCCAxWgAwIBAgIUExU12XrdIPy6vZAtk9toGh2jbEwDQYJKoZIhvcNAQEN
BQAwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEwLjEw
MDAeFw0yMjA1MTEyMTAyMzFaFw0yMzA1MTEyMTAyMzFaMC0xExARBgNVBAoMcm1v
LnN0cm1temkxZjAUBgNVBAMMDWNSdXN0ZXIyY2EgdjAwggIiMA0GCSCqGSIb3DQEBA
QUAA4ICDwAwggIKAoICAQDh6bYIudhZQ1/rR9IgsB7pZqTvtRiN0vzmnZPdtVtT
q71NLytPqpR6uuCIrhpuR0CPb++Rvjp2QrWgXr5VWBktT1MLk8WzDfX3+qxd5xC8
B00EKneBZkhohxBdb0co8ipxDpQAFTy+SeXhuR0d5vwLEuh30JeZMEUfTcNfUbv0

```

```
J/IHUIGeDmhK//DumQE79z3vflC2EcQgenMo0VoBy4ooQ2o4B7Y3p1XHStvt6h
1am30rSA+p3nKskrMDDpNKadHtmCrlwI/rZZBFYb7DTdUpi69New3TEMRXGG3dMdk
YYTdKN0zkB5BTvRx5FC6GX+cz/Uq3Snx1SmWB1DT+2n1n1wzVAgbNdsW4HiDUIdI
FBJyQDqWTH9e7aUv3Rz1rT4c995YBTfh1Jdvq5mzneMf61ab7iZow1hGYQLRRC5y
v8iTyCwHd7EEGf/tjGrJ/s5nWPgGv/DE0g95/UvTRz9dZUWRwHCFANd0LaFW/HdF
qkhuivZ0KNXqfr7zxnCw/F+0408+vcR43HKUTwId7vq1+F+Egjt69U5pDF4sh6ep
SgLTHoCGd/bekq5HHkrylCoty+ZU9EEWp4fQD+wN3RzGxJ080AA3RjkqsXmHbd5e
aX1nhDB68mWp0HFuJ6YciNBBX1C/2HhDeR7PiMD9Zj0/7A3UHZj4hHXcS0oCnSW7
mwIDAQABo0UwQzAdBgNVHQ4EFgQU6yQK1Z+FEJyMkjsPxmHERps1vgwEgYDVR0T
AQH/BAGwBgEB/wIBADA0BgNVHQ8BAf8EBAMCAQYwDQYJKoZIhvcNAQENBQADggIB
AEfcrS4I2xsbTuULMtH1OGLgv7Mo+aJ80s+vCE+MvSMVrsSvs1VnigzE6aSvi7Ys
TTpstmAhiF0cEEqldRa5GcG6Az6NWlbskZXfftojWtjnZevkuRnn/xICdizX+mj4
A3wL/GOVpTAWVUa5+1Uh1AzFwhBw5kDvMxHyQhmpet98ptxNpj5n9cHSWwJpjXl
boNil+y5ka4raWGa6gE0E0lwmLyS5pjOWCTCTD2MvldNakYPMqObVPE4DNia4qa1
hux0yxd151KNBc7yVgQ1Fa7ZD+rF1a6aa6GwvwAKYNoxd7VW7fmZBSckpuWer9+R
YCVvgE2a4vLnc5zLfw0fhjqazSiIx0PMEkHx1ZTriVg0GVZ8beU+I9BxUQsJyJU
S4z9UaHexmYu/YRAQXK0Dw1xhqqR6oW2+CXyrtUvzN6kamFh8jN3AKf4PKA+TmjL
maW0M7FVp+0Erne59hBcZhKG0QYx4AkjCwKc1RwDBxXcBTcmXduDFeGzLub0napJ
Uczo2zURQ7L6qPew9Guh001dnGp+kgi8T8kt/DniMvQBWDK3GvFi0A5mVjLQqMHQ
HvAPzshx7S11045hepGK4fxQMcwAHw6c1V3j10R8RHh7bck1d5mJ5Nh/BjZhk/LK
N5K1fwoek0QSVAXQfnX1YtJfrHfz5+TYx0NnYTcgX6fE
-----END CERTIFICATE-----
```

c) Get the Kafka username and password

```
oc extract secret/maskafkauser -n kafka --keys=sasl.jaas.config --to=-
```

d) Get the SASL mechanism.

```
oc get Kafka.kafka.strimzi.io maskafka -n kafka -o
jsonpath='{.spec.kafka.listeners[0].authentication}' | jq -r .
```

Sample output:

```
{
  "type": "scram-sha-512"
}
```

Note: The SASL Mechanism is SCRAM-SHA-512.

Customer-managed **Deploying IBM Maximo Monitor**

By using Maximo Monitor, you can visualize current and historical trend data for your devices and assets on customizable dashboards.

Before you begin

The following components must be installed and configured before you deploy Maximo Monitor:

- [IBM Db2 Warehouse database](#)
- [“Apache Kafka” on page 23](#)
- [MongoDB](#)

The following components can either be deployed before or as part of deploying Maximo Monitor:

“IoT tool” on page 76

The IoT tool provides device connectivity, data filtering and mapping, and device management tools needed for the Maximo Monitor application.

About this task

The following steps are specific to the Maximo Monitor application and are part of the overarching [application deployment process](#). Select your application update method. To later change from channel subscription versioning to manual versioning, you delete and then redeploy the application. After you activate the application, you must grant users [access](#) to it.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Catalog** and then click the **Monitor** card. Click **Continue**.

Note: If insufficient AppPoints are available to deploy this application, you can still complete the application configuration. The application is automatically deployed when the required number of AppPoints are available.

2. Select your application update method.

- a) Select an update method.

To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.

To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.

- b) Subscribe to a channel by selecting a version from the list.

For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.

- c) Click **Subscribe to channel <version>**

3. On the **Deploy** Monitor application page, click **Show advanced settings** and set **System managed** to off.

- Choose from developer, small, or medium. For more information about workload size, see the [Detailed System Requirements](#).
- Click **Deploy**.

The estimated deployment time is an estimate of the time that it takes to configure and deploy the application. The time includes both processing and configuration.

On the **Applications** page for Monitor, you can monitor the deployment status. If the deployment fails, delete and then redeploy the application. Deployment is complete when the **Application** card displays the **Monitor is ready** message and the **Activate** button is displayed. Click **Activate**.

What to do next

The Asset Data Dictionary component is installed automatically when a consuming application, such as Maximo Monitor, is installed. The Asset Data Dictionary is a repository of data and metadata that facilitates data sharing and synchronization between applications in the suite. If you plan to synchronize data between Maximo Monitor and Maximo Manage, complete configure the Asset Data Dictionary. For more information, see [Asset Data Dictionary implementation](#) in the Maximo Manage product documentation.

Related concepts

[IBM Maximo Monitor](#)

IBM Maximo Monitor is an application in Maximo Application Suite. By using Maximo Monitor, you can visualize current and historical trend data for devices and assets in customizable dashboards.

Related tasks

[Activating IBM Maximo Monitor](#)

Before you can grant users access and start working with the application, you must activate Maximo Monitor. You can activate the application after the deployment is complete.

[Updating Maximo Monitor](#)

Customer-managed

Deploying IBM Maximo Predict

Maximo Predict can use historical and recent asset performance data to correlate performance factors that predict asset degradation or failure. Other types of data that can be correlated include maintenance records, inspection reports, and environmental data. Maximo Predict uses artificial intelligence to optimize predictive model accuracy.

Before you begin

The following components must be installed and configured before you deploy Maximo Predict:

- [IBM Watson Studio](#)
- [“Watson Machine Learning” on page 20](#)
- [“Watson OpenScale” on page 20](#)
- [“Spark” on page 21](#)

Note: Watson OpenScale is not required in Maximo Application Suite 9.0.4 or later.

The following components must be installed before the deployment of Maximo Predict, but can be configured before or after deployment:

- [Maximo Monitor](#)
- [Maximo Health](#)

About this task

Select your application update method. To later change from channel subscription versioning to manual versioning, you must first delete and then redeploy the application.

Note: If insufficient AppPoints are available to deploy this application, you can still complete the application configuration. The application is automatically deployed when the required number of AppPoints are available.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Catalog** and click the **Predict** tile. Click **Continue**.
2. Select your application update method.
 - a) Select an update method.

To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.

To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.
 - b) Subscribe to a channel by selecting a version from the list.

For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
 - c) Click **Subscribe to channel <version>**
3. On the **Deploy** Predict application page, verify that Db2 Warehouse is configured. Maximo Predict can use the Db2 Warehouse instance that is configured at the system scope. For more information about configuring Db2 Warehouse, see [Db2 Warehouse](#).
4. Click **Deploy**.

On the **Applications** page for Predict, you can monitor the deployment status. If the deployment fails, delete and then redeploy the application. Deployment is complete when the **Application** card displays the **Monitor is ready** message and the **Activate** button is displayed. Click **Activate**.

What to do next

To later change from channel subscription versioning to manual versioning, you must first delete and then redeploy the application.

Related concepts

[IBM Maximo Predict](#)

IBM Maximo Predict is an application in Maximo Application Suite. By using Maximo Predict, you can leverage your historical and near real-time asset performance data, maintenance records, inspection

reports, and environmental data to correlate performance factors that predict asset degradation or failure. Maximo Predict also uses artificial intelligence to optimize predictive model accuracy.

Related tasks

[Activating IBM Maximo Predict](#)

Before you can grant users access and start working with Maximo Predict, you must activate the application. You can activate Maximo Predict after the deployment is complete.

[Updating Maximo Predict](#)

[Spark](#)

Related information

[Getting started with Maximo Predict](#)

Deploying Maximo Real Estate and Facilities in Maximo Application Suite

Starting in IBM Maximo Application Suite 9.1, you can deploy and activate IBM Maximo Real Estate and Facilities to make it available for use in IBM Maximo Application Suite. Before deployment, you must prepare the database. After deployment, you must create the mandatory initial FACILITIESADMIN user for the first login and to set up your administrator users to configure the application.

A system administrator can do the following tasks:

- Prepare the application database for Maximo Real Estate and Facilities.
- Deploy, configure, and activate Maximo Real Estate and Facilities, see [“Deploying and activating Maximo Real Estate and Facilities”](#) on page 379.
- Create the mandatory initial FACILITIESADMIN user in Maximo Application Suite, see [“Administering Maximo Real Estate and Facilities users”](#) on page 385.

Important: First log in to Maximo Real Estate and Facilities with the mandatory initial FACILITIESADMIN user and set up other users. If you try to log in to Maximo Real Estate and Facilities with any other user before the FACILITIESADMIN user, you get a blank screen and an error message that says the user is invalid, see [“Administering Maximo Real Estate and Facilities users”](#) on page 385.

- Complete any remaining post-deployment configuration and administration tasks, see [“Configuring Maximo Real Estate and Facilities after deployment”](#) on page 384.

Related concepts

[Application database](#)

To deploy Maximo Health, Maximo Real Estate and Facilities, or Maximo Manage, a database instance must be configured and running. The applications support Db2, Db2 Warehouse, Microsoft® SQL Server, or Oracle Database. If Maximo Health is deployed as part of Maximo Manage, the two applications share a database.

Preparing to deploy

To prepare for deployment, review workload sizes and deployments, and complete tasks such as configuring the database for Maximo Real Estate and Facilities. If you are deploying in multiple languages, review the supported languages that are available.

If you intend to use IBM Maximo Monitor for Workplace Analytics, you must also deploy IBM Maximo Monitor, see [“Deploying IBM Maximo Monitor”](#) on page 370. For more information, see [Maximo Monitor for Maximo Real Estate and Facilities](#).

System requirements

To deploy in Maximo Application Suite, your environment must meet the hardware and software requirements for Maximo Real Estate and Facilities and Maximo Application Suite.

If you intend to use IBM Maximo Monitor for Workplace Analytics, you must also deploy IBM Maximo Monitor, see [“Deploying IBM Maximo Monitor”](#) on page 370.

For more information about system requirements, see [IBM Maximo Application Suite system requirements](#) and [“Application-specific requirements for Maximo Real Estate and Facilities”](#) on page 187.

For detailed system requirements, see [Software Product Compatibility Reports](#).

Workload sizes and deployments

You can choose a small, medium, or large workload size for your deployment. The Maximo Real Estate and Facilities workload size determines how agents are deployed across pods and the number of pods that are deployed in the Maximo Application Suite/Red Hat OpenShift Container Platform cluster.

The main difference between the workload sizes is how the agents are deployed across pods. Review how the pods are deployed for each size to see how the workload is distributed.

For more information about agents, see [Business process agents](#).

For more information about customizing workloads, see [“Customizing workloads in Maximo Real Estate and Facilities”](#) on page 385.

Pods common to all sizes

appserver (UI)

The application server is always present and, if needed, you can add more `appserver` pods by configuring the CR. The `dataimportagent` agent runs on each `appserver` pod.

- Pod1 (default)
- Pod2
- ...
- Podn

Dedicated workflow agents

Dedicated workflow agents are custom agents that you can create to handle dedicated workflows that require more resources. Each dedicated workflow agent runs on its own pod. If needed, you can configure them in the `FacilitiesWorkspace` CR. Do not scale by increasing the number of replicas for `dwfagents`. Scale by adding multiple individual dedicated workflow agents in the CR. For more information, see [“Configuring dedicated Maximo Real Estate and Facilities workflow agents”](#) on page 386.

- Pod1 - `dwfagent-mycustomagent1`
- Pod2 - `dwfagent-mycustomagent2`
- ...
- Podn - `dwfagent-mycustomagentn`

Agent pods for small

For a small workload size, agents are deployed as follows:

Agents

- Pod1 - `multiagents`
 - `extendedformulaagent`
 - `formularecalcagent`
 - `incomingmailagent`
 - `objectmigrationagent`
 - `objectpublishagent`
 - `maintenanceagent`
 - `reportqueueagent`
 - `scheduleragent`
 - `wfagent`
 - `wffutureagent`

- wfnotificationagent
- reservesmtpagent
- dataconnectagent

Agent pods for medium

For a medium workload size, agents are deployed as follows:

Agents

- Pod1 - multiagents
 - extendedformulaagent
 - formularecalcagent
 - incomingmailagent
 - objectmigrationagent
 - objectpublishagent
 - maintenanceagent
 - scheduleragent
 - wffutureagent
 - wfnotificationagent
 - reservesmtpagent
 - dataconnectagent
- Pod2 - wfagent
- Pod3 - reportqueueagent

Agent pods for large

For a large workload size, agents are deployed as follows:

Agents

- Pod1 - extendedformulaagent
- Pod2 - formularecalcagent
- Pod3 - incomingmailagent
- Pod4 - objectmigrationagent
- Pod5 - objectpublishagent
- Pod6 - maintenanceagent
- Pod7 - scheduleragent
- Pod8 - wffutureagent
- Pod9 - wfnotificationagent
- Pod10 - wfagent
- Pod11 - reportqueueagent
- Pod12 - reservesmtpagent
- Pod13 - dataconnectagent

Preparing your Maximo Real Estate and Facilities database

Before you can deploy IBM Maximo Real Estate and Facilities, you must configure your database and determine how your database is encrypted. You can prepare IBM Db2 in the cluster or a supported external database outside the cluster.

If you are installing by using the Maximo Application Suite CLI, Db2 can be installed and prepared in the cluster during the installation, see [“Standard installation with IBM Maximo Application Suite CLI” on page 218](#). You must manually install and prepare external databases.

For more information about AES encryption for database passwords, see [AES encryption](#).

Preparing your external Db2 database

Prepare your external IBM Db2 database by using the provided the database preparation scripts.

Before you begin

Download the external Db2 scripts from the [IBM Maximo Application Suite CLI GitHub](#).

Procedure

1. Install your Db2 database by using the provider's documentation.
2. Configure the Db2 instance by creating and running the following script as the instance owner, for example db2inst1:

```
./db2configinst.sh <instance_name> <port> <db2_installation_directory>
```

For example:

```
./db2configinst.sh db2inst1 50001 /home/db2inst1/sqllib
```

3. Create the Db2 database by creating and running the following script as the instance owner:

```
./db2createdb.sh <db_name> <instance_name> US <db2_installation_directory>  
<linux_user_who_owns_db>
```

For example:

```
./db2createdb.sh mrefdb db2inst1 US /home/db2inst1/sqllib mref
```

4. Connect to the Db2 database by creating and running the following script as the instance owner:

```
db2 connect to $DB_NAME
```

Ensure that you replace \$DB_NAME with the required value.

5. Create the Db2 table space instance by creating and running the following script as the instance owner:

```
db2 -tvf create-ts.sql
```

6. Set up SSL and extract the certificate by creating and running the following script as the instance owner:

```
./ssl-setup.sh <instance_name>
```

Preparing your external Oracle Database

Prepare your external Oracle Database by using the provided the database preparation scripts.

Before you begin

Download the external Oracle Database scripts from the [IBM Maximo Application Suite CLI GitHub](#).

About this task

For more information about preparing your Oracle Database for installation, see:

- [SSL With Oracle JDBC Thin Driver - White Paper](#)

- [HOW TO: Setting up Encrypted Communications Channels in Oracle Database](#)
- [Configuring the Oracle Database server](#)

Procedure

1. Install your Oracle Database by using the provider's documentation.
2. Run the `create-ts.sql` command to create the table spaces.

Ensure that you update the `DATAFILE` path based on your Oracle installation.

3. Create the database schema and user.

- a) Run the following command:

```
alter session set "_ORACLE_SCRIPT"=true;
```

- b) Then, run the `createuser.sql` command.

Ensure that you update the `$dbuser$` and `$dbpassword$` values based on your Oracle Database installation.

4. Set up SSL and extract the certificate.

- a) Run the following commands, substituting your Oracle Database installation directory:

```
export ORACLE_SID=orcl
export ORACLE_HOME=/u01/app/oracle/product/19.0.0/dbhome_1
export PATH=/u01/app/oracle/product/19.0.0/dbhome_1/bin:$PATH
```

- b) Locate the following files, depending on your Oracle Database version:

For Oracle 19c:

<i>Table 27. Oracle 19c</i>	
File name	Path
<code>listener.ora</code>	<code>\$ORACLE_HOME/network/admin</code>
<code>sqlnet.ora</code>	<code>\$ORACLE_HOME/network/admin</code>
<code>xdbwallet</code>	<code>/u01/app/oracle/admin/orcl/ xdb_wallet</code>

- c) Update `listener.ora`:

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP)(HOST = <hostname>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPRO1521))
      (ADDRESS = (PROTOCOL = TCPS)(HOST = <hostname>)(PORT = 5500))
    )
  )

SSL_VERSION = 1.2
SSL_CLIENT_AUTHENTICATION = FALSE

WALLET_LOCATION =
  (SOURCE =
    (METHOD = FILE)
    (METHOD_DATA =
      (DIRECTORY = /u01/app/oracle/admin/orcl/xdb_wallet)
    )
  )
```

- d) Update `sqlnet.ora`:

```
SQLNET.AUTHENTICATION_SERVICES= (TCPS, NTS)
NAMES.DIRECTORY_PATH= (TNSNAMES)

SSL_VERSION = 1.2
SSL_CLIENT_AUTHENTICATION = FALSE
```

```

WALLET_LOCATION =
(SOURCE =
(METHOD = FILE)
(METHOD_DATA =
(DIRECTORY = /u01/app/oracle/admin/orcl/xdw_wallet)
)
)

```

e) Restart Listener Control by running the following commands:

```

lsnrctl stop
lsnrctl start
lsnrctl status

```

Ensure that the listener shows that the services have started their instances and are in ready state.

If the listener shows no services, run the following command:

```

dbstart $ORACLE_HOME

```

f) Extract the certificate by running the following commands:

```

orapki wallet display -wallet /u01/app/oracle/admin/orcl/xdw_wallet
orapki wallet export -wallet /u01/app/oracle/admin/orcl/xdw_wallet -dn 'CN=orcl' -cert /
home/oracle/CA.cert
cat /home/oracle/CA.cert

```

Preparing your external Microsoft SQL Server database

Prepare your external Microsoft SQL Server database by using the provided the database preparation scripts.

Before you begin

Download the external Microsoft SQL Server scripts from the [IBM Maximo Application Suite CLI GitHub](#).

About this task

For more information about preparing your database for installation, see [SQL Docs: Encrypting Connections to SQL Server on Linux](#).

Procedure

1. Install your Microsoft SQL Server database by using the provider's documentation.
2. Create the database and user by running the following command:

```

CREATE DATABASE "mref" collate SQL_Latin1_General_CP1_CI_AS;
use mref;
CREATE LOGIN puriuser
    WITH PASSWORD = 'puripass',
    DEFAULT_DATABASE = mref,
    CHECK_POLICY = OFF;
exec sp_changedbowner 'puriuser','puriuser';

```

3. Set up SSL and extract the certificate by running the following script as a root user:

Note: This step enables SSL at the instance level. If other databases in the instance do not require SSL to be enabled, consider moving the IBM Maximo Real Estate and Facilities database into its own separate instance.

```

./ssl-setup.sh

```

Deploying and activating Maximo Real Estate and Facilities

As a system administrator, you deploy and activate IBM Maximo Real Estate and Facilities in Maximo Application Suite. Select the application version, configure the application, and point to the prepared IBM Maximo Real Estate and Facilities database to deploy and activate the application.

Before you begin

- Confirm that the Maximo Real Estate and Facilities database is configured and you have its information available, including the correct JDBC URL in the correct format.

About this task

First, deploy Maximo Real Estate and Facilities in Maximo Application Suite by selecting the application version. This process usually takes a few minutes to complete.

Then, configure and activate Maximo Real Estate and Facilities for use in Maximo Application Suite. After you confirm activation, the main deployment process starts. This process takes a couple of hours or more to finish.

Deploying Maximo Real Estate and Facilities

You deploy Maximo Real Estate and Facilities in the Maximo Application Suite user interface.

Before you begin

Prepare the Maximo Real Estate and Facilities database before you deploy.

You must have the following information available:

- **Customer-managed** Confirm the correct pricing plan to select based on the expected number of users.
- The database credentials that are needed to configure the JDBC connection. Confirm that the Maximo Real Estate and Facilities database is configured and you have its information available, including the JDBC URL in the correct format.

Procedure

1. On the **Suite administration** page, click **Catalog** and on the **Applications** tab, select the **Real Estate and Facilities** tile.
2. On the **Setup** tab, review the information. Click **Continue**.
3. Select your application update method and subscription channel.

Note: For production Maximo Real Estate and Facilities deployments, set **Automatic approval** to **Off**. If you choose an automatic upgrade strategy, required downtime might occur before you have a chance to take preparatory action, such as reviewing changes or backing up the database.

With manual updates, you must trigger the updates yourself. Review the changes, run backups of the Maximo Real Estate and Facilities configuration and custom resource definitions, schedule the update, and communicate the scheduled downtime to users.

- a) Select an update method.

To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.

To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.

- b) Subscribe to a channel by selecting a version from the list.
For example, select channel **9.1.x**.
- c) Click **Subscribe to channel <version>**

Wait a few minutes while the subscription to the channel is established. The **Deploy Facilities** page is displayed.

4. **Customer-managed**

The Base pricing plan is selected by default. If needed, click **Change plan** to review the plans and choose one of the following options. The appropriate number of AppPoints is consumed:

Limited

Up to 10,000 users or room panels can use Reserve.

Base

10,001-100,000 users or room panels can use Reserve.

Advanced

More than 100,000 users or room panels can use Reserve.

5. Click **Deploy**, and on the **Confirmation** page, click **Begin Deployment**.

You can view the deployment progress on the **Real Estate and Facilities** page.

What to do next

Configure and activate Maximo Real Estate and Facilities, see [“Configuring and activating Maximo Real Estate and Facilities”](#) on page 380.

Customer-managed

Configuring and activating Maximo Real Estate and Facilities

Before you can grant users access and start working with the application, you must configure and activate IBM Maximo Real Estate and Facilities. You can activate Maximo Real Estate and Facilities after the deployment is complete.

Before you begin

Important:

Before you activate Maximo Real Estate and Facilities, ensure that a default storage class is set in the Red Hat OpenShift cluster or specify a storage class for logs and user files. If you don't specify a storage class for logs and user files in the custom resource, the operator uses the default storage class that is set in the cluster and if it is not set, the deployment fails.

Do not change the storage class after you start activation. After the storage class is set, you cannot change it without losing data.

Confirm that the Maximo Real Estate and Facilities database is configured and you have its information available, including the JDBC URL in the correct format.

If you need to integrate with external servers over TLS, you must provide CA certificates for the servers. Have the CA certificates details ready.

Procedure

1. Create a secret with the default name format of <workspaceId>-facilities-vs--sn to hold the value for AES encryption. This value must be available during activation. For more information, see [AES encryption](#).
2. In Maximo Application Suite, from the side navigation menu, click **Applications** and then click **Real Estate and Facilities**. On the Real Estate and Facilities page, click **Activate**.
3. Set **Advanced settings** to on to see the settings for Real Estate and Facilities.
4. Configure the database connection information for Maximo Real Estate and Facilities. The Real Estate and Facilities configuration scope is workspace-application, which means that the configuration is set for and can be used with a single application in the default workspace.
 - a) In the **Dependencies and Integrations** section, on the **Database connection** tile, select **View**.
 - b) On the **Database connection** page, click **Edit**.

- c) In the **JDBC connection information** section, select the **SSL-enabled** checkbox and specify the following information.

Note: Only secure sockets layer (SSL) connections to the database are supported.

Specify the JDBC connection URL, username, password to connect to the database as shown in the following examples.

- For IBM Db2 database:

```
jdbc:db2://<hostname>:<port>/<database name>;sslConnection=true;sslVersion=TLSv1.2;
```

For example:

```
jdbc:db2://169.61.55.213:50010/mref428;sslConnection=true;sslVersion=TLSv1.2;
```

By default, if a value is not specified, the database schema value is the database username in capital letters.

You can optionally specify a custom database schema value that doesn't match the username, such as MREFDB, in capital letters.

```
jdbc:db2://<hostname>:<port>/<database name>;sslConnection=true;sslVersion=TLSv1.2;currentSchema=<SCHEMA_NAME>;
```



Warning: When you specify a custom schema, ensure that the custom database schema value does not match the database user name or the Db2 driver fails. If you want the username and database schema values to match, do not set *currentSchema=<SCHEMA_NAME>*.

For example:

```
jdbc:db2://169.61.55.213:50010/mref428;sslConnection=true;sslVersion=TLSv1.2;currentSchema=MREFDB;
```

- For Oracle Database:

```
jdbc:oracle:thin@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(HOST=<hostname>)(PORT=<port number>)(CONNECT_DATA=(SID=<Service ID>))))
```

The protocol for SSL is TCPS.

Service names are not supported.

For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(HOST=169.61.55.213)(PORT=5500))(CONNECT_DATA=(SID=orashift)))
```

- For Microsoft SQL Server database:

```
jdbc:jtds:sqlserver://<host>:<port>;databaseName=<database name>;tds=<jTDS version>;SendStringParametersAsUnicode=false;prepareSQL=2;ssl=authenticate
```

jTDS is a required driver for Maximo Real Estate and Facilities and must be specified in the JDBC URL.

For example:

```
jdbc:jtds:sqlserver://169.61.55.214:2022;databaseName=justice4;tds=8.0;SendStringParametersAsUnicode=false;prepareSQL=2;ssl=authenticate
```

Additional driver options (optional)

Typically, you can specify JDBC options as part of the URL for the database. However, in some cases you might want to specify extra JDBC options. You can add multiple options separated by a semicolon (;).

For example, your URL might exceed the maximum length that is allowed. Or, you might want to configure a JDBC option that cannot be included in the connection URL. However, you cannot specify the same JDBC option in both the URL and the additional driver options. If you do, JDBC driver errors can cause the connection to fail.

Security: Custom certificates

If you are using certificates from a nontrusted certificate authority or if default trusted certificate authorities are not automatically used, add the custom certificate here.

If you chose to use an SSL-enabled database connection, click **Add +** to display the fields to include in your database certificate.

In the **Alias** field, specify an alias name to identify the certificate, for example, DB2WHcert. Each alias name must be unique within the workspace.

In the **Certificate content** field, copy and paste your certificate in the format that is mentioned in the field content. After you copy and paste the text into the field, including the BEGIN CERTIFICATE, and END CERTIFICATE text, click **Confirm**.

Click **Retrieve** to retrieve certificates from the server. You can retrieve a Privacy Enhanced Mail (PEM) certificate for your database. The file must be a Base-64 encoded X.509 file. You do not need to retrieve a private key. For more information, see the documentation for your database. The suite automatically downloads and saves any CA certificates that are configured with the server. The certificates are not validated, so you must verify them and remove any unwanted certificates for the server.

d) Click **Save**. The **Database connection** page is closed.

Note: Click **Save and Select** if you do not want to wait for the database connection verification after you complete the fields.

e) Click the **Database connection** tile to verify your database connection. Expand the **Status** section that is loading in the **Configuration Scope: Workspace-application** section to display some tiles. Click **Select** when the **Status** icon displays the **Ready** message ready.

5. By default, the system manages advanced settings and the recommended values are set. To review or modify the default values, set **Advanced settings** to off.

6. To modify the workload size, in the **Deployment** section, set **System Managed** to off.

Choose **Small**, **Medium**, or **Large**. The default value is **Small**. For more information, see [“Workload sizes and deployments”](#) on page 374.

7. To configure persistent storage, in the **Storage** section, set **System Managed** to off.

Log files storage access mode

Supported access modes are ReadWriteOnce and ReadWriteMany. The default is ReadWriteOnce.

Log files storage class

The storage class can vary based on your environment. For example, `ibmc-file-gold-gid` for IBM Cloud, or `ocs-storagecluster-cephfs`.

Note: The Red Hat OpenShift Container Platform cluster must be equipped with a StorageClass that can grant read/write permission to the Linux root group and to support Kubernetes ReadWriteMany or ReadWriteOnce access mode.

To view a list of available storage classes in your cluster, run the following Red Hat OpenShift console command:

```
oc get storageclasses
```

Log files storage size

Specify a size in GB based on your requirements for logs. The default is 30.

User files storage access mode

Supported access modes are `ReadWriteOnce` and `ReadWriteMany`. The default is `ReadWriteOnce`.

User files storage class

The storage class can vary based on your environment. For example, `ibmc-file-gold-gid` for IBM Cloud, or `ocs-storagecluster-cephfs`.

Note: The Red Hat OpenShift Container Platform cluster must be equipped with a `StorageClass` that can grant read/write permission to the Linux root group and to support Kubernetes `ReadWriteMany` or `ReadWriteOnce` access mode.

To view a list of available storage classes in your cluster, run the following Red Hat OpenShift console command:

```
oc get storageclasses
```

User files storage size

Specify a size in GB based on your requirements for logs. The default is 50.

8. To connect to external systems, in the **Imported certificates** section, set **System Managed** to off.

Click **Add** to add trusted certificates for external sites into the truststore in WebSphere Application Server Liberty. Specify the certificate alias and copy and paste the contents of the certificate. Each alias name must be unique within the workspace. For more information, see [“Adding trusted certificates” on page 387](#).

9. To configure Liberty application server, in the **Server configuration** section, set **System Managed** to off.

Liberty Server XML extensions

By default, no extensions are specified. You can specify a sequence of XML elements to be added to the `server.xml` file of the WebSphere Application Server Liberty as shown in the following example:

```
<logging consoleLogLevel="AUDIT"
traceSpecification="*=audit:com.ibm.ws.security.saml.*=all" />
```

The XML elements are added and the `<workspaceId>-facilities-lexml--sn secret` is created.

Maximum connection pool size

The default value for the maximum connection pool size of the database is 100 per pod.

10. To modify the timeout for Maximo Real Estate and Facilities routes, in the **Network** section, set the **System Managed** to off.

The default value is 600s. The supported values for the time unit are microseconds (us), milliseconds (ms), seconds (s), minutes (m), hours (h), or days (d). For more information about routes, see [“Network considerations” on page 385](#).

11. Click **Start activation**.
12. On the **Confirmation** page, click **Confirm**.

What to do next

Verify that IBM Maximo Real Estate and Facilities was deployed in Suite administration.

Then, you must create the mandatory initial `FACILITIESADMIN` administrator user in Maximo Application Suite administration and give them access to Maximo Real Estate and Facilities, [“Administering Maximo Real Estate and Facilities users” on page 385](#).

If you intend to use IBM Maximo Monitor for Workplace Analytics, you must also deploy IBM Maximo Monitor, see [“Deploying IBM Maximo Monitor” on page 370](#).

Deployment and activation status information

You can see status information for deployment and activation on the Real Estate and Facilities page. The Facilities operator maintains a status subresource for each instance that it manages. This status includes both Kubernetes conditions and reason codes.

Conditions

Running

True while the operator is running the reconciliation for the custom resource (CR) instance.

Successful

True when the reconciliation cycle for the CR instance finishes without errors. The operator then waits for the next reconciliation action, either the reconcile period, dependent watches triggers, or the resource is updated.

Failed

True if an error occurs during the reconciliation run for the CR instance, and the error message from the error that caused this condition is reported. The error message is the raw output from the run for reconciliation. If the failure is intermittent, the situation can often be resolved when the operator reruns the reconciliation loop.

Initialized

True when all resources are created, including all the application pods.

Ready

True when all application pods reach the Ready condition.

Reason codes

FacilitiesApp

- PodTemplatesInvalid - Validation of podTemplates keys failed.
- CertificateNotReady - Secret is not available.
- TruststoreNotReady - Suite truststore is not ready.

FacilitiesWorkspace

- AppserverNotReady - The Appserver StatefulSet has not finished its initialization.
- MinimumReplicasNotAvailable - The minimum required number of replicas for StatefulSet is not met.
- StorageError - Unable to create or modify StorageClass.
- DWFAgentsNotReady - The dedicated workflow agents have not finished initialization.
- AgentsNotReady - One or more agents have not finished initialization.
- BindingFailure - Facilities did not correctly bind to IBM Maximo Application Suite.
- CompatibilityCheckFailure - The supported Maximo Real Estate and Facilities and Maximo Real Estate and Facilities versions do not match.
- CompatibilityCheckFailure - The supported Manage Foundation and Maximo Real Estate and Facilities versions do not match.
- DataInitializationPending - Data Initialization not yet finished. This can take up to a few hours, depending on your network speed and database storage type.
- DataInitializationFailed - Data Initialization has failed.

Configuring Maximo Real Estate and Facilities after deployment

Complete the following post deployment steps for IBM Maximo Real Estate and Facilities.

For more information about setting up Maximo Real Estate and Facilities, see [Configuring and Administering](#) in the Maximo Real Estate and Facilities documentation.

Administering Maximo Real Estate and Facilities users

Maximo Real Estate and Facilities users are created and authenticated in Maximo Application Suite. User data is automatically synchronized to Maximo Real Estate and Facilities.

1. **Important:** The first login to Maximo Real Estate and Facilities must be as the mandatory initial FACILITIESADMIN administrator user.

To get started, you must create the mandatory initial FACILITIESADMIN administrator user, and other Maximo Real Estate and Facilities administrator users, in Maximo Application Suite administration and give them access to Maximo Real Estate and Facilities, see [“Creating users in Maximo Application Suite 9.1”](#) on page 784.

2. Log in to Maximo Real Estate and Facilities as the FACILITIESADMIN user to set up the Maximo Real Estate and Facilities administrator users. Maximo Real Estate and Facilities administrator users can then grant users more granular application permissions, or administrator permissions, as needed in Maximo Real Estate and Facilities administration. For more information, see [Administering user access and permissions](#).

For more information about administering users, see the following topics:

- [Configuring authentication in Maximo Application Suite](#)
- [Understanding application points](#)

Network considerations

Maximo Real Estate and Facilities exposes routes and an SMTP server endpoint.

Maximo Real Estate and Facilities run on JVMs on Red Hat OpenShift pods.

Routes

The main Maximo Real Estate and Facilities route for users is the `appserver` route, which can connect to one or more `appserver` pods based on the workload size.

You can find the route in the **Routes** section of the Red Hat OpenShift web console. All other routes that you can see point to individual pods. Use these routes for any pod-specific configuration in the Maximo Real Estate and Facilities Administrator Console, which you can use to update properties in the `TRIRIGAWEB.properties` file.

All Maximo Real Estate and Facilities routes have a default timeout of 600 seconds (600s). You can change this in Suite administration. The supported values for the time unit are microseconds (us), milliseconds (ms), seconds (s), minutes (m), hours (h), or days (d).

SMTP server endpoint

An SMTP server endpoint is needed only for the Maximo Real Estate and Facilities Reserve and Microsoft Exchange integration. The supported ports are 1025 and 587 for the SMTP server endpoint.

Customizing workloads in Maximo Real Estate and Facilities

You can customize workloads in Maximo Real Estate and Facilities by using the standard suite process. The following information is specific to Maximo Real Estate and Facilities.

Workload scaling

Regardless of the deployment size, the following workload scaling is supported by configuring the Custom Resource (CR):

- The UI application server workload is scalable both horizontally and vertically. That is, you can have multiple UI application server workloads, and also increase CPU and memory resources allocation.

- All other workloads are scalable vertically. You can configure different CPU and memory resources allocation for individual agents, including dedicated workflow agents.

For more information about customizing workloads and the supported pods for Maximo Real Estate and Facilities, see [“Customizing workloads”](#) on page 648.

Deployment size

If needed, you can change the workload size for an existing installation in Suite administration. For example, you can change from small to medium.

Important: To avoid any service disruption, make sure that all pods in the existing FacilitiesWorkspace CR are in the Ready state before you change the workload size.

Configuring dedicated Maximo Real Estate and Facilities workflow agents

You can create dedicated workflow agents for specific Maximo Real Estate and Facilities workflows by using the optional `spec.settings.dwfagents` property in the FacilitiesWorkspace CR. Each dedicated workflow agent runs on its own pod.

- Restrict each user to a single agent. Sharing users across multiple dedicated workflow agents is not recommended because you want all workflows to be executed by one process.
- Using groups is possible but not recommended. Ensure that all users are added to the group before you create a dedicated workflow agent. If new users are added to the group, the dedicated workflow agent pod must be restarted to pick up new changes.
- You can scale dedicated workflow agent pods with `podTemplates` vertically, but not horizontally. Do not increase the number of replicas for `dwfagents`. To add more pods, add multiple individual dedicated workflow agents in the FacilitiesWorkspace CR.

Use the following syntax to create dedicated workflow agents.

```
apiVersion: apps.mas.ibm.com/v1
kind: FacilitiesWorkspace
metadata:
  name: my-facilities
  namespace: ibm-mas-facilities
spec:
  settings:
    ...
    dwfagents:
      - name: dwfagent-mycustomagent1
        members:
          - class: user
            name: myUserName1
          - class: user
            name: myUserName2
          - class: group
            name: myGroupName1
          - class: group
            name: myGroupName2
      - name: dwfagent-mycustomagent2
        members:
          - class: user
            name: myUserName3
      - name: dwfagent-mycustomagent3
        members:
          - class: group
            name: myGroupName3
    ...
```

- Each item in the `spec.settings.dwfagents` array represents a dedicated workflow agent.
 - The `spec.settings.dwfagents[i].name` string is the name of the dedicated workflow agent. The name must have the `dwfagent-` prefix.
 - The `spec.settings.dwfagents[i].members` array is a list of users or groups that the workflow agent is dedicated to.
- If the item represents a user, set the `spec.settings.dwfagents[i].members[j].class` property to `user`.

- If the item represents a group, set the `spec.settings.dwfagents[i].members[j].class` property to `group`.
- In both cases, the `spec.settings.dwfagents[i].members[j].name` property is the name of the user or group.
- The `members` array of each of the included items cannot be empty. At least one user or group must be assigned to the dedicated workflow agent. Each `spec.settings.dwfagents[i].name` must have a unique name and be a part of the dedicated workflow agent pod.

Known limitations

Deleting a dedicated workflow agent pod from Maximo Real Estate and Facilities does not update the UI (Agents Page).

Adding trusted certificates

To enable secure communication with an external server, administrators can add trusted certificates for external sites into the truststore in WebSphere Application Server Liberty. The list of aliases and certificates represents a complete chain of trust for the host. The certificates are necessary only if the host presents a certificate that is signed by a certificate authority that is not already trusted by Maximo Application Suite.

Adding trusted certificates

You can add certificates through the user interface when you configure for deployment. The format of the certificate must be PEM.

Adding trusted certificates in the user interface

Procedure

Complete the following steps to add certificates.

For more information about certificates, see [“Certificate management” on page 588](#).

- Log in to IBM Maximo Application Suite as a system administrator.
- From **Suite administration**, select **Workspaces** from the side navigation menu and then select the workspace that you want to configure.
- On the **Overview** tab for the workspace, select the Maximo Real Estate and Facilities tile.
- On the **Real Estate and Facilities** panel, click **Actions** and select **Update configuration**.
- In the **Imported certificates** row on the **Update Real Estate and Facilities configuration** window, click the **Edit** icon.
- In the **Imported certifications** section, set **System managed** to off and click **Add** to add a certificate.
- Specify the certificate alias and copy and paste the contents of the certificate.
Each alias name must be unique within the workspace.
- Click **Confirm** to save your changes.
- Select one of the following options:
 - If you are configuring a new deployment, click **Activate**.
 - If you are updating a deployment, click **Apply changes**.

Changing database credentials

If you change the username and password of your database after you deploy IBM Maximo Real Estate and Facilities, you must update the database connection in Maximo Application Suite. Updating ensures synchronization of the change to the new credentials and prevents database connection errors.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. From the side navigation menu, click **Workspace** and then click the Maximo Real Estate and Facilities tile.
3. On the Maximo Real Estate and Facilities details page, click **Actions** and then select **Update configuration**.
4. Click the **Edit** icon for **Database connection**.
5. In the **Dependencies and integrations** section, on the **Database connection** tile, click **View** and then click **Edit**.
6. Enter the new database username and password and save your changes.
7. To update the deployment with the new database credentials, click **Apply changes**.

The FacilitiesWorkspaces custom resource

The following parameters are in the FacilitiesWorkspace custom resource.

The parameter values in the following table are provided as examples. Update the parameter values based on your configuration:

Parameter	Value
<code>spec.binding.jdbc</code>	Specify workspace-application.
<code>spec.deployment.size</code>	Specify small, medium, large based on your environment.
<code>spec.settings.vaultSecret.secretName</code>	Specify the name of the vault secret. The secret value must be encoded in Base64. For example, <code><workspaceId>-facilities-vs--sn</code>
<code>spec.settings.routes.timeout</code>	Specify the timeout value for routes. The default is 600s.
<code>spec.settings.libertyExtensionXML.secretName</code>	Specify a sequence of XML elements to be added to the <code>server.xml</code> file of the WebSphere Application Server Liberty. The XML elements are added and the <code><workspaceId>-facilities-lxml--sn</code> secret is created.
<code>spec.settings.imagePullPolicy</code>	By default, the pull policy is set to <code>ifNotPresent</code> and the image is pulled only if it is not present locally. For more information, see the kubernetes documentation .

Parameter	Value
spec.settings.storage.log.class	<p>The storage class can vary based on your environment. For example, <code>ibmc-file-gold-gid</code> for IBM Cloud, or <code>ocs-storagecluster-cephfs</code>.</p> <p>Note: The Red Hat OpenShift Container Platform cluster must be equipped with a <code>StorageClass</code> that can grant read/write permission to the Linux group and to support Kubernetes <code>ReadWriteMany</code> or <code>ReadWriteOnce</code> access mode.</p> <p>If the storage class is not specified in the custom resource, the operator uses the default storage class set in the Red Hat OpenShift cluster. If a default storage class is not set in the Red Hat OpenShift cluster, PVC provisioning fails. Check whether the storage class is set in the cluster before you activate Maximo Real Estate and Facilities without setting the storage class.</p> <p>To view a list of available storage classes in your cluster, run the following Red Hat OpenShift console command:</p> <pre data-bbox="862 884 1474 961">oc get storageclasses</pre>
spec.settings.storage.log.mode	Supported access modes are <code>ReadWriteOnce</code> and <code>ReadWriteMany</code> . The default is <code>ReadWriteOnce</code> .
spec.settings.storage.log.size	Specify a size based on your requirements for logs. The default is 30.

Parameter	Value
spec.settings.storage.userfiles.class	<p>The storage class can vary based on your environment. For example, <code>ibmc-file-gold-gid</code> for IBM Cloud, or <code>ocs-storagecluster-cephfs</code>.</p> <p>Note: The Red Hat OpenShift Container Platform cluster must be equipped with a <code>StorageClass</code> able to grant read/write permission to the <code>Linux</code> root group and to support Kubernetes <code>ReadWriteMany</code> or <code>ReadWriteOnce</code> access mode.</p> <p>If the storage class is not specified in the custom resource, the operator uses the default storage class set in the Red Hat OpenShift cluster. If a default storage class is not set in the Red Hat OpenShift cluster, PVC provisioning fails. Check whether the storage class is set in the cluster before you activate Maximo Real Estate and Facilities without setting the storage class.</p> <p>To view a list of available storage classes in your cluster, run the following Red Hat OpenShift console command:</p> <pre>oc get storageclasses</pre>
spec.settings.storage.userfiles.mode	Supported access modes are <code>ReadWriteOnce</code> and <code>ReadWriteMany</code> . The default is <code>ReadWriteOnce</code> .
spec.settings.storage.userfiles.size	Specify a size based on your requirements for user files. The default is 50.
spec.settings.dwfagents	Specify the name of the dedicated workflow agents and a list of users and groups that the workflow agent is dedicated to. For more information, see “Configuring dedicated Maximo Real Estate and Facilities workflow agents” on page 386.
spec.settings.db.maxconnpoolsize	Maximum number of physical connections for the application server database connection pool. The default is 100 and the minimum is 100.

```

apiVersion: apps.mas.ibm.com/v1
kind: FacilitiesWorkspace
metadata:
  name: <instanceId>-<workspaceId>
  namespace: mas-<instanceId>-facilities
spec:
  bindings:
    jdbc: workspace-application
  settings:
    deployment:
      size: small
    vaultSecret:
      secretName: <workspaceId>-facilities-vs--sn
    routes:
      timeout: 600s
    libertyExtensionXML:
      secretName: <workspaceId>-facilities-lexml--sn
    imagePullPolicy: IfNotPresent
    storage:
      log:

```

```

class: ibmc-file-gold-gid
mode: ReadWriteMany
size: 30
userfiles:
  class: ibmc-file-gold-gid
  mode: ReadWriteMany
  size: 50
dwfagents:
- name: dwfagent-mycustomagent1
  members:
  - class: user
    name: myUserName1
  - class: user
    name: myUserName2
- name: dwfagent-mycustomagent2
  members:
  - class: user
    name: myUserName3
db:
  maxconnpoolsize: 100

```

Customer-managed **Deploying IBM Maximo Visual Inspection**

By using deep learning AI, Maximo Visual Inspection can identify production defects and monitor assets for potential disruptions. The application requires graphical processing units (GPUs) to conduct deep learning model training.

Before you begin

Complete the following tasks to set up a GPU-accelerated node:

1. [Enabling GPU passthrough](#)
2. [Installing the NVIDIA operator](#)

If you are using Amazon Web Services, you need to complete the following task:

- [Adding a GPU worker node to a Red Hat OpenShift cluster on AWS](#)

About this task

By deploying and activating the Maximo Visual Inspection application, you make it available for use in Maximo Application Suite.

Customer-managed **Enabling GPU passthrough**

For Red Hat OpenShift clusters that are provisioned on hypervisors, such as VSphere, Citrix, and KVM, you can use a GPU passthrough and assign GPU resources directly to specific virtual machines (VMs).

About this task

For Red Hat OpenShift clusters that are provisioned on hypervisors, such as VSphere, Citrix, and KVM, two methods can be applied to expose the underlying GPU hardware to the Red Hat OpenShift cluster. The first is to install a GPU virtualization driver on the hypervisor so that the physical GPUs can be virtualized and its resources are shared among the VMs through the hypervisor. The second method is to enable GPU passthrough and assign GPU resources directly to specific VMs, thus bypassing the hypervisor. The second method, as it applies to the VSphere hypervisor, is described in this task.

Note:

- Installation of Red Hat OpenShift on bare metal machines without a hypervisor are not subject to this documentation.
- This documentation does not apply to AWS. For more information about enabling GPU passthrough on AWS, see [“Adding a GPU worker node to a Red Hat OpenShift cluster on AWS” on page 394](#).

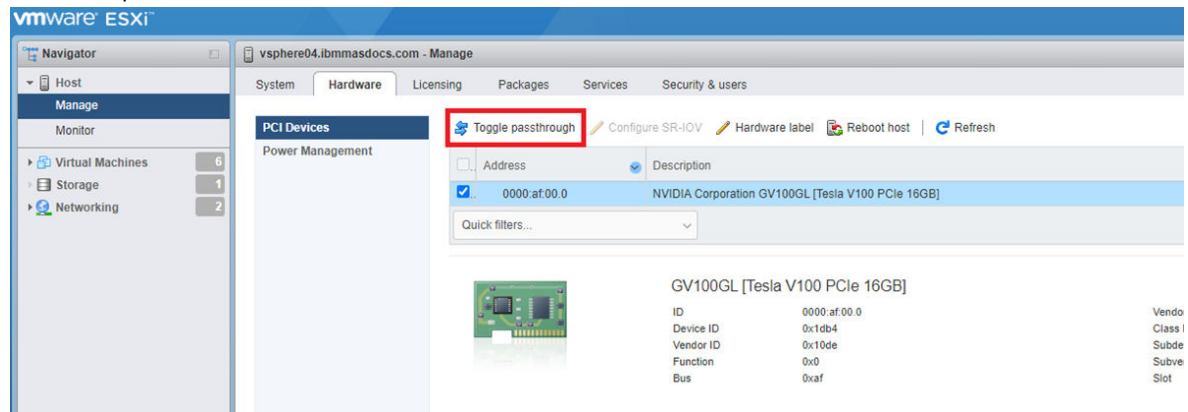
To enable GPU passthrough and assign GPU resources, see the following NVIDIA documentation that correlates to your specific VM:

- [Virtual GPU Software R510 for Linux with KVM Release Notes](#)
- [Virtual GPU Software R510 for Citrix Hypervisor Release Notes](#)
- [Virtual GPU Software R510 for VMware vSphere Release Notes](#)

Procedure

1. Enable GPU passthrough in the VSphere hypervisor.
 - a) From the **Navigator** on the VSphere browser-based console, click **Host > Manage**.
 - b) On the **Hardware** tab, click **PCI Devices** and search for the GPU devices.
 - c) Click the checkbox next to all wanted GPU devices to enable passthrough and click **Toggle passthrough**.

For example:

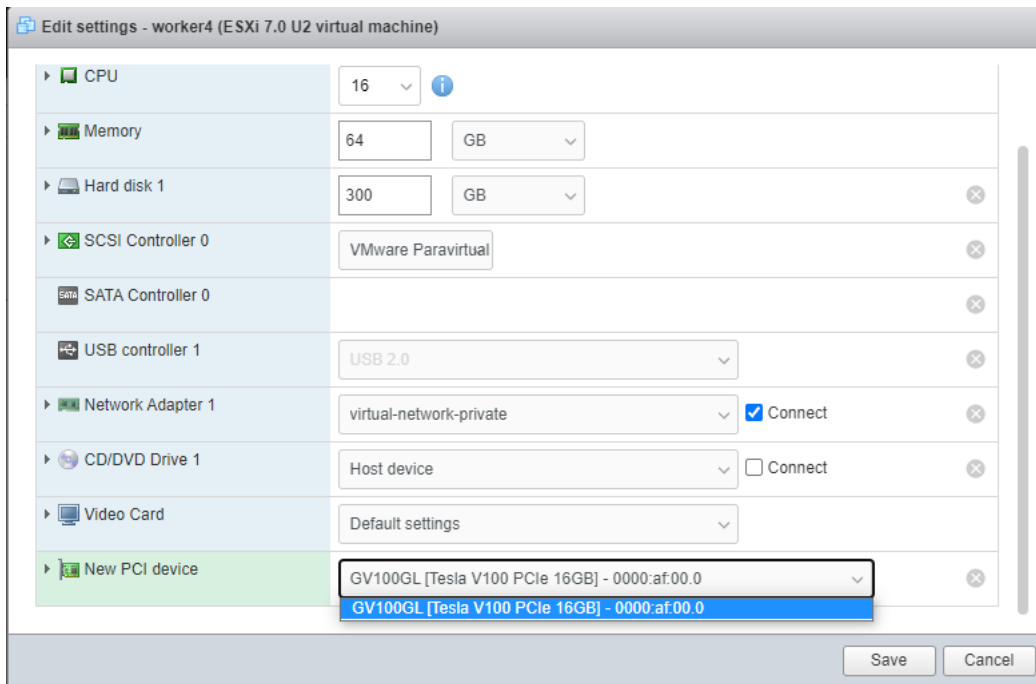


Note:

Now that GPU passthrough is enabled on the GPU devices, the next steps assign the physical GPU devices to VMs that become GPU-enabled compute nodes in the OpenShift Container Platform cluster.

2. Add or modify an existing compute node VM and assign the GPU to the compute node.
3. Edit the VM settings and expand **Memory**. Ensure that the **Reserve all guest memory** checkbox is enabled.
4. From the VM settings, select **Add other device > PCI device**.
5. From the **New PCI device** setting, select the GPU devices from the drop-down menu to assign to the VM. Click **Save**.

For example:



6. Click **VM Options** on the **Edit** settings page.

7. Expand **Boot Options** and ensure that **EFI** is selected from the **Firmware** drop-down menu.

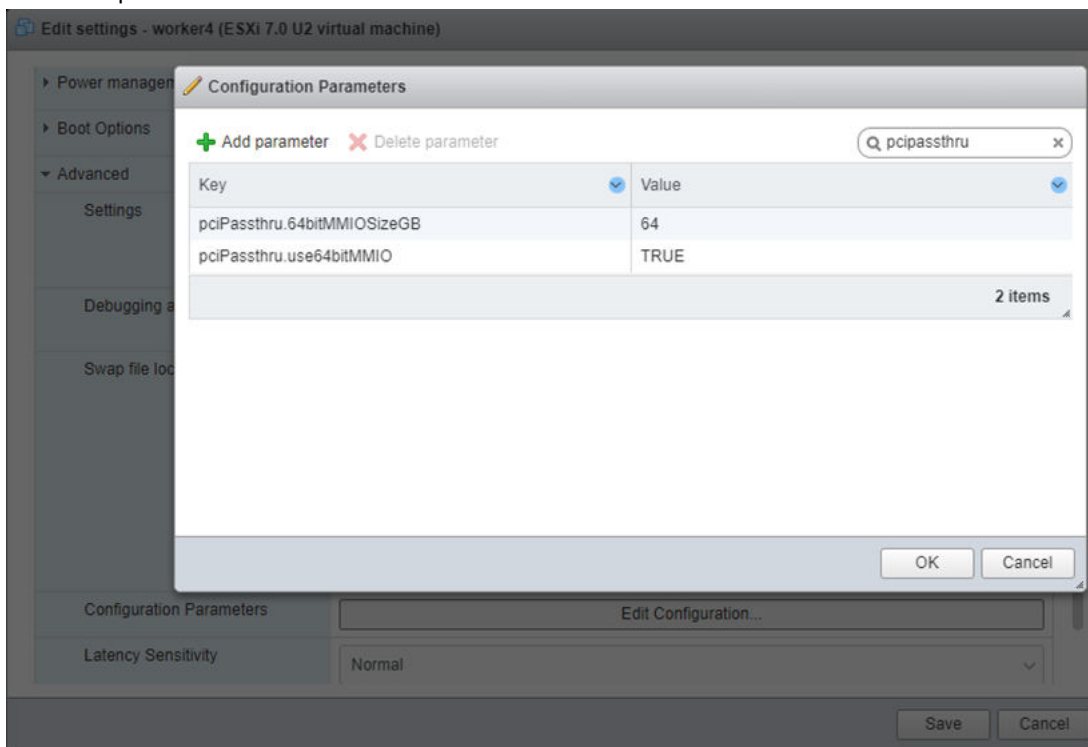
Note: Booting from EFI is a requirement. If booting from BIOS is configured for the VM, the GPU fails to pass through.

8. From **VM Options**, click **Advanced**. Click **Edit Configuration**.

9. Add the following configuration parameters to the compute node VM. Click **Save**.

- **pciPassthru.64bitMMIOSizeGB** to 64
- **pciPassthru.use64bitMMIO** to TRUE

For example:



10. The VM is now assigned the specified GPU resources. Power on the VM and run the following command to confirm that the GPU device is visible to the Guest OS:

```
[core@worker4 ~]$ lspci | grep NVIDIA
13:00.0 3D controller: NVIDIA Corporation GV100GL [Tesla V100 PCIe 16GB] (rev a1)
[core@worker4 ~]$
```

What to do next

Next, install the NVIDIA operator to automate the management of all NVIDIA software components. For more information about installing the NVIDIA operator, see [“Installing the NVIDIA operator” on page 394](#)

Customer-managed **Installing the NVIDIA operator**

The NVIDIA GPU Operator uses the operator framework within Kubernetes to automate the management of all NVIDIA software components that are needed to provision the GPU.

Before you begin

For more information about the NVIDIA GPU Operator, see the following information on the NVIDIA website: [GPU Operator on OpenShift Overview](#).

Complete the following tasks, dependent on your software needs:

- For AWS, see: [“Adding a GPU worker node to a Red Hat OpenShift cluster on AWS” on page 394](#).
- **Note:** For AWS clusters, check if entitlements are necessary in instances that the client is on Red Hat OpenShift 4.9.8 or previous versions.
- For on-premises, see: [“Enabling GPU passthrough” on page 391](#).

Procedure

Install the NVIDIA operator by using the Red Hat OpenShift Container Platform web console and complete the following steps on the NVIDIA site: [Installing the NVIDIA GPU Operator](#).

Tip: This task can also be done by using the following Ansible role: `nvdiia_gpu`. For more information about the Ansible role, see [“IBM Maximo Application Suite installation with Ansible collection” on page 276](#).

What to do next

Next, complete deployment by following the steps in [“Deploying IBM Maximo Visual Inspection” on page 391](#).

Customer-managed **Adding a GPU worker node to a Red Hat OpenShift cluster on AWS**

Before you begin

Ensure that you have the following requirements:

- A GPU worker node is added to the Red Hat OpenShift cluster. . Current AWS Maximo Application Suite BYOL offerings do not include nodes with GPU.
- A control shell.

The control shell can be the boot node. Locate the control shell in the EC2 dashboard after the list of all instances. If the boot node is in a stopped state, restart the instance. Connect to this instance as the EC2-user.

Tip: You can use Visual Studio Code to remotely connect to the boot node, but it is not necessary.

```
oc jq
```

- The appropriate EC2 GPU instance is selected and has sufficient availability in the region where the Maximo Application Suite instance is installed.

Obtain this information from the instance type page that is located after the EC2 service in the AWS console.

For example, if you deployed the Maximo Application Suite instance in the `us-east-1` region, go to the EC2 instance type page for that region by navigating to the [AWS](#) website. The instance type page details the compute, networking, storage, accelerators, and pricing information. The networking section details the availability zones.

About this task

For more information about the processes in this task, see:

- [AWS Recommended GPU Instances](#)
- **Note:** AWS offers EC2 instances that come with GPUs. Use `p3.2xlarge` as the EC2 instance type for MVI.
- [Install & use GPU on AWS](#)
- [Creating a machine set on AWS](#)

Procedure

1. In the control shell, log in as `masocpuser` (or `kubeadmin`).
2. Switch to the `openshift-machine-api` namespace.

```
oc project openshift-machine-api
```

Note:

If the namespace is not switched, use the `-n` flag and provide `openshift-machine-api` as an argument in the succeeding steps.

3. List the available machine sets in the cluster.

```
oc get machineset -o name
```

4. Select an appropriate machine set as a template for the new GPU worker node's YAML custom resource. Pick a machine set that is located in the same availability zone as the GPU EC2 instance type to use to create the new node.
For example, if `p3.2xlarge` is available in `us-east-1b`, pick a machine set that has `us-east-1b` as part of its name.
5. Assign a variable for the template machine set name.
For example,

```
SOURCE_MACHINESET=machine set.machine.openshift.io/masocp-4kyowr-mm5b5-worker-us-east-1b
```

6. Copy the source machine set's custom resource to a new file.

```
oc get -o json $SOURCE_MACHINESET | jq -r > source-machineset.json
```

Note: The file `source-machineset.json` is created in the current folder.

7. Define variables to use for later.

```
OLD_MACHINESET_NAME=$(jq '.metadata.name' -r source-machineset.json)
```

```
NEW_MACHINESET_NAME=${OLD_MACHINESET_NAME}/worker/worker-gpu}
```

- Change the `instanceType` and if needed, change the number of replicas. Delete some metadata and copy the resulting code into a new file `gpu-machineset.json`. This file is used to create the new machine set with the GPU.

```
jq -r '.spec.template.spec.providerSpec.value.instanceType = "p3.2xlarge"
| .spec.replicas = 1
| del(.metadata.selfLink)
| del(.metadata.uid)
| del(.metadata.creationTimestamp)
| del(.metadata.resourceVersion)
| source-machineset.json > gpu-machineset.json'
```

- Change the machine set name in `gpu-machineset.json`.

```
sed -i "s/${OLD_MACHINESET_NAME}/${NEW_MACHINESET_NAME}/g" gpu-machineset.json
```

- Run the **diff** command to check changes.

```
diff -Naur source-machineset.json gpu-machineset.json
```

For more information, see [Install & use GPU on AWS](#).

- Check the value for `availabilityZone` (found under `spec.template.spec.providerSpec.value.placement`). Ensure that the new instance type (`p3.2xlarge`) has the same availability zone, or you can omit the availability key-value pair from the JSON file. If not, an error is displayed after you create the machine set. For more information, see the troubleshooting section at the end of this task.
- Create a machine set:

```
oc create -f gpu-machineset.json
```

Example output

```
machineset.machine.openshift.io/masocp-4kyowr-mm5b5-worker-gpu-us-east-1b created
```

- Verify that the machine set is created.

```
oc get machineset
```

Example output

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
masocp-4kyowr-mm5b5-worker-gpu-us-east-1b	1	1			10s
masocp-4kyowr-mm5b5-worker-us-east-1a	3	3	3	3	7d8h
masocp-4kyowr-mm5b5-worker-us-east-1b	2	2	2	2	7d8h
masocp-4kyowr-mm5b5-worker-us-east-1c	2	2	2	2	7d8h
masocp-4kyowr-mm5b5-workerocs-us-east-1a	1	1	1	1	7d7h
masocp-4kyowr-mm5b5-workerocs-us-east-1b	1	1	1	1	7d7h
masocp-4kyowr-mm5b5-workerocs-us-east-1c	1	1	1	1	7d7h

Note: The output shows that the new GPU node was created but is not ready and available yet.

- Get the list of machines to show the status:

```
oc get machine
```

Example output				
NAME	AGE	PHASE	TYPE	REGION
masocp-4kyowr-mm5b5-master-0		Running	m5.2xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-master-1		Running	m5.2xlarge	us-east-1
us-east-1b	7d8h			
masocp-4kyowr-mm5b5-master-2		Running	m5.2xlarge	us-east-1
us-east-1c	7d8h			
masocp-4kyowr-mm5b5-master-3		Running	m5.2xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-master-4		Running	m5.2xlarge	us-east-1
us-east-1b	7d8h			
masocp-4kyowr-mm5b5-worker-gpu-us-east-1b-nrr4n		Provisioning	p3.2xlarge	us-east-1
us-east-1b	22s			
masocp-4kyowr-mm5b5-worker-us-east-1a-kx449		Running	m5.4xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1a-nn72q		Running	m5.4xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1a-p5nqf		Running	m5.4xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1b-7r5wz		Running	m5.4xlarge	us-east-1
us-east-1b	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1b-94khr		Running	m5.4xlarge	us-east-1
us-east-1b	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1c-fvv52		Running	m5.4xlarge	us-east-1
us-east-1c	7d8h			
masocp-4kyowr-mm5b5-worker-us-east-1c-rsnwf		Running	m5.4xlarge	us-east-1
us-east-1c	7d8h			
masocp-4kyowr-mm5b5-workerocs-us-east-1a-hwb4m		Running	m5.4xlarge	us-east-1
us-east-1a	7d7h			
masocp-4kyowr-mm5b5-workerocs-us-east-1b-979w8		Running	m5.4xlarge	us-east-1
us-east-1b	7d7h			
masocp-4kyowr-mm5b5-workerocs-us-east-1c-85ktb		Running	m5.4xlarge	us-east-1
us-east-1c	7d7h			

When the machine set is done provisioning, the output for `oc get machineset` is similar to the following example:

Example output						
NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE	
masocp-4kyowr-mm5b5-worker-gpu-us-east-1b	1	1	1	1	3m38s	
masocp-4kyowr-mm5b5-worker-us-east-1a	3	3	3	3	7d8h	
masocp-4kyowr-mm5b5-worker-us-east-1b	2	2	2	2	7d8h	
masocp-4kyowr-mm5b5-worker-us-east-1c	2	2	2	2	7d8h	
masocp-4kyowr-mm5b5-workerocs-us-east-1a	1	1	1	1	7d7h	
masocp-4kyowr-mm5b5-workerocs-us-east-1b	1	1	1	1	7d7h	
masocp-4kyowr-mm5b5-workerocs-us-east-1c	1	1	1	1	7d7h	

b) Run the `oc get machine` command. The output indicates that the machine is provisioned:

Example output				
NAME	AGE	PHASE	TYPE	REGION
...				
masocp-4kyowr-mm5b5-master-3		Running	m5.2xlarge	us-east-1
us-east-1a	7d8h			
masocp-4kyowr-mm5b5-master-4		Running	m5.2xlarge	us-east-1
us-east-1b	7d8h			
masocp-4kyowr-mm5b5-worker-gpu-us-east-1b-nrr4n		Provisioned	p3.2xlarge	us-east-1
us-east-1b	107s			
masocp-4kyowr-mm5b5-worker-us-east-1a-kx449		Running	m5.4xlarge	us-east-1
us-east-1a	7d8h			
...				

Note: You can also check the Red Hat OpenShift console, by clicking **Compute > Nodes** or click **Compute > Machinesets**.

What to do next

To verify that the process is completed successfully, or in instances that errors occur, ensure that you run the commands in the `openshift-machine-api` namespace.

Next, run the command `oc create -f <machine set custom resource>` (Step “12” on page 396). The output always indicates that the machine is created. However, if there is a failure in creating the machine, the machine set is not ready and available. Running `oc get machine` can immediately indicate the failure:

Run the `oc get machine` command. The output indicates that the machine is provisioned:

Example output					
NAME	AGE	PHASE	TYPE	REGION	
masocp-qxkempl-wh7px-master-0		Running	m5.2xlarge	us-east-1	us-
east-1a	18h				
masocp-qxkempl-wh7px-master-1		Running	m5.2xlarge	us-east-1	us-
east-1b	18h				
masocp-qxkempl-wh7px-master-2		Running	m5.2xlarge	us-east-1	us-
east-1c	18h				
masocp-qxkempl-wh7px-worker-gpu-us-east-1a-5z7sd		Failed			
				4s	
masocp-qxkempl-wh7px-worker-gpu-us-east-1a-nhldx		Failed			
				20s	
masocp-qxkempl-wh7px-worker-us-east-1a-h2c8g		Running	m5.4xlarge	us-east-1	us-
east-1a	18h				
masocp-qxkempl-wh7px-worker-us-east-1a-p7mt9		Running	m5.4xlarge	us-east-1	us-
east-1a	18h				
masocp-qxkempl-wh7px-worker-us-east-1b-4rlrq		Running	m5.4xlarge	us-east-1	us-
east-1b	18h				
masocp-qxkempl-wh7px-worker-us-east-1b-dhv6g		Running	m5.4xlarge	us-east-1	us-
east-1b	18h				
masocp-qxkempl-wh7px-worker-us-east-1c-ks85p		Running	m5.4xlarge	us-east-1	us-
east-1c	18h				
masocp-qxkempl-wh7px-workerocs-us-east-1a-9r6pj		Running	m5.4xlarge	us-east-1	us-
east-1a	17h				
masocp-qxkempl-wh7px-workerocs-us-east-1b-p9psl		Running	m5.4xlarge	us-east-1	us-
east-1b	17h				
masocp-qxkempl-wh7px-workerocs-us-east-1c-94d7q		Running	m5.4xlarge	us-east-1	us-
east-1c	17h				

To see the reason for the failure, run `oc describe machine <machine name>` or `oc describe machineset <machineset name>` and check the error message that is listed after **Status** or **Events**:

```

Status:
Conditions:
  Last Transition Time: 2022-05-26T15:20:25Z
  Message: Instance has not been created
  Reason: InstanceNotCreated
  Severity: Warning
  Status: False
  Type: InstanceExists
Error Message: error launching instance: Your requested instance type (p3.2xlarge)
is not supported in your requested Availability Zone (us-east-1a). Please retry your request by
not specifying an Availability Zone or choosing us-east-1b, us-east-1c, us-east-1d, us-east-1f.
Error Reason: InvalidConfiguration
Last Updated: 2022-05-26T15:20:26Z
Phase: Failed
Provider Status:
Conditions:
  Last Probe Time: 2022-05-26T15:20:26Z
  Last Transition Time: 2022-05-26T15:20:26Z
  Message: error launching instance: Your requested instance type (p3.2xlarge)
is not supported in your requested Availability Zone (us-east-1a). Please retry your request by
not specifying an Availability Zone or choosing us-east-1b, us-east-1c, us-east-1d, us-east-1f.
Reason: MachineCreationFailed
Status: False
Type: MachineCreation
Events:
  Type      Reason      Age      From      Message
  ----      -
Warning FailedCreate 52s (x2 over 53s) awscontroller masocp-qxkeml-wh7px-worker-gpu-us-
east-1a-5z7sd: reconciler failed to Create machine: failed to launch instance: error launching
instance: Your requested instance type (p3.2xlarge) is not supported in your requested
Availability Zone (us-east-1a). Please retry your request by not specifying an Availability Zone
or choosing us-east-1b, us-east-1c, us-east-1d, us-east-1f.

```

In this case, you can delete the machine set:

```
oc delete machineset <machineset name>
```

Edit the `availabilityZone` value in the custom resource and rerun `oc create -f <customresource.json>`. Monitor the creation of the machine set and machines by using the commands that are listed in step “13” on page 396. For any other types of errors, delete the machine set, edit the custom resource, and re-create the machine set by using the edited custom resource JSON file.

Customer-managed **Creating and applying YAML files for deploying IBM Maximo Visual Inspection**

YAML files contain information that is related to security constraints and access permissions, such as cluster roles and cluster role bindings. When you create and apply YAML files with Maximo Visual Inspection, you can view important information such as graphical processing unit (GPU) usage statistics.

Before you begin

Check that you have the following resources for creating and updating YAML files:

- A text editor or your preferred application to create the YAML files.
- [Red Hat OpenShift CLI](#) to run the commands.

Note: Your environment might already be set up to run commands. For example, if you are using Maximo Application Suite on Amazon Web Services, you can use the boot node as the command shell.

Procedure

1. Open a command shell on the Red Hat OpenShift cluster.
2. Connect as the cluster administrator by using the `oc login` command.
3. Create a project in the Red Hat OpenShift cluster and name it using the following format: `mas-<instanceId>-visualinspection`

Note:

If you are using Maximo Application Suite on Amazon Web Services, change *InstanceId* to the following format: *mas-ClusterUniqueString*.

4. Create a file that is named `customsccl.yaml` and paste in the following text:

```
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities:
- CHOWN
- DAC_OVERRIDE
- FOWNER
- FSETID
- KILL
- SETGID
- SETUID
- SETPCAP
- NET_BIND_SERVICE
- NET_RAW
- SYS_CHROOT
allowedUnsafeSysctls: null
apiVersion: security.openshift.io/v1
defaultAddCapabilities: null
fsGroup:
  type: RunAsAny
  ranges:
  - max: 65535
    min: 1
groups: []
kind: SecurityContextConstraints
metadata:
  annotations:
    kubernetes.io/description: "This policy is the most restrictive for IBM Maximo Visual
Inspection."
  name: ibm-mas-visualinspection-scc
readOnlyRootFilesystem: false
requiredDropCapabilities:
- ALL
runAsUser:
  type: MustRunAsRange
  uidRangeMax: 65535
  uidRangeMin: 0
selinuxContext:
  type: RunAsAny
seccompProfiles: null
supplementalGroups:
  type: MustRunAs
  ranges:
  - max: 65535
    min: 1
users: []
volumes:
- configMap
- downwardAPI
- emptyDir
- persistentVolumeClaim
- projected
- secret
```

5. Create a file that is named `clusterrole.yaml` and paste in the following text:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: ibm-mas-visualinspection-clusterrole
rules:
- apiGroups:
  - 'security.openshift.io'
  resources:
  - 'securitycontextconstraints'
  resourceNames:
  - 'ibm-mas-visualinspection-scc'
```



```

verbs:
- use
- apiGroups:
- ""
resources:
- nodes
- pods
verbs:
- list

```

6. Create a project in the Red Hat OpenShift cluster and name it using the following format: mas-
<instanceId>-visualinspection.

For Maximo Application Suite on Amazon Web Services, the instanceId has the following format: mas-<ClusterUniqueString>.

7. Create a file that is called clusterrole_binding.yaml. Paste in the following text and replace <MVI_deployment_namespace> with the name of the project that you created.

```

kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: ibm-mas-visualinspection-clusterrolebinding
subjects:
- kind: ServiceAccount
  name: ibm-mas-visualinspection-operator
  namespace: <MVI_deployment_namespace>
roleRef:
  kind: ClusterRole
  name: ibm-mas-visualinspection-clusterrole
  apiGroup: rbac.authorization.k8s.io

```

8. To apply the files that you created, run the following commands:

```

oc apply -f customssc.yaml
oc apply -f clusterrole.yaml
oc apply -f clusterrole_binding.yaml

```

Customer-managed **Completing the deployment of Maximo Visual Inspection**

By default, the application is not available to any users until it is activated and users are granted access.

Before you begin

The persistent volume claim (PVC) is automatically created when you complete the deployment of the Maximo Visual Inspection application. To specify the configuration details of the PVC for deployment, check that you have the following resources:

- A storage class that supports the ReadWriteMany access mode. For more information, see [ReadWriteMany](#)

If the storage class that you specify does not support the ReadWriteMany access mode or the class name does not exactly match an available Red Hat OpenShift storage class name, the deployment fails.

- A storage class that is supported by your cloud vendor.

The default storage class is `ibmc-file-gold` and is suitable for deployments to IBM Cloud. However, some cloud vendors do not support the default storage class.

To see a list of available storage classes, use the following command in the Red Hat OpenShift console:

```
oc get sc
```

Procedure

1. In Maximo Application Suite, open the **Suite Administration Catalog** page and then click **Visual Inspection**.

- Optional: Select the number of IBM Maximo Visual Inspection Edge licenses that you want to use and then click **Continue**.

By using IBM Maximo Visual Inspection Edge, you can call models that were trained in Maximo Visual Inspection to complete inspections from edge devices.

Note: When an insufficient number of AppPoints are available for deploying, you can continue to complete the application configuration. The application is automatically deployed when the sufficient number of AppPoints are available.

- Optional: To keep the application up to date by subscribing to the upgrade channel, click **Channel subscription**.

When new versions are available, they are added to the channel and updated in your suite instance either automatically or after manual approval.

- Optional: To keep the application up to date by selecting the version that you want to manually deploy, click **Manual**.

If you want to deploy IBM Maximo Visual Inspection version 8.5.0, complete the remaining steps in this procedure. If you want to deploy an earlier version, complete the deployment steps for that version.



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription”](#) on page 482.

- Select an update method.
To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.
To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.
 - Subscribe to a channel by selecting a version from the list.
For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
 - Click **Subscribe to channel <version>**
- Optional: On the **Deploy application** page, configure the IBM Maximo Visual Inspection Edge add-on.
For more information, see [“Deploying IBM Maximo Visual Inspection Edge”](#) on page 408.
 - On the **Deploy application** page, click **Show advanced settings** and then set **System managed** to off.
 - For the PVC configuration, specify the storage class and the storage size.

To specify the minimum storage size of 100 gibibytes, enter 100 Gi.

If you are using the `ocs-storagecluster-cephfs` storage class, complete the following steps as a workaround.

Initialize the PVC:

- Add `anyuid scc` to the Maximo Visual Inspection service account by entering the following command:

```
oc adm policy add-scc-to-user anyuid system:serviceaccount:
{{ mas_app_namespace }}:ibm-mas-visualinspection-operator
```

When the application is ready for deployment:

- Apply the custom security context constraints (SCC).
- Remove `anyuid scc` from the Maximo Visual Inspection service account by entering the following command:

```
oc adm policy remove-scc-from-user anyuid system:serviceaccount:
{{ mas_app_namespace }}:ibm-mas-visualinspection-operator
```

8. Apply the changes and click **Deploy**

On the **Applications** page for Visual Inspection, you can monitor the deployment status. If the deployment fails, delete and then redeploy the application. Deployment is complete when the **Application** card displays the Monitor is ready message and the **Activate** button is displayed. Click **Activate**.

If the deployment fails, delete and redeploy the application.

What to do next

[“Activating IBM Maximo Visual Inspection” on page 469](#)

Customer-managed **Deploying IoT tool**

The IoT tool provides device connectivity, data filtering and mapping, and device management tools. Maximo Monitor Maximo Monitor requires the IoT tool.

Customer-managed **Installing Apache Kafka for IoT tool**

Apache Kafka provides a buffer for messages that are sent to and received from external interfaces. Apache Kafka is needed only in instances where Maximo Manage interfaces with external systems.

What to do next

Configure Apache Kafka Suite parameters. For more information, see [“Apache Kafka” on page 23](#).

Customer-managed **Deploying IoT tool**

The IoT tool provides device connectivity, data filtering and mapping, and device management tools and is needed by Maximo Monitor.

Before you begin

The following components must be installed and configured before you deploy the IoT tool:

- [“Db2 Warehouse” on page 21](#).

Note: For Maximo Application Suite on AWS (BYOL) offerings, an instance is already created and the `jdbccfg` custom resource is already configured at the System scope.

- [“MongoDB” on page 21](#).

Note: For Maximo Application Suite on AWS (BYOL) offerings, MongoDB is installed and configured for you.

- [Apache Kafka for IoT tool](#).

About this task

These deployment steps are completed in Maximo Application Suite.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Catalog**. Next, select the **Tools** tab and then click the **IoT** card.
2. Click **Continue**.
3. Select your application update method.

- a) Select an update method.
To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.
To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.
 - b) Subscribe to a channel by selecting a version from the list.
For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
 - c) Click **Subscribe to channel <version>**
4. Configure advanced configurations.
On the Deploy application screen, if you want to customize your settings for storage class and deployment size, click **Show advanced settings**. Set the **System managed** toggle for the settings that you want to configure manually.
 5. Select a storage class and size for the Persistent Volume Claims (PVCs) which is used to persist messaging data.
 6. Choose from developer, small, or medium.
For more information about available and tested storage classes, and for deployment size guidance, see the Monitor and IoT section of the [Maximo Application Suite system requirements](#) document.
 7. Click **Deploy**.
 8. Verify the deployment configuration and click **Begin deployment**.
The estimated deployment time is an estimate of the time that it takes to configure and deploy the application. The time includes processing and configuration.

Monitor the environment. Check that the namespace `mas-<instanceid>-iot` is created successfully. Check the pods to ensure that no error conditions are present and wait for all pods to be ready.
 9. When the deployment is complete, click **Activate**.

What to do next

[“Deploying IBM Maximo Monitor” on page 370](#)

Related concepts

IoT tool

The IoT tool is an add-on to Maximo Application Suite. By using the IoT tool, you can access device management tools and configure device connectivity, data filtering, and data mapping.

Customer-managed **Deploying industry solutions**

The industry solutions that are available for your installed Maximo Application Suite version can be configured for use with your environment.

About this task

Note: The industry solutions that are described in the proceeding topics apply to industry solutions that can be deployed in Maximo Application Suite, excluding industry solutions deployed in Maximo Manage. In Maximo Application Suite, from the side navigation menu, click **Catalog** and then select **Industry solutions**. For more information about industry solutions that are included with Maximo Manage, see [“Deployment of industry solutions and add-ons” on page 326](#).

Related concepts

[Maximo Application Suite Industry solutions](#)

On the Industry solutions tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove industry solutions.

Deploying IBM Maximo Health and Predict - Utilities

Maximo Health and Predict - Utilities supports maintenance, operations, and performance of assets and networks for energy and utility companies.

Before you begin

Deploy Maximo Health before you deploy Maximo Health and Predict - Utilities. For more information, see [Deploy Maximo Health](#)

Note: Complete all steps to deploy and activate Maximo Health and enable asset investment optimization. If Maximo Health is not deployed as a part of Maximo Manage, do not complete the system integration or use the data loader until after Maximo Health and Predict - Utilities is deployed and activated. Maximo Health and Predict - Utilities uses the same database as Maximo Health.

The following components must be configured before you deploy Maximo Health and Predict - Utilities:

IBM Watson Studio on Cloud Pak for Data.

Create at least one project and add the notebook asset type to the project. For more information, see [Installing Watson Studio](#).

Maximo Scheduler Optimization.

Maximo Scheduler Optimization is required for asset investment optimization, which you enable as part of deploying the Maximo Health dependency. You do not need to configure the `optimization.mofapi` or `optimization.mofui` system properties after you deploy and activate Maximo Scheduler Optimization. However, if you enable asset investment optimization and you are using the limited version of Maximo Scheduler Optimization, all other active models in that Maximo Scheduler Optimization instance are disabled.

If you load data into Maximo Health and Predict - Utilities by using App Connect from Maximo Application Suite 8.8, App Connect 12.0.4 or later, it is supported. The App Connect dashboard must use an App Connect Enterprise Production license. Install App Connect in the Red Hat OpenShift Container Platform cluster. For more information, see [Configuring IBM App Connect](#).

If you use App Connect to load data, in the App Connect **Dependencies** section,, click **Configure**. Then, specify the App Connect dashboard URL and user.

Note:

If App Connect is used to load data, before you deploy Maximo Health and Predict - Utilities, configure App Connect.

Set **spec.UseCommonServices** to **false**. Ensure that there are no API keys or credentials in the App Connect dashboard. If any API keys or credentials exist, delete them from App Connect, enable authentication, and add them back. For more information, see [Editing the settings for a deployed integration server](#) and [Configuration reference](#).

After deployment and activation are complete for Maximo Health and Predict - Utilities, to enable the authentication option, set the **spec.UseCommonServices** to **true**.

About this task

Note:

Starting in Maximo Application Suite 8.11, Maximo Health and Predict - Utilities is no longer available as a separate industry solution. The information that is provided is applicable only to Maximo Application Suite 8.10 and earlier versions. For more information, see [“Upgrading IBM Maximo Application Suite” on page 473](#). Before you upgrade to Maximo Application Suite 8.11, deactivate and delete Maximo Health and Predict - Utilities.

The following steps are specific to the Maximo Health and Predict - Utilities industry solution. Complete the predeployment and deployment steps before you activate the industry solution. After you activate the industry solution, you must grant users access to it.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Catalog** and then select **Industry solutions**. Click the tile.
2. On the Maximo Health and Predict - Utilities catalog page, verify the information.
3. Click **Continue**.

Note: If insufficient AppPoints are available to deploy this application, you can still complete the application configuration. The application is automatically deployed when the required number of AppPoints are available.

4. Determine how you want to keep the industry solution up to date.



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription”](#) on page 482.

5. Click **Deploy**. After you click **Deploy**, the deployment process starts in Red Hat OpenShift. The estimated deployment time is an estimate of the time that it takes to configure and deploy the industry solution. The time includes both processing and configuration.

What to do next

Track the deployment process on the **Details** page or as an administrator, log in to the Red Hat OpenShift web console and from the Red Hat OpenShift cluster, in the navigation menu, click **Workloads > Pods**. Ensure that you are reviewing either all projects or only the Maximo Health and Predict - Utilities, which is named `mas-instance name-hutilities`. The Maximo Health and Predict - Utilities operator supports single namespace deployment.

Login to the Red Hat OpenShift cluster. Next, check the Maximo Health and Predict - Utilities deployment status by running the Red Hat OpenShift client command line.

```
oc get hutilitiesapp -n {{mas-instance name-hutilities}}
```

The following output shows an example of the status:

NAME	VERSION	STATUS	AGE
devtest	8.5.0	Ready	2d1h

Maximo Health and Predict - Utilities uses the Ready condition reason code of the custom resource to indicate progress. The following reason codes are used:

Ready

The Health and Predict Utilities workspace or application Operator custom resource is up to date and ready to use.

CompatibilityCheckFailure

The operator version is not compatible with Maximo Application Suite or the Maximo Health version.

AppPointReservationFailed

Failed to reserve the app points.

As the industry solution is deployed, the following processes happen in the following order in the Red Hat OpenShift:

1. The `ibm-mas-hutilities-operator-string` pod is created, and the status changes to **Running**.
2. A `instance name -hutilities-entymgr-ws-string` pod is created. You can open the logs for these pods to monitor the deployment. The status changes to **Running**.
3. Three `instance name -hutilities-model-engine-string` pods are created, and the status changes to **Running**.

The strings are randomly generated sets of numbers and letters.

After the status of all pods is **Running**, the industry solution is ready to be activated. Activate Maximo Health and Predict - Utilities to configure a workspace and give users access to the industry solution.

Related concepts

[Maximo Application Suite Industry solutions](#)

On the Industry solutions tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove industry solutions.

Related tasks

[Activating Maximo Health and Predict - Utilities](#)

Related information

[Getting started with Maximo Health and Predict - Utilities](#)

Customer-managed

Configuring IBM MRO Inventory Optimization in Maximo Application Suite

In version 8.11 and earlier, MRO Inventory Optimization is a stand-alone but linked product that requires an externally purchased license. To give users access to MRO Inventory Optimization in Maximo Application Suite, you can add MRO Inventory Optimization as an external launcher.

Before you begin

Install MRO Inventory Optimization in Maximo Application Suite before configuration.

For more information, see the [MRO Inventory Optimization product page](#)

MRO Inventory Optimization is not used with any other Maximo Application Suite component.

About this task

Note:

Starting in Maximo Application Suite 9.0, MRO Inventory Optimization is no longer available to be added as an externally configured application and must be accessed by the dedicated URL. The information that is provided is applicable to Maximo Application Suite 8.11 and earlier versions. If MRO Inventory Optimization is configured as an external launcher and you are upgrading to Maximo Application Suite 9.0, you must remove MRO Inventory Optimization before you can complete the upgrade.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Workspace** and then select the **External launchers** tab.
2. Click **Edit** and select MRO Inventory Optimization.
3. Enter the solution portal URL of your MRO Inventory Optimization environment.

Important: The Solution portal URL must be properly formed, including the initial `http://` or `https://`.

4. Save your changes.

After MRO Inventory Optimization is successfully configured, you can access the product login page from the Suite navigator.

What to do next

To remove MRO Inventory Optimization, you can delete the URL for MRO Inventory Optimization on the **External launchers** tab. When the product URL is removed, users can no longer access MRO Inventory Optimization from the suite navigator.

Related concepts

[IBM MRO Inventory Optimization](#)

Customer-managed **Deploying add-ons**

Add-ons extend the capabilities of Maximo Application Suite, its applications, and its tools. Add-ons that are available for your installed Maximo Application Suite version can be configured for use with your environment.

Customer-managed **Deploying IBM App Connect**

With App Connect, you can connect applications and data from existing systems and modern technologies across all their environments.

Before you begin

Install and configure App Connect. You need the application dashboard URL to complete the configuration.

For more information, see the [IBM App Connect documentation](#).

Maximo Application Suite supports cloud-based and on-premises App Connect:

Cloud-based

For more information, see [IBM App Connect on Cloud](#) and [Getting started with App Connect](#).

On-premises

Follow the procedure that is outlined in the [Maximo Application Suite download document](#) to access and download the software and the related installation and quick start documentation.

Important: The Maximo Health and Predict - Utilities industry solution requires an on-premises instance of IBM App Connect. For Maximo Application Suite 8.8 and 8.9, App Connect 12.0.4 and 12.0.5 are supported. For Maximo Application Suite 8.10, App Connect 12.0.6 and 12.0.7 are supported. Also, the dashboard must use an App Connect Enterprise Production license. Install IBM App Connect in a Cloud Pak for Data cluster. For more information, see [Applying service to IBM App Connect Enterprise](#).

To install IBM App Connect Enterprise that is needed to install the IBM Maximo Health and Predict - Utilities, see the ansible role [appconnect](#).

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Configurations**, and click **App Connect**.
2. Click **Configure**.
3. Configure **App Connect**:
 - a) Add the dashboard URL for your instance of App Connect.
For example, the format is `http://dash-1204-xxx.cloud`. The dashboard URL must be correct.
Do not add extra `/`.
4. Click **Save**.

Related concepts

[IBM® App Connect Enterprise](#)

[IBM App Connect](#)

App Connect is an add-on to Maximo Application Suite. By using App Connect, you can connect applications and data from existing systems and modern technologies across all their environments.

Customer-managed **Deploying IBM Maximo Visual Inspection Edge**

By using Maximo Visual Inspection Edge, you can call or deploy multiple models that were trained in Maximo Visual Inspection Edge to conduct inference operations from edge devices.

Before you begin

Before you can configure IBM Maximo Visual Inspection Edge, you must deploy IBM Maximo Visual Inspection. For more information, see [Deploying IBM Maximo Visual Inspection](#).

On the system that you want to install Maximo Visual Inspection Edge, ensure that all hardware, platform, and software prerequisites are met. For more information, see [Planning for IBM Maximo Visual Inspection Edge](#).

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Configurations** and in the **Other** section, select **Visual Inspection Edge**.
2. Click **Configure**.
3. On the **Parameters** tab, describe the location and purpose of the edge deployment.
For example, enter a location of *Detroit car assembly plant* and a name of *Check gearbox assembly*
4. Click **Save**.
5. Install Maximo Visual Inspection Edge on edge devices.
 - a) Install IBM Maximo Visual Inspection Edge. For more information, see [Installing and uninstalling IBM Maximo Visual Inspection Edge](#).

Customer-managed

Deploying IBM Parts Identifier

Parts Identifier is an IBM Cloud service that can be used as an add-on with IBM Maximo Manage and IBM Maximo Mobile. It enables technicians to search for and identify industrial parts on a mobile device.

Before you begin

Before you can configure the Parts Identifier add-on in Maximo Application Suite, purchase a license for Parts Identifier from IBM. After you purchase a license, credentials are provided to use to in the configuration in Maximo Application Suite.

About this task

The Parts Identifier application is not integrated with Maximo Application Suite. It is used as a stand-alone but linked product that requires an externally purchased license.

Note: Starting in Maximo Application Suite 8.11, Parts Identifier is no longer available. If Parts Identifier is deployed and active in your environment and you are upgrading to Maximo Application Suite 8.11, you must deactivate and delete Parts Identifier before you can complete the upgrade.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Configurations** and then in the **Other** section, click **Parts Identifier**.
2. Click **Configure**.
3. In the **Username** and **Password** fields, enter the credentials that you received from Humai.
4. Click **Save** to complete the configuration.

What to do next

After you configure the Parts Identifier add-on in Maximo Application Suite, you must complete the following steps:

- Connect Maximo Manage to the Humai cloud service to enable Maximo Manage to send the parts data to the service.
- Send images of Maximo Manage parts to Humai to train an AI model to identify images. For more information about these steps, see [Configuring the Parts Identifier integration with Maximo Manage](#).

Related concepts

IBM Parts Identifier

IBM Parts Identifier is an IBM Cloud service that can be used as an add-on to IBM Maximo Manage and IBM Maximo Mobile in Maximo Application Suite. By using Parts Identifier, you can search for and identify industrial parts from a mobile device.

Customer-managed **Deploying IBM Maximo Optimizer**

By using Maximo Optimizer, you can automate efficient decisions for long-range plans, schedules, and the dispatch of resources for asset maintenance, helping to balance competing objectives and constraints.

Before you begin

Note: If Maximo Scheduler Optimization is installed and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization. Then, you can complete the upgrade and deploy Maximo Optimizer.

Before you deploy Maximo Optimizer, install and configure MongoDB.

About this task

The following steps are specific to Maximo Optimizer and are part of the overarching deployment process. Select your application update method. To later change from channel subscription versioning to manual versioning, you delete and then redeploy the application. After you activate the application, you must grant users access to it.

Procedure

1. In Maximo Application Suite, from the side navigation menu, click **Catalog** and then on the **Add-ons** tab, click the **Optimizer** tile.
2. On the Optimizer deployment page, select a pricing plan for deployment.
 - Optimizer
 - Optimizer Limited
3. Click **Continue**.
4. Select your application update method.
 - a) Select an update method.

To subscribe to automatic updates, set **Automatic approval** to **On**. When new application updates are available, they are added to the channel and automatically updated in your Maximo Application Suite instance.

To subscribe to manual updates, set **Automatic approval** to **Off**. When new application updates are available, you receive a notification, and you can manually approve the updates.
 - b) Subscribe to a channel by selecting a version from the list.

For example, select channel 8.x.x, 8.x, 9.x.x, or 9.x.
 - c) Click **Subscribe to channel <version>**
5. On the **Deploy** application page, confirm that MongoDB is configured.
6. Click **Deploy** and then click **Begin deployment**. The estimated deployment time is an estimate of the time that it takes to configure and deploy the application. The time includes both processing and configuration.

What to do next

After deployment, Maximo Optimizer is not immediately available. Activate Maximo Optimizer and grant users access.

Related concepts

IBM Maximo Optimizer

IBM Maximo Optimizer is an add-on to Maximo Application Suite. By using Maximo Optimizer, you can automate efficient decisions for long-range planning, scheduling, and dispatching of resources for asset maintenance while balancing competing objectives and constraints.

Related tasks

Activating IBM Maximo Optimizer

Before Maximo Optimizer is available for use, you must activate Maximo Optimizer. Activating the application does not automatically grant your users access to the application.

Modifying optimization system properties

To run optimization, system properties for Maximo Optimizer are automatically enabled during deployment and activation. If needed, you can modify the system properties in Maximo Manage.

Upgrading from Maximo Scheduler Optimization to Maximo Optimizer

In Maximo Application Suite 8.8, you use Maximo Optimizer instead of Maximo Scheduler Optimization. If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade.

Customer-managed

Modifying optimization system properties

To run optimization, system properties for Maximo Optimizer are automatically enabled during deployment and activation. If needed, you can modify the system properties in Maximo Manage.

About this task

You might want to adjust the system properties in certain instances. For example, in Maximo Application Suite, if Maximo Optimizer is deployed in a test instance, you can change the system properties to configure a development system to use a test environment.

The API key and URL that are needed for optimization are located in the Maximo Optimizer application. To view the URL and API key for the optimization project, in Maximo Optimizer, click **Projects** and expand **Rest API URL**. Modify the global value for the following system properties and then run optimization.

System Property	Global Value
optimization.mofapi.apikey	The API key for the Maximo GS/GA/GSLP optimization project.
optimization.mofapi.baseurl	<code>https://{instanceid}-{workspaceid}-api.mas-{instanceid}-optimizer.svc</code>
optimization.mofui.url	<code>https://{workspaceId}.optimizer.{masdomain}.com</code>

Procedure

1. In Maximo Manage, from the side navigation menu, click **System Configuration > Platform Configuration > System Properties**.
2. In the **optimization.mofapi.apikey** system property, add the API key for your optimization project.
 - a) Click the **Open Filter** icon, search for **optimization.mofapi.apikey**, and then expand the property.
 - b) In the **Global Value** field, enter the API key for the **Maximo GS/GA/GSLP** optimization project.
 - c) From the **Common Actions** menu, click **Save Property**.
 - d) Select the checkbox for the property name that you updated and then click **Live Refresh**.
 - e) In the **Live Refresh** window, confirm the changes that you made and click **OK** to refresh.

3. In the **optimization.mofapi.baseurl** system property, specify `https://{instanceid}-{workspaceid}-api.mas-{instanceid}-optimizer.svc` for the global value.
4. In the **optimization.mofui.url** system property, specify `https://{workspaceId}.optimizer.{masdomain}.com` for the global value.

Related concepts

IBM Maximo Optimizer

IBM Maximo Optimizer is an add-on to Maximo Application Suite. By using Maximo Optimizer, you can automate efficient decisions for long-range planning, scheduling, and dispatching of resources for asset maintenance while balancing competing objectives and constraints.

Related tasks

Deploying IBM Maximo Optimizer

By using Maximo Optimizer, you can automate efficient decisions for long-range plans, schedules, and the dispatch of resources for asset maintenance, helping to balance competing objectives and constraints.

Maximo AI Service and AI features in Maximo Manage

You can enable AI features for Maximo Manage by deploying Maximo AI Service. Maximo AI Service is an integrated add-on for Maximo Application Suite 9.1. Maximo AI Service enables access to the AI assistant, field value recommendations, including problem code recommendations for work orders, locating similar work orders, and AI recommendations in Reliability Strategies.

Note:

The AI broker, which was introduced in 9.0, is replaced with Maximo AI Service as of 1 August 2025. To continue using the features that were enabled by the AI broker after that time, you must uninstall any instance of the broker and then deploy and use Maximo AI Service 9.1. You can deploy Maximo AI Service 9.1 with Maximo Application Suite 9.0 or 9.1. If Maximo AI Service is deployed with Maximo Application Suite 9.0, you can use only the AI features that were included in Maximo Application Suite 9.0. For more information about uninstalling the AI broker, see [Uninstalling the AI broker](#).

Note: Maximo AI Service is supported by an integration with IBM watsonx. Maximo AI Service 9.1 includes a limited use license to watsonx.ai™ and incurs an additional AppPoint cost. For more information, see [Licensing in Maximo Application Suite 9.1](#).

Maximo AI Service overview

Maximo AI Service is the integration hub that facilitates communication between Maximo Manage and watsonx AI systems or services.

Maximo AI Service facilitates the following tasks:

- Manages configuration, training, and retraining AI models and retains data during training.
- Delegates inferencing jobs to watsonx AI or to a local embedded runtime.
- Completes health checks of the AI model runtime and the individual models.

Maximo AI Service supports multitenancy. Model inferencing and training supports only the English language. Data that is not in English cannot be processed as part of inferencing or used to generate output. To enable Maximo AI Service in production, development, and testing environments, you must enable Maximo AI Service individually in each environment type. Maximo AI Service requires a unique tenant ID per environment. You cannot track AppPoint usage for Maximo AI Service in Maximo Application Suite.

Maximo AI Service also enables some AI features in Maximo IT. For more information, see [Integrating AI with Maximo IT](#).

Data privacy and security

When used with Maximo AI Service, watsonx.ai does not retain inference data or results for inferencing. Data sent to watsonx.ai is encrypted in transit by using TLS 1.2. Exchanges with watsonx.ai are stateless and not stored. No data is retained by watsonx.ai remote services.

For MCC and PCC models, model training occurs local to Maximo Application Suite as part of Maximo AI Service, which is hosted in Red Hat OpenShift. After model training, the trained model runs locally in Red Hat OpenShift. Production data for inferencing for PCC models is not sent to the cloud except for data that the administrator selects for a preprocessing synthetic data generation step that leverages watsonx.ai.

For more information about security and privacy practices, see [IBM Trust Center](#) and [IBM Office of Privacy and Responsible Technology](#).

Enabling AI features by using Maximo AI Service on-premises

You can deploy Maximo AI Service on-premises with Maximo Application Suite 9.0 or 9.1. For example, you can deploy Maximo AI Service 9.1 on-premises on Maximo Application Suite 9.0.

The following list contains high-level steps to enable AI features by using Maximo AI Service 9.1 on-premises:

1. If you installed the AI broker, uninstall the broker and MariaDB. For more information, see [Uninstalling the AI broker](#).
2. If you started a deployment for Maximo AI Service 9.1 and have MariaDB as part of that deployment, uninstall that version of MariaDB. Earlier versions of Maximo AI Service required MariaDB.
 - a. Open Red Hat OpenShift web console.
 - b. From the side navigation, click **Home > Projects**.
 - c. Search for the mariadb project name.
 - d. For the project, click the three-dot menu and then click **Delete Project**.
3. Deploy Maximo AI Service.

To deploy Maximo AI Service, you must first set up and configure the prerequisite software, including watsonx.ai. You can then complete the deployment by using a CLI or Ansible collection, connect Maximo AI Service to Maximo Manage, and then verify that Maximo AI Service is running and connected. For more information, see [Deploying Maximo AI Service](#).

4. Create AI configurations for the AI features that you want to enable.

You create AI configurations in Maximo Manage in the AI configuration application. Each AI feature that you want to enable requires its own configuration.

The following table contains the available AI features, associated model template name, the required Granite™ model, and links to procedures that describe how to set up the AI configuration, including preparing data.

Feature	Model template	Model	Steps
Problem code recommendations for work orders	pcc	Granite 3.0 8B Instruct	“Enabling recommended problem codes for Work orders” on page 446
Field value recommendations	mcc	Granite 3.2 8B Instruct	Enabling field value recommendations
AI assistant	nl2oslc	Granite 3.2 8B Instruct	Enabling the assistant

Feature	Model template	Model	Steps
Locating similar work orders	similarity	Granite 3.0 8B Instruct	“Enabling locating of similar work orders” on page 449
AI recommendations for asset boundary and failure list in Reliability Strategies	fmea	Granite 3.2 8B Instruct	“Enabling AI recommendations in Reliability Strategies” on page 455

Note: Model training for mcc and pcc models can use significant CPU resources. Ensure that your Red Hat OpenShift cluster can handle the load. Inferencing occurs locally in the cluster and does not consume as significant CPU resources as training.

To determine what version of Maximo AI Service you are using, in the Red Hat OpenShift web console, in the navigation menu, click **Operators > Installed operators** and locate the Maximo AI Service operator. The version is listed. If the operator does not exist, Maximo AI Service is not installed.

Enabling AI features by using Maximo AI Service SaaS

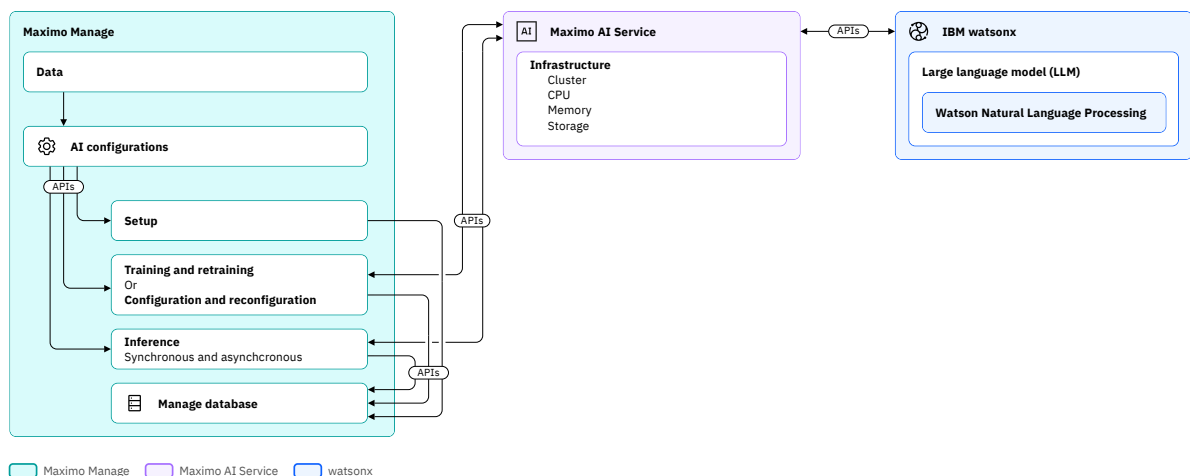
To deploy Maximo AI Service SaaS, contact your IBM representative. The following table describes the steps that you must complete depending on how Maximo Application Suite and Maximo AI Service SaaS are deployed.

Maximo AI Service SaaS cannot be used with watsonx.ai on-premises.

Architecture overview

Maximo Application Suite includes a foundation for integrating AI into Maximo Manage by using the Maximo Manage extensible framework.

[Open image in new tab.](#)



Metadata coordinates the configuration of AI models for the Maximo Manage AI configuration framework. Metadata also binds AI models to Maximo Manage application objects and manages the inferencing process. The AI configuration framework enables Maximo Manage applications to communicate with Maximo AI Service.

Maximo AI Service uses watsonx large language models (LLMs) to enable AI features.

Note: Maximo AI Service 9.1 includes a limited use license to watsonx.ai and incurs an additional AppPoint cost. For more information, see [Licensing in Maximo Application Suite 9.1](#).

The following list describes how Maximo Manage integration components generally support AI features.

Object structures

Documents the Maximo Manage objects' architecture, defines the API set that is used to configure the model for the specified Maximo Manage business object, and defines permissions for users and groups who are authorized to use the AI features.

Object structure query definitions

Osclause queries that are used to select a training data set and control the inference set or configuration set. The osclause query contains an SQL where clause, without the WHERE keyword, that is based on the root object of the object structure.

Object structure query templates

Defines the schema that is used for training, inference, or configuration requests.

Endpoints

Communicates to external services, such as the Maximo AI Service APIs. The Maximo AI Service API is a Java application that enables the flow of data between Maximo Application Suite and watsonx.ai.

Invocation channels

Defines the processing logic and mapping of inbound and outbound data, which enables Maximo Manage to process responses to and from Maximo AI Service. Invocation channels also prepare training and inference data from Maximo Manage for the selected object structure and object structure query template.

Maximo Manage system properties are also used to connect Maximo Manage to Maximo AI Service.

Components for each AI feature

An *AI configuration* is type of record that you create that contains settings for your AI feature, such as the required model template that's available from Maximo AI Service, and where you can start and monitor any required training or inferencing. Each model template has unique data requirements and not all require training and inferencing. You manage AI configurations in the AI configuration application in Maximo Manage.

AI configurations are available by default for finding similar work orders, the AI assistant, and problem code recommendations for work orders. These AI configurations use default integration components that do not have to be configured. To enable recommending field values other than problem codes, in AI configuration and all of the required integration components must be created or selected for use. To enable AI recommendations in Reliability Strategies, you must also create an AI configuration, but the AI feature does not require integration components.

Locate similar work orders

This feature uses the similarity model template and the following default integration components:

- MXAPIWODETAIL object structure with the following values configured:
 - WOSIMILARITYCFG query template for configuration
 - WOSIMILARITYINF query template for inferencing
 - WOSIMILARITYCFGFILTER query definition for configuration
 - WOSIMILARITYINFFILTER query definition for inferencing
- AIBROKERAPI endpoint
- AIWOSIMILARITYCFG invocation channel for configuration, with the following values configured:
 - com.ibm.tivoli.maximo.ai.AITrainReqExit request processing class
 - com.ibm.tivoli.maximo.ai.AITrainRespExit response processing class
- AIWOSIMILARITYINF invocation channel for inferencing, with the following values configured:
 - com.ibm.tivoli.maximo.ai.AIINFReqExit request processing class

- `com.ibm.tivoli.maximo.ai.AIINFRespExit` response processing class

Recommend problem codes in Work orders

This feature uses the pcc model template and the following default integration components:

- MXAPIWODETAIL object structure with the following values configured:
 - WOPROBLEMTRAIN query template for training
 - WOPROBLEMINF query template for inferencing
 - AITRAINFILTER query definition for training
 - AIINFERENCEFILTER query definition for inferencing
- AIBROKERAPI endpoint
- AITRAINWOPROBLEMCODE invocation channel for training with the following values configured:
 - `com.ibm.tivoli.maximo.ai.AITrainReqExit` request processing class
 - `com.ibm.tivoli.maximo.ai.AITrainRespExit` response processing class
- AIINFWOPROBLEMCODE invocation channel for inferencing with the following values configured:
 - `com.ibm.tivoli.maximo.ai.AIINFReqExit` request processing class
 - `com.ibm.tivoli.maximo.ai.AIINFRespExit` response processing class

Recommended field values

This feature uses the mcc model template. You can reuse the following integration components:

- AIBROKERAPI endpoint
- `com.ibm.tivoli.maximo.ai.AIINFReqExit` request processing class
- `com.ibm.tivoli.maximo.ai.AIINFRespExit` response processing class
- `com.ibm.tivoli.maximo.ai.AITrainReqExit` request processing class
- `com.ibm.tivoli.maximo.ai.AITrainRespExit` response processing class

The following integration components must be created or selected for use:

- An object structure that contains the following values:
 - Query template for training
 - Query template for inferencing
 - Query definition for training
 - Query definition for inferencing
- An invocation channel for inferencing
- An invocation channel for inferencing

For more information about creating these components, see [“Preparing required components for field value recommendations”](#) on page 443.

Enable an AI assistant

This feature uses the assistant model template. The assistant uses object structures and the AIBROKERAPI end point. The administrator can select which object structures the assistant uses. For more information, see [“Enabling the AI assistant”](#) on page 436.

Generate AI recommendations in Reliability Strategies

This assistant uses the fmea model template. The AI feature for Reliability Strategies does not require any integration components.

Cron tasks for training and inferencing

A cron job for scheduling model training and a cron job for inference requests are available by default.

The AITRAINJOB crontask determines the frequency at which the training process runs for all eligible AI configurations. By default, the crontask runs every five minutes. The AIINFJOB crontask determines the

frequency at which the inference process runs for all eligible AI configurations. By default, the crontask runs every hour.

The **AICONFIGS** parameter controls which AI configurations are processed. By default, the **AICONFIGS** crontask parameter is set to *, which indicates that the crontask processes all active and ready AI configurations. The parameter also verifies that the model status is ready.

For the cron job to retrieve an AI configuration for training, the configuration must be active and have a train request pending. A train request is initiated by, in the AI configuration application, in a configuration, selecting **Actions** and then selecting **Train** or **Re-train**.

For the cron job to retrieve an AI configuration for inferencing, the name of the configuration must be listed in the **AICONFIGS** parameter for the inference cron job. If the **AICONFIGS** parameter has a value of *, inferencing for an AI configuration will occur automatically after training.

To run inferencing or training for only specific AI configurations, you can create other instances of the associated crontask and **AICONFIGS** parameter and then specify a single AI configuration for each or specify a comma-separated list of the AI configurations.

Ensure that values do not overlap. Overlapping value can negatively affect system performance.

For example, the **AICONFIGS** parameter for the AIINFJOB crontask uses the * value, which means inferencing is run automatically for all active and ready AI configurations. You create another crontask, AICONFIG, that runs inferencing for a specific AI configuration. AIINFJOB crontask and the AICONFIG crontask run inferencing for the same AI configuration, which extends the inferencing process unnecessarily.

Uninstalling AI broker for Maximo Manage

Before you can deploy Maximo AI Service, you must uninstall any existing instances of the AI broker.

About this task

If you also have existing versions of Suite License Service, IBM Data Reporter Operator , Db2, MinIO, or Amazon Web Services S3, you are not required to delete those versions. You can use your existing versions with Maximo AI Service. However, if you used Db2 with Amazon Web Services S3 and the AI broker, you must create new tables to use with Maximo AI Service.

Procedure

1. Open Red Hat OpenShift web console.
2. On the dashboard, uninstall the AI broker operators.
 - a) From the side navigation, click **Operators > Installed Operators**.
 - b) For the following operators, click the three-dot menu and then click **Uninstall Operator**.
Uninstalling these operators also uninstalls all associated custom resource definitions (CRD).
 - Authorino Operator
 - IBM Maximo AiBroker
 - IBM Truststore Manager
 - Open Data Hub Operator
 - Red Hat OpenShift Pipelines
 - Red Hat OpenShift Serverless
 - Red Hat OpenShift Service Mesh 2
3. Delete the mas-*{instance_id}*-aibroker project.
Deleting the project also removes all associated resources, such as secrets, deployments, ConfigMaps, and inventory.
 - a) From the side navigation, click **Home > Projects**.
 - b) For the mas-*{instance_id}*-aibroker project, click the three-dot menu and then click **Delete project**.

4. On the dashboard, delete any aibroker-*{tenant_name}* tenant projects.
 - a) From the side navigation, click **Administration > CustomResourceDefinition**.
 - b) Search for the AiBrokerWorkspace definition name.
 - c) Select the definition and then on the **Instances** tab, delete all listed instances.
5. Delete any existing MariaDB project.
 - a) From the side navigation, click **Home > Projects**.
 - b) Search for the mariadb project name.
 - c) For the project, click the three-dot menu and then click **Delete Project**.

Deploying Maximo AI Service

You can use Maximo AI Service SaaS or on-premises. To deploy Maximo AI Service on-premises, you initiate the deployment by using the CLI or Ansible collection, connect Maximo AI Service to Maximo Manage by using system properties, and then import the Maximo AI Service certificate.

Before you begin

Before you begin, review the entire process for enabling AI features. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#).

You must also set up prerequisites. The following list describes the prerequisites:

- Install or set up an account for watsonx.ai. You can use watsonx.ai on-premises or SaaS. You can use an existing instance or install or sign up for a new instance.
 - For more information about installing watsonx.ai on-premises, see [“Installing watsonx.ai on-premises for Maximo AI Service” on page 420](#). The cluster that watsonx.ai is installed in requires a GPU.
 - For watsonx.ai SaaS, no installation is required, but you must set up an account for watsonx.ai. For more information, see [Step 1: Sign up for watsonx](#). You must complete only step 1.
- At a minimum, for a single user and tenant, you must have three primary nodes that have eight CPUs with 32 GB memory each and six secondary nodes that have four CPUs with 32 GB memory each. Additional users, tenants, and workloads require more resources.
- Red Hat OpenShift 4.16 or later
- Suite License Service. You can use an existing instance or install a new instance as part of the deployment for Maximo AI Service.
- IBM Data Reporter Operator . You can use an existing instance or install a new instance as part of the deployment for Maximo AI Service.
- One of the following object storage providers:
 - MinIO. You can use an existing instance or install a new instance as part of the deployment for Maximo AI Service. Install MinIO in the same cluster as Maximo AI Service.
 - Amazon Web Services S3. The buckets that are used for Maximo AI Service must have unique names, and all typical dependencies of Amazon Web Services S3 must be deployed.
- IBM Db2. You can use an existing instance or install a new instance as part of the deployment for Maximo AI Service.

About this task

As part of the deploying, you use system properties to establish a connection from Maximo Manage to the Maximo AI Service. This connection enables Maximo Manage to communicate with the services or runtimes that are hosted within the Maximo AI Service cluster.

The following steps describe how to deploy Maximo AI Service on-premises. For more information about Maximo AI Service SaaS, see [“Enabling AI features by using Maximo AI Service SaaS” on page 414](#).

Procedure

1. Install Maximo AI Service by using the Maximo Application Suite CLI or the Ansible collection.

For more information about deploying by using Ansible, see [Installing Maximo AI Service with the Ansible collection](#).

For more information about deploying by using the CLI in interactive mode, see [Installing Maximo AI Service from the CLI in interactive mode](#).

For more information about deploying by using the CLI in noninteractive mode, see [Installing Maximo AI Service from the CLI in noninteractive mode](#).

2. Retrieve the required values for connecting Maximo Manage to Maximo AI Service.

You must retrieve the Maximo AI Service API key, URL, and tenant ID. All of these values are in the . You must have these values to complete step 3.

- a) Retrieve the API key.

- i) In Red Hat OpenShift web console, click **Home** > **Projects** and search for the `aiservice-{instance_id}` project name.
- ii) Open the `aiservice-{instance_id}` project and in the **Inventory** section, click **number Secrets**.
- iii) Click the **aiservice-user----apikey-secret** secret and in the **Data** section, click **Reveal values**. Save this value on your local machine.

- b) Retrieve the URL.

- i) In Red Hat OpenShift web console, click **Networking** > **Routes** and search for the `aibroker` route.
- ii) Open the route.
- iii) Copy the value that is in the **Location** field.
- iv) On your local machine, save that value and append the following path to the end of the URL: `/ibm/aibroker/service/rest/api/v1`

- c) Retrieve the tenant ID.

- i) In Red Hat OpenShift web console, click **Administration** > **CustomResourceDefinition** and search for the `AIServiceTenant` definition name.
- ii) Open the definition.
- iii) On the **Instances** tab, copy the value in the **Name** column that corresponds to the system that you are setting up. For example, if you are deploying Maximo AI Service in your production environment, copy the **Name** value for your production tenant.

3. Configure the Maximo Manage system properties.

- a) In Maximo Manage, open the System Properties application.

- b) Search for and then select and add global values for the following properties. Use the values that you retrieved in step 2.

- **mxo.int.aibrokerapikey**. The value is the Maximo AI Service API key.
- **mxo.int.aibrokerapiurl**. The value is the Maximo AI Service URL.
- **mxo.int.aibrokertenantid**. The value is the Maximo AI Service tenant ID.

- c) After you edit each property, in the **Common Actions** menu, click **Save Property**.

- d) In the **Common Actions** menu, click **Live Refresh**.

4. Create and import the `ca.cert` file for Maximo AI Service.

- a. In Red Hat OpenShift web console, **Home** > **Projects** and search for the `mas-{instanceid}-broker` project.

- b. Open the project.

- c. From the side navigation menu, click **Workloads** > **Secrets**.

- d. Select the `{instanceid}-public-aibroker-tls` secret.
 - e. Copy the content that is in the `ca.crt` field. The value for the `ca.crt` field is the content for the `ca.cert` file.
 - f. On your local machine, paste the content into an empty `.txt` file.
 - g. Save the file as `ca.cert`.
 - h. Import the certificate. For more information, see [Adding trusted certificates](#).
5. Verify that Maximo AI Service is running and connected to Maximo Manage.
- To verify that Maximo AI Service is running and connected, you can run the following command:

```
curl -X 'GET' \ 'https://{hostname}/ibm/aibroker/service/rest/api/v1/health' \ -H 'accept: */*'
```

If Maximo AI Service is running, the following output is returned:

```
{ "max_number_of_tenant": "<number>", "kmodel": "running", "healthy": true, "version": "<version>", "status": { "KMODELS": { "healthy": true }, "DB2": { "healthy": true } } }
```

Alternately, you can check the status of Maximo AI Service in Maximo Manage in the AI configuration application. In the AI configuration application, click **AI Service Health**. If the status is running, Maximo AI Service is ready to use.

What to do next

If another status or an error is displayed, you can access more details in the Red Hat OpenShift logs. For more information about troubleshooting, see [“Troubleshooting Maximo AI Service and AI features”](#) on page 458.

If Maximo AI Service is available and running, you can start setting up your AI configurations. Each configuration represents a feature that you want to enable, for example, problem code recommendations or the AI assistant.

The following table contains the available AI features, associated model template name, the required Granite model, and links to steps to set up the AI configuration.

Feature	Model template	Model	Steps
Problem code recommendations for work orders	pcc	Granite 3.0 8B Instruct	“Enabling recommended problem codes for Work orders” on page 446
Field value recommendations	mcc	Granite 3.2 8B Instruct	Enabling field value recommendations
AI assistant	nl2oslc	Granite 3.2 8B Instruct	Enabling the assistant
Locating similar work orders	similarity	Granite 3.0 8B Instruct	“Enabling locating of similar work orders” on page 449
AI recommendations for asset boundary and failure list in Reliability Strategies	fmea	Granite 3.2 8B Instruct	“Enabling AI recommendations in Reliability Strategies” on page 455

Installing watsonx.ai on-premises for Maximo AI Service

Before you can deploy Maximo AI Service on-premises, you must install watsonx.ai on-premises or SaaS.

Before you begin

- Review the online product documentation for IBM Software Hub, especially planning, system requirements, and storage requirements. For more information, see the following topics:
 - [Planning](#)
 - [System requirements](#)
 - [Storage requirements](#)
- Review the online product documentation for installing watsonx.ai.

For more information, see [Preparing to install IBM watsonx.ai](#). Install the following models according to which AI features you want to use.

Template	Feature	Model to install
mcc	Recommend field values.	Granite 3.2 8B Instruct
pcc	Recommend problem codes in work orders.	Granite 3.0 8B Instruct
similarity	Locate similar work orders.	Granite 3.0 8B Instruct
nl2oslc	Enable the assistant.	Granite 3.2 8B Instruct
fmea	Generate AI recommendations in Reliability Strategies.	Granite 3.2 8B Instruct

About this task

The following steps apply to watsonx.ai on-premises. For watsonx.ai SaaS, no installation is required, but you must set up an account for watsonx.ai. For more information, see [Step 1: Sign up for watsonx](#).

Procedure

1. Install IBM Software Hub.

For more information, see [Installing IBM Software Hub](#).

2. Install watsonx.ai in full version mode.

For more information, see [Installing watsonx.ai](#).

Ensure that you install watsonx.ai on a different cluster than the Maximo AI Service cluster. The cluster that watsonx.ai is installed on requires a GPU. The Maximo AI Service cluster does not require a GPU.

Do not install watsonx.ai in lightweight engine mode.

3. Use the **oc patch** command to update the **watsonxaiifm** custom resource.

Installing Maximo AI Service through the Maximo Application Suite CLI

You can install Maximo AI Service by using the Maximo Application Suite CLI in interactive or noninteractive mode.

Installing Maximo AI Service from the CLI in interactive mode

You can install Maximo AI Service by using the Maximo Application Suite CLI in interactive mode.

Procedure

1. Copy the Maximo Application Suite CLI license to a local folder, for example, `/home/{user}/git/devops-configs`.

For more information about obtaining the license, see [Setting up IBM Maximo Application Suite](#).

2. Open a command line and run the following command.

```
docker run -v /home/{user}/git/devops-configs:/tmp -ti --rm --pull always quay.io/ibmmas/cli:master
```

3. Ensure that the installation process is successful. The following text is an example of the installation process. Provide values during the installation process and follow the prompts. `***` denotes a variable.

```
[***-]$ docker run -v /home/***git/***:/tmp -ti --rm --pull always quay.io/ibmmas/cli:master
Trying to pull quay.io/ibmmas/cli:master...
Getting image source signatures
Copying blob 095405e15023 done |
Copying blob 671f5ec6b9c1 done |
Copying blob 4fa67d0346c5 done |
Copying blob 706d1e2e240b done |
Copying blob 0df697abec85 skipped: already exists
Copying blob 11ea16df3746 done |
Copying blob 9b4960ee63a1 done |
Copying config 7e92c566f5 done |
Writing manifest to image destination
IBM Maximo Application Suite CLI Container ***

https://github.com/ibm-mas/ansible-devops
https://github.com/ibm-mas/cli

MAS Management:
- mas install to install a new MAS instance
- mas update to apply a new catalog update
- mas upgrade to upgrade an existing MAS install to a new release
- mas must-gather to perform must-gather against the target cluster
- mas uninstall to uninstall a MAS instance
- mas configtool-oidc to configure OIDC integration
Disconnected Install Support:
- mas setup-registry to setup a private container registry on an OCP cluster
- mas teardown-registry to delete a private container registry on an OCP cluster
- mas mirror-images to mirror container images required by MAS to a private registry
- mas configure-airgap to configure a cluster to use a private registry as a mirror
OpenShift Cluster Management:
- mas provision-aws to provision an OCP cluster on AWS
- mas provision-roks to provision an OCP cluster on IBMCloud Red Hat OpenShift Service (ROKS)
- mas provision-rosa to provision an OCP cluster on AWS Red Hat OpenShift Service (ROSA)
- mas provision-fyre to provision an OCP cluster on IBM DevIT Fyre (internal)
AI Service (Standalone) Management:
- mas aiservice-install to install a new AI Service instance

[ibmmas/cli:***]mascli$ mas aiservice-install

IBM Maximo Application Suite Admin CLI v***
Powered by https://github.com/ibm-mas/ansible-devops/ and https://tekton.dev/

AI Broker 9.0 Deprecation Notice

Maximo AI Broker (introduced with MAS 9.0) has been replaced with Maximo AI Service as of Aug 1 2025
To continue using the features that were enabled by the AI broker after that time, you must deploy and use Maximo AI Service 9.1:
- Maximo AI Service 9.1 is compatible with both Maximo Application Suite 9.0 and 9.1 releases
- If Maximo AI Service is deployed with Maximo Application Suite 9.0, you can use only the AI features that were included in Maximo Application Suite 9.0

Note: Maximo AI Service 9.1 includes a limited-use license to watsonx.ai and incurs an additional AppPoint cost

1) Set Target OpenShift Cluster
Server URL: ***
Login Token: ***
Disable TLS Verify? [y/n] y

2) IBM Maximo Operator Catalog Selection
The catalog you choose dictates the version of everything that is installed, with Maximo
```

Application Suite this is the only version you need to remember; all other versions are determined by this choice.
Older catalogs can still be used, but we recommend using an older version of the CLI that aligns with the release date of the catalog.

- Learn more: <https://ibm-mas.github.io/cli/catalogs/>

Supported Catalogs:

- ***

Select catalog ***

Catalog Details

Catalog Image: ***
Catalog Digest: ***
AI Service Releases: 9.1
MongoDb: ***

Two types of release are available:

- GA releases of Maximo Application Suite are supported under IBM's standard 3+1+3 support lifecycle policy.
- 'Feature' releases allow early access to new features for evaluation in non-production environments and are only supported through to the next GA release.

	9.1	
AI Service	***	

Select release ***

3) License Terms

To continue with the installation, you must accept the license terms:

- ***

Do you accept the license terms? [y/n] y

4) Configure Storage Class Usage

Maximo Application Suite and its dependencies require storage classes that support ReadWriteOnce (RWO) and ReadWriteMany (RWX) access modes:

- ReadWriteOnce volumes can be mounted as read-write by multiple pods on a single node.
- ReadWriteMany volumes can be mounted as read-write by multiple pods across many nodes.

Storage provider auto-detected: NFS Client

- Storage class (ReadWriteOnce): nfs-client
- Storage class (ReadWriteMany): nfs-client

Use the auto-detected storage classes? [y/n] y

5) Configure AppPoint Licensing

By default the AI Service instance will be configured to use a cluster-shared License, this provides a shared pool of AppPoints available to all MAS & AI Service instances on the cluster.

Alternatively you may choose to install using a dedicated license only available to this AI Service instance.

1. Install AI Service with Cluster-Shared License (AppPoints)
2. Install AI Service with Dedicated License (AppPoints)

SLS Mode 1

License file ***

Contact e-mail address ***

Contact first name ***

Contact last name ***

IBM Data Reporter Operator (DRO) Namespace ***

6) Configure IBM Container Registry

IBM entitlement key ***

7) Configure AI Service Instance

Instance ID restrictions:

- Must be 3-12 characters long
- Must only use lowercase letters, numbers, and hyphen (-) symbol
- Must start with a lowercase letter
- Must end with a lowercase letter or a number

Instance ID ***

8) Configure Operational Mode

Maximo Application Suite can be installed in a non-production mode for internal development and testing, this setting cannot be changed after installation:

- All applications, add-ons, and solutions have 0 (zero) installation AppPoints in non-production installations.
- These specifications are also visible in the metrics that are shared with IBM and in the product UI.

1. Production

2. Non-Production Operational Mode 1

9) AI Service Settings

9.1) Storage Configuration

AI Service requires object storage for pipelines, tenants, and templates. You can either install MinIO in-cluster or connect to external storage.

```
Install Minio? [y/n] y
minio root username ***
minio root password ***
```

10) AI Service Tenant Settings

AI Service will reserve AppPoints for a fixed period of time based on the values you enter:
Tenant entitlement type ***
Entitlement end date (YYYY-MM-DD) ***

11) WatsonX Integration

This CLI section configures the integration between the AI Service and IBM watsonx.ai. AI Service uses watsonx for model deployment and inferencing.

The WatsonX API key must be a **platform API key** associated with a user that has at least:

- **Editor permission** for the project
- **Viewer permission** for the space

You can generate this key by following IBM's documentation: https://www.ibm.com/docs/en/watsonx/w-and-w/2.2.0?topic=tutorials-generating-api-keys#api-keys__platform__title__1

The endpoint URL is your WatsonX Machine Learning service URL. It can be found in the watsonx.ai

documentation: <https://cloud.ibm.com/apidocs/watsonx-ai-cp/watsonx-ai-cp-2.2.0#endpoint-url>

The project ID refers to your specific watsonx.ai project where your ML models and assets are stored.

```
Watsonxai api key ***
Watsonxai machine learning url ***
Watsonxai project id ***
```

12) RSL Integration

RSL (Reliable Strategy Library) connects to strategic asset management via STRATEGIZEAPI.

RSL URL: ***

Org ID: Get from MAS Manage > System Properties > 'mxe.rs.rslorgid'

Token: Use your IBM entitlement key (same as MAS installation)

Note: Future versions will auto-configure these from MAS Manage.

```
RSL url ***
ORG Id of RSL ***
Token for RSL ***
Does the RSL API use a self-signed certificate?? [y/n] y
RSL CA certificate (PEM format) ***
```

13) Configure MongoDB

The installer can setup mongoce in your OpenShift cluster (available only for amd64) or you may choose to configure MAS to use an existing mongodb

Create MongoDB cluster using MongoDB Community Edition Operator? [y/n] y

MongoDb namespace mongoce

14) Non-Interactive Install Command

Save and re-use the following script to re-run this install without needing to answer the interactive prompts again

```
export IBM_ENTITLEMENT_KEY=***
mas aiservice-install --mas-catalog-version *** --ibm-entitlement-key $IBM_ENTITLEMENT_KEY \
  --aiservice-instance-id "***" \
  --storage-class-rwo "***" --storage-class-rwx "***" \
  --storage-pipeline "***" --storage-accessmode "***" \
  --license-file "***" \
  --uds-email "***" --uds-firstname "***" --uds-lastname "***" \
  --dro-namespace "***" \
  --mongodb-namespace "***" \
  --aiservice-channel "***" \
  --s3-accesskey "***" \
  --s3-secretkey "***" \
  --s3-host "***" \
  --s3-port "***" \
  --s3-ssl "***" \
  --s3-region "***" \
  --s3-bucket-prefix "***" \
  --s3-tenants-bucket "***" \
```



```

--s3-templates-bucket "***" \
--watsonxai-apikey "***" \
--watsonxai-url "***" \
--watsonxai-project-id "***" \
--minio-root-user "***" \
--minio-root-password "***" \
--tenant-entitlement-type "***" \
--tenant-entitlement-start-date "***" \
--tenant-entitlement-end-date "***" \
--rsl-url "***" \
--rsl-org-id "***" \
--rsl-token "Bearer ***" \
--rsl-ca-crt "***" \
--accept-license --no-confirm

```

15) Review Settings

Connected to:

- https://console-openshift-console.apps.***

15.1) Pipeline Configuration

```

Service Account ..... Default
Image Pull Policy ..... Default
Skip Pre-Install Healthcheck ..... No

```

15.2) OpenShift Container Platform

```

Worker Node Architecture ..... ***
Storage Class Provider ..... ***
ReadWriteOnce Storage Class ..... ***
ReadWriteMany Storage Class ..... ***
Certificate Manager ..... ***
Cluster Ingress Certificate Secret ..... Default

```

15.3) Maximo Operator Catalog

```

Catalog Version ..... ***

```

15.4) IBM Container Registry

```

IBM Entitled Registry ..... ***
IBM Open Registry ..... ***

```

15.5) AI Service

```

Release ..... ***
Instance ID ..... ***
Environment Type ..... ***

```

15.6) AI Service Tenant Entitlement

```

Entitlement Type ..... ***
Start Date ..... ***
End Date ..... ***

```

15.7) S3 Configuration

```

Minio Root Username ..... ***

Host ..... ***
Port ..... ***
SSL Enabled ..... ***
Region ..... ***
Bucket Prefix ..... ***
Templates Bucket Name ..... ***
Tenants Bucket Name ..... ***

```

15.8) IBM WatsonX

```

URL ..... ***
Project ID ..... ***

```

15.9) RSL

```

URL ..... ***
Organization ID ..... ***

```

15.10) IBM Data Reporter Operator (DRO) Configuration

```

Contact e-mail ..... ***
First name ..... ***
Last name ..... ***
Install Namespace ..... ***

```

15.11) IBM Suite License Service

```

Namespace ..... ***
Subscription Channel ..... ***
IBM Open Registry ..... ***
License File ..... ***

```

15.12) MongoDB

```

Type ..... ***
Install Namespace ..... ***

```

```

15.13) IBM Db2 Univeral Operator Configuration
Action ..... install
Install Namespace ..... Default
Subscription Channel ..... Default

Please carefully review your choices above, correcting mistakes now is much easier than
after the install has begun
Proceed with these settings? [y/n] y

16) Launch Install
If you are using storage classes that utilize 'WaitForFirstConsumer' binding mode choose
'No' at the prompt below
Wait for PVCs to bind? [y/n] y
 OpenShift Pipelines Operator is installed and ready to use
 Namespace is ready (***)
 MAS CLI image deployment test completed
 Latest Tekton definitions are installed (***)
 PipelineRun for apmdevops install submitted

View progress:
***

```

Installing Maximo AI Service from the CLI in noninteractive mode

You can install Maximo AI Service by using the Maximo Application Suite CLI in noninteractive mode.

Before you begin

Installing Maximo AI Service through the Maximo Application Suite CLI in noninteractive mode requires a script that contains values. You generate this script by running the installation program and providing values to add to the script. Generate the script and then complete the following procedure.

Procedure

1. Copy the Maximo Application Suite CLI license to a local folder, for example, `/home/{user}/git/devops-configs`.

For more information about obtaining the license, see [Setting up IBM Maximo Application Suite](#).

2. Open a command line and run the following command.

```
docker run -v /home/{user}/git/devops-configs:/tmp -ti --rm --pull always quay.io/ibmmas/cli:master
```

3. Ensure that the installation process is successful. The following text is an example of the installation process. Provide values during the installation process and follow the prompts. `***` denotes a variable.

```

[*** ~]$ docker run -v /home/***/git/***/:/tmp -ti --rm --pull always quay.io/ibmmas/cli:master
Trying to pull quay.io/ibmmas/cli:master...
Getting image source signatures
Copying blob 5c9621f2a073 skipped: already exists
Copying blob fe0a337f348a skipped: already exists
Copying blob 2686ea3b22ef skipped: already exists
Copying blob 47f5b092a7aa skipped: already exists
Copying blob 76feb2c4dd7f skipped: already exists
Copying blob 8bc213e3816c skipped: already exists
Copying blob 0df697abec85 skipped: already exists
Copying config 3bafaf4d51 done
Writing manifest to image destination
IBM Maximo Application Suite CLI Container ***

https://github.com/ibm-mas/ansible-devops
https://github.com/ibm-mas/cli

MAS Management:
- mas install to install a new MAS instance
- mas update to apply a new catalog update
- mas upgrade to upgrade an existing MAS install to a new release
- mas must-gather to perform must-gather against the target cluster
- mas uninstall to uninstall a MAS instance
- mas configtool-oidc to configure OIDC integration
Disconnected Install Support:
- mas setup-registry to setup a private container registry on an OCP cluster

```

```

- mas teardown-registry to delete a private container registry on an OCP cluster
- mas mirror-images to mirror container images required by MAS to a private registry
- mas configure-airgap to configure a cluster to use a private registry as a mirror
OpenShift Cluster Management:
- mas provision-aws to provision an OCP cluster on AWS
- mas provision-roks to provision an OCP cluster on IBMCloud Red Hat OpenShift Service (ROKS)
- mas provision-rosa to provision an OCP cluster on AWS Red Hat OpenShift Service (ROSA)
- mas provision-fyre to provision an OCP cluster on IBM DevIT Fyre (internal)
AI Service (Standalone) Management:
- mas aiservice-install to install a new AI Service instance

[ibmmas/cli:***]mascli$ mas aiservice-install

export IBM_ENTITLEMENT_KEY="***"
mas aiservice-install --mas-catalog-version *** --ibm-entitlement-key $IBM_ENTITLEMENT_KEY \
--aiservice-instance-id "***" \
--storage-class-iwo "***" --storage-class-iwx "***" \
--storage-pipeline "***" --storage-accessmode "***" \
--license-file "***" \
--uds-email "***" --uds-firstname "***" --uds-lastname "***" \
--dro-namespace "***" \
--mongodb-namespace "***" \
--aiservice-channel "***" \
--s3-accesskey "***" \
--s3-secretkey "***" \
--s3-host "***" \
--s3-port "***" \
--s3-ssl "***" \
--s3-region "***e" \
--s3-bucket-prefix "***" \
--s3-tenants-bucket "***" \
--s3-templates-bucket "***" \
--watsonxai-apikey "***" \
--watsonxai-url "***" \
--watsonxai-project-id "***" \
--minio-root-user "***" \
--minio-root-password "***" \
--tenant-entitlement-type "***" \
--tenant-entitlement-start-date "***" \
--tenant-entitlement-end-date "***" \
--rsl-url "***" \
--rsl-org-id "***" \
--rsl-token "Bearer ***" \
--rsl-ca-crt "***" \
--accept-license --no-confirm

```

What to do next

On your Red Hat OpenShift dashboard, click **Pipeline** > **Pipelines** and locate the aiservice-install pipeline. When the pipeline run is complete, Maximo AI Service is installed.

Installing Maximo AI Service by using an Ansible collection

You can install Maximo AI Service by using an Ansible collection.

Before you begin

Complete the following prerequisites:

- Gather the values for the required environment variables. For more information, see [Environment variables](#).
- Ensure that Git is installed on your system.
- Ensure that a code editor is installed on your system.
- Create a GitHub account or identify an existing account that you own.
- Connect your code editor to your GitHub account. For more information, see your code editor's documentation.

About this task

For general information about Ansible, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

The Ansible playbook for Maximo AI Service can be run against any Red Hat OpenShift cluster regardless of the cluster type.

The following table describes the Ansible roles' mapping to tasks.

Role	Tasks
aiservice_odh	<ul style="list-style-type: none">• Install Red Hat OpenShift serverless operator• Install Red Hat OpenShift service mesh operator• Install Authorino operator• Install Open Data Hub operator• Create DSCInitialization instance• Create Data Science cluster• Create Data Science Pipelines application
aiservice_kmodels	<ul style="list-style-type: none">• Install Kmodel controller• Install istio• Install Kmodel store• Install Kmodel watcher
aiservice	<ul style="list-style-type: none">• Install Maximo AI Service API application• Create and delete Maximo AI Service API key• Create and delete Amazon Web Services S3 API key
aiservice_tenant	<ul style="list-style-type: none">• Create config for Suite License Service• Create config for Reliability Strategies• Create config for Data Reporter Operator• Create and delete Amazon Web Services S3 API key• Create and delete watsonx API key• Create Maximo AI Service tenant

Procedure

1. Clone the ansible-devops repository, which is at the following link: <https://github.com/ibm-mas/ansible-devops>. Do not use the CLI container. You must clone the repository.

You can clone the repository by using the following command or manually clone the repository in your code editor. For more information on how to clone GitHub repositories in your code editor, see your code editor's documentation and the documentation for your operating system. Different operating systems can require unique steps for cloning a repository from GitHub.

```
git clone https://github.com/ibm-mas/ansible-devops.git
```

2. Run the following command to build and install the Ansible collection.

```
cd git/ansible-devops/ibm/mas_devops/  
ansible-galaxy collection build
```

```
ansible-galaxy collection install ibm-mas_devops-*.tar.gz --ignore-certs --force
rm ibm-mas_devops-*.tar.gz
```

3. Run the following command to set up the virtual environment.

Do not install the environment globally.

```
python3 -m venv ~/venv/
source ~/venv/bin/activate
python3 -m pip install --upgrade pip
python3 -m pip install ansible junit_xml pymongo xmljson jmespath kubernetes==12.0.1
openshift==0.12.1 boto3 requests mas-devops
```

4. For your Red Hat OpenShift cluster, set the environment variables.

For more information about the variables and what values to specify, see [Environment variables](#). The following content is an example structure for how you set the environment variables. The example does not include example values for the variables.

```
# General environment variables
export IBM_ENTITLEMENT_KEY=${MAS_ENTITLEMENT_KEY}
export MAS_CATALOG_VERSION="<value>"
export MAS_CONFIG_DIR="<value>"
export ENVIRONMENT_TYPE="<value>"
export AISERVICE_INSTANCE_ID="<value>"
export MAS_ENTITLEMENT_KEY="<value>"
export MAS_ENTITLEMENT_USERNAME="<value>"
export AISERVICE_CHANNEL="9.1.x"
export MAS_AISERVICE_TENANT_ENTITLEMENT_TYPE="standard"
#Storage provider environment variables
export INSTALL_MINIO="<value>"
export MINIO_ROOT_PASSWORD="<value>"
export AISERVICE_S3_ACCESSKEY="<value>"
export AISERVICE_S3_SECRETKEY=${MINIO_ROOT_PASSWORD}
export AISERVICE_S3_HOST="<value>"
export AISERVICE_S3_SSL="<value>"
export AISERVICE_S3_PROVIDER="<value>"
export AISERVICE_S3_PORT="<value>"
export AISERVICE_S3_REGION="<value>"
export AISERVICE_S3_PIPELINES_BUCKET="<value>"
export AISERVICE_S3_TENANTS_BUCKET="<value>"
export AISERVICE_S3_TEMPLATES_BUCKET="<value>"
export AISERVICE_S3_BUCKET_PREFIX="<value>"
export AISERVICE_S3_REGION="<value>"
export AISERVICE_S3_ENDPOINT_URL="<value>"
#Watsonx environment variables
export AISERVICE_WATSONXAI_APIKEY="<value>"
export AISERVICE_WATSONXAI_URL="<value>"
export AISERVICE_WATSONXAI_PROJECT_ID="<value>"
#Data Reporter Operator environment variables
export INSTALL_DRO="<value>"
export DRO_CONTACT_EMAIL="<value>"
export DRO_CONTACT_FIRSTNAME="<value>"
export DRO_CONTACT_LASTNAME="<value>"
export AISERVICE_DRO_TENANT_ID="<value>"
export AISERVICE_DRO_SECRET_NAME="<value>"
export AISERVICE_DRO_API_KEY="<value>"
export AISERVICE_DRO_URL="<value>"
export AISERVICE_DRO_CA_CERT="<value>"
#Suite License Service environment variables
export INSTALL_SLS="<value>"
export INSTALL_MONGO="<value>"
export SLS_MONGODB_CFG_FILE=${MAS_CONFIG_DIR}/mongo-mongoce.yml
export SLS_LICENSE_ID="<value>"
export SLS_LICENSE_FILE="<value>"
export AISERVICE_SLS_SUBSCRIPTION_ID="<value>"
export AISERVICE_SLS_SECRET_NAME="<value>"
export AISERVICE_SLS_REGISTRATION_KEY="<value>"
export AISERVICE_SLS_URL="<value>"
export AISERVICE_SLS_CA_CERT="<value>"
#Reliability Strategy Library environment variables
export RSL_URL="<value>"
export RSL_ORG_ID="<value>"
export RSL_TOKEN="<value>"
#Db2 environment variables
export INSTALL_DB2="<value>"
export AISERVICE_DB2_USERNAME="<value>"
export AISERVICE_DB2_PASSWORD="<value>"
export AISERVICE_DB2_JDBC_URL="<value>"
```

```
export AISERVICE_DB2_SSL_ENABLED="value"
export AISERVICE_DB2_CA_CERT="value"
```

5. Run the following command to run the Ansible playbook.

Replace all variables with your values.

```
oc login --token=token --server=https://server:port
oc projects
ansible-playbook playbooks/oneclick_add_aibroker.yml
```

6. Create a tenant.

When you create a tenant, a custom resource is also created that contains the configuration that is related to that tenant. Each custom resource is customized to a specific tenant. You can create more than one custom resource. The tenant reuses the specific environmental variables unless the values are defined by the associated dependency.

Run the following command to create the tenant. Replace all variables with your values.

For more information about the variables and what values to specify, see [Environment variables](#).

```
export TENANT_ACTION="install"
export AISERVICE_INSTANCE_ID="value"
export AISERVICE_TENANT_ID="value"
export AISERVICE_NAMESPACE="aiservice-value"
export ROLE_NAME="aiservice_tenant"

# Data Reporter Operator variables
export AISERVICE_DRO_SECRET_NAME="value"
export AISERVICE_DRO_API_KEY="value"
export AISERVICE_DRO_URL="value"
export AISERVICE_DRO_CA_CERT="value"

# Suite License Server variables
export AISERVICE_SLS_SECRET_NAME="value"
export AISERVICE_SLS_REGISTRATION_KEY="value"
export AISERVICE_SLS_URL="value"
export AISERVICE_SLS_CA_CERT="value"

# Reliability Strategy Library variables (optional)
export RSL_URL="value"
export RSL_ORG_ID="value"
export RSL_TOKEN="Bearer value"
export RSL_CA_CERT="value"

# Watsonx.ai variables
export AISERVICE_WATSONXAI_APIKEY="value"
export AISERVICE_WATSONXAI_URL="value"
export AISERVICE_WATSONXAI_PROJECT_ID="value"

# TENANT
export AISERVICE_SLS_SUBSCRIPTION_ID="value"
export AISERVICE_TENANT_ENTITLEMENT_START_DATE="value"
export AISERVICE_TENANT_ENTITLEMENT_END_DATE="value"

ansible-playbook playbooks/run_role.yml
```

What to do next

If any updates are made to the playbooks, you must use the following command to rebuild the collection:

```
cd git/ansible-devops/ibm/mas_devops/
ansible-galaxy collection build
ansible-galaxy collection install ibm-mas_devops-*.tar.gz --ignore-certs --force
rm ibm-mas_devops-*.tar.gz
```

Environment variables

To deploy Maximo AI Service by using the Ansible collection, you must set certain environment variables.

<i>Table 32. General environment variables</i>	
Variable	Description
IBM_ENTITLEMENT_KEY	Specify <code>\$\$MAS_ENTITLEMENT_KEY</code> .

Table 32. General environment variables (continued)

Variable	Description
IBM_ENTITLEMENT_USERNAME	Specify your IBM entitlement username.
MAS_CATALOG_VERSION	Specify the latest Maximo Application Suite operator catalog version for amd64. Use the <i>version-amd64</i> format. For more information see IBM Maximo Operator Catalog .
MAS_CONFIG_DIR	Specify the configuration directory that you created.
ENVIRONMENT_TYPE	Specify production or non-production.
AISERVICE_INSTANCE_ID	Specify your installation namespace for your Maximo AI Service. Any value can be the namespace.
MAS_ENTITLEMENT_KEY	Specify your IBM entitlement key. For more information, see “Obtaining your IBM Entitlement key from the IBM Entitled Registry” on page 203.
MAS_ENTITLEMENT_USERNAME	Specify your IBM entitlement user for IBM Container Registry.
AISERVICE_CHANNEL	Specify 9.1.x.
AISERVICE_TENANT_ENTITLEMENT_TYPE	Specify standard, essentials, or premium.

Table 33. Storage provider environment variables

Variable	Instructions
INSTALL_MINIO	Specify whether to install MinIO. Default value is <code>false</code> . Set to <code>true</code> if you want to install MinIO. If set to <code>true</code> , you must specify values only for the <code>MINIO_ROOT_PASSWORD</code> and <code>AISERVICE_S3_ACCESSKEY</code> variables. If you set to <code>false</code> because you already have an instance of MinIO installed, you must specify values for all of the following storage provider variables unless otherwise specified. If you plan to use Amazon Web Services S3 instead of MinIO, set <code>INSTALL_MINIO</code> to <code>false</code> .
MINIO_ROOT_PASSWORD	Specify your MinIO password. Do not specify a value for this variable if you are using Amazon Web Services S3.
AISERVICE_S3_ACCESSKEY	Specify your storage provider access key.
AISERVICE_S3_SECRETKEY	Specify <code>\$_{MINIO_ROOT_PASSWORD}</code> . Do not specify a value for this variable if you are using Amazon Web Services S3.
AISERVICE_S3_HOST	For MinIO, specify <code>http://minio-service.minio.svc.cluster.local</code> . For Amazon Web Services S3, specify your unique storage host.

Table 33. Storage provider environment variables (continued)

Variable	Instructions
AISERVICE_S3_REGION	If you are using MinIO, do not specify a value. If you are using Amazon Web Services S3, specify your storage provider region for your Amazon Web Services S3 instance.
AISERVICE_S3_PROVIDER	Specify minio or aws.
AISERVICE_S3_SSL	Specify your storage SSL value, either true or false.
AISERVICE_S3_PORT	For MinIO, specify 9000. For Amazon Web Services S3, specify your unique port.
AISERVICE_S3_PIPELINES_BUCKET	For MinIO, specify km-pipelines. For Amazon Web Services S3, specify your unique bucket name for pipelines.
AISERVICE_S3_TENANTS_BUCKET	For MinIO, specify km-tenants. For Amazon Web Services S3, specify your unique bucket name for tenants.
AISERVICE_S3_TEMPLATES_BUCKET	For MinIO, specify km-templates. For Amazon Web Services S3, specify your unique bucket name for templates.
AISERVICE_S3_BUCKET_PREFIX	Specify a prefix that has fewer than 6 characters.
AISERVICE_S3_REGION	For MinIO, send an empty string. For Amazon Web Services S3, specify your unique Amazon Web Services S3 region.

Table 34. watsonx environment variables

Variable	Instructions
AISERVICE_WATSONXAI_APIKEY	Specify your watsonx platform API key. The user who is associated with the API key must have a minimum of editor permission for the project and viewer permission for the space. For more information about how to create the API key, see Generating API keys for authentication .
AISERVICE_WATSONXAI_URL	Specify your watsonx endpoint URL. For more information about how to get the URL, see Introduction to IBM watsonx.ai software .
AISERVICE_WATSONXAI_PROJECT_ID	Specify your watsonx project ID. For more information, see Finding your project ID .

Table 35. Data Reporter Operator environment variables

Variable	Instructions
INSTALL_DRO	Specify whether to install Data Reporter Operator . Default value is <code>false</code> . Set to <code>true</code> if you want to install Data Reporter Operator . If set to <code>true</code> , do not set values for the other Data Reporter Operator variables. If you have an existing instance of Data Reporter Operator , set to <code>false</code> and then set values for the other Data Reporter Operator variables.
DRO_CONTACT_EMAIL	Specify a valid email.
DRO_CONTACT_FIRSTNAME	Specify a given name. You can specify any person in your organization.
DRO_CONTACT_LASTNAME	Specify a surname. You can specify any person in your organization.
AISERVICE_DRO_TENANT_ID	Specify your namespace ID.
AISERVICE_DRO_SECRET_NAME	Specify the secret name. You can retrieve the secret name from Red Hat OpenShift web console.
AISERVICE_DRO_API_KEY	Specify the API key. To locate the API key, complete the following steps: <ol style="list-style-type: none"> 1. In Red Hat OpenShift web console, click Home > Projects and search for the <code>aiservice-instance_id</code> project. 2. Select the <code>aiservice-instance_id</code> project and in the Inventory section, click <i>number</i> Secrets. 3. Select the dro-token secret. 4. In the Data section, click Reveal values. The value is the API key.
AISERVICE_DRO_URL	Specify <code>https://ibm-data-reporter-redhat-marketplace.cluster_domain</code>
AISERVICE_DRO_CA_CERT	Specify the certificate file content. To locate the certificate content, complete the following steps: <ol style="list-style-type: none"> 1. In Red Hat OpenShift web console, click Home > Projects and search for the <code>aiservice-instance_id</code> project.. 2. Select the <code>aiservice-instance_id</code> project and in the Inventory section, click <i>number</i> Secrets. 3. Select the dro-certificates secret 4. In the Data section, click Reveal values. The value is the certificate content.

Table 36. Suite License Service environment variables

Variable	Instructions
INSTALL_SLS	Specify whether to install Suite License Service. The default value is <code>false</code> . You already have Suite License Service installed if you have an existing deployment of Maximo Application Suite. Set to <code>true</code> only if you are deploying Maximo AI Service as part of a new deployment of Maximo Application Suite. If <code>true</code> , do not set values for the other Suite License Service variables. If <code>false</code> , set values for all Suite License Service variables unless otherwise noted.
INSTALL_MONGO	Specify whether to install MongoDB. The default value is <code>false</code> . MongoDB is installed as part of Suite License Service. If set to <code>false</code> , you must specify a value for the <code>SLS_MONGODB_CFG_FILE</code> variable.
SLS_MONGODB_CFG_FILE	If <code>INSTALL_MONGO</code> is <code>false</code> , specify <code>\${MAS_CONFIG_DIR}/mongo-mongoce.yml</code> . If <code>INSTALL_MONGO</code> is <code>true</code> , do not set a value for this variable.
SLS_LICENSE_ID	Specify your license ID. For more information about the entitlement license, see “Installing Maximo Application Suite and Maximo Manage” on page 906 .
SLS_LICENSE_FILE	Specify the directory that stores the <code>authorized_entitlement.lic</code> file. You must create the directory, for example, <code>/Users/name/Documents/GitHub/devops-configs/config/authorized_entitlement.lic</code> . For more information about the entitlement license, see “Installing Maximo Application Suite and Maximo Manage” on page 906 .
AIService_SLS_Subscription_ID	Specify the subscription ID.
AIService_SLS_Secret_Name	Specify the secret name. You can retrieve the secret name from the Red Hat OpenShift web console,
AIService_SLS_Registration_Key	Specify the registration key. To locate the key, complete the following steps: <ol style="list-style-type: none"> 1. In the Red Hat OpenShift web console, click Home > Projects and search for the <code>ibm-sls</code> project.. 2. Select the ibm-sls project and in the Inventory section, click number Pods. 3. Select the <code>sls-api-licensing</code> pod and on the Environment tab, in the Single values section, copy the value for REGISTRATION_KEY. That registration key value is the value for the <code>AIService_SLS_Registration_Key</code> variable.

Table 36. Suite License Service environment variables (continued)

Variable	Instructions
AISERVICE_SLS_URL	Specify the URL.
AISERVICE_SLS_CA_CERT	Specify the certificate file content. To locate the content, complete the following steps: <ol style="list-style-type: none"> 1. In the Red Hat OpenShift web console, click Home > Projects and search for the <code>ibm-sls</code> project.. 2. Select the ibm-sls project and in the Inventory section, click number Secrets. 3. Select the sls-cert-ca secret and in the Data section, for ca.cert, click Reveal values. The value is the certificate content.

Table 37. Db2 environment variables

Variable	Instructions
INSTALL_DB2	Specify whether to install Db2. If the default value is <code>false</code> . If <code>false</code> , which means you have an existing version of Db2, set values for the other Db2 variables. If <code>true</code> , do not set values for the other Db2 variables.
AISERVICE_DB2_USERNAME	Specify the username.
AISERVICE_DB2_PASSWORD	Specify the password.
AISERVICE_DB2_JDBC_URL	Specify the JDBC URL.
AISERVICE_DB2_SSL_ENABLED	Specify the SSL value.
AISERVICE_DB2_CA_CERT	Specify the certificate.

Table 38. Reliability Strategy Library environment variables

Variable	Instructions
RSL_URL	Specify <code>https://api.rsl-service.suite.maximo.com/api/v3/vector/query</code>
RSL_ORG_ID	Specify the Reliability Strategy Library organization ID. The organization ID is located in the mxe.rs.rslogid system property in the System Properties application in Maximo Manage.
RSL_TOKEN	Specify the Reliability Strategy Library token. The token format is <code>Bearer token</code> , where <code>token</code> is your container library entitlement key.

Note: You must specify the Reliability Strategy Library variables if you want to configure and use the AI recommendations in Reliability Strategies. If you do not want to enable that feature, the variables are optional. For more information, see [Enabling AI recommendations in Reliability Strategies](#).

Administering AI features

As a system administrator, you can enable AI features by setting up AI configurations in the AI configuration application in Maximo Manage. In Maximo Application Suite 9.1, you can enable an AI assistant, field value recommendations, including problem code recommendations for work orders, locating similar records, and AI recommendations in Reliability Strategies.

Before you begin

If you are deploying Maximo AI Service on-premises, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412.](#)

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

About this task

An *AI configuration* is type of record that you create that contains settings for your AI feature, such as the required model template that's available from Maximo AI Service, and where you can start and monitor any required training or inferencing. Each model template has unique data requirements and not all require training and inferencing. You manage AI configurations in the AI configuration application in Maximo Manage.

Adding AI configurations

Use the AI configuration application in Maximo Manage to configure AI features for Maximo Manage.

Enabling the AI assistant

You can enable an AI assistant for Maximo Manage, Maximo Health, Operational Dashboard, and asset dashboards. The assistant uses the nl2oslc model template.

Before you begin

If you are using Maximo AI Service on-premises, before you begin, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412.](#)

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

In both cases, optionally, you can complete prompt tuning. *Prompt tuning* is a process in which you specify parameters, also known as *prompts*, for the model to consider. Each prompt is context that the assistant uses when answering questions or requests. Complete prompt tuning only if you want to alter how the assistant responds to certain questions.

If you are enabling your organization's first deployment of the AI assistant, you might skip prompt tuning until you enabled and completed testing of the assistant. If the assistant can answer questions accurately for your organization, you do not need to complete prompt tuning. For more information, see [“Prompt tuning for the AI assistant” on page 437.](#)

The quality of the assistant's answers is directly dependent on the quality of data for assets, work orders, and service requests.

About this task

Starting in 9.1, the assistant can answer questions about only assets, work orders, and service requests. In the Object Structures application, the **Use for assistant** checkbox is selected in the object structures that are related to these data types. If you do not want the assistant to answer questions about those data types, you can clear the checkboxes. The checkbox is available for all object structures, but the assistant can use only those object structures that have the checkbox selected by default.

Procedure

1. In Maximo Manage, open the AI configuration application and then in the table, click the **ASSISTANT** AI configuration.

If you delete the default configuration and must add another configuration for the assistant, ensure that you use the nl2oslc model template and that you select the latest version of the template.

2. Review the content in the **Edit AI configuration** dialog.

- a) Click **Actions > Edit**.

- b) In the **Template version** field, ensure that the latest version is selected. To view all versions, click the **Lookup** icon.

- c) Edit the quick starters.

Quick starters are suggested requests or questions that appear in the welcome message when the assistant is opened by one of your users. Quick starters can enable users to explore the assistant and its features more efficiently. You can use the default quick starters or specify your own. Quick starters cannot be customized to certain roles or users. Specify requests or questions that most of your users might want and that the assistant can answer based on the selected object structures.

- d) Click **OK**.

3. Click **Actions > Activate**.

Activating the AI configuration indicates that the AI configuration is prepared.

4. From the **Actions > Create model**

Creating the model enables the assistant on the user interface. The assistant appears in the lower right of the interface for Maximo Manage, Maximo Health, Operational Dashboard, and asset dashboards.

You can drag and drop the icon to different areas on the user interface, but the assistant chat window always opens on the right side of the page.

What to do next

Access to the AI assistant chat window is managed through security groups. In Security Groups, open a security group and on the **Applications** tab, in the Applications table, in the **Description** column, search for assistant. In the Options for assistant table, as a system administrator, you can grant or remove access to the Maximo Assistant option.

Users must also have access to work orders, service request, and assets. Security groups can grant users access to the chat window, but if users do not have access to data, the assistant cannot answer questions. Users can only receive answers that contain data that they have access to.

Later, if you want to edit the quick starters or other configuration settings, you must deactivate the model first. Click **Actions > Deactivate**.

The assistant does not currently include a feedback mechanism for answers. Conversation history cannot be saved. The assistant does not support integrations with external systems.

For more information about using the assistant, see [Using the assistant](#).

Prompt tuning for the AI assistant

As a system administrator, you can affect how the AI assistant responds to questions or requests. For example, you can clarify data that is specific to your organization.

To complete prompt tuning, you specify parameters, also known as *prompts*, for the model to consider. Each prompt is the context that the assistant uses when answering questions or requests. Prompts are not rules that the assistant follows. To provide an answer, the assistant uses any prompts that are specified for tuning and the user's question or request. Data quality and availability can also affect responses.

Prompt tuning can include basic prompts and template prompts. A *basic prompt* is a natural language statement that is used by the model directly. A *template prompt* is a prompt that contains variables and is used to generate multiple, similar prompts. You might specify a template prompt if you want a prompt to

apply to data under various conditions. You might specify a basic prompt if you have an exact attribute or value that you want the model to consider.

Complete prompt tuning only if you want to alter how the assistant responds to certain questions. If you are enabling your organization's first deployment of the AI assistant, you might skip prompt tuning until you have enabled and completed testing of the assistant. If the assistant can answer questions accurately for your organization, you do not need to complete prompt tuning.

Basic prompts

The following list contains example prompts. You can use these prompts for prompt tuning or write your own prompts.

- 'Site ID' is typically the name of a site, place, or location
- 'Asset condition' implies **assethealth**
- High priority work order implies `wopriority < 3`
- 'SR' is the same as 'service request'
- Estimated total cost of a work order implies the summation of `estlabcost`, `estmatcost`, `stservcost`, and `esttoolcost`
- 'Equipment' is the same as 'asset'
- 'Downtime' or 'outage hours' implies `totdowntime`

Template prompts

When you write template prompts, use the following format:

term1{*object*[*where*].*attributefrom*} implies *term2*'{*attributeto*}'

term1

A keyword.

term2

A keyword.

object

The name of the referenced object.

where

A where clause. This value is optional.

attributefrom

The attribute of the object that is mapped from.

attributeto

The attribute of the object that is mapped to.

Do not specify more than one object in a prompt. Ensure that both *attributefrom* and *attributeto* are from the same object.

For example, the following prompt is a template prompt: `Failures related to {failurecode[failurecode in (select failurecode from failurelist where type is null)].description} implies failurecode = '{failurecode}'`

The prompt enables failure code descriptions to be used primarily to understand and apply failure codes. To tie a failure code to its description, you can write one individual, basic prompt per failure code, or you can use a template prompt to automatically generate one prompt per failure code.

term1 is `Failures related to`, and *term2* is `failurecode`.

object is `failurecode`. A period (.) is used to separate the object from the attribute. If a period is not used, the value is treated as an attribute.

where is `failurecode in (select failurecode from failurelist where type is null)`.

attributefrom is `description`, and *attributeto* is `failurecode`.

The following prompts are examples of prompts that might be generated from the template prompt:

- Failures related to IT implies failurecode = 'IT'
- Failures related to Pipe Failures implies failurecode = 'PIPES'
- Failures related to CONVEYOR LINE FAILURES implies failurecode = 'CONVEYOR'

Procedure

To complete prompt tuning, complete the following steps:

1. In Maximo Manage, in the Database Configuration application, from the **More actions** menu, click **Prompt tuning for the assistant**. Some prompts are provided by default. If you do not want to use these prompts, deactivate them by clearing the **Activate** checkbox. Do not edit the default prompts.
2. Click the **New Row** icon.
3. In the **Name** field, specify a name for the prompt.
4. Optionally, in the **Description** field, specify a description for the prompt.
5. Optionally, in the **Object** field, select the object that the prompt relates to. For example, if the prompt is used for sites, select an object structure for sites. Selecting an object for your prompt can enable the assistant to better understand the prompt and apply it as you intend.
6. If the prompt is a template prompt, select the **Template** checkbox.
7. In the **Prompt** field, specify your prompt.
8. Select the **Activate** checkbox. When a basic prompt is activated, it is sent to the assistant as metadata. When the template prompt is activated, those generated prompts are sent to the assistant as metadata.
9. Click **OK**.

Enabling field value recommendations

You can use the AI configuration application to enable AI-recommended values for fields. The field value recommendations use the mcc model template. This configuration does not support problem code recommendations in work orders. That AI feature requires a separate configuration.

Before you begin

If you are deploying Maximo AI Service on-premises, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#).

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

In both cases, prepare the required data and components. For more information, see [“Preparing required components for field value recommendations” on page 443](#).

About this task

After recommendations are enabled, the recommendation feature is accessible to all users who have access to the base record. For example, if you enable this feature for work orders, all users with access to work orders can see the feature.

To configure AI-recommended problem codes for work orders, you must enable another type of AI configuration. For more information, see [“Enabling recommended problem codes for Work orders” on page 446](#).

Procedure

1. In Maximo Manage, in the AI configuration application, click **Add configuration**.
2. In the **Name** field, specify a name

3. Optional: In the **Description** field, specify a description.
4. In the **Template** field, select the mcc template.
5. In the **Template version** field, select the latest template version.
6. In the **Object structure** field, select the object structure that you used in the invocation channels and the training and inference filters.
7. Optional: In the **Target object path** field, select the hierarchy path for the child object.
This value is required only when the inference is on an attribute of a child object and not the main object of the object structure. When you select child objects, during training and inferencing, the target object path maps the hierarchy path in the object structure. The hierarchy path supports only the immediate child to the root object.
8. In the **Attribute** field, specify the specific object structure attribute that represents the field for which you want to generate recommended values. Select the same attribute as you specified in the invocation channel request template for training.
9. Optional: In the **Target description** field, specify the name of the attribute that contains the description for the target attribute.
For example, if the target attribute is for work types, specify the description field object for work types. By selecting the target description, the model can more effectively recommend values.
10. In the **Training invoke channel** field, select the invocation channel for training.
11. In the **Training filter** field, select the training filter.
12. In the **Inference invoke channel** field, select the invocation channel for inferencing.
13. In the **Inference filter** field, specify the inferencing filter.
14. Click **Create**.
15. In the AI configurations table, select the AI configuration that you created.
16. Set up arguments.
Arguments control some aspects of the models output, such as the threshold for acceptable values.
 - a) Click **Actions** > **Set arguments**.
 - b) Specify a value for the arguments. The following table describes the arguments. You must set a value for the features argument.

Argument key	Description	Type	Default value
features	<p>Specify description-related features that you specified in the query template for the object structure. Specify a comma-separated list.</p> <p>If the target attribute description is sourced from a related object, specify only the attribute description.</p> <p>To specify a description that is a child attribute of multiple related objects, use the following notation: <i>relationship_name.attribute_name*</i></p>	string	No default value
score_threshold	<p>A score that measures how recommendable a value is. Values that have scores above the threshold are considered for recommendation.</p> <p>If you are setting up this type of recommendation for the first time, you might choose to set the score threshold to a lower value. A lower value increases the likelihood of any output from the model, although it does typically decrease desirable output, especially if the training filter does not contain diverse or adequate amounts of data.</p>	integer	0.5

c) Click **Save**.

17. Optional: In the **Edit AI configuration** dialog, specify information about your model.

a) Click **Actions > Edit**.

b) In the Additional details for AI explained section, provide any information that is specific to your organization and relevant to your end users to help them understand the model and its output.
An **AI** icon is located alongside your model's output. Your users can click the **AI** icon and then access any information you specify in this section alongside other general model information that's provided by IBM. You can complete the AI configuration process, review the **AI** icon and its content in context, and then edit this section later as needed.

c) Click **Save**.

18. Click **Actions > Check data requirement**.

When you check the data requirements, the training data is reviewed to determine if the data contains enough detail. If the data check fails, you must add or improve the quality of data in your training filter. For more information, see [“Preparing required components for field value recommendations”](#) on page 443.

19. Click **Actions > Activate**

Activating the AI configuration indicates that the AI configuration is prepared and the model is ready to be trained.

20. Optional: Change the frequency of the training process

Training is controlled on a crontask schedule. The AITRAINJOB crontask initiates training for all eligible AI configurations. By default, the crontask runs every five minutes.

For more information, see [“Cron tasks for training and inferencing”](#) on page 416.

If you want training to run sooner, you can edit the cron task schedule.

a) In Maximo Manage, in the Cron Task Setup application, open the AITRAINJOB cron task.

b) In the Cron Task Instances table, for the WOAI instance, in the **Schedule** field, change the value.

Wait a few minutes before continuing to the next step.

21. Click **Actions > Train model**

Training begins when the AITRAINJOB crontask runs. Training can take a few hours.

For conceptual information about training, see [“Model training overview”](#) on page 455.

You can monitor training in the **Model training log** table or in the **Model status** dialog.

The **Model training log** table is in the AI configuration that you created and contains step-by-step updates for the training process, but you must refresh the page to see updates. To refresh the page, click **Refresh**.

To access the **Model status** dialog, in the AI configuration, click **Actions > Check model status**.

The model accuracy score is a measure of how the model performs on the training data. The score represents the amount of values that the model recommended that it considers reasonable to be the correct or best value. The closer to 1, the more accurate the output likely is in context of the data on which the model was trained. If the model was trained on data that was not complete or diverse but the score threshold is low, the accuracy score might be high but the output is not accurate.

If training fails, you can complete some troubleshooting steps. For more information, see [“Troubleshooting Maximo AI Service and AI features”](#) on page 458.

What to do next

Inferencing is controlled on a crontask schedule and is initiated immediately after training. The AIINFJOB crontask initiates inferencing for all eligible AI configurations. By default, the crontask runs every hour. If inferencing is not run promptly after training is completed, you can change the crontask schedule.

1. In Maximo Manage, in the Cron Task Setup application, open the AIINFJOB cron task.

2. In the Cron Task Instances table, for the WOAI instance, in the **Schedule** field, change the value.

After inferencing is complete, locate the field that contains the AI recommendations. Confirm that the recommendations appear as expected. For recommendations to be enabled, the base record must be included in the inference filter.

If you want to retrain the model on new data in the same filters or you want to edit the arguments, make the changes and then in the AI configuration, click **Actions > Re-train model**.

If you need to change other configuration settings, you must deactivate the configuration first. In the AI configuration, click **Actions > Deactivate**, edit the configuration, activate the model again, and then click **Actions > Train**.

Preparing required components for field value recommendations

Before you can enable the AI configuration for field recommendations, you must prepare your training and inference data and create training and inference filters.

About this task

If you are enabling problem code recommendations in work orders, you must complete another task. For more information, see [“Enabling recommended problem codes for Work orders”](#) on page 446.

Procedure

1. Decide what data you will use for training and prepare that data.

You must choose at least 20 records for training. 10 of those records must contain one possible value type that you want AI to recommend. The other 10 records must contain another value type. For example, if you wanted AI to recommend work types, choose 10 work orders that contain one work type and 10 work orders that contain another work type.

To decide what records to choose, consider the following best practices:

- Use unique records. The training process filters out duplicate records of the same type, for example, duplicate purchase orders.
- Use diverse records of the same type. Ideally, the records include a range of descriptions and address a range of problems but are all the same type of record, for example, all purchase orders.
- Use records that have accurate instances of the values.
- Ensure that each possible recommended value has an accurate description and details.
- Ensure that each record has an accurate description.
- Although the minimum is 20 records, the larger, more diverse, and accurate your training filter data is, the more likely the model can accurately recommend values. The ideal training filter contains 20 - 50 records per possible value.

2. Create the training and inference filters.

Filters, also known as *query definitions*, determine what records are used for training or inferencing and what records have the value recommendation feature enabled. Create one filter for training and one filter for inferencing. All of the data that is included in the filter is used for the filter's respective process.

These filters are defined as query definitions for an object structure. Select an object structure that represents the records that you want recommended values for, such as MXAPIWODETAIL. You use the same object structure for both filters, and then use the query definitions to specify a subset of data for each filter.

- a. In Maximo Manage, in the Object structures application, open the object structure.
- b. In the More actions menu, click **Query definition**.
- c. Click the plus (+) icon.
- d. In the **Query Type** field, select **osclause**.
- e. In the **Query Clause Name** field, specify a name for the filter.

Ensure that the name indicates whether the filter is for training or inferencing and what AI configuration it is associated with.

- f. In the **Query Clause** field, specify a WHERE clause that retrieves the records that you want to use for training or inferencing.

The following text is an example of a query clause that retrieves all EM and CM work orders.

```
worktype in ('EM','CM')
```

For the training filter, specify a clause that retrieves the records that you chose in step 1.

For the inference filter, specify a clause that returns all of the records for which you want to enable the recommendation feature. To improve system performance, set an age limitation on the records. For example, if the training filter retrieves 20 EM and CM work orders, configure the inference filter to retrieve all CM and EM work orders in the last 30 days. The following text is an example of a query clause for an inference filter.

```
worktype in ('EM', 'CM') and status in (select value from synonymdomain where domainid = 'WOSTATUS' and maxvalue in ('WAPPR')) and reportdate > CURRENT DATE - 30 DAYS
```

- g. Click the **Is Public?** checkbox.
 - h. Click **OK**.
 - i. Repeat the steps to add the other filter.
3. Determine the attribute names for the following values. You must have the attribute names to complete step 4.
 - Each feature, or field, that can be compared to determine recommended values, for example, **description** and **description_longdescription**.
 - The unique attribute ID for the object for which inferencing is completed. For example, WORKORDERID is the unique attribute ID for the WORKORDER object.
 - The attribute name for the field that the model is recommending a value for. For example, if you want the model to recommend work types, determine the attribute name for work types.
 - The description attribute for the attribute name.
 4. Create query templates for training and inferencing.

For the object structure that the filters use, add a query template for training and a query template for inferencing. Query templates contain attributes. For AI configurations, the template defines the JSON structure for training data. Maximo AI Service uses the template to fetch the data.

- a) In Maximo Manage, in the Object Structures application, open the object structure.
- b) In the **More actions** menu, click **Query Template**.
- c) In the Query Templates for table, click **New Row**.
- d) In the **Description** field, specify a description that indicates that the query template is used for training or inferencing and what AI configuration it is associated with.
- e) In the **Page Size** field, specify 1,000 for training or 10 for inferencing.
- f) Select the **Default Projection?** checkbox.
- g) In the **Query Template Attributes for** table, add a row for each required attribute.

For the training, the table must contain the following attributes:

- Each feature, or field, that can be compared to determine recommended values. Specify a comma-separated list.

To specify a feature that is a child attribute of a single related object, use the following notation:
relationship.attribute

To specify a feature that is a child attribute of multiple related objects, use the following notation:
rel.relationship{attribute}

- The unique attribute ID for the object for which inferencing is completed. For example, WORKORDERID is the unique attribute ID for the WORKORDER object. You determine the object as part of preparing the data. For more information, see [“Preparing required components for field value recommendations”](#) on page 443.
- The field that the model is recommending a value for. For example, if you want the model to recommend work types, specify the attribute name for work types.
- The description attribute for the attribute name. Use the following format:
{attribute}_description

If the target description is sourced from a related object, use the following format:
relationship.description--attribute_description

For more information about the notation, see [“Query template JSON examples”](#) on page 456.

For inferencing, the table must contain the features and the unique attribute ID for the object.

5. Create an invocation channel for training.

Invocation channels facilitate training and inferencing for the AI configuration. For more information about creating channels, see [Creating invocation channels](#).

- In Maximo Manage, in the Invocation Channels application, search for and select the AITRAINWOPROBLEMCODE channel.
- In the **More actions** menu, click **Duplicate Invocation Channel**.
- In the **Invocation Channel** field, specify a unique name and description. Ensure that the name and description indicates that the channel is used to train AI models and which AI configuration it is used for.
- In the **Endpoint** field, ensure AIBROKERAPI is specified. If you edit any end point properties, ensure that the **Override** checkbox is selected.
- In the **Request Object Structure** field, select the object structure that you used for the inference and training filters. The object structure, in the invocation channel, is used to generate the compressed file for training and to fetch the data details for inferencing. Use the same object structure for both channels.
- In the **Request Template** field, select the query template for training that you created for the object structure.
- In the **Request Processing Class** field, ensure that `com.ibm.tivoli.maximo.ai.AITrainReqExit` is specified.

Note: When you add an AI configuration, you can choose to specify a hierarchy path. If a hierarchy path is specified, the `AITrainReqExit` processing class transforms the request JSON to make the child object that is referred by the hierarchy path in to the root object of the JSON.

For example, if an object structure contains PO as an object and POLINE as a child, and the hierarchy path is set to PO/POLINE, then the `AITrainReqExit` request processing class uses POLINE as the root object and removes the PO parent in the transformed JSON. The following text is an example structure:

```
PO
POLINE
POLINE
```

The following text is an example of an updated structure:

```
POLINE
POLINE
```

- In the **Response Processing Class** field, ensure that `com.ibm.tivoli.maximo.ai.AITrainRespExit` is specified.

6. Create an invocation channel for inferencing.

- a) In Maximo Manage, in the Invocation channels application, search for and select the AIINFWOPROBLEMCODE channel.
- b) In the **More actions** menu, click **Duplicate Invocation Channel**.
- c) In the **Invocation Channel** field, specify a unique name and description. Ensure that the name and description indicates that the channel is used for inferencing for AI models.
- d) In the **Endpoint** field, ensure AIBROKERAPI is specified. If you edit any end point properties, ensure that the **Override** checkbox is selected.
- e) In the **Request Object Structure** field, specify the object structure that you used for the inference filter.
- f) In the **Request Template** field, select the query template for inferencing that you created for the object structure.
- g) In the **Request Processing Class** field, ensure that `com.ibm.tivoli.maximo.ai.AIINFReqExit` is specified.
- h) In the **Response Processing Class** field, ensure that `com.ibm.tivoli.maximo.ai.AIINFRespExit` is specified.

Enabling recommended problem codes for Work orders

You can use the AI configuration application to enable AI-recommended problem codes for Work orders. The problem code recommendation AI feature uses the pcc model template.

Before you begin

If you are using Maximo AI Service on-premises, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage”](#) on page 412.

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

In both cases, prepare the required data. For more information, see [“Preparing required data for problem code recommendations”](#) on page 448.

About this task

The WOPROBLEMCODE AI configuration is available by default. The configuration uses preconfigured Maximo Manage integration components, such as an object structure and invocation channels. You do not have to configure these components to set up the AI configuration. Later, if you want to fine-tune how data is transferred or managed, you might alter some aspects of the components at your own digression. For more information about the components, see [Integration components for recommended problem codes in Work orders](#).

After problem code recommendations are enabled, the recommendation feature is accessible to all users who have access to Work orders.

Procedure

1. In Maximo Manage, in the AI configuration application, select the **WOPROBLEMCODE** configuration.
2. Review the **Edit AI configuration** dialog.
 - a) Click **Actions > Edit**.
 - b) In the **Template version** field, ensure that the latest version is selected. To view all versions, click the **Lookup** icon.
 - c) In the Additional details for AI explained section, provide any information that is specific to your organization and relevant to your end users to help them understand the model and its output.

An **AI** icon is located alongside your model's output. Your users can click the **AI** icon and then access any information that you specify in this section alongside other general model information

that's provided by IBM. You can enable the AI configuration, view the **AI** icon and its content in context, and then edit this section later as needed.

- d) Click **Save**.
3. Optional: Set up arguments.
 - a) Click **Actions > Set arguments**
 - b) Change the default value for the **score_threshold** argument.

A numerical score is used to measure how recommendable each problem code is for a given work order. Problem codes that have scores above the threshold are considered for recommendation.

If you are setting up problem code recommendations for the first time, you might choose to set the threshold to a lower value. A lower value increases the likelihood of any output from the model, although it does typically decrease desirable output, especially if the training filter does not contain diverse or adequate amounts of data.
 - c) Click **Save**.
4. Click **Actions > Check data requirement**.

When you check the data requirements, the training data is reviewed to determine whether the data contains enough problem codes and work orders. If the data check fails, you must add or improve the quality of data in your training filter. For more information, see [“Preparing required data for problem code recommendations”](#) on page 448.
5. Click **Actions > Activate**

Activating the AI configuration indicates that the AI configuration is prepared and the model is ready to be trained.
6. Optional: Change the frequency of the training process.

Training is controlled on a crontask schedule. The AITRAINJOB crontask initiates training for all eligible AI configurations. By default, the crontask runs every five minutes.

For more information, see [“Cron tasks for training and inferencing”](#) on page 416.

If you want training to run sooner, you can edit the cron task schedule.

 - a) In Maximo Manage, in the Cron Task Setup application, open the AITRAINJOB cron task.
 - b) In the Cron Task Instances table, for the WOAI instance, in the **Schedule** field, change the value.
 - c) Wait a few minutes before continuing to the next step.
7. Click **Actions > Train model**

Training begins when the AITRAINJOB crontask runs. Training can take a few hours.

You can monitor training in the **Model training log** table or in the **Model status** dialog.

The **Model training log** table is in the AI configuration that you created and contains step-by-step updates for the training process, but you must refresh the page to see updates. To refresh the page, click **Refresh**.

To access the **Model status** dialog, in the AI configuration, click **Actions > Check model status**.

The model accuracy score is a measure of how the model performs on the training data. The score represents the amount of values that the model recommended that it considers reasonable to be the correct or best value. The closer to 1, the more accurate the output likely is in context of the data on which the model was trained. If the model was trained on data that was not complete or diverse but the score threshold is low, the accuracy score might be high but the output is not accurate.

If training fails, you can complete some troubleshooting steps. For more information, see [“Troubleshooting Maximo AI Service and AI features”](#) on page 458.

For conceptual information about training, see [“Model training overview”](#) on page 455.

What to do next

Inferencing starts automatically after training. Inferencing is controlled on a crontask schedule. The AIINFJOB crontask initiates inferencing for all eligible AI configurations. By default, the crontask runs every hour. If inferencing is not running promptly after training is completed, you can change the crontask schedule.

1. In Maximo Manage, in the Cron Task Setup application, open the AIINFJOB cron task.
2. In the Cron Task Instances table, for the WOAI instance, in the **Schedule** field, change the value.

After inferencing is complete, check that problem code recommendations are enabled for work orders. For problem code recommendations to be enabled, the work order must be included in the inference filter. For more information about accessing the problem code recommendations, see [Using AI recommendations in work orders](#).

If you want to retrain the model on new data in the same filters or you want to edit the arguments, make the changes and then in the AI configuration, click **Actions > Re-train model**.

If you need to change other configuration settings, you must deactivate the configuration first. In the AI configuration, click **Actions > Deactivate**, edit the configuration, activate the model again, and then click **Actions > Train**.

Preparing required data for problem code recommendations

Before you can enable the AI configuration for problem code recommendations, you must prepare your training and inference data, review your training and inference filters, and manually select the work orders to include in training.

Before you begin

Review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#).

About this task

To prepare the data and components, you must complete the following high-level steps:

1. Review the default training and inferencing filters. The training filter determines what work orders can be selected for training. The inference filter determines what work orders are used to test the model after training and also what work orders have problem code recommendation enabled.
2. In Work orders, select which work orders are used for training. Select work orders that are included in the training filter.

Procedure

1. Review the training and inference filters.

Filters determine what work orders can be used for training, retrieve all of the work orders that are used for inferencing, and determine what work orders have the problem code recommendation feature enabled. The filters are defined as query definitions for the MXAPIWODETAIL object structure. For problem code recommendations, the object structure and the query definitions are available by default, but you can edit the query definition to change what work orders can be included in training or inferencing.

- a) In Maximo Manage, in the Object structures application, open the MXAPIWODETAIL object structure.
- b) In the More actions menu, click **Query definition**.
- c) In the Queries to be assigned table, click the **Filter** icon and locate the AITRAINFILTER query, which is used for training, and the AIINFERENCEFILTER filter, which is used for inferencing. Select a filter and then in the **Query Clause** field, review the WHERE clause and change as needed for your use case. For the AITRAINFILTER filter, if ai_usefortraining=1 is removed from the

Query Clause field, you can manually select work orders for inclusion in training but the work orders are not added to the training data.

The following text is an example of the training filter.

```
worktype in ('EM', 'CM') and ai_usefortraining = 1
```

The following text is an example of the inference filter. You do not manually select work orders for the inference filter. The filter retrieves all qualifying work orders. For a work order to have problem code recommendations enabled, the work order must be included in the inference filter.

```
worktype in ('EM', 'CM') and status in (select value from synonymdomain where domainid = 'WOSTATUS' and maxvalue in ('WAPPR')) and reportdate > CURRENT DATE - 30 DAYS and ai_usefortraining = 0
```

d) Click **OK**.

2. Select work orders for training.

From the work orders that are included in the training filter, you must select at least 20 work orders for training. 10 of those work orders must contain one problem code. The other 10 work orders must contain another problem code. To decide what work orders to choose, consider the following best practices:

- Use unique work orders. The training process filters out duplicate work orders.
- Use diverse work orders. Ideally, work orders include a range of descriptions and address a range of problems.
- Use work orders that have accurate problem codes.
- Ensure that each problem code has an accurate description. You can edit problem code descriptions in the Failure Codes application.
- Ensure that each work order has an accurate description. You can edit work order descriptions in either the Work orders or Work Order Tracking applications.
- Although the minimum is 20 work orders, the larger, more diverse, and accurate your training filter data is, the more likely the model can accurately recommend problem codes. The ideal training filter contains 20 - 50 work orders per problem code.

a) In Maximo Manage, open Work orders.

b) Select work orders to include in the training filter. Ensure that the work orders qualify for the training filter that you reviewed in step 1.

To select individual work orders, open a work order and then in the **Actions** menu, click **Add to AI training model**. A tag appears in the work order to indicate that the work order is added to the filter. Ensure that you save the work order.

To select work orders in bulk, in the Work orders table, first ensure that the Problem code column is added. To add the column, click **Manage columns** and then select the column in the list. After the column is added, click the **Filter** icon and then filter the work orders by the problem code that you want to include in training. Select the work order check boxes and then click **Add to AI training model**.

Note: The **Add to AI training model** action is applicable to only the training filter for only problem code recommendations. The action is available for all work orders, but if the query clause for the training filter does not include a work order, selecting the action for that work order causes the tag to appear on the work order but does not add the work order to the training data. If a work order qualified for the filter and was selected but then later was disqualified from the filter, the tag remains on the work order.

Enabling locating of similar work orders

You can use the AI configuration application to AI to find work orders that are similar to another work order. For example, if applied to work orders, you can find work orders that are similar to another work order. The similar work orders feature uses the similarity model template.

Before you begin

If you are using Maximo AI Service on-premises, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#).

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

In both cases, you can optionally review the data that is used for configuration and inferencing. For more information, see [“Preparing required data for problem code recommendations” on page 448](#).

About this task

The WOSIMILARITY AI configuration is available by default. The configuration uses preconfigured Maximo Manage integration components, such as an object structure and invocation channels. You do not have to configure these components to set up the AI configuration. Later, if you want to fine-tune how data is transferred or managed, you might alter some aspects of the components at your own digression. For more information about the components, see [Integration components for locating similar work orders](#).

The AI configuration for finding similar work orders does not require training but does require some automated configuration processes after activation and prior to inferencing.

Procedure

1. In Maximo Manage, in the AI configuration application, in the table, click the **WOSIMILARITY** AI configuration.

If you delete the default configuration and must add another configuration for the same feature, ensure that you use the similarity model template and that you select the latest version of the template.

2. In the **Edit AI configuration** dialog, review the default settings.

- a) Click **Actions > Edit**.

- b) Optional: In the **Related Records** field, specify a WHERE clause to limit what work orders can be reviewed for similarity.

For example, you can limit the search to work orders that are for a certain asset or of a certain type. If a work order does not meet the criteria of the WHERE clause, you can trigger the similarity feature from that work order, but no similar work orders are returned.

In Maximo Manage, in the System properties application, you can configure the `mxe.ai.comparetolimit` property to determine the maximum number of work orders that can be reviewed for similarity.

- c) In the **Template version** field, ensure that the latest version is selected. To view all versions, click the **Lookup** icon.

- d) Optional: In the Additional details for AI explained section, provide any information that is specific to your organization and relevant to your end users to help them understand the model and its output.

An **AI** icon is located alongside your model's output. Your users can click the **AI** icon and then access any information that you specify in this section alongside other general model information that's provided by IBM. You can complete the AI configuration process, view the **AI** icon and its content in context, and then edit this section later as needed.

- e) Click **Save**.

3. Set up arguments.

- a) Click **Actions > Set arguments**.

- b) For each of the arguments, review any default value and change as needed. Specify values for the **date_field** and **features** arguments. The following table describes all of the arguments.

<i>Table 40. Arguments for similarity model template</i>			
Argument key	Description	Type	Default value
date_field	A date field in work orders that is used to determine similarity. If you change the value for this parameter, you must also change the query templates for similarity in the MXAPIWODETAIL object structure. For more information, see Integration components for locating similar work orders .	string	changedate
features	Fields to compare to determine how similar work orders are. If you change the value for this parameter, you must also change the query templates for similarity in the MXAPIWODETAIL object structure. For more information, see Integration components for locating similar work orders .	string	description, description_longdescription
history_days	The maximum age in days of work orders that can be included for inferencing. This value is calculated from the base work order's date field value.	integer	30
max_similarity_records	The maximum number of similar work orders that can be returned. Note: In the System properties application, you can configure the <code>mxe.ai.comparetolimit</code> property to determine the maximum number of work orders that can be reviewed for similarity.	integer	10

<i>Table 40. Arguments for similarity model template (continued)</i>			
Argument key	Description	Type	Default value
retention_days	The maximum age in days from the current date for work orders that can be included for configuration. Work orders can be at most 1000 days old, but the higher this value, the longer it can take to complete inferencing and render results. This value also determines the age of work orders that can be similar. For example, a work order with a change date of today can return similar work orders that have a change date of up to 90 days ago.	integer	90
similarity_threshold	The minimum similarity score that work orders must have to be included.	integer	70

c) Click **Save**.

4. Click **Actions > Activate**.

Activating the AI configuration indicates that the AI configuration is prepared and the model is ready to be processed by the cron tasks. Activation also automatically runs configuration and inferencing.

5. Click **Actions > Create model**.

What to do next

Inferencing is controlled on a crontask schedule. The AIINFJOB crontask initiates inferencing for all eligible AI configurations. By default, the cron task runs every hour. If inferencing does not start promptly after the model is activated, you can edit the cron task schedule:

1. In Maximo Manage, in the Cron Task Setup application, open the AIINFJOB cron task.
2. In the Cron Task Instances table, for the WOAI instance, in the **Schedule** field, change the value.

For more information about the cron task, see [“Cron tasks for training and inferencing”](#) on page 416.

If you want to reconfigure the model on new data in the same filters or you want to edit the arguments, make any changes and then in the AI configuration, click **Actions > Reconfigure model**.

If you need to change other configuration settings, you must first deactivate the configuration. In the AI configuration, click **Actions > Deactivate**, edit the configuration and then activate the configuration again. Activating the similarity AI configuration automatically starts reconfiguration and inferencing.

Access to the similarity feature is controlled through a signature option for the associated object structure. By default, the MXAPIWODETAIL object structure is used. In the Security Groups application, select the security group to which you want to grant access. On the **Object Structures** tab, in the Object

Structures table, select the object structure that you associated with the similarity AI configuration. In the Options for table, for the View similar work order option, select the **Grant Access** checkbox.

For more information about using this feature to locate similar work orders, see [Viewing similar work orders](#).

Reviewing data for configuration and inferencing for finding similar work orders

The AI configuration for finding similar work orders is available by default. The configuration includes default invocation channels and filters, but you can create your own configuration and inference filters.

Before you begin

Review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#).

About this task

The AI configuration for finding similar work orders does not require training but does require some automated configuration processes after activation and before inferencing.

To optionally prepare the data, complete the following high-level steps.

1. Review the configuration and inference filters.
2. Review the query templates for configuration and inferencing.
3. Review the invocation channels for configuration and inferencing.

Procedure

1. Review the configuration and inference filters.

Your filters, also known as *query definitions*, determine what work orders can be used for configuration or inferencing and what work orders have the similarity feature enabled. One filter exists for configuration data and one filter exists for inferencing data. All of the data that is included in the filter is used for the filter's respective process.

The default filters are defined as query definitions for the MXAPIWODETAIL object structure.

- a. In Maximo Manage, in the Object structures application, open the MXAPIWODETAIL object structure.
- b. In the More actions menu, click **Query definition**.
- c. In the Queries to be assigned table, click the **Filter** icon and locate the WOSIMILARITYCFGFILTER query, which is used for configuration, and the WOSIMILARITYINFFILTER query, which is used for inferencing. Select a filter and then in the **Query Clause** field, review the WHERE clause and change as needed for your use case.

The following text is an example of the configuration filter. Ensure that the filter retrieves a diverse set of work orders, that the work orders have sufficient data, such as detailed descriptions, and that the work orders have the values that you want to use to determine similarity.

```
status in (select value from synonymdomain where domainid = 'WOSTATUS' and maxvalue not in ('CAN', 'CLOSE'))
```

The following text is an example of the inference filter. The inference filter is meant to retrieve all of the work orders for which you want to enable the similarity AI feature to be initiated from. To improve system performance, set an age limitation on the work orders, for example, all work orders in the last 30 days.

```
status not in (select value from synonymdomain where domainid = 'WOSTATUS' and maxvalue in ('CAN', 'CLOSE')) and (worktype != 'PM' or worktype is null)
```

- d. Click **OK**.

If you choose to create new filters, ensure that in the **Query Type** field, select **osclause** and ensure that the **Is Public?** checkbox is selected.

2. Review the query templates for configuration and inferencing.
 - a) In Maximo Manage, in the Object Structures application, open the MXAPIWODETAIL object structure.
 - b) In the **More actions** menu, click **Query Template**.
 - c) In the **Query Templates for** table, click the **Filter** icon and locate the WOSIMILARITYCFG template, which is used for configuration, and the WOSIMILARITYINF template, which is used for inferencing.

By default, the templates contains the following attributes:

- changedate
- description
- description_longdescription
- workorderid

If you choose to create new templates, ensure the **Default Projection?** checkbox is selected and include the following attributes in the **Query Template Attributes for** table:

- Each feature, or field, that can be compared to determine similarity, for example, **description** and **description_longdescription**. Specify a comma-separated list.

To specify a feature that is a child attribute of a single related object, use the following notation:
relationship.attribute

To specify a feature that is a child attribute of multiple related objects, use the following notation:
rel.relationship{attribute}

- The unique ID for the object structure that you used for the configuration and inference filters.

d) Click **OK**.

3. Review the invocation channels for configuration and inferencing.

Invocation channels facilitate configuration and inferencing for the AI configuration. For more information about creating channels, see [Creating invocation channels](#).

- a) In Maximo Manage, in the Invocation channels application, search for and select either the AIWOSIMILARITYINF channel, which is used for inferencing, or the AIWOSIMILARITYCFG channel, which is used for configuration.
- b) Review the values.

Note: When you add an AI configuration, you can choose to specify a hierarchy path. If a hierarchy path is specified, the `AITrainReqExit` processing class transforms the request JSON to make the child object that is referred by the hierarchy path in to the root object of the JSON.

For example, if an object structure contains PO as an object and POLINE as a child, and the hierarchy path is set to PO/POLINE, then the `AITrainReqExit` request processing class uses POLINE as the root object and removes the PO parent in the transformed JSON. The following text is an example structure:

```
PO
POLINE
POLINE
```

The following text is an example of an updated structure:

```
POLINE
POLINE
```

If you choose to create your own invocation channels, in the **Endpoint** field, ensure AIBROKERAPI is specified. If you edit any end point properties, ensure that the **Override** checkbox is selected. You can

reuse the response and request processing classes. The **Request Template** field contains the query template, or filter, that you specified for the associated object structure.

Enabling AI recommendations in Reliability Strategies

You can use the AI configuration application to enable AI recommendations for asset boundary and failure lists in Reliability Strategies. The AI recommendations use the fmea model template.

Before you begin

If you are using Maximo AI Service on-premises, review the entire process for enabling AI features and complete all prerequisites, including deploying Maximo AI Service. For more information, see [“Maximo AI Service and AI features in Maximo Manage” on page 412](#). If you deploy Maximo AI Service by using the Ansible collection, ensure that as part of the deployment, you specify values for the Reliability Strategy Library variables.

If you are using Maximo AI Service SaaS, ensure that you have specified values for the Maximo Manage system properties. For more information, review your welcome letter.

Procedure

1. In the AI configuration application, click **Add configuration**.
2. In the **Name** field, specify a name
3. Optional: In the **Description** field, specify a description.
4. In the **Template** field, select the fmea template.
5. In the **Template version** field, select the latest template version.
6. Click **Create**.
7. In the AI configurations table, select the configuration that you created.
8. Click **Actions > Activate**.

Activating the AI configuration indicates that the AI configuration is prepared.

9. Click **Actions > Create model**.

What to do next

In Reliability Strategies, select and edit a strategy. Ensure that the **AI** icon is located with the asset boundary details and in the failure list. For the failure list, you can use AI to generate content from the hierarchy view only.

Model training overview

When you train a model, you give the model data to use to teach itself how to generate the output. Not all models require training and inferencing.

Training a model

In the AI configurations application, in an AI configuration, in the **Actions** menu, use the **Train** action if the model has not been trained before or if the model has been trained but you want to apply new parameters or to discard the original model but maintain the model parameters. After a model is created, any change in the metadata, such as arguments, template version, or training and inference parameters, are not reflected in the model unless the AI model is trained again. Training does not require any downtime for the model.

When you initiate training for a model, the following processes occur:

1. The Maximo Manage AI framework uses the training query and the configured object structure or query template to collect the training data from the target object. The framework paginates the data based on the configured page size in the training query template. The JSON data pages are written to the directory that is configured in the **mxe.int.airoot** system property. The default directory is **/airoot**, which is available by default for every Maximo Manage pod. The AI framework compresses

the data with the AI model metadata during only the training process and not during retraining. The compressed file is uploaded to Maximo AI Service.

2. Maximo AI Service takes over the process and works asynchronously to train the model.

You can monitor training in the **Model training log** or **Logs** table or in the **Model status** dialog.

Model training log table or **Logs** table is in the AI configuration and contains step-by-step updates for the training process, but you must refresh the page to see updates. To refresh the page, click **Refresh**.

To access the **Model status** dialog, in an AI configuration, click **Actions > Check model status**.

- When the **State** field value is Ready to serve (running), a value exists in the **Model accuracy score (between 0 - 1)** field, but the **Ready** field value is false, training is done but inferencing is not done.
- When the **State** field value is Ready to serve (running), a value exists in the **Model accuracy score (between 0 - 1)** field, and the **Ready** field value is true, training and inferencing are done.

Retraining a model

Use the **Re-train model** action to train the model on the existing filter that contains new data. Retraining can result in a more accurate model. Retraining requires downtime for the model. During retraining, the AI framework uses only the object data, and uploads the compressed data to Maximo AI Service.

Query template JSON examples

To enable an AI configuration, you must configure invocation channels for training and inferencing. The channels define the JSON structure of the data for training and inferencing.

The following text is an example of JSON for a feature list.

```
{
  "description": "Cannot start pump motor",
  "description_longdescription": "Suspect electrical problems, need to inspect power, wiring,
and the motor itself, including the capacitor, rotor, and bearings",
  ....
}
```

The following text is an example of JSON for a feature list that contains a feature that is a child attribute of a related object.

```
{
  "description": "Cannot start pump motor",
  "description_longdescription": "Suspect electrical problems, need to inspect power, wiring,
and the motor itself, including the capacitor, rotor, and bearings",
  "asset": {
    "description": "Secondary pump used for irrigation in sector 1"
  },
  ....
}
```

The following text is an example of JSON for a feature list that contains a feature that is a child attribute of multiple related objects.

```
{
  "description": "Cannot start pump motor",
  "description_longdescription": "Suspect electrical problems, need to inspect power, wiring,
and the motor itself, including the capacitor, rotor, and bearings",
  "woactivity": [
    {
      "description": "Probe the power supply line to this device"
    },
    {
      "description": "Verify the health of the rotor shaft"
    }
  ],
  ....
}
```


The following text is an example of JSON for training. The JSON contains feature attributes, the target inference attribute, the inference attribute's corresponding description, and the unique ID.

The feature attributes are `description` and `description_longdescription`.

The target inference attribute is `failurecode`.

The inference attribute's corresponding description is `failurecode_description`.

The unique ID is `workorderid`.

```
{
  "description": "Cannot start pump motor",
  "description_longdescription": "Suspect electrical problems, need to inspect power, wiring,
and the motor itself, including the capacitor, rotor, and bearings",
  "workorderid": 1234567,
  "failurecode": "CE",
  "failurecode_description": "Combustion Engine"
}
```

Upgrading Maximo AI Service

You upgrade Maximo AI Service by using a CLI.

Procedure

Complete the upgrade process. The following text is an example of the upgrade process. Provide values during the upgrade process and follow the prompts. `***` denotes a variable.

```
docker run -v /home/***/git/devops-configs:/tmp -ti --rm --pull always quay.io/ibmmas/cli:master
oc login ***

[ibmmas/cli:***]mascli$ mas update
IBM Maximo Application Suite Update Manager (v***)
Powered by https://github.com/ibm-mas/ansible-devops/ and https://tekton.dev/

IBM Maximo Application Suite Admin CLI v***
Powered by https://github.com/ibm-mas/ansible-devops/ and https://tekton.dev/

1) Set Target OpenShift Cluster
Already connected to OCP Cluster:
***

Proceed with this cluster? [y/n] y

2) Review Installed Catalog
The currently installed Maximo Operator Catalog is IBM Maximo Operators (***).
icr.io/cpopen/ibm-maximo-operator-catalog:***

3) Review MAS Instances
No MAS instances were detected on the cluster (***/v1 API is not available)

4) Review AI Service Instances
The following AI Service instances are installed on the target cluster and will be affected by
the catalog update:
- ***

5) Select IBM Maximo Operator Catalog Version
Select MAS Catalog
1) ***
2) ***
3) ***
Select catalog version 1

6) Dependency Update Checks
 IBM Watson Discovery is not installed
 IBM Watson Openscale is not installed
 IBM Certificate-Manager is not installed
 IBM User Data Services is not installed
 Grafana Operator v4 is not installed
 MongoDB CE is already installed at version ***
 1 Db2uClusters (***) in namespace '***' will be updated
 *** is not available in the cluster
 IBM Cloud Pak for Data is not installed
```

```

7) Review Settings
Connected to:
- https://console-openshift-console.***

7.1) IBM Maximo Operator Catalog
  Installed Catalog ..... ***
  Updated Catalog ..... ***

7.2) Supported Dependency Updates
  IBM Db2 ..... All Db2uCluster instances in ***
  MongoDB CE ..... All MongoDBCommunity instances in ***
  Apache Kafka ..... No action required
  IBM Cloud Pak for Data ..... No action required

7.3) Required Migrations
  IBM Certificate-Manager ..... No action required
  IBM User Data Services ..... No action required
  Grafana *** Operator ..... No action required

Please carefully review your choices above, correcting mistakes now is much easier than after
the update has begun
Proceed with these settings? [y/n] y

8) Launch Update
 OpenShift Pipelines Operator is installed and ready to use
 Namespace is ready (***)
 Latest Tekton definitions are installed (***)
 PipelineRun for MAS update submitted

View progress:
https://console-openshift-console.apps.***

[ibmmas/cli:master]mascli$

```

Troubleshooting Maximo AI Service and AI features

You can complete steps to troubleshoot Maximo AI Service and AI configurations in Maximo Manage.

Troubleshoot Maximo AI Service status

To check the status of Maximo AI Service in Maximo Manage, open the AI configuration application and then click **AI Service Health**. If the status is running, Maximo AI Service is ready and available.

In the **AI Service Health** dialog, the following errors or issues can occur. You can view more details for each of the errors in the aibroker-api pod logs. For more information, see [“Accessing logs for Maximo AI Service in Red Hat OpenShift web console”](#) on page 463.

<i>Table 41. Maximo AI Service health errors</i>	
Error	Troubleshooting steps
Failed to connect to server	<p>This error occurs when one or more system properties in Maximo Manage contain an incorrect value. For more information, see Configure the Maximo Manage system properties.</p> <p>The incorrect property might be for the API key, URL, or tenant ID, or it might be the mxe.int.airoot system property. The mxe.int.airoot value is the root URL or directory path for the AI integration within Maximo Manage. Ensure that the value is the location where resources or services are hosted. The default value is /airoot.</p> <p>To edit the mxe.int.airoot property, in Maximo Manage, in the System Properties application, search for and select the property, edit the value, and then save your changes.</p>

Table 41. Maximo AI Service health errors (continued)

Error	Troubleshooting steps
Certificate error	This error occurs if the Maximo AI Service certificate must be imported or was imported incorrectly. For more information, see Create and import the ca.cert file for Maximo AI Service .
BMXAA1482E - The response code received from the HTTP request from the endpoint is not successful Internal Server Error	<p>This error typically occurs if the watsonx.ai API key is not properly set. To determine the root problem of this error, you must check the Maximo AI Service log for the Maximo AI Service project. For more information, see “Accessing logs for Maximo AI Service in Red Hat OpenShift web console” on page 463.</p> <p>If the Provided API key could not be found error appears in the aibroker-api logs, you must update the watsonx.ai API key. To update the key, in Red Hat OpenShift web console, from the side navigation, click Workloads > Secrets and select <code>aiservice-<i>{instance_id}</i>---wx-secret</code>. Update the value for that secret to another API key. For more information about generating a watsonx.ai API key, see watsonx environment variables.</p>

Alternatively, you can check the status of Maximo AI Service by running the following command:

```
curl -X 'GET' \ 'https://{hostname}/ibm/aibroker/service/rest/api/v1/health' \ -H 'accept: */*'
```

If Maximo AI Service is available and running, the following output is returned:

```
{"max_number_of_tenant": "<number>", "kmodel": "running", "healthy": true, "version": "<version>", "status": {"KMODELS": {"healthy": true}, "DB2": {"healthy": true}}}
```

Troubleshooting the AI assistant

If the AI assistant cannot accurately answer questions, alter the prompts for prompt tuning, which can help the assistant better understand your organization's data. For more information about prompt tuning, see [“Prompt tuning for the AI assistant” on page 437](#).

The following table describes errors that can occur in the assistant's chat window.

Table 42. AI assistant errors

Error	Troubleshooting steps
There is a problem	<p>This general error might indicate any of the following issues:</p> <ul style="list-style-type: none"> • Timeout • Connectivity problems • The AI configuration for the assistant is not configured or activated <p>To determine if the issue is connectivity, review the status of Maximo AI Service. For more information, see “Troubleshoot Maximo AI Service status” on page 458.</p>
Invalid datasource query	<p>This error might indicate any of the following issues:</p> <ul style="list-style-type: none"> • watsonx.ai API key is not properly set. • The question or request that the user submitted is not supported by the assistant or the user does not have access to data. • Maximo AI Service cannot be connected to <p>If the question or request is supported by the assistant, in the AI configuration application, click AI Service Health and review the status. If the status is in error, Maximo AI Service connectivity is likely the root of the error. For more information, see “Troubleshoot Maximo AI Service status” on page 458.</p>

Troubleshooting data requirements

The following errors can occur when you check the data requirement for an AI configuration. To check a data requirement, in the AI configuration application, in an AI configuration, click **Actions > Check data requirement**.

Table 43. Data requirement errors

Error	Troubleshooting steps
<p>BMXAA10199E: Insufficient labels in training data. Training data should contain at least two labels.</p>	<p>This error can occur if the training data is not diverse or complete enough. To troubleshoot this error, complete the following steps:</p> <ul style="list-style-type: none"> • Ensure that the records and the value that the model is trying to recommend have sufficient descriptive details. For example, if the model is recommending problem codes, ensure that each problem code has a description, and ensure that the work orders in the training data have complete descriptions and include adequate examples of the value. For more information, see the enablement documentation for the AI configuration that you are enabling. • If you are enabling problem code recommendations, ensure that you have selected which work orders are included in the training filter. For more information, see “Preparing required data for problem code recommendations” on page 448. <p>After any data improvements, you must retrain the model. In the AI configuration application, in your AI configuration, click Actions > Re-train model</p>

Troubleshooting model training and inferencing

Some AI configurations require training and inferencing. You can review the **Model status** dialog to track the status of training and inferencing. To access the dialog, in the AI configurations application, in the AI configuration, click **Actions > Check model status**. For more information about the different statuses, see [“Model training overview”](#) on page 455.

If the **Model status** dialog indicates that training or inferencing failed, before you retry training, complete the following troubleshooting tasks:

- Review the training logs. The logs are located in the AI configuration, in the **Logs** or **Model training log** table.
- Review the status of Maximo AI Service. Training or inference will fail if Maximo AI Service is not available. To review the status of Maximo AI Service, in the AI configurations application, click **AI Service Health**. If the status is **running**, Maximo AI Service availability is not causing training or inferencing to fail.

The following table describes errors that can occur during training or inferencing.

Table 44. Training or inferencing errors

Error	Troubleshooting steps
<p>Training failed. Reason: score <i>score number</i> is below the expected <i>score_threshold threshold</i>. The <i>value</i> which underperformed are {'code'}</p> <p>Resolution: please try with a lower threshold or provide a dataset with better description for the underperforming <i>value</i>.</p>	<p>This error occurs if the score threshold is too high for output to be generated, which can occur when training data is not diverse enough or is incomplete. If the model cannot recommend values with enough accuracy, the training fails.</p> <p>To troubleshoot this error, you can complete the following tasks:</p> <ul style="list-style-type: none"> • Improve your training data. Ensure that the value that the model is trying to recommend has sufficient descriptive details. For example, if the model is recommending problem codes, ensure that each problem code has a description, and ensure that the work orders in the training data have complete descriptions and include adequate examples of the value. For more information, see the enablement documentation for the AI configuration that you are enabling. <p>After any data improvements, you must retrain the model. In the AI configuration application, in your AI configuration, select Actions > Re-train model</p> <ul style="list-style-type: none"> • Lower the score threshold. Setting the threshold to a lower value will decrease the likelihood of desirable output, but it will likely increase the likelihood of generating some output and completing training without failures. To lower the threshold, complete the following steps: <ol style="list-style-type: none"> 1. Note the current score threshold. The score threshold is included in the error message. 2. In the AI configurations application, in your AI configuration, select Actions > Deactivate 3. Select Actions > Set arguments. 4. In the Value field, specify a value that is between 0 and 1 and is lower than your original threshold. 5. Click Save. 6. Select Actions > Activate. 7. Select Actions > Train.
<p>unknown error</p>	<p>To determine why this error occurred, check the logs in the Red Hat OpenShift web console pods for the model ID. The model ID is listed in the AI configuration. For more information, see “Accessing logs for Maximo AI Service in Red Hat OpenShift web console” on page 463.</p>

To check the status of the inferencing servers, complete the following steps:

1. In Red Hat OpenShift web console, click **Administration > CustomResourceDefinition**.
2. Search for `isvc` and then open the inference services resource.

3. On the **Instances** tab, select each instance and then on the **YAML** tab, review the contents of the file. The file contains the status of the inference servers.

If you are troubleshooting training for the similarity model template, you can also review the persistent volume claims for model training to determine if training is in progress.

1. In the Red Hat OpenShift web console, click **Home** > **Projects** and search for the `aiservice-
{instance_id}-user` project name.
2. Open the project and then in the side navigation, click **Storage** > **PersistentVolumeClaims**

When model training is underway, persistent volume claims are listed on the tab. If nothing is listed, no training is actively occurring.

Troubleshooting output availability

If the output for an AI configuration, such as problem code recommendations, is not available after the model is activated, complete the following troubleshooting steps:

- Ensure that the data that you are reviewing is included in the inference filter. If a model requires inferencing, do have the feature enabled, the data must be included in the inference filer.
- Ensure that you have specified the correct features for the features argument, if the argument is used for the model.
- Ensure that you have access to view the feature. Some AI features require specific permissions. For more information, review the documentation for enabling the AI feature. The following list contains links to the documentation:
 - [“Enabling recommended problem codes for Work orders” on page 446](#)
 - [Enabling field value recommendations](#)
 - [Enabling the assistant](#)
 - [“Enabling locating of similar work orders” on page 449](#)
 - [“Enabling AI recommendations in Reliability Strategies” on page 455](#)
-

Accessing logs for Maximo AI Service in Red Hat OpenShift web console

As a general troubleshooting task, you can access logs for various Maximo AI Service operators and pods in Red Hat OpenShift web console.

About this task

As a general troubleshooting task, you can also download and review the object storage log files. For more information, review the documentation for your storage provider.

Procedure

1. Review the logs for the Maximo AI Service project, which includes the logs for the `km-controller`, `km-store`, `km-watcher`, and `aibroker-api` pods.
 - a) In Red Hat OpenShift web console, click **Home** > **Projects** and search for the `aiservice-
{instance_id}` project name.
 - b) On the **Logs** tab, review the log for errors.
 - c) Review the logs for the `km-controller`, `km-store`, `km-watcher`, and `aibroker-api` pods.
 - i) For the `aiservice-
{instance_id}` project, on the **Details** tab, in the **Inventory** section, select **number Pods**.
 - ii) Review the status of the pods and ensure that they are all listed as running. For the `km-controller`, `km-store`, `km-watcher`, and `aibroker-api` pods, open the pods and then on the **Logs** tab, review the log for errors.

2. Review the training and inference pods.
 - a) Return to the projects table by clicking **Home > Projects** and then search for and select the **aiservice-*{instance_id}*-user** project, which is the tenant project. Trainings and inferencing runs within the tenant project.
 - b) In the **Inventory** section, select **number Pods**.
 - c) For each pipeline or predictor pod, open the pod and on the **Logs** tab, review the logs for errors.
If a pod name contains *pipeline*, the pod is used for training. If the pod name contains *predictor*, the pod is used for inferencing.

Customer-managed Accelerators

Accelerators are solutions that are provided by IBM and partners of IBM that extend Maximo Application Suite capabilities. These accelerators are hosted on the [Maximo Application Suite accelerators solutions page](#).

Related concepts

[Maximo Application Suite accelerators](#)

Accelerators are solutions that complement or extend Maximo Application Suite capabilities or experience or accelerate time to value.

Activating and deactivating accelerators

After the deployment process is completed, the accelerator is not immediately available in your environment. Before you can grant users access and start working with the accelerator, you must activate the accelerator.

Entitled accelerators are not immediately available. You have to add an accelerator into your environment from the Suite administration page by selecting it from the Catalog and then clicking **Add to Suite**.

After you have deployed an accelerator, you can track the activation status.

Activating an accelerator from the details page

The accelerators details page displays the status and information about the accelerator. You can activate or deactivate an accelerator from its details page.

Procedure

1. Log in to Maximo Application Suite as an administrator.
2. From the **Suite administration** menu, click **Suite**.
3. Select the **Accelerators** tab.

The Accelerators page contains the list of entitled Maximo Application Suite accelerators purchased.

4. Select an accelerator from the catalog.
5. From the accelerator details page, select **Activate** from the actions menu.
You can alternatively deactivate the accelerator if it is currently active.
6. Confirm that you want to activate the accelerator.

Activating an accelerator from the accelerators list page

The accelerators list page displays accelerators that you are entitled to activate in your environment. You can activate or deactivate an accelerator from the accelerators list page.

Procedure

1. Log in to Maximo Application Suite as an administrator.
2. On the **Suite administration** page, from the side navigation menu, click **Suite**.

3. Select the **Accelerators** tab.

The Accelerators page contains the list of entitled Maximo Application Suite accelerators purchased on Red Hat Marketplace.

4. From the ellipses menu of an accelerator listed, select **Activate**.

You can alternatively deactivate the accelerator if it is currently active.

5. Confirm that you want to activate the accelerator.

Customer-managed **Activating applications**

After the deployment process is completed, the application is not immediately available in your environment. Before you can grant users access and start working with the application, you must activate the application.

Important: Before you activate the application, complete any required post-deployment steps and ensure that any prerequisite application is activated or in the activation process. If you do not complete these steps before you activate the application, it might result in unexpected behavior that might require you to deactivate and redeploy the application.

Customer-managed **Activating IBM Maximo Collaborate**

After the deployment is complete, activate Maximo Collaborate to make it available in the Suite. Then, grant users access to it.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Before you begin

For more information about deploying Maximo Collaborate, see [“Deploying IBM Maximo Collaborate” on page 292](#).

About this task

Note:

Starting in Maximo Application Suite 9.0, Watson Discovery, which is used to support the query and diagnose functions, is no longer available as a dependency in Maximo Assist. If Maximo Assist is already deployed and activated with Watson Discovery, and you are upgrading to Maximo Application Suite 9.0, before you can complete the upgrade, you must contact IBM Support to help with the manual removal of Watson Discovery.

Note:

Starting in Maximo Application Suite 9.0, voice inspections are no longer available in Maximo Assist. If voice inspections are enabled and you are upgrading to Maximo Application Suite 9.0, this feature is automatically removed during the upgrade.

Procedure

1. On the **Collaborate** application page, click **Activate**.

If the application deployment is not complete, the **Activate** button is not enabled.

2. Click **Show advanced settings**.

3. Configure Watson Discovery by entering the following information for your Watson Discovery instance. Save your configurations.

Note: This step is applicable for Maximo Assist deployed in Maximo Application Suite 8.11 and earlier.

- Enter the URL of the Watson Discovery API.

Note: To access the URL, log in to the IBM Cloud Pak for Data portal and launch the Watson Discovery. The URL is located in the **Access Information Page**.

- Enter the Watson Discovery username.
 - Enter the Watson Discovery password.
4. Optional: To enable the voice inspection feature in Maximo Application Suite 8.11 or earlier so that inspectors can complete hands-free inspections on a mobile device, configure the **Watson Text To Speech, Watson Speech To Text, and Watson Assistant** settings. Contact your IBM representative for the IBM Cloud Watson service instance and API key information.

Note: This step is applicable for Maximo Assist deployed in Maximo Application Suite 8.11 and earlier.

- a) For each setting, set **System managed** to on and add your custom settings.
- b) Enter your Watson service instance in the **Endpoint** field.
- c) Specify the API key for Watson services.

Option	Description
If you are configuring Watson services for the first time during the activation.	Enter the API key for your Watson service.
If you are updating an existing configuration for Watson services or configuring Watson services after Maximo Assist is activated.	Click Replace secret .

If you do not edit the three configuration settings, the voice inspection feature is not enabled by default.

5. Click **Activate**.

On the confirmation page, click **Activate**. The estimated deployment time is an estimate of the time that it takes to activate the application. You can track the activation process on the Collaborate application page.

Activated applications are available from the Suite navigator and at specific URLs. For more information, see [Maximo Collaborate](#).

Note: It can take a few hours to complete the activation. If the application was recently activated, it might appear in the Suite navigator but might not be ready for use.

6. Validate the activation.

- You can validate the activation status on the Collaborate application page: `https://admin.{{mas_domain}}/applications/collaborate`
- You can also check the Collaborate workspace status by using the Red Hat OpenShift client command prompt.

```
oc login <OCP_cluster>
oc get collaborateworkspace -n {{Collaborate_project}}
```

```
NAME VERSION STATUS AGE masdev-main 9.1.0 Ready 73d
```

What to do next

After you deploy and activate Maximo Collaborate, complete the post-activation tasks.

As an application administrator, you can now give new Maximo Collaborate users [access to the application](#). You can also configure automatic refresh for work orders. For more information, see [IBM Support](#).

Related concepts

[IBM Maximo Collaborate](#)

Related tasks

[Deploying IBM Maximo Collaborate](#)

[Updating Maximo Collaborate](#)

Related information

[Getting started](#)

[Maximo Assist](#)

Configuring the push notification service

After you deploy and activate Maximo Collaborate, configure the push notification service in Maximo Application Suite so that the Collaborate feature works with your mobile service.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Procedure

1. Generate an API key.
 - a) Open the following URL:
<https://pushnotification.suite.maximo.com/pushnotification/auth>
 - b) Click **Sign in or up**.
 - c) Log in by using your IBMid or create an IBMid .
 - d) After you log in, enter a name and click **Generate API key**.
 - e) Save the API key for configuration.
2. As a system administrator in Maximo Application Suite, from the side navigation menu, click **Configurations** and then click **Push Notification**
3. Click **Edit**.
4. In the **Bundle identifier** field, specify `com.ibm.iot.maximo.mobile`.
5. Optional: Set **Configure with the nonproduction account** to **Production**.
6. In the **URL** field, specify the following:
<https://pushnotification.suite.maximo.com/pushnotification/sendMessage>
7. In the **API key** field, specify the API key that you generated.
8. Save the changes.

Customer-managed **Activating IBM Maximo Monitor**

Before you can grant users access and start working with the application, you must activate Maximo Monitor. You can activate the application after the deployment is complete.

Procedure

To activate Maximo Monitor, on the Monitor application page, click **Activate**.

It can take up to 15 minutes to activate Maximo Monitor. If the **Activate** button is disabled, the deployment process is not complete.

Access activated applications from the Maximo Application Suite Suite navigator and specific URLs for Maximo Monitor and for the IoT tool. For more information, see [“Maximo Application Suite application URLs” on page 77](#).

What to do next

As an application administrator, you can continue with the following tasks:

- Grant your users access to the application. Activating the application makes it available but does not automatically grant your users access to the application. Administrators with user management administration can grant users access. For more information, see [grant users access to the application](#).
- For more information about other getting started, see [Getting started with Maximo Monitor](#).
- Your data is stored in IBM Db2 Warehouse. For more information about how to encrypt your data in Maximo Monitor, see [Encryption of data at rest and Encryption of data in transit](#). For more information about how to back up and restore your data in Maximo Monitor, see [Backing up data](#).

Related tasks

[Deploying IBM Maximo Monitor](#)

By using Maximo Monitor, you can visualize current and historical trend data for your devices and assets on customizable dashboards.

Customer-managed

Activating IBM Maximo Predict

Before you can grant users access and start working with Maximo Predict, you must activate the application. You can activate Maximo Predict after the deployment is complete.

Before you begin

Before you activate Maximo Predict, ensure that the Maximo Health and Maximo Monitor applications are deployed. For more information about deploying Maximo Health, see [Deploy Maximo Health](#). For more information about deploying Maximo Monitor, see [Deploy Maximo Monitor](#).

Watson Studio must also be installed before you deploy Maximo Predict, but it can be configured before or during activation of Maximo Predict. The Watson Studio URL, username, and password are the IBM Cloud Pak for Data URL, username, and password.

If you are using the explainability service, configure Watson Studio. For more information, see [explainability service](#)

Procedure

1. If **System managed** is set to off for the explainability service, configure the settings for the log, such as the level and the batch scoring limit. Configure the batch scoring limit in the **WML batch scoring limit** field.
2. Click **Activate**.
3. On the confirmation page, click **Activate**. Activation is complete when the **Workspace** card displays the Ready message and a green checkmark.

What to do next

Review [Getting started as an application administrator](#).

Related concepts

[IBM Maximo Predict](#)

IBM Maximo Predict is an application in Maximo Application Suite. By using Maximo Predict, you can leverage your historical and near real-time asset performance data, maintenance records, inspection reports, and environmental data to correlate performance factors that predict asset degradation or failure. Maximo Predict also uses artificial intelligence to optimize predictive model accuracy.

Related tasks

[Deploying IBM Maximo Predict](#)

Maximo Predict can use historical and recent asset performance data to correlate performance factors that predict asset degradation or failure. Other types of data that can be correlated include maintenance records, inspection reports, and environmental data. Maximo Predict uses artificial intelligence to optimize predictive model accuracy.

[Updating Maximo Predict](#)

Related information

[Getting started with Maximo Predict](#)

Customer-managed

Activating IBM Maximo Visual Inspection

Before you can grant users access and start working with Maximo Visual Inspection, you must activate the application. You can activate Maximo Visual Inspection after the deployment is complete.

Procedure

On the **Deploy** Maximo Visual Inspection application page, click **Activate**. When the application deployment is complete, the **Activate** button is enabled.

Activated applications are available from the suite catalog and at specific URLs. For more information, see [“Maximo Application Suite application URLs” on page 77](#).

Note: Activating an application does not automatically grant your users access to the application. You need to grant users access.

What to do next

After you deploy and activate Maximo Visual Inspection, you can complete the following tasks:

- [Grant users access to Maximo Visual Inspection](#).
- [Train and work with models](#).
- [Optional: Configure the IBM Maximo Visual Inspection Edge add-on](#).

For more information about getting started in Maximo Visual Inspection, see [Getting Started steps](#).

Customer-managed

Activating IBM Maximo Optimizer

Before Maximo Optimizer is available for use, you must activate Maximo Optimizer. Activating the application does not automatically grant your users access to the application.

About this task

When the application deployment is complete, the **Activate** button is enabled.

Procedure

1. On the Maximo Optimizer details page, click **Activate**.
2. To view or specify the **Execution service** configuration settings, click **Show advanced settings**. The system defines these settings by default. If your environment requires different values, you can set **System managed** to off to enable the fields and add your custom settings.
3. Click **Activate**. The estimated time is an estimate of the time that it takes to activate the application. You can track the activation process on the Optimizer details page.
4. Optional: Disable the configuration option **Allow model customization artifacts** for optimization models, which is enabled by default.
 - a) Open **Suite Administration**.
 - b) Select **Workspaces**.
 - c) In the **Applications** tab, select **Optimizer**.
 - d) Click the **Actions** button and from the Actions menu select **Update configuration**.
 - e) Within the **Configurations** subsection, select **Model customization artifacts**.
 - f) Disable the follow 2 toggle buttons.
 - **System managed**

- **Allow model customization artifacts**

What to do next

Grant your users access to the application. For more information about granting user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier”](#) on page 796. Starting in Maximo Application Suite 9.1, you grant your users access to the application by assigning them to security groups. For more information, see [“Creating users in Maximo Application Suite 9.1”](#) on page 784.

After activation, Maximo Optimizer is available from the Suite navigator and at the following URL:
`https://<workspace_id>.optimizer.<mas_domain>/.`

Related concepts

[IBM Maximo Optimizer](#)

IBM Maximo Optimizer is an add-on to Maximo Application Suite. By using Maximo Optimizer, you can automate efficient decisions for long-range planning, scheduling, and dispatching of resources for asset maintenance while balancing competing objectives and constraints.

Related tasks

[Deploying IBM Maximo Optimizer](#)

By using Maximo Optimizer, you can automate efficient decisions for long-range plans, schedules, and the dispatch of resources for asset maintenance, helping to balance competing objectives and constraints.

Customer-managed

Activating Maximo Health and Predict - Utilities

Note:

Starting in Maximo Application Suite 8.11, Maximo Health and Predict - Utilities is no longer available as a separate industry solution. The information that is provided is applicable only to Maximo Application Suite 8.10 and earlier versions. For more information, see [“Upgrading IBM Maximo Application Suite”](#) on page 473. Before you upgrade to Maximo Application Suite 8.11, deactivate and delete Maximo Health and Predict - Utilities.

After deployment is complete, the **Activate** button is enabled. The industry solution might appear in the Suite navigator but might not be ready to use immediately. Activating the industry solution does not automatically grant your users access to the industry solution. As an administrator, you must grant users access to it.

Before you begin

Important:

Maximo Application Suite supports Cloud Pak for Data 4.6 with Python 3.10.

Note:

1. Specify the root URL for Watson Studio.
2. Specify the Watson Studio username.
3. Specify the Watson Studio password.
4. Click **Show advanced settings** to configure the Watson Studio project ID. To retrieve the project ID, open your project and then copy the project ID from the page URL.
5. Maximo Health and Predict - Utilities uses the most recent version of Python to deploy notebooks. If you need to specify the Python version manually, complete the following steps.
 - a. In the Red Hat OpenShift console, open the `mas-<instanceID>-hputilities` project.
 - b. From the side navigation menu, select **Workloads > ConfigMaps** and then click **Create ConfigMap**.
 - c. Replace the Python version, for example, type `Python 3.10, apiVersion v1, kind ConfigMap, metadata, name hputilities-watsonstudio-project-envname, data, PYTHON_ENV_NAME Python version`.

6. Save your changes.

Procedure

1. On the Maximo Health and Predict - Utilities page, click **Activate**.
2. On the Health and Predict - Utilities details page, configure your Watson Studio instance or use the system-managed options. If **System managed** is set to on, the project ID is `<instanceID>-<workspaceID>-hutilities`
3. Click **Activate**.
4. On the confirmation page, click **Activate**. The estimated deployment time is an estimate of the time that it takes to activate the industry solution. You can track the activation process on the Maximo Health and Predict - Utilities details page or review the logs for the `instance name-hutilities-entitymgr-ws-string pod`.

Login to the Red Hat OpenShift cluster and then check the Maximo Health and Predict - Utilities deployment status by running the Red Hat OpenShift client command line.

```
oc get hutilitiesworkspace -n mas-devtest-hutilities
```

The following output shows an example of the status:

NAME	VERSION	STATUS	AGE
devtest-masdev	8.5.0	Ready	2d2h

After the industry solution is activated, it is available in the Suite navigator and at specific URLs.

What to do next

For next steps, review [Getting started as an application administrator](#).

Related tasks

[Deploying IBM Maximo Health and Predict - Utilities](#)

Maximo Health and Predict - Utilities supports maintenance, operations, and performance of assets and networks for energy and utility companies.

Completing post-activation steps for IBM Maximo Health and Predict - Utilities

After activation is complete for Maximo Health and Predict - Utilities, configure the API key and verify that the components for Maximo Health and Predict - Utilities are configured.

Before you begin

If Cloud Pak for Data is using Security Assertion Markup Language (SAML), you must retrieve the API key before you configure it. Complete the following steps to retrieve the API key:

1. Log in to Watson Studio as an administrator.
2. Click your user icon and then click **Profile and settings**.
3. Click the API key.
4. Copy the API key value.

Procedure

1. If Cloud Pak for Data is using Security Assertion Markup Language (SAML), you must also configure an API key for Watson Studio. To configure the API, complete the following steps:
 - a) As an application administrator, in Maximo Health and Predict - Utilities, click **Application administration**.
 - b) In the System Properties application, open the `hutilities.watsonstudio.apikey` and specify the Watson Studio project administrator's API key.
 - c) Save the property.

- d) Select the checkbox for the property and then click **Live Refresh** and complete the live refresh.
2. Add the API key for the administrative user and configure the EXTENGINEAPI endpoint. Complete the following steps to create and configure the API key:
 - a) In Maximo Health and Predict - Utilities, from the navigation menu, select **Application administration** to take you to Maximo Manage.
 - b) From the side navigation menu, click **Integration > API Keys**.
 - c) Click **Add API Key**.
 - d) In the **User** field, select the user ID for the administrative user, ensure that the value in the **Token expiry in minutes** field is -1. Click **Create**.
 - e) In the **API key card** click **Copy key**.
 - f) From the Start Center, open the End Points application.
 - g) Open the EXTENGINEAPI endpoint and in the Properties for End-Point table, for the PASSWORD property, in the **Encrypted Value** field, specify the administrative user's API key.
3. Verify that the components for Maximo Health and Predict - Utilities are configured.
 - a) In Maximo Health and Predict - Utilities, in the End Points application, for the EXTENGINEAPI endpoint, verify that the property value is the model engine URL.
 - b) Verify that your Watson Studio project contains data assets, notebooks, and jobs.
 - c) In Maximo Health and Predict - Utilities, in the System Properties application, verify that a value is specified for each of the following properties:
 - hutilities.health.endpoint
 - hutilities.watsonstudio.endpoint
 - hutilities.watsonstudio.username
 - hutilities.watsonstudio.password
 - hutilities.watsonstudio.projectid
 - d) If you are going to use the data loader, for App Connect, verify that the data loader bar and configuration are uploaded.

To verify that the data loader bar is uploaded, on the App Connect dashboard, click **Bar Files** and ensure that a version `ibm-mas-hutilities-<new_version>-string-data-loader` bar exists. To verify that the data loader configuration is uploaded, on the App Connect dashboard, click **Configurations**. Ensure that the following configurations are listed:

 - ibm-hutilities-string-healthurl
 - ibm-hutilities-string-keystore.p12
 - ibm-hutilities-string-sts-serverconf
 - ibm-hutilities-string-https
 - ibm-hutilities-string-users
 - ibm-hutilities-string-store-password
 - ibm-hutilities-string-sts-trustcert

What to do next

For more information about next steps, see [Getting started as an application administrator](#).

Customer-managed

Deactivating and deleting applications

By deleting an application, you remove access to it from the environment. To reinstate it, you can redeploy the application. Application data and metadata for the deleted application is not deleted but is retained in the corresponding repositories, such as MongoDB and Db2.

About this task

Important: When you deactivate or delete an application, access is revoked for all users. If you later redeploy the application, you must grant the users access again. An application must be deactivated before it can be deleted.

Important: If you need to keep data from your persistent volumes, take backups before you delete or deactivate an application. When an application is deleted or deactivated, the respective CRs and all resources that belong to the operator are deleted.

Procedure

1. On the **Suite administration** page, from the side navigation menu, click **Applications**.

2. Select the application that you want to deactivate or delete.

3. On the application details page, select **Actions > Deactivate**

4. Verify that you want to deactivate the application and then click **Deactivate**.

The application is no longer available in the Suite navigator. If an AppPoint cost is associated with the application, it is returned to your license AppPoint pool.

Important: If your intent is only to deactivate the application, stop here.

5. On the application overview page, select **Actions > Delete application**.

6. Verify that you want to delete the application and then click **Delete**.

The application is no longer available in **Applications** in the catalog. To reinstate it, you must redeploy the application.

Related tasks

Upgrading from Maximo Scheduler Optimization to Maximo Optimizer

In Maximo Application Suite 8.8, you use Maximo Optimizer instead of Maximo Scheduler Optimization. If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade.

Upgrading

You can upgrade Maximo Application Suite to keep up to date with new versions and updates to applications. You can also upgrade from Maximo Asset Management to Maximo Manage in Maximo Application Suite.

Customer-managed

Upgrading IBM Maximo Application Suite

Note: The Maximo Application Suite upgrade policy supports n-1 versions within a cluster. This means that you can upgrade directly from the version immediately preceding the current one. For example, if the current version of Maximo Application Suite is 9.1, you can upgrade directly from 9.0 in a cluster.

You can upgrade in IBM Maximo Application Suite by subscribing to a channel. Before you upgrade, you must complete several prerequisite tasks to ensure that your Maximo Application Suite environment meets the upgrade requirements.

For information about upgrading application and add-ons, see [“Application and add-ons upgrade prerequisites”](#) on page 475

To upgrade Maximo Application Suite and its applications, you can use the methods in the following table. The upgrade method that is used is determined by your chosen installation method.

Upgrade method	Maximo Application Suite upgrades	Application upgrades
Channel subscription	<ul style="list-style-type: none"> Your Maximo Application Suite instance is automatically kept up to date with the latest version that is available in the Maximo Application Suite operator's subscription channel. Optionally, you can reconfigure the subscription to ensure that upgrades require manual approval before they begin. This upgrade method is used if you installed Maximo Application Suite from the IBM Operator catalog in your Red Hat OpenShift cluster. 	<ul style="list-style-type: none"> Your application is automatically kept up to date with the latest version that is available in the application operator's subscription channel. You are notified when an update is available. If you turned on the Automatic approval switch when you deployed the application, it is automatically updated to the most current version in the channel. If you turned off this switch, you must manually approve the update.
<p>Manual</p> <p>Manual upgrade is applicable only to Maximo Application Suite 8.9 or earlier.</p>	<ul style="list-style-type: none"> You manually keep Maximo Application Suite up to date with the latest version. You are notified when a new version is available. You are responsible for ensuring compatibility between versions. You use this upgrade method if you installed Maximo Application Suite by downloading the software from IBM Passport Advantage and running the installer script. 	<ul style="list-style-type: none"> You manually update your applications to new versions. You are notified when new application versions are available. You are responsible for ensuring application version compatibility.

Before you begin

For information about upgrading application and add-ons, see [“Application and add-ons upgrade prerequisites”](#) on page 475



Attention: IBM App Connect and Cloud Pak for Data do not support odd-numbered Red Hat OpenShift Container Platform versions. If Maximo Collaborate, Maximo Health and Predict - Utilities, or Maximo Predict is deployed, you must use even-numbered Red Hat OpenShift Container Platform versions.

Procedure

- Use the upgrade process that applies to your upgrade method.
 - For channel subscription upgrades, use [this process](#).
 - In Maximo Application Suite 8.9 and earlier, for manual upgrades, use [this process](#).
- Update applications

When the Maximo Application Suite upgrade is complete, updated applications are made available. For each application, an update is automatically deployed only if you chose the **Subscription** upgrade method and the **Automatic approval** option when you deployed the application.

To update your remaining deployed Maximo Application Suite applications, use the [application update process](#). You must complete any pre-update, update, and post-update steps before you start to use the updated applications.

What to do next

If you encounter issues during or after the upgrade, use the troubleshooting information in this documentation to help you resolve the issues. For more information, see [Troubleshooting upgrade issues](#).

Related tasks

[Upgrading from Maximo Scheduler Optimization to Maximo Optimizer](#)

In Maximo Application Suite 8.8, you use Maximo Optimizer instead of Maximo Scheduler Optimization. If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade.

Application and add-ons upgrade prerequisites

Before you upgrade, you must consider whether the IBM Maximo Application Suite applications and add-ons are available for upgrade. If the applications or add-ons are no longer available, you must deactivate and delete those applications and add-ons.

Upgrading to 9.1

If you are upgrading to Maximo Application Suite 9.1, review the following prerequisites:

Stand-alone Maximo Health no longer a suite application

Maximo Health is no longer available as a stand-alone suite application but remains an add-on in Maximo Manage.

If Maximo Health is deployed as a stand-alone suite application and you are upgrading to 9.1, you must add Maximo Health as an add-on in Maximo Manage 9.0. Then, you upgrade to 9.1. For more information, see [Upgrading Health stand-alone 9.0 to Manage with Health 9.1](#)

IBM Maximo Assist renamed IBM Maximo Collaborate

Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate. If Maximo Assist is deployed in Maximo Application Suite 9.0 or earlier and you upgrade to Maximo Application Suite 9.1, the name is automatically changed in the user interface.

Note: In the Maximo Application Suite documentation, Maximo Assist is now referred to as Maximo Collaborate.

Upgrading to 9.0

If you are upgrading to Maximo Application Suite 9.0, review the following prerequisites:

MRO Inventory Optimization

Starting in Maximo Application Suite 9.0, MRO Inventory Optimization is no longer available to be added as an externally configured application. If MRO Inventory Optimization is configured as an external launcher and you are upgrading to Maximo Application Suite 9.0, you must remove MRO Inventory Optimization before you can complete the upgrade.

To remove MRO Inventory Optimization as an external launcher, you delete the solution portal URL for MRO Inventory Optimization on the **External launcher** page. When the product URL is removed, users can no longer access from the suite navigator. For more information, see [“Configuring external launchers”](#) on page 644.

Support for MongoDB 5.0 and 6.0

MongoDB 5.0 and 6.0 are supported in Maximo Application Suite.

If you are upgrading from an earlier version of MongoDB to MongoDB 5.0 or 6.0, consider the following requirements:

- Before you upgrade MongoDB, create a backup.
- If you are using a cloud-hosted instance of MongoDB, such as on IBM Cloud, refer to the documentation for that cloud host for more information about how to support and upgrade MongoDB.
- If MongoDB is hosted in another way, such as on premises, refer to the upgrade documentation for MongoDB for that host.
- If you are using the MongoDB community operator with automation that is provided by Maximo Application Suite, refer to the [Maximo Application Suite Ansible documentation](#).

Upgrading to 8.11

If you are upgrading to Maximo Application Suite 8.11, review the following prerequisites:

IBM Parts Identifier

Starting in 8.11, the IBM Parts Identifier add-on is no longer available. If Parts Identifier is installed, and you are upgrading to 8.11, you must deactivate and delete Parts Identifier before you can complete the upgrade.

IBM Maximo Health and Predict - Utilities

Starting in 8.11, the Maximo Health and Predict - Utilities add-on is no longer available. If Maximo Health and Predict - Utilities is installed, and you are upgrading to 8.11, you must deactivate and delete Maximo Health and Predict - Utilities before you can complete the upgrade.

If you want to use your existing health models or data loader of Maximo Health and Predict - Utilities, in the Integrations and dependencies section, configure the Watson Studio project ID and the dashboard URL for App Connect that was used for Maximo Health and Predict - Utilities. For more information, see [Deploying Maximo Health](#).

Note: The Maximo Models for Electrical Distribution accelerator replaces Maximo Health and Predict - Utilities in Maximo Application Suite 8.11.

Upgrading to 8.10

If you are upgrading to Maximo Application Suite 8.10, review the following prerequisites:

- Starting in Maximo Application Suite 8.10, the manual deployment for applications and by using installation script for Maximo Application Suite are discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method, and subscribe to the latest channel.

For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription” on page 482](#).

- The installation method that you use determines how Maximo Application Suite is upgraded. If you install the Maximo Application Suite manually by downloading and running an installer script, you upgrade the Maximo Application Suite the same way. Depending on your chosen approval strategy, the upgrade begins automatically or after you approve it.

Upgrading to 8.9

If you are upgrading to Maximo Application Suite 8.9, review the following prerequisites:

Service Binding Operator

Uninstall Service Binding Operator (SBO) because it is not a dependency from Maximo Application Suite 8.9. Therefore, before uninstalling SBO from the cluster, ensure your Maximo Application Suite instances (core and all applications) are upgraded to the latest Maximo Application Suite available fixes.

To uninstall, in Red Hat OpenShift operators, go to **Operators > Installed operators**, and search for Service Binding Operator 1.0.1 and click **Uninstall Operator**.

Cloud Pak for Data

Optionally upgrade to IBM Cloud Pak for Data 4.5.x. For more information, see [Installing IBM Cloud Pak for Data](#).

Red Hat OpenShift

Optionally upgrade your Red Hat OpenShift cluster to 4.10. For more information, see [Updating clusters, worker nodes, and cluster components](#).

Maximo Safety

In Maximo Application Suite 8.9, Maximo Safety is not longer available. If Maximo Safety is deployed and active in your environment and you are upgrading to Maximo Application Suite 8.9, you must deactivate and delete Maximo Safety before you can complete the upgrade.

IBM Maximo Optimizer

If Maximo Optimizer is deployed in your environment and you are upgrading to Maximo Application Suite 8.9, Maximo Optimizer must be version 8.2.2, or higher before you can complete the upgrade.

Maximo Health and Predict - Utilities

If Maximo Health and Predict - Utilities is installed, optionally upgrade the IBM App Connect dashboard. For more information, see [“Updating Maximo Health and Predict - Utilities” on page 491](#).

Upgrading to 8.8

If you are upgrading to Maximo Application Suite 8.8, review the following prerequisites:

Maximo Scheduler Optimization

If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade. For more information, see [“Upgrading from Maximo Scheduler Optimization to Maximo Optimizer” on page 493](#).

Customer-managed

Upgrading IBM Maximo Application Suite by using the channel subscription method

Note: The Maximo Application Suite upgrade policy supports n-1 versions within a cluster. This means that you can upgrade directly from the version immediately preceding the current one. For example, if the current version of Maximo Application Suite is 9.1, you can upgrade directly from 9.0 in a cluster.

Channel subscription upgrades of IBM Maximo Application Suite involve setting the subscription channel for your installed Maximo Application Suite operator, ensuring that the upgrade is approved, and updating the applications.

Before you begin

Before you can upgrade the Maximo Application Suite, you must complete the [prerequisite steps](#), such as deleting incompatible applications, installing the IBM Certificate Manager service, and upgrading your Cloud Pak for Data System.

Note: If IBM Maximo Manage is installed and the Anywhere component is enabled as part of the Maximo Manage 8.4 installation, and when you upgrade to Maximo Manage 8.5. Anywhere is automatically removed as a component.



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription” on page 482](#).

About this task

To upgrade your Suite instance, you set the subscription channel for your installed Maximo Application Suite operator in Red Hat OpenShift and ensure that the upgrade is approved.

Procedure

- In the Red Hat OpenShift web console, in your installed Suite operator, set the subscription channel, such as to either 8.x.x or 8.x.

For example, set the subscription channel to either **8.9.x** or **8.x**.

For more information, see [Upgrading installed Operators](#) in the Red Hat OpenShift documentation.

The upgrade will begin after it is approved. The approval strategy that is configured in your Maximo Application Suite operator's subscription determines how upgrades are approved.

If the strategy is set to **Manual**, you can manually approve the upgrade in the **Subscription** tab. If the strategy is set to **Automatic**, the upgrade is approved and completes automatically.

If the automatic approval strategy is set to **On**, when new application updates are available, they are added to the channel and updated in your Maximo Application Suite instance automatically. If the strategy is set to **Off**, when new application updates are available, you receive a notification, and you can manually approve the updates.

What to do next

When the upgrade is complete, you can log in to the Maximo Application Suite **Administration** page and continue with the upgrade process by [updating applications](#).

Related tasks

[Setting version lock in Maximo Application Suite](#)

Upgrading Maximo Application Suite by using static catalog

Note: The Maximo Application Suite upgrade policy supports n-1 versions within a cluster. This means that you can upgrade directly from the version immediately preceding the current one. For example, if the current version of Maximo Application Suite is 9.1, you can upgrade directly from 9.0 in a cluster.

You can upgrade your Maximo Application Suite by using static catalogs. If you are using the dynamic catalog, then updates are automatically applied as soon as they are released.

Upgrade is the act of switching a Maximo Application Suite installation to a new subscription channel. Upgrade is distinct from an update in which new versions are available on existing subscription channels. New features are available from new subscription channels, while updates to existing functions (including security updates) and bug fixes are available within existing subscription channels.

Before you begin

You can upgrade to a Maximo Application Suite version already supported by the `ibm-operator-catalog` CatalogSource currently installed in the cluster.

- If you are using the static catalog and not updated to a catalog that includes the Maximo Application Suite version you want to upgrade to, then you must first update your Maximo Application Suite.
- If you are using the dynamic catalog, then no action is required. The catalog source automatically updates to the latest Maximo Application Suite version.

Procedure

If you are using a static catalog, first update your catalog that includes the Maximo Application Suite version. You can then upgrade your Maximo Application Suite.

Note: Mirroring the images is a simple but time consuming process.

1. Optional: Prepare for a disconnected update by mirroring the container images.

Use the following three modes:

- **direct** - mirrors images directly from the source registry to your private registry
- **to-filesystem** - mirrors images from the source to a local directory
- **from-filesystem** - mirrors images from a local directory to your private registry.

Run the **mirror-images** command to mirror images.

```
docker run -ti --rm --pull always quay.io/ibmmas/cli mas mirror-images
```

Note: Ensure you select the Maximo Application Suite version that is installed in the cluster. If multiple Maximo Application Suite instances are running in the cluster for different versions, then you must run the command multiple times to ensure that you mirror the content for all Maximo Application Suite versions that are used in the cluster.

When prompted, set the target registry to mirror the image, select the Maximo Application Suite Operator catalog to mirror, and select the subset of content that you want to mirror. You can choose to mirror everything from the catalog, or control exactly what is mirrored to your private registry to reduce the time and bandwidth that is used to mirror the images, and reduce the storage requirements of the registry.

You can run the **mirror-images** command in noninteractive mode.

```
mas mirror-images \  
-m direct \  
-d /mnt/local-mirror/ \  
-H myprivateregistry.com -P 5000 -u $REGISTRY_USERNAME -p $REGISTRY_PASSWORD \  
-c v8-221129-amd64 -C 8.9.x --mirror-core --mirror-iot --mirror-optimizer --mirror-manage \  
--ibm-entitlement $IBM_ENTITLEMENT_KEY \  
--redhat-username $REDHAT_USERNAME --redhat-password $REDHAT_PASSWORD \  
--no-confirm
```

For more information, see [Mirror images](#).

2. Update Maximo Application Suite.

Run the **mas update** command and choose the catalog for update. This command updates the operator catalog that is installed in your cluster, and the Operator Lifecycle Manager (OLM) automatically updates all installed operators to the newest version available on the current subscription channel.

```
docker run -ti --rm --pull always quay.io/ibmmas/cli mas update
```

Important: You must select a newer catalog than what is already in use. Updating to an older static catalog is not supported.

You can run the **update** command in noninteractive mode.

```
mas update -c v8-221129-amd64 --no-confirm
```

3. Upgrade Maximo Application Suite.

Run **mas upgrade** and choose the Maximo Application Suite instance to upgrade. The upgrade automatically detects the installed release, and upgrades to the next available release.

```
docker pull quay.io/ibmmas/cli  
docker run -ti --rm quay.io/ibmmas/cli mas upgrade
```

Note: Ensure you select the Maximo Application Suite release that you want to upgrade to instead of the one you are currently using. For example, if you are upgrading from Maximo Application Suite 8.8 to Maximo Application Suite 8.9, use `-C 8.9.x`.

You can run the **mas upgrade** command in noninteractive mode:

```
mas upgrade -i inst1 --no-confirm
```

Related concepts

[Prerequisites for installing](#)

Before you begin, ensure that your environment meets the prerequisites by downloading, and installing the software and interfaces that you use to install IBM Maximo Application Suite.

[Planning to install in disconnected environment](#)

Upgrading IBM Maximo Application Suite manually



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription”](#) on page 482.

Manual upgrades of IBM Maximo Application Suite involve downloading the latest version from IBM Passport Advantage, running the installer to upgrade your environment, and updating your applications in Maximo Application Suite.

By completing the following upgrade process, the core Maximo Application Suite systems are upgraded and restarted. These systems include management APIs, the user registry, the Maximo Application Suite administration page, and the Maximo Application Suite navigator.

Before you begin

Before you can upgrade the Maximo Application Suite, you must perform the following tasks:

1. Complete the [prerequisite steps](#), such as deleting incompatible applications, installing the IBM Cloud Certificate Manager service, and upgrading your Cloud Pak for Data system.
2. Ensure that the workstation on which you run the Maximo Application Suite installer has the [Java Runtime Environment](#) installed. You need this component to accept the Maximo Application Suite licence.

Procedure

1. Download the Maximo Application Suite software installer.
Follow the procedure that is outlined in the [Maximo Application Suite download document](#) or the instructions that are provided by your IBM representative to access and download the software installer for the version of the product that you want to upgrade to.
2. Extract the contents of the compressed installer file to your workstation.
3. Log in to the Red Hat OpenShift web console as an administrator.
4. Copy the login token.
 - a. Click your user account icon and select **Copy Login Command**.
 - b. Copy the login token.
5. In your local command line, paste in the login token and press enter to log in to the Red Hat OpenShift cluster.

```
oc login --token=kiaj2_jkoasunJljsLdsqdsa787asd --server=https://
api.myopenshiftcluster.com:6443
```

Note: You can also use a valid token for a service account if you are using a service account for the installation. Maximo Application Suite.

6. Locate and run the `install-mas.sh` file.

The `install-mas.sh` file is part of the Maximo Application Suite V8.N for Multiplatform package that you downloaded from Passport Advantage.

```
./install-mas.sh -i <instance_name> -d <mas_domain> -c <cluster_issuer>
```

Where:

- `<instance_name>` is the Red Hat OpenShift instance that you want to upgrade.
- `<mas_domain>` is the domain name for your environment.
- `<cluster_issuer>` is the cluster issuer that was specified during the initial installation. If no cluster issuer was specified, remove the `-c <cluster_issuer>` part of the command.

7. Accept the license terms.

The Maximo Application Suite license terms are displayed. To continue with the upgrade, you must accept the license terms.

8. Monitor the upgrade.

The installer downloads the Maximo Application Suite container images from IBM Entitled Registry. Depending on your cluster's network speed, this download might take up to 30 minutes.

As the upgrade progresses, verify the successful upgrade of each component in your command shell. For each installation step, debug log information, including warnings and errors, and the step result are displayed.

```
....
....
Installing ibm-mas operator
-----
namespace/mas-appconnccost-core unchanged
Warning: oc apply should be used on resource created by either oc create --save-config or oc
apply
customresourcedefinition.apiextensions.k8s.io/appconnects.addons.mas.ibm.com configured

Warning: oc apply should be used on resource created by either oc create --save-config or oc
apply
customresourcedefinition.apiextensions.k8s.io/mviedges.addons.mas.ibm.com configured
Warning: oc apply should be used on resource created by either oc create --save-config or oc
apply
customresourcedefinition.apiextensions.k8s.io/bascfgs.config.mas.ibm.com configured
....
....
```

Tip: You can monitor the Maximo Application Suite initialization from the Red Hat OpenShift user interface under **Workloads > Pods**. Filter by namespace to locate the Maximo Application Suite pods.

9. Review the upgrade summary.

After the successful completion of the Maximo Application Suite upgrade, the following information is displayed:

- Administration dashboard URL. For example: `https://admin.<mas_domain>`

What to do next

Important: The complete upgrade process might take around 15 minutes, if most of the pods are replaced. The pod replacement continues even after the upgrade summary is displayed.

Ensure that the upgrade is complete before you continue. You can monitor the Pods in **Workloads > Pods**, or by using the following command from the Red Hat OpenShift command line:

```
oc get pods -n mas-<instance_name>-core
```

You can also run the following command to verify that the upgrade completed successfully:

```
oc get Suite -n mas-<instanceId>-core <instanceId>
```

If the upgrade is successful, this command displays output that is similar to the following text:

NAME	VERSION	STATUS	SYSTEMDATABASEREADY	BASINTEGRATIONREADY
SLSINTEGRATIONREADY	AGE			
appconnccost	8.7.0-pre.m3dev87	Ready	Ready	Ready
Ready	4d1h			

When the upgrade is complete, you can log in to the Suite administration page and continue with the upgrade process by [updating applications](#).

Converting IBM Maximo Application Suite from manual deployment to channel subscription

If you are using IBM Maximo Application Suite 8.9 or earlier and are upgrading to Maximo Application Suite 8.10, you must run the Operator Lifecycle Manager (OLM) conversion script to use a subscription method and subscribe to the latest channel.

The manual deployment of Maximo Application Suite core and its applications is discontinued starting in Maximo Application Suite 8.10.

Procedure

1. On your local machine, run the command to pull the image and initiate the container.
For example, use Docker to pull the image.

```
docker run -ti --rm --pull always quay.io/ibmmas/cli
```

Alternatively, you can use Podman or Skopeo.

2. On the command line inside the container, run the **oc login** command.
You can access your Red Hat OpenShift cluster by using the **oc** command directly from a terminal on the client machine that oc was installed on.
 - a) In the Red Hat OpenShift web console, click your login name and select **Copy login command**.
 - b) Click **View token**.
 - c) In the **Login with this token** field, copy the entire content that has the token and paste it on the command line inside the Docker container.
3. Export the required environment variables.

```
export MAS_INSTANCE_ID=<YourMASInstanceID>
```

```
export IBM_ENTITLEMENT_KEY=<YourIBMENTitlementKey>
```

4. Optional: Configure the upgrade strategy for the subscriptions.
By default, the upgrade strategy is set to Manual.

```
export MAS_UPGRADE_STRATEGY=Automatic
```

5. Run the **ansible-playbook** command with the **convert_to_olm** role.

```
ansible-playbook ibm.mas_devops.oneclick_convert_to_olm
```

For more information, see [convert_to_olm](#).

Any supported applications that are deployed in the cluster are converted to a channel subscription. Any applications that aren't deployed are not converted.

Upgrading to Maximo Application Suite 9.1

When you upgrade to Maximo Application Suite 9.1, either Maximo Manage or the foundation service is required. The upgrade path to Maximo Application Suite 9.1 depends on your current installation and configuration.

Maximo Application Suite 9.1 includes the foundation service bundle that provides common functionality, such as a unified authorization services and one common navigation for all suite applications. For more information, see [“Foundation service” on page 136](#).

Upgrading from Maximo Application Suite 9.0 or earlier to 9.1

You can upgrade Maximo Application Suite by using a channel subscription. For more information, see [“Upgrading IBM Maximo Application Suite by using the channel subscription method” on page 477](#).

If Maximo Manage is deployed

If Maximo Manage 9.0 or earlier is deployed and you are upgrading to 9.1, your upgrade includes Maximo Manage 9.1. The **cron**, **mea**, **report**, and **ui** server bundles for Maximo Manage are updated. The foundation service is also added as an additional bundle. Alternatively, if you choose to deploy only the **all** server bundle, which combines these bundle types, then the foundation service is included in the **all** bundle.

If other suite applications are deployed without Maximo Manage

If version 9.0 or an earlier version of other suite applications is deployed, you can upgrade by using the CLI.

Before you upgrade, you are asked whether to also deploy Maximo Manage.

If you select Yes, then Maximo Manage is also included in the upgrade, which provides full Manage capabilities and the Manage industry solutions and add-ons that you also choose. For more information, see [“Deploying Maximo Manage in Maximo Application Suite”](#) on page 298.

If you select No, then the foundation service is installed, which provides the common functionality, such as the unified authorization service and one common navigation for all suite applications. For more information, see [“Foundation service”](#) on page 136. The Manage namespace is created and the Manage tile is added to the catalog in the user interface.

To complete the upgrade, you must also provide a database that supports Maximo Manage. For more information, see [“Preparing your database for deployment”](#) on page 301

If stand-alone Maximo Health 9.0 or earlier is deployed

Starting in Maximo Application Suite 9.1, Maximo Health is no longer available as a stand-alone suite application but is still available as an add-on in Maximo Manage.

If Maximo Health 9.0 or earlier is deployed as a stand-alone suite application, you must add Maximo Health as an add-on in Maximo Manage 9.0 and then upgrade to 9.1. For more information, see [Upgrading Health stand-alone 9.0 to Manage with Health 9.1](#).

Upgrading from Maximo Asset Management 7.6.1.3 or earlier to Maximo Application Suite 9.1

If you are upgrading from Maximo Asset Management 7.6.1.3 or earlier to Maximo Manage in Maximo Application Suite 9.1, then the **cron**, **mea**, **report**, and **ui** server bundles for Maximo Manage are installed. The foundation service bundle is also included. Alternatively, if you choose to deploy only the **all** server bundle, which combines these bundle types, then the foundation service is included in the **all** bundle.

For more information, see [“Upgrading from Maximo Asset Management to IBM Maximo Manage”](#) on page 494.

Related concepts

[Foundation service](#)

Customer-managed **Updating applications**

When new versions of applications become available as part of a Maximo Application Suite upgrade, you can update your deployed applications.

About this task

The application update steps differ depending on the [upgrade method](#) for the application:

- For a deployment method, you must upgrade the application from the Maximo Application Suite catalog.
- For a subscription method, the application is automatically updated after your approval.

Updating IBM Maximo Manage

When new versions of IBM Maximo Manage become available as part of a Maximo Application Suite upgrade, you can update your deployed application as the application administrator.

Procedure

- Updating Maximo Manage in a channel subscription
 - a) From the **App switcher**, select **Suite administration**.
 - b) In the **Suite administration** navigation menu, select **Applications**.
 - c) From the list of applications, select the **Manage** application.
 - d) In the **Manage** application management page, select **Actions** > **Administer versions**.
 - e) In the **Change application upgrades** window, expand the drop-down list for the **Select new channel** field, and select a channel.
 - f) Click **Apply changes**. A confirmation dialog is displayed about establishing subscription to the update channel. After subscription to the update channel, the deployment and activation starts. The following **Status** tiles are updated in the **Application details** section. For example,

Operator	Deployed
Application	Ready
Application	Running
 - g) Select a Reduce Downtime state. If you don't want the application update to proceed with the default state for reduced system downtime, then you must select another Reduce Downtime state. Complete the following steps. For more information about the functions, see [“Reducing system downtime” on page 323](#).
 - a. In the **Manage** application management page, click **Actions** > **Update configuration**.
 - b. In the **Update Manage configuration** window, go to the **Activation configuration** section, and click **MAS Development workspace details page**.
 - c. In the **Manage Workspace** details page, click **Actions** > **Update configuration**.
 - d. In the **Update Manage configuration** window, go to the **Activation configuration** section, and click the edit icon next to **Database connection**.
 - e. Click **Show advanced settings**.
 - f. Within **Advanced settings**, go to the **Database** section, and go to the **Reduce Downtime** property.
 - g. In the **Reduce Downtime** category, select one of the states.
 - h. In the **Failure Control** category, select a value.
 - i. Click **Apply changes** to save and apply the configuration changes. The configuration change applies to future upgrades.
 - j. Start the activation process. In the confirmation window **Activate Manage application in workspace**, click **Confirm**.

What to do next

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated. The update time is based on network speeds and processing performance. Therefore, communicate with users about any system downtime due to an update.

If the update fails, retry the update with another Reduce Downtime state.

Updating IBM Maximo Health

When new versions of IBM Maximo Health become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

Procedure

- Updating Maximo Health in a channel subscription
 - a) From the Application management page, click **Update available** for the Maximo Health application.

Information about the latest version is displayed.

- Click **Select Maximo Health** to continue.
- You can also click **Available versions** and then select a different supported version to deploy.
- To deploy a version that is not listed, such as a specific fix, click **Deploy another version** and then manually enter the version number in the search field.

Note: Fix version information is provided by your IBM representative. A complete list of application versions is also available in the product documentation.

- b) Click **Deploy** to start the deployment process.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

- Manually updating Maximo Health



Attention: Starting in Maximo Application Suite 8.10, the manual deployment for applications that use the installation script for Maximo Application Suite is discontinued. To upgrade Maximo Application Suite and its applications, you must run a conversion script to use a subscription method and subscribe to the upgraded channel. For more information, see [“Converting IBM Maximo Application Suite from manual deployment to channel subscription”](#) on page 482.

- a) From the Suite catalog, from the **Applications** tab, on the **Health** tile, click **Update available**. Information about the latest version is displayed. From the **Version** field, you can select a different supported version to deploy. To deploy a version that is not listed, such as a specific fix, from the **Version** field, click **Other version** and then in the **Other version number** field, specify the version number. Fix version information is provided by your IBM representative.
- b) Click **Deploy version** to start the deployment process. The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

What to do next

As an application administrator, you can now complete the following tasks:

- Check the system status. For more information, see [Checking the system status](#).
- For other getting started steps in Maximo Health, review [Getting started for application administrators](#).

Related concepts

[IBM Maximo Health](#)

IBM Maximo Health is an application in Maximo Application Suite. By using Maximo Health, you can improve your asset's reliability by understanding asset health and taking action. You can review your assets' performance and condition indicators, such as the last failure date and the maintenance-to-replacement ratio (MRR), and take action by creating work orders and service requests. You can use work

queues to improve the quality of your asset's details and related data. You can also configure scoring for assets' health, criticality, and risk.

[Maximo Health and Predict - Utilities](#)

Related tasks

[Deploying IBM Maximo Health](#)

Improve the reliability of your assets by proactively monitoring and managing asset health by using Maximo Health.

Related information

[Getting started with Maximo Health and Predict - Utilities](#)

Customer-managed **Updating Maximo Monitor**

When new versions of Maximo Monitor become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

About this task

The following steps are specific to the Maximo Monitor application and are part of the overarching application update process.

Procedure

1. From the **Application management** page, click **Update available** for the Maximo Monitor application. Information about the latest version is displayed.

Click **Select Maximo Monitor** to continue.

You can also click **Available versions** and then select a different supported version to deploy.

Note: To deploy a version that is not listed, such as a specific fix, click **Deploy another version** and then manually enter the version number in the search field.

Fix version information is provided by your IBM representative. A complete list of application versions is also available in the product documentation.

2. Click **Deploy** to start the deployment process.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

What to do next

The Asset Data Dictionary component in Maximo Manage is installed automatically when a consuming application, such as Maximo Monitor, is installed. The Maximo Monitor is a repository of data and metadata in Maximo Manage that facilitates data sharing and synchronization between Maximo Manage and other applications in the suite. If you are upgrading to Maximo Monitor and you plan to synchronize data between Maximo Monitor and Maximo Manage, complete the configuration steps in the [Asset Data Dictionary implementation](#) section of the Maximo Manage product documentation.

For other getting started steps in Maximo Application Suite, review the [Getting started with Maximo Monitor](#).

Related tasks

[Deploying IBM Maximo Monitor](#)

By using Maximo Monitor, you can visualize current and historical trend data for your devices and assets on customizable dashboards.

Updating Maximo Real Estate and Facilities

When new versions of Maximo Real Estate and Facilities Platform become available as part of a Maximo Application Suite upgrade, you can update your deployed Maximo Real Estate and Facilities Platform.

Procedure

- Updating Maximo Real Estate and Facilities in a channel subscription
 - a) From the Application management page, click **Update available** for the Maximo Real Estate and Facilities application.

Information about the latest version is displayed.

 - Click **Select Maximo Real Estate and Facilities** to continue.
 - You can also click **Available versions** and then select a different supported version to deploy.
 - To deploy a version that is not listed, such as a specific fix, click **Deploy another version** and then manually enter the version number in the search field.

Note: Fix version information is provided by your IBM representative. A complete list of application versions is also available in the product documentation.
 - b) Click **Deploy** to start the deployment process.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.
- Manually updating Maximo Real Estate and Facilities
 - a) From the Suite catalog, from the **Applications** tab, on the **Real Estate and Facilities** tile, click **Update available**. Information about the latest Maximo Real Estate and Facilities Platform version is displayed. From the **Version** field, you can select a different supported version to deploy. To deploy a version that is not listed, such as a specific fix, from the **Version** field, click **Other version** and then in the **Other version number** field, specify the version number. Fix version information is provided by your IBM representative.
 - b) Click **Deploy version** to start the deployment process. The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During deployment of the new Maximo Real Estate and Facilities Platform version, users might temporarily lose access to applications while they are being updated.

What to do next

As an application administrator, you can now complete the following tasks:

- Check the system status. For more information, see [Checking the system status](#).
- Apply Maximo Real Estate and Facilities Applications updates, if needed depending on your Maximo Real Estate and Facilities Applications version.

Updating Maximo Predict

When new versions become available, you can update Maximo Predict.

Before you begin

Ensure that Maximo Health is updated to the same version.

About this task

If automatic updates are not configured for the channel subscription, you must complete a few steps to update the application.

Procedure

1. To update Maximo Predict in a channel subscription, complete the following steps:

- a) In the side navigation menu, under **Suite** application, select the **Applications** page from **Administration > Suite** module.
- b) Click **Update available** for **Predict**.

Information about the latest version is displayed. **Multiple updates available** is displayed instead of the **Update available** when both a version and channel update are available.

- c) Select a different supported version to deploy and apply your changes.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

2. To manually update Maximo Predict, complete the following steps:

- a) In the side navigation menu, under **Suite** application, select the **Applications** page from **Administration > Suite** module.
- b) Select **Predict**, click **Update available**. Information about the latest version is displayed.
- c) From the **Version** drop-down menu, you can select a different supported version to deploy.
- d) Optional: To deploy a version that is not listed, such as a specific fix, from the **Version** drop-down menu, select **Other version**. In the **Other version number** field, specify the version number. The fix version information is provided by your IBM representative.
- e) Click **Deploy version** to start the deployment process. The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

You can also manually upgrade and edit the configuration from the **Overview** tab for the application. To upgrade, from the **Actions** menu, click **Application versions**.

Related concepts

IBM Maximo Predict

IBM Maximo Predict is an application in Maximo Application Suite. By using Maximo Predict, you can leverage your historical and near real-time asset performance data, maintenance records, inspection reports, and environmental data to correlate performance factors that predict asset degradation or failure. Maximo Predict also uses artificial intelligence to optimize predictive model accuracy.

Related tasks

Deploying IBM Maximo Predict

Maximo Predict can use historical and recent asset performance data to correlate performance factors that predict asset degradation or failure. Other types of data that can be correlated include maintenance records, inspection reports, and environmental data. Maximo Predict uses artificial intelligence to optimize predictive model accuracy.

Activating IBM Maximo Predict

Before you can grant users access and start working with Maximo Predict, you must activate the application. You can activate Maximo Predict after the deployment is complete.

Related information

[Getting started with Maximo Predict](#)

Customer-managed **Updating Maximo Collaborate**

When new versions of Maximo Collaborate become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

Before you begin

The Maximo Collaborate application requires no pre-update steps.

About this task

The following steps are specific to the Maximo Collaborate application and are part of the overarching application update process.

Procedure

1. To update Maximo Collaborate in a channel subscription, complete the following steps:
 - a) In the side navigation menu, under **Suite** application, select the **Applications** page from **Administration > Suite** module.
 - b) Click **Update available** for **Collaborate**.

Information about the latest version is displayed. **Multiple updates available** is displayed instead of the **Update available** when both a version and channel update are available.
 - c) Select a different supported version to deploy and apply your changes.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.
2. To manually update Maximo Collaborate, complete the following steps:
 - a) In the side navigation menu, under **Suite** application, select the **Applications** page from **Administration > Suite** module.
 - b) Select **Collaborate**, click **Update available**. Information about the latest version is displayed.
 - c) From the **Version** drop-down menu, you can select a different supported version to deploy.
 - d) Optional: To deploy a version that is not listed, such as a specific fix, from the **Version** drop-down menu, select **Other version**. In the **Other version number** field, specify the version number. The fix version information is provided by your IBM representative.
 - e) Click **Deploy version** to start the deployment process. The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

What to do next

The Maximo Collaborate application requires no post-update steps.

For other getting started steps in Maximo Collaborate, see [Getting started with Maximo Collaborate](#).

Related concepts

[IBM Maximo Collaborate](#)

Related tasks

[Deploying IBM Maximo Collaborate](#)

[Activating IBM Maximo Collaborate](#)

[Rolling back Maximo Collaborate](#)

Related information

[Getting started](#)

[Maximo Assist](#)

Customer-managed **Updating Maximo Visual Inspection**

As part of a Maximo Application Suite upgrade, a new version of Maximo Visual Inspection is available for update. The following steps are specific to Maximo Visual Inspection and are part of the overarching application update process.

Before you begin

Ensure that you complete the pre-update steps relevant to your version of Maximo Visual Inspection before you update the application.

Updating to version number	Updating from version number	Pre-update steps
8.2.0	8.1.0	None
8.3.0	8.2.0 and earlier versions	This update is not supported. To move to version 8.3.0, complete the following steps: <ol style="list-style-type: none">1. Back up your data.2. Delete the earlier version.3. Install version 8.3.0.4. Import your data.
8.4.0	8.2.0 and earlier versions	Create and apply YAML files in your Red Hat OpenShift cluster. For more information about creating and applying YAML files, see Deploying Maximo Visual Inspection
8.5.0	8.4.0 and earlier versions	This update is not supported. To move to version 8.5.0, complete the following steps: <ol style="list-style-type: none">1. Back up your data.2. Delete the earlier version.3. Install version 8.5.0.4. Import your data.
8.6.0	8.5.0	None

About this task

Complete the following steps in Maximo Visual Inspection.

Procedure

1. Open the **Applicant management** page.
2. In the **Visual Inspection** tile, select **Update available**.
3. Select the version that you want to update to.
4. Click **Deploy version** to start the deployment process.
5. Optional: To update to the most recent version, complete steps 5-8 of [Deploying Maximo Visual Inspection](#)
6. Optional: To update to an earlier version, complete the deployment steps for that version.

Important: The estimated update time is based on typical network speeds and processing performance. During deployment, you might temporarily lose access to the application.

Customer-managed

Updating Maximo Health and Predict - Utilities

Note:

Starting in Maximo Application Suite 8.11, Maximo Health and Predict - Utilities is no longer available as a separate industry solution. The information that is provided is applicable only to Maximo Application Suite 8.10 and earlier versions. For more information, see [“Upgrading IBM Maximo Application Suite” on page 473](#). Before you upgrade to Maximo Application Suite 8.11, deactivate and delete Maximo Health and Predict - Utilities.

When new versions of Maximo Health and Predict - Utilities become available as part of a Maximo Application Suite upgrade, you can update your deployed application.

The update process that you complete is determined by the update method that the industry solution uses.

Before you begin

Important:

IBM Maximo Application Suite supports Cloud Pak for Data 4.6 with Python 3.10.

Ensure that the following dependencies are also updated:

- Maximo Health.
- App Connect.
 - If you want to update Maximo Health and Predict - Utilities to a new version, you need to check the Maximo Health and Predict - Utilities supported App Connect version. From IBM Maximo Application Suite v8.8, App Connect 12.0.4 or later is supported. If the old version App Connect is not supported anymore, you must first upgrade App Connect to a later version by completing the following steps:
 1. Upgrade the App Connect operator. For more information, see [Upgrading the IBM App Connect Operator](#).
 2. If you customized the App Connect instance, review the [Upgrade considerations for IBM App Connect instances](#).
 3. Upgrade the App Connect instance. For more information, see [Upgrading your instances](#).
 4. Upgrade your integration servers to a newer version. You also need to update the corresponding license for this version. For more information, see [Licensing reference for IBM App Connect Operator](#).
 5. Ensure that before you set **Use Common Service** to false, there is no apikey or credentials in IBM App Connect dashboard. If the apikey or credentials already exist, delete them from IBM App Connect and add them back after enabling authentication. For more information, see [Editing the settings for a deployed integration server and Configuration reference](#).
 - If you have a new App Connect dashboard URL, you must also configure the new URL. For more information, see [Configuring IBM App Connect](#).

- IBM Watson Studio on Cloud Pak for Data.

About this task

If automatic updates are not configured for the channel subscription, you must complete a few steps to update the application.

Procedure

- To update Maximo Health and Predict - Utilities in a channel subscription, complete the following steps:
 - On the **Suite administration** page, select **Applications**.
 - Click **Update available** for .

Information about the latest version is displayed. **Multiple updates available** is displayed instead of the **Update available** when both a version and channel update are available.
 - Select a different supported version to deploy and apply your changes.

The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.
- To manually update Maximo Health and Predict - Utilities, complete the following steps:
 - On the **Suite administration** page, select **Applications** and for **Health and Predict - Utilities**, click **Update available**. Information about the latest version is displayed.
 - From the **Version** drop-down menu, you can select a different supported version to deploy.
 - Optional: To deploy a version that is not listed, such as a specific fix, from the **Version** drop-down menu, select **Other version**. In the **Other version number** field, specify the version number. The fix version information is provided by your IBM representative.
 - Click **Deploy version** to start the deployment process. The application update process might require software downloads and more configuration steps. The estimated update time is an estimate based on typical network speeds and processing performance.

Important: During the deployment of the new versions, users might temporarily lose access to the applications while they are being updated.

You can also manually upgrade and edit the configuration from the **Overview** tab for the application. To upgrade, from the **Actions** menu, click **Application versions**.
- With Cloud Pak for Data 4.6, the IBM Runtime 22.1 on Python 3.9 works well. If you want to use IBM Runtime 22.2 on Python 3.10 for existing notebook, stop its kernel and then change the environment.
 - To stop the kernel, click the overflow menu and then click **Stop Kernel**.
 - To change the environment, click the overflow menu and then click **Change Environment** . Select **IBM Runtime 22.2 on Python 3.10 (1 vCPU and 2 GB RAM)** and then click **Associate**.
- Existing notebooks that are deployed or edited must be updated to use the latest versions of Python that is supported by Cloud Pak for Data. Refer to the following steps to edit the notebooks:
 - Update the healthlib file name in cell 1


```
for f in ['IBM-Transformers-Tap-Changers-DGA-4.0.0.cfg', 'healthlibv4-1.0.0-cp38-cp38-linux_x86_64.whl']:
```
 - Update healthlib file name in cell 3


```
for f in ['IBM-Transformers-Tap-Changers-DGA-4.0.0.cfg', 'healthlibv7-1.0.0-cp310-cp310-linux_x86_64.whl']:s
```

```
!pip install healthlibv4-1.0.0-cp38-cp38-linux_x86_64.whl
```

```
!pip install healthlibv7-1.0.0-cp310-cp310-linux_x86_64.whl
```

c) Update healthlib in cell 6

```
import healthlibv4 as healthlib
from healthlibv4 import context
from healthlibv4 import maximo_function
```

```
import healthlibv7 as healthlib
from healthlibv7 import context
from healthlibv7 import maximo_function
```

d) Update cell 9

```
healthlib.set_asset_query(select="assetid,assetnum,siteid,orgid,location,installdate,assetmeter{lastreading,metername},assetspec{alnvalue,numvalue,assetattrid}", page_size=100)
```

```
healthlib.set_asset_query(select="assetid,assetnum,siteid,orgid,location,installdate,status,assetmeter{lastreading,metername},assetspec{alnvalue,numvalue,assetattrid}", page_size=100)
```

e) Update cell 10

```
healthlib.set_location_query(select="locationsid,location,siteid,orgid,installdate,locationmeter{lastreading,metername},locationspec{alnvalue,numvalue,assetattrid},rel.activeasset{assetid,assetnum,siteid,orgid,installdate,rel.assetspec{alnvalue,numvalue,assetattrid},rel.assetmeter{lastreading,metername}}", page_size=100)
```

```
healthlib.set_location_query(select="locationsid,location,siteid,orgid,installdate,status,locationmeter{lastreading,metername},locationspec{alnvalue,numvalue,assetattrid},rel.activeasset{assetid,assetnum,siteid,orgid,installdate,rel.assetspec{alnvalue,numvalue,assetattrid},rel.assetmeter{lastreading,metername}}", page_size=100)
```

f) To save the file, click **File > Save**. To save the version, click **File > Save Version**.

5. On the **Job tab**, you can edit the job running environment

- a) Click **Edit configuration**
- b) Select **IBM Runtime 22.2 on Python 3.10**
- c) Click **Next**, then **Review and Save** to the edit to the job.

Customer-managed **Upgrading from Maximo Scheduler Optimization to Maximo Optimizer**

In Maximo Application Suite 8.8, you use Maximo Optimizer instead of Maximo Scheduler Optimization. If Maximo Scheduler Optimization is deployed in your environment and you are upgrading to Maximo Application Suite 8.8, you must uninstall Maximo Scheduler Optimization before you can complete the upgrade.

Before you begin

Back up customized model data in Maximo Scheduler Optimization that you want to migrate to Maximo Optimizer by recording or taking a screen capture of the following model information:

- Model project
- Model name
- Model version
- Model type
- Entry class name

Also, download the model artifact files and record the model type of the model or extension.

Note: You do not need to back up default models, such as Graphical Assignment (GA), Graphical Scheduling (GS). You also do not need to back up the Graphical Scheduling Large Projects (GSLP) models and Asset Investment Optimization (AIO) models that are used in Maximo Health. These models are automatically migrated when Maximo Optimizer is activated in Maximo Application Suite 8.8.

About this task

You can upgrade without migrating any historical optimization jobs. If you want to back up your historical optimization jobs with the attachments into your local file directory, you can use a script that is provided by IBM. To request the script, you open a case with IBM Support.

Procedure

1. In Maximo Application Suite 8.7, deactivate and delete Maximo Scheduler Optimization 8.1.x.
2. Upgrade from Maximo Application Suite 8.7 to 8.8.
3. Deploy Maximo Optimizer 8.2.
4. Restore the customized models by manually creating them in the Optimizer administration dashboard.

Related concepts

[IBM Maximo Optimizer](#)

IBM Maximo Optimizer is an add-on to Maximo Application Suite. By using Maximo Optimizer, you can automate efficient decisions for long-range planning, scheduling, and dispatching of resources for asset maintenance while balancing competing objectives and constraints.

Related tasks

[Deploying IBM Maximo Optimizer](#)

By using Maximo Optimizer, you can automate efficient decisions for long-range plans, schedules, and the dispatch of resources for asset maintenance, helping to balance competing objectives and constraints.

[Upgrading IBM Maximo Application Suite](#)

[Deactivating and deleting applications](#)

By deleting an application, you remove access to it from the environment. To reinstate it, you can redeploy the application. Application data and metadata for the deleted application is not deleted but is retained in the corresponding repositories, such as MongoDB and Db2.

Upgrading from Maximo Asset Management to IBM Maximo Manage

Earlier known as IBM Maximo Asset Management, Maximo Manage is part of IBM Maximo Application Suite and has enhanced and new features.

Notice:

To download a PDF version of this entire section, click [Upgrading from Maximo Asset Management to IBM Maximo Manage PDF document](#).

The PDF content is updated regularly. To ensure that you use the latest version, print the guide when you are ready to upgrade to Maximo Manage in Maximo Application Suite.

Watch a video to understand the changes in the new Upgrading from Maximo Asset Management to Maximo Manage documentation.

Maximo Manage in Maximo Application Suite overview

IBM Maximo Asset Management is now called Maximo Manage. Maximo Manage is one of many applications that are included within Maximo Application Suite.

Maximo Application Suite is an integrated suite of applications that is built on Red Hat OpenShift environment to provide multi-cloud portability, including support for Hybrid Cloud or on-premises deployments. For more information, see [IBM Maximo Application Suite architecture](#).

Maximo Application Suite also contains other applications such as IBM Maximo Health, IBM Maximo Predict, and IBM Maximo Monitor.

Maximo Application Suite is based on the Red Hat OpenShift deployment model. Red Hat OpenShift is a platform-as-a-service system that is built around containers and uses container orchestration that is provided by Kubernetes. It is designed to integrate well with IBM Cloud Pak for Data components.

After you upgrade to Maximo Manage, user licenses are managed in Maximo Application Suite.

- In IBM Maximo Application Suite, subscription term or perpetual license is managed through AppPoints.
 - Contact IBM account team about AppPoints and IBM Maximo Application Suite license file.
 - Buy a pool of AppPoints.
 - The same AppPoints can be applied to any application in IBM Maximo Application Suite.
 - For more information, see [“Administering licenses and AppPoints usage” on page 813](#)

Available user entitlements:

- Limited
- Base
- Premium
- Self-Service

What's changed in Maximo Manage

In previous versions, Maximo Asset Management was an independent product. Now, Maximo Manage is part of Maximo Application Suite. Many processes such as licensing and user management that were managed at the product-level are now managed at the suite level.

Changes in Maximo Manage

Changes in licensing model

Maximo Application Suite uses a different licensing model and uses AppPoints to track application usage, runtime, and user access. AppPoints are allocated in your organization as defined by your license entitlement. You can configure your environment to enforce the AppPoint entitlement.

Changes in deployment and architecture

Installation, configuration, deployment, and upgrade are done by the Red Hat OpenShift operator in Maximo Application Suite. The entire deployment is based on Red Hat OpenShift. For more information, see [Upgrading to Maximo Application Suite](#).

Changes due to containerization

Containerization is the packaging of software code with just the operating system (OS) libraries and dependencies that are required to run the code to create a single lightweight executable package that is infrastructure-agnostic. System properties, server bundles, integration, and customization process when you upgrade from Maximo Asset Management to Maximo Manage are different because of containerization.

System properties

The Maximo properties file in Maximo Asset Management is replaced by the System properties application in Maximo Manage, through which you can set system properties in the Maximo Manage user interface. System properties include global properties that apply to all the server instances that use a common database, including a clustered environment. For more information, see [Setting system properties](#).

- The bootstrap properties, such as database username and password and encryption keys, are set in IBM Maximo Application Suite during deployment or in a custom resource in Red Hat OpenShift. These properties are applied to all workload deployments.

- The bundle-level properties for workloads are applied to the specific server and set in IBM Maximo Application Suite and stored as a ConfigMap file in Red Hat OpenShift. ConfigMap files contain configuration data that the Red Hat OpenShift pods consume.

For more information, see [Configuring Maximo Manage for deployment](#).

- Use the System Properties application in Maximo Manage to set other properties. These properties are applied to all the workloads. For example, if you change the administrative password, you must update the `mxadminPasswd` property.

Server bundles

The Maximo Manage application can be deployed in one or more workloads called server bundles. A server bundle isolates the workload processes so that they can be independently managed. Server bundles can be independently scaled and managed based on your needs. For more information, see [Server bundle overview](#).

Integration

An integration framework integrates data with other applications, within your enterprise or with external systems. The framework includes predefined content and a toolkit to extend the predefined content to new integration points. It also enables message providers and an abstraction of message queuing features making Maximo Manage independent of messaging models like JMS or Kafka.

Customization process

Maximo Archiving is not supported in Maximo Application Suite. Maximo Manage uses customization archives for using customized Java classes or database scripts. Maximo Asset Management 7.6.1.2 included customization archives. However, Maximo Manage has enhanced the customization process further. For more information, see [Customizing the application](#).

Changes in technology

RMI replaced by REST API

You use REST API instead of Remote Method Invocation (RMI) for interactions with the product from custom extensions or external applications. Maximo Application Suite does not support RMI.

For more information, see [REST API as replacement for RMI](#).

Message queues

If you are using Service Integration (SI) buses, you must migrate to Apache Kafka or any other supported JMS provider. For more information, see [Enabling data import and export through message queues](#).

Changes in application server

You do not need to install or migrate IBM WebSphere Application Server unless you used any of its features outside Maximo Asset Management. IBM WebSphere Application Server Liberty is embedded in the Maximo image and deployed automatically by Maximo Manage. You can configure other server configurations in the Maximo Application Suite user interface. For more information, see [Configuring the application server](#).

Changes in authentication and user management

In Maximo Asset Management 7.6, user authentication is configured in the application itself or at the application server level. For example, WebSphere Application Server Network Deployment.

In Maximo Manage, user authentication is configured at the Maximo Application Suite level. The upgrade process migrates the existing users to the **Users** application in Maximo Application Suite where you can view and edit their details. The users are synchronized from Maximo Application Suite to Maximo Manage by using a cron task. Security groups are configured in Maximo Manage. For more information, see [Administering users and user access](#).

Impact of upgrading from Maximo Asset Management to Maximo Manage

Upgrading from Maximo Asset Management to Maximo Manage has an impact on users and administrators.

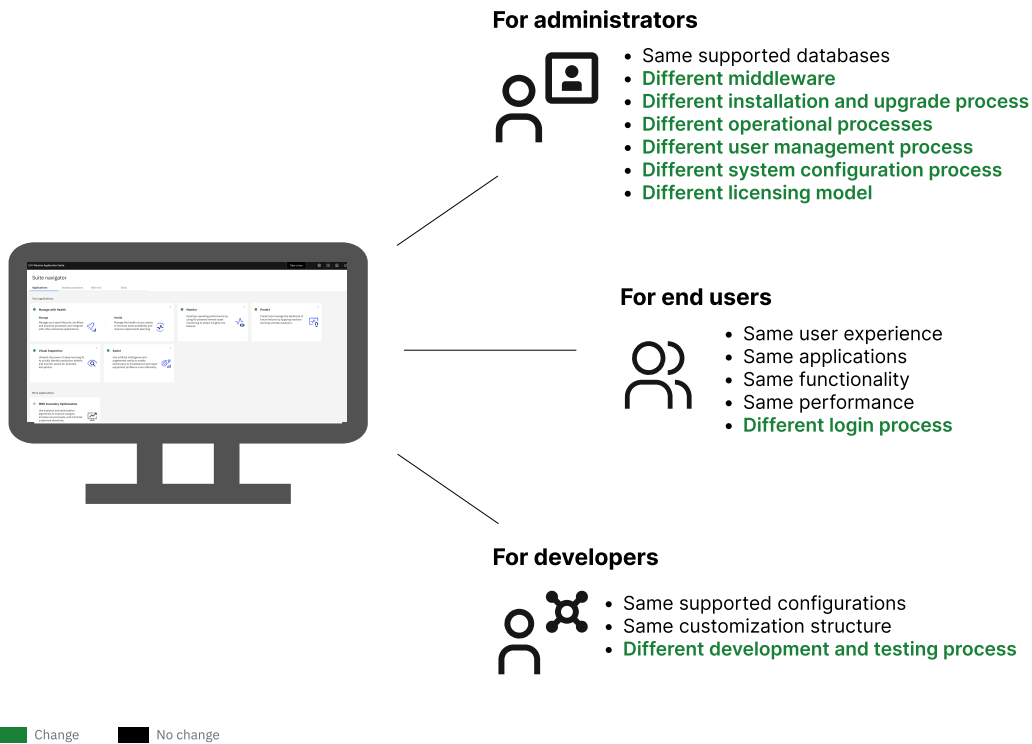


Figure 2. Impact of upgrading to Maximo Application Suite for different users, administrators, and developers

Changes for administrators

Maximo Manage supports the same databases as Maximo Asset Management. The following aspects are different when you upgrade to Maximo Application Suite:

Middleware implementation

Maximo Application Suite is deployed on Red Hat OpenShift Container Platform. For more information, see [IBM Maximo Application Suite technical overview](#).

Installation and configuration

Maximo Application Suite installation and configuration process includes customer-managed or IBM managed installations. You can choose from multiple platforms and environments. For more information, see [Installing Maximo Application Suite](#).

User management

User management process differs in that users are created and managed at the suite-level and are made available for application-level access with application-specific role assignments such as administrator or user. Users who are assigned a Maximo Manage entitlement, must be provided security group permissions in Maximo Manage to define the Maximo Manage applications, options, and data that the user can access. For more information, see [Administering users and users access](#) and [Security and user management](#).

System configuration

In Maximo Manage, the System Configuration module contains the Platform Configuration module and the Migration module. You use the applications in the Platform Configuration module to perform numerous tasks, such as managing systems properties and domains. You use the applications in the Migration module to migrate configuration content from one environment to another. For more information, see [System configuration module](#).

Upgrade process

You can upgrade Maximo Application Suite automatically or manually. For more information see, [Upgrading IBM Maximo Application Suite](#).

Changes for users

Maximo Application Suite users have a different login process.

Login process

The login process differs based on whether you have a customer-managed, IBM managed, or SaaS instance of Maximo Application Suite. For more information, see [Getting started](#).

Changes for developers

If your role entails developing and testing Maximo Manage code locally, you are provided with more flexibility and can set up a local development environment. For more information see, [Setting up a local Maximo Manage development environment](#).

What's supported in Maximo Application Suite

Products

IBM Maximo Anywhere and IBM Maximo Mobile

Maximo Anywhere is supported until Maximo Application Suite 8.8. Maximo Mobile has all connected and disconnected functions in one application. The initial version of Maximo Mobile includes Assist, Technician, and Inspections. For more information, see [Managing Maximo Mobile](#).

Industry solutions and add-ons

All industry solutions and add-ons that were a part of Maximo Asset Management are supported in Maximo Application Suite. For more information, see [Applications, industry solutions, add-ons, and tools](#).

Other products

IBM Maximo Scheduler, IBM Maximo Calibration and IBM Maximo Linear Asset Manager are now a part of Maximo Manage. IBM Maximo for Life Sciences is covered in Maximo Calibration. For more information, see [Maximo Manage components](#).

SAP and Oracle integration

SAP and Oracle integrations supported in Maximo Manage through connectors.

For more information, see:

- [Preparing Oracle and SAP Connector before upgrade](#)
- [Configuring Oracle connector after upgrade](#)
- [Configuring SAP connector after upgrade](#)

Third-party add-ons

Third-party changes that exist in Maximo Asset Management are migrated to Maximo Manage only if the changes can be extracted completely in a customization archive. For more information, see [Migrating customizations in a customization archive](#).

Features

File system

For users that use the file system mount, the same file system mount can be configured in Maximo Application Suite through volume mount in Red Hat OpenShift. If the mount point is the same as in Maximo Asset Management, no change is needed for Maximo Manage configuration. For more information, see [Configuring persistent volume claims](#).

Object storage or S3

No change is needed to migrate attached documents. A public certificate for the object storage server might need to be imported. For more information, see [Configuring attached documents](#).

Reporting

Business Intelligence and Reporting Tools (BIRT) supported. For more information, see [Report migration](#).

Migration Manager

All pending migration packages must be migrated before the upgrade. Configuration options are provided during deployment for persistent volume mount for file-based packages. Maximo Anywhere APIs are available. Migration Manager automates creating and deploying packages. For more information, see [Migration Manager](#).

Language support

All languages that were supported in Maximo Asset Management are supported in Maximo Manage. For more information, see [Language support](#).

What's not changed

Database

There are no significant changes in the database for Maximo Manage. Your Maximo Asset Management database is upgraded to the Maximo Manage database when you activate it.

For more information, see [Activating Maximo Manage](#).

User experience

- User interface and usability
- Applications
- Functions
- Performance

For more information, see [Maximo Manage overview](#).

Architecture and deployment models

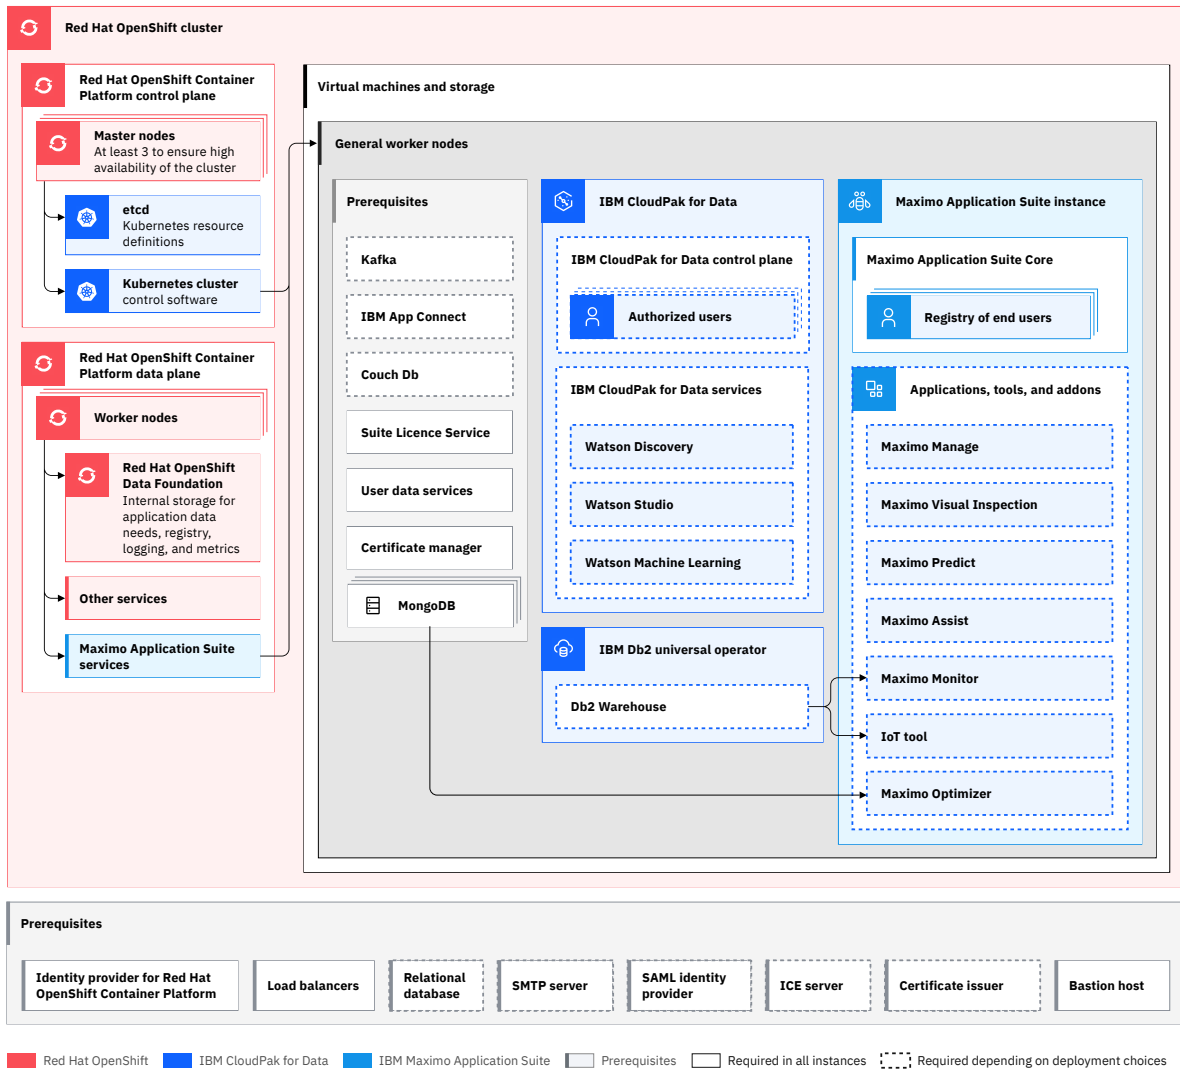
Maximo Application Suite architecture differs significantly from Maximo Asset Management because Maximo Application Suite is deployed on Red Hat OpenShift Container Platform. The shift in technology results in different deployment models for IBM Maximo Manage and Maximo Asset Management.

IBM Maximo Application Suite architecture

IBM Maximo Application Suite helps drive operational resiliency and reliability. With expanded access to enterprise asset management, Maximo Manage, Maximo Health, Maximo Predict, Maximo Monitor, and other applications, your team can reach across the enterprise to unify operations and maintain business continuity.

Maximo Application Suite has the following capabilities:

- Integrated suite of applications. For example, Maximo Manage, Maximo Health, Maximo Predict, and Maximo Monitor.
- Simplified licensing model.
- Hybrid-cloud deployment. Maximo Application Suite can be deployed on premises or in a public cloud.
- Comprehensive view of your assets.



- Maximo Application Suite includes applications, such as Maximo Manage, Maximo Monitor, Maximo Health, Maximo Predict, Maximo Visual Inspection, and Maximo Collaborate.
- IBM Cloud Pak for Data can run on your Red Hat OpenShift cluster. It offers AI-infused services for business and IT operations, development, data science, and management.
 - Note:** If you intend to integrate Maximo Manage with other application suites, you must install IBM Cloud Pak for Data. Otherwise, install only the IBM Db2 Warehouse stand-alone operator that can be one of the supported database options for Maximo Manage.
- MongoDB can exist in a cluster, cloud, or outside a cluster.
 - User, application, and entitlement metadata, such as OpenID Connect (OIDC) registration and user management, is stored in MongoDB. OIDC is an identity layer on the OAuth 2.0 protocol. OpenID Provider (OIDP) is an identity provider that is also used for user authentication.
- IBM Event Streams is a fully managed Apache Kafka-As-A-Service Platform for Cloud. The Maximo Manage application in the suite can be configured to use Kafka service.
- Maximo Manage support for all three databases, such as IBM Db2, Oracle Database, and Microsoft SQL Server.
- Maximo Manage can be configured to use cloud services, such as Cloud Object Storage, block, or file storage.

Deployment differences between Maximo Asset Management and Maximo Manage

Maximo Asset Management deployment is different from Maximo Manage deployment due to the change to Red Hat OpenShift Container Platform.

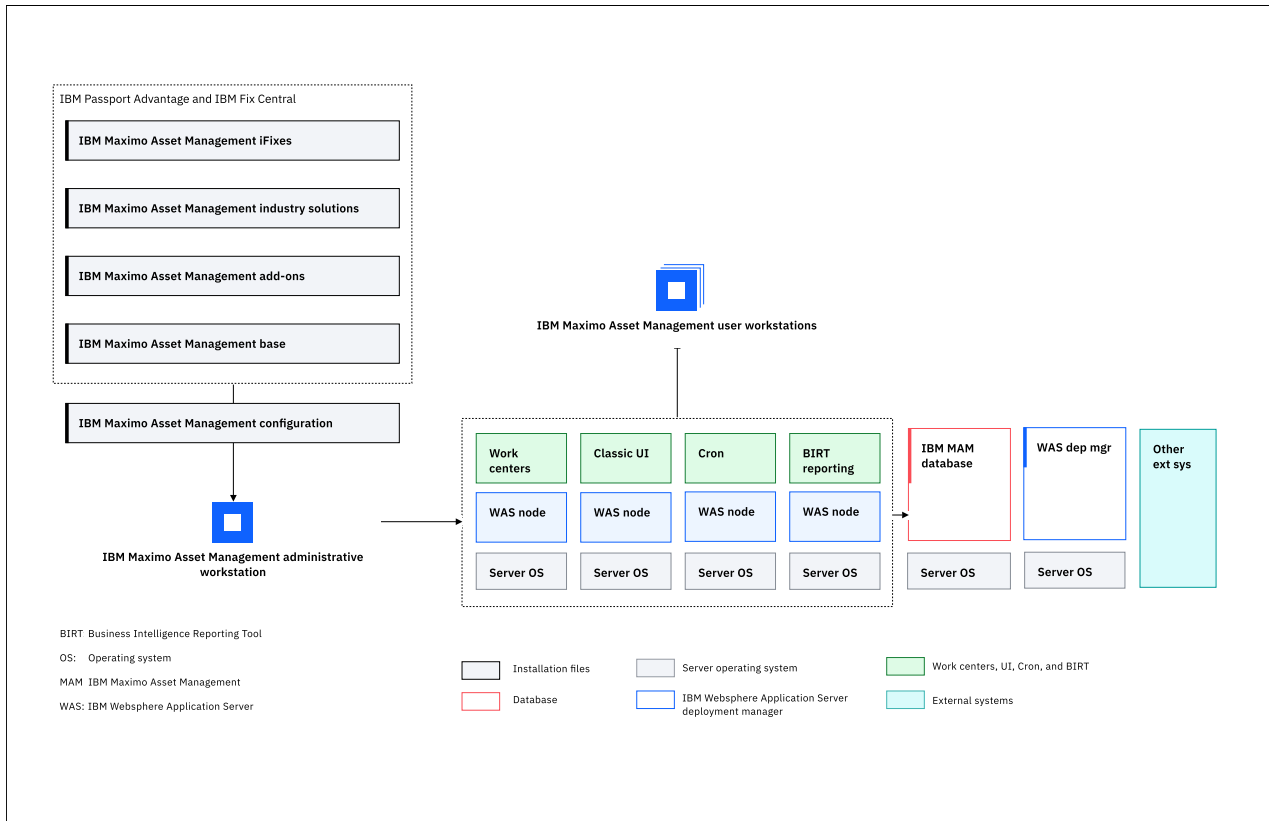


Figure 3. Maximo Asset Management 7.6 deployment architecture

Maximo Asset Management deployment

Maximo Asset Management 7.6 has an administrative workstation where the deployment software is downloaded. Deployment files are generated on the administrative workstation and deployed on an application server. The Maximo database and other external systems are run on separate servers.

In Maximo Asset Management deployment, the software is downloaded on the administrative workstation. The ear and war files are generated on the administrative workstation and deployed on an application server. In this deployment example, the cluster is running multiple servers for Classic UI, Work Centers, Cron, and BIRT reporting to distribute the load in a WebSphere Application Server Network Deployment. The Maximo Asset Management database and other external systems for integration are running on separate servers.

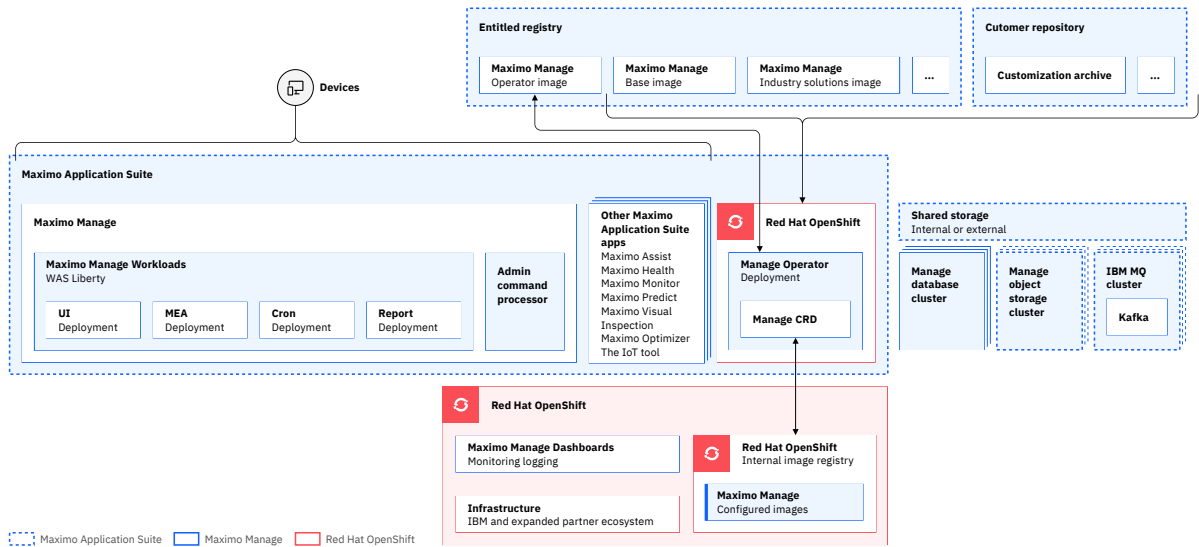
Maximo Manage deployment in Red Hat OpenShift Container Platform

When you deploy Maximo Manage in Red Hat OpenShift Container Platform, the Maximo Manage operator pulls the images from the IBM Entitled Registry and, if any customization exists, pulls the customization archive from the customer repository. The operator builds a Maximo Manage administrative image and configures images or workloads, such as UI, Cron, BIRT reporting.

The configured images and workloads are deployed to the Maximo Application Suite containers.

The built images are stored in the Red Hat OpenShift internal image registry repository.

The following diagram shows an example of a Maximo Manage application deployment in Red Hat OpenShift Container Platform:



Red Hat OpenShift

Red Hat OpenShift contains the infrastructure layer, and an internal image registry to store Maximo Manage configured images and services. Maximo Manage dashboards are provided in Red Hat OpenShift. You can use the dashboards to review logs and monitor Red Hat OpenShift and other applications that are deployed in Maximo Application Suite.

Images

The images for the Maximo Manage operator, Maximo Manage base, industry solutions, and add-ons are provided by the Entitled Registry. These images contain software application classes, deployment descriptors, XML, and scripts.

Maximo Manage operator

The Maximo Manage operator is similar to an installer. The operator pulls the images from the Entitled Registry. This process is similar to downloading software from Passport Advantage and Fix Central.

Maximo Manage workloads or server bundles

Maximo Manage supports All, UI, Cron, Report, and Maximo Enterprise Adapter workloads. You can configure the workloads during the application deployment. In the diagram, the UI, Cron, Report, and Maximo Enterprise Adapter workloads are deployed.

Customization archive

The customization archive is stored in the customer repository and can be accessed from Maximo Application Suite by using HTTPS.

Admin command processor

The admin command processor pod runs Maximo Manage tools, such as the integrity checker.

Other Maximo Application Suite applications

You can deploy other Maximo Application Suite applications, such as Maximo Collaborate, Maximo Health, Maximo Predict, and Maximo Monitor.

Maximo Manage database

You can deploy the Maximo Manage database, such as IBM Db2 Warehouse, in a cluster.

Maximo Manage object storage

You can configure IBM Cloud Object Storage for document storage. For example, you can configure doc-links to use IBM Cloud Object Storage for storing documents.

Maximo Manage IBM MQ or Kafka

You can configure a messaging provider, such as JMS, IBM MQ, or Kafka, in a cluster.

Deployment from operator

The following steps describe how Maximo Manage is deployed by using the operator:

1. Create a custom resource (CR).

An administrator selects, configures, and deploys the Maximo Manage application. The deployment creates a CR. A CR contains a user-entered configuration for the application, for example, the name, version, number of pods, database type and connection, type of workloads, and location of customization archive.

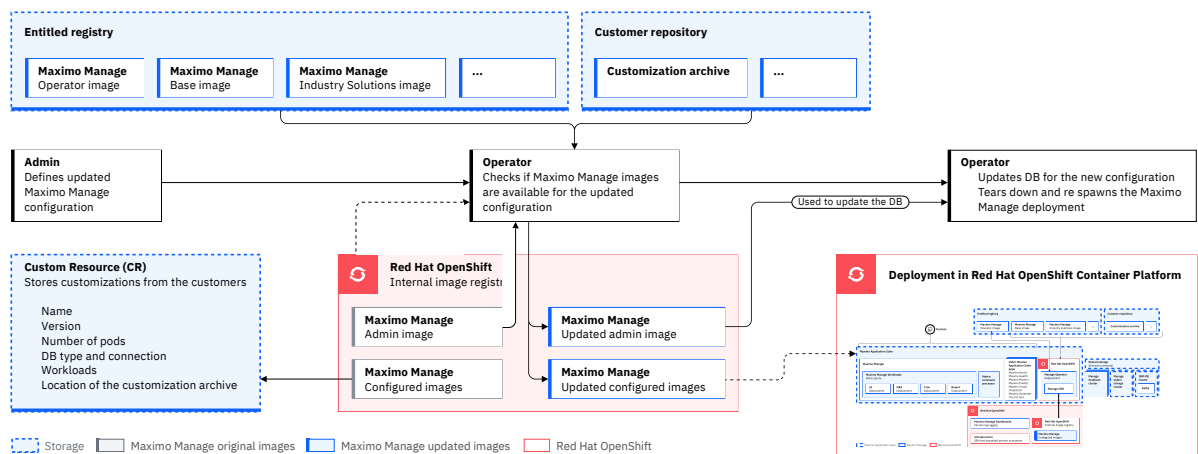
2. Create the image.

The Maximo Manage operator pulls the images from the Entitled Registry and, if any customization exists, pulls the customization archive from the customer repository. The operator deploys the industry solutions or add-on images and the customization archive over the Maximo Manage base image to create the final images. The operator also validates for the dependency matrix. The operator creates an Maximo Manage administrative image and the Maximo Manage configured images or workloads. The final images are stored in the image registry repository.

3. Update the configuration.

If the Maximo Manage database does not exist, the operator installs the Maximo Manage administrator image. If the database does exist, the operator upgrades the Maximo Manage database for a new configuration and restarts the Maximo Manage deployment. Maximo Manage configured images or workloads are deployed to containers in Maximo Application Suite.

The following diagram shows how the operator is used to deploy Maximo Manage.



Maximo Manage deployment model

In Maximo Manage deployment is done in Red Hat OpenShift clusters.

The following diagram shows a sample Maximo Manage deployment model.

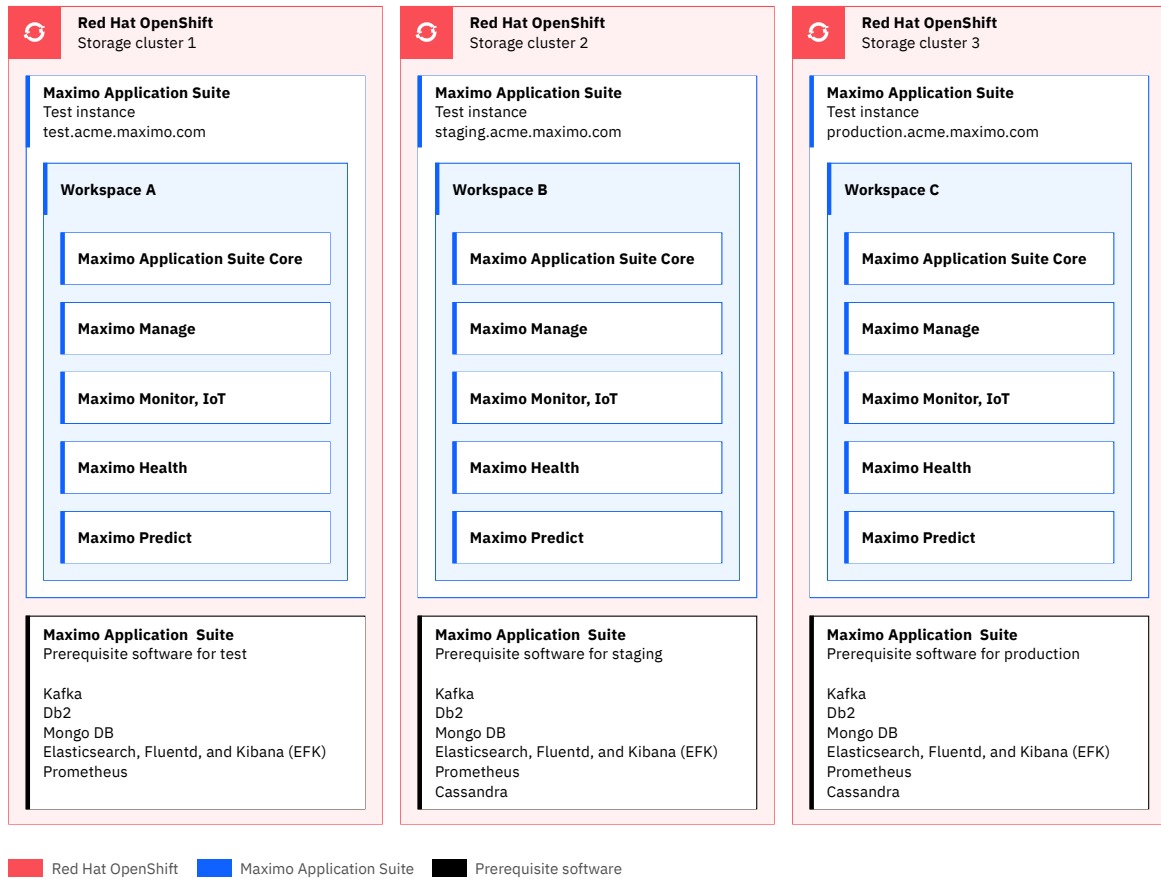


Figure 4. Maximo Manage deployment model sample

- Each cluster can run multiple Maximo Application Suite instances.
- Each instance runs its own set of pods and runtime code.
- All instances within the cluster can share a license pool.
- Each instance has its own workloads.
- Worker nodes capacity management is done per instance.
- Prerequisites stack is defined at the instance level.
- Each instance has an Red Hat OpenShift Data Foundation (ODF) storage cluster, that provides the necessary storage classes for all applications in Maximo Application Suite.

Deployment options for Maximo Application Suite

One of the main changes when you upgrade from Maximo Asset Management to Maximo Application Suite is the migration to Red Hat OpenShift Container Platform. Maximo Application Suite allows multiple ways to install and configure it.

All the applications, tools, add-ons, and utilities in Maximo Application Suite are based on three technology tiers.

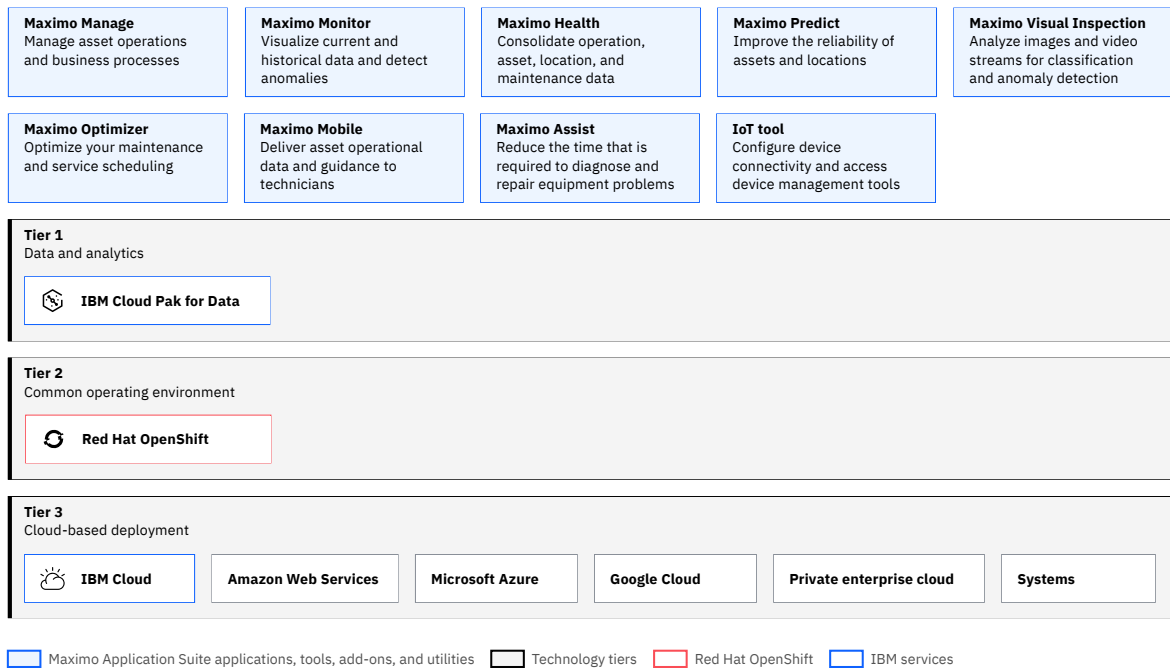


Figure 5. Maximo Application Suite

- Tier 1 for data and analytics
 - IBM Cloud Pak for Data is a data and Artificial Intelligence (AI) platform with a data fabric that makes all data available for AI and analytics.
 - IBM Watson Studio enables data scientists, developers, and analysts to build, run, and manage AI models, and optimize decisions on IBM Cloud Pak for Data.
 - IBM Watson Machine Learning provides tools and services to build, train, and deploy Machine Learning models.
- Tier 2 for a common operating environment uses Red Hat OpenShift Container Platform, which implements containerization by using Red Hat OpenShift clusters.
- Tier 3 is a cloud-based deployment to select deployment options best suited for your organization.

Note: The deployment options in the table are limited to options supported by IBM. For a full list of Maximo Application Suite installation paths, see [Supported installation paths](#).

Deployment	Procurement	Provisioning and operation	Benefits
On-premises - customer-managed	Customers buy Maximo Application Suite and use their own infrastructure.	Customers provision, manage, and operate the full technology stack.	Maximum operational flexibility
On-premises - hybrid-managed	Customers procure and manage infrastructure and application and avails PaaS services from IBM	IBM manages PaaS	PaaS services managed on both on-premises and hyperscalers

Note: The deployment options in the table are limited to options supported by IBM. For a full list of Maximo Application Suite installation paths, see [Supported installation paths](#).

(continued)

Deployment	Procurement	Provisioning and operation	Benefits
Hyperscalers – customer-managed IBM or Amazon Web Services or Microsoft Azure	Customers buy software from IBM and infrastructure from hyperscalers.	Customers run IBM-provided automation scripts to deploy Maximo Application Suite on hyperscalers' cloud.	<ul style="list-style-type: none"> • Simplifies procurement and deployment • Allows customers to select their hyperscalers. • Flexibility for customers to manage and operate their environment.
	Customers buy software and infrastructure from hyperscalers.	Customers manage and operate both software and infrastructure.	
SaaS – IBM-managed	Customers buy a single part that includes software, infrastructure, and operations from either standard IBM sales channels or Amazon Web Services marketplace.	IBM provisions, manages, and operates customers' Maximo Application Suite environment on Amazon Web Services cloud by using IBM's cloud account.	<ul style="list-style-type: none"> • Reduced time-to-value • Reduced operational costs • Allows customers to focus on their business priorities.

Related information

[Getting started with Maximo Application Suite as a Service](#)

Planning your upgrade schedule

Upgrading involves planning, preparing, executing multiple activities and tasks, and troubleshooting. Planning enough time for all necessary activities helps ease the process.

Plan enough time to complete each phase of the upgrade process:

- Determining when you can upgrade
- Planning the upgrade
- Performing pre-upgrade tasks. For more information, see [Before you upgrade](#).
- Upgrading in a test environment
- Troubleshooting your test upgrade
- Upgrading your production environment

Plan time into your schedule to perform the upgrade in a test environment. You can perform a test upgrade to test and troubleshoot your upgrade to avoid additional downtime in your production environment.

You can also schedule sufficient time to train administrators and users to use Maximo Application Suite and Maximo Manage. Team members need to understand the capabilities of the new software to participate in the upgrade planning process.

Training and courses

Maximo Application Suite is based on a different technology platform than Maximo Asset Management. IBM training can help you to upgrade your skills to manage the upgrade process, administer, and operate Maximo Application Suite and Maximo Manage.

Recommended skills for Maximo Manage administrators

Maximo Manage administrators perform the following tasks:

- Install and configure software.
- Understand Maximo Application Suite tasks, such as creating and modifying records.
- Understand relational database concepts, such as views and joins.
- Understand the Maximo Application Suite database and data relationships
- Construct Structured Query Language (SQL) statements.
- Understand the SQL syntax for your database.
- Set Maximo Application Suite properties for proper configuration.
- Define security privileges for users and groups.

Recommended skills for Maximo Manage users and developers


Maximo Manage users, depending on their access and entitlements could be customizing applications or using workflow for business processes. Users might benefit from training in the following areas:

- Customizing existing applications.
- If you are making changes or improvements to your business processes, training on the new processes.

References

To learn about	See
Red Hat OpenShift Container Platform	Certification <ul style="list-style-type: none">• Red Hat Certified System Administrator (RHCSA) certification path for new learners• Red Hat Certified Specialist in OpenShift Administration certification path for experienced learners• Containerization and RHOCP essentials for Maximo Application Suite and Sterling solutions for an introduction on containerization and Red Hat OpenShift Container Platform essentials• Maximo Application Suite & Red Hat OpenShift Container Platform Deployment Technical Essentials for information on Maximo Application Suite and Red Hat OpenShift Container Platform installation.• Maximo Application Suite deployment overview for an overview on deploying Maximo Application Suite with an interactive flow chart and a series of demonstrations.
Maximo Mobile	Self-paced training Getting started with Maximo Mobile v2 to learn about a mobile solution that keeps technicians connected and your organization productive.

To learn about	See
Maximo Manage	<p>Self-paced training</p> <ul style="list-style-type: none"> • Maximo Application Suite - Manage: Introduction for an overview of the core functional areas of Maximo Manage and an introduction to the application relationships and overall usage. • Maximo Application Suite - Manage: Core Data Setup to learn how to create a new organization and site in Maximo Manage. • Maximo Application Suite - Manage: Users and Security to learn how to create new users in Maximo Manage and manage labor, crafts, and calendar records.
IBM Cloud	<p>If you have an Red Hat OpenShift deployment on IBM Cloud , Amazon Web Services, Microsoft Azure, or Google Cloud, you can deploy IBM Cloud Pak for Data on your cluster. You can also run Cloud Pak for Data on your private, on-premises cluster. For more information, see Overview of IBM Cloud Pak for Data.</p>
Automate Maximo Application Suite installation	<p>It is possible to automate some of the manual steps of installing Maximo Application Suite and its components, using Ansible collection roles. For more information see, IBM Maximo Application Suite installation with Ansible collection.</p>
MongoDB	<p>Maximo Application Suite uses MongoDB for its data dictionary and local user management. Your MongoDB instance can run in the Red Hat OpenShift cluster or external to it. For more information, see Installing MongoDB.</p>
IBM Suite License Service	<p>IBM Suite License Service provides features for managing virtualized environments and measuring license utilization. Suite License Service discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and generates audit reports. Each report provides you with different information about your infrastructure, for example the computer groups, software installations, and the content of your software catalog. For more information, see “Suite License Service” on page 7.</p>
IBM Data Reporter Operator	<p>An operator that accepts events and transforms them into reports that are submitted to the Data Service of the IBM Metrics Operator. For more information, see “Data Reporter Operator ” on page 7.</p>

To learn about	See
IBM Cloud Pak for Business Automation	IBM Cloud Pak for Business Automation assembles certified software from the IBM Automation Platform for Digital Business on multiple cloud infrastructures. A private cloud vendor can be used as an enabling layer with a user interface and command line to limit access to members of an enterprise and partner networks. For more information, see Overview .
Migrating from an existing Maximo Asset Management implementation	Migrating from an existing Maximo Asset Management implementation, from Maximo SaaS Flex or on premises, see Migration from SaaS Flex or On-Premise .
Selenium Automation Framework V3	<p>Maximo Selenium Automation Framework V3 hosts test scripts for Maximo Asset Management 7.6.1.0 and earlier versions. For more information, see Maximo Selenium Automation Framework V3.</p> <p> Warning: The Maximo Selenium automation framework tool is not officially supported by IBM.</p>
IBM Maximo Test Automation Framework	As part of Maximo Application Suite, the IBM Maximo Test Automation Framework can be utilized to validate or re-validate Maximo Manage processes and capabilities for a given release based on approved and certified configurations from IBM. The Test Automation Framework consists of a series of validation test scripts that encompass asset and work management business processes. These test scripts contain information that can be built upon to develop and document a manufacturer's policies and procedures according to the implementation and use of Maximo Manage application software. These scripts can be used as quality assurance test cases to validate information systems. They are updated for every Maximo Application Suite long term supported release. For more information, see IBM Maximo Test Automation Framework .

Process overview

The migration process consists of an initial test deployment, which is followed by the production deployment after successful testing. Install and configure Maximo Application Suite, create a customization archive if you find it necessary, and finally deploy and activate Maximo Manage.

Note: If possible, deploy the Maximo Application Suite to multiple nonproduction environments for testing purpose before deploying the same on production environment.

The migration process supports the following elements of your Maximo Asset Management systems:

- All data
- All customizations. You must create a customization archive to store any specific changes, such as Java classes, XML files, and database scripts. You create the customization archive in a location accessible to IBM Maximo Application Suite during deployment. The structure of the customization archive is the same as the Maximo Asset Management folder structure.

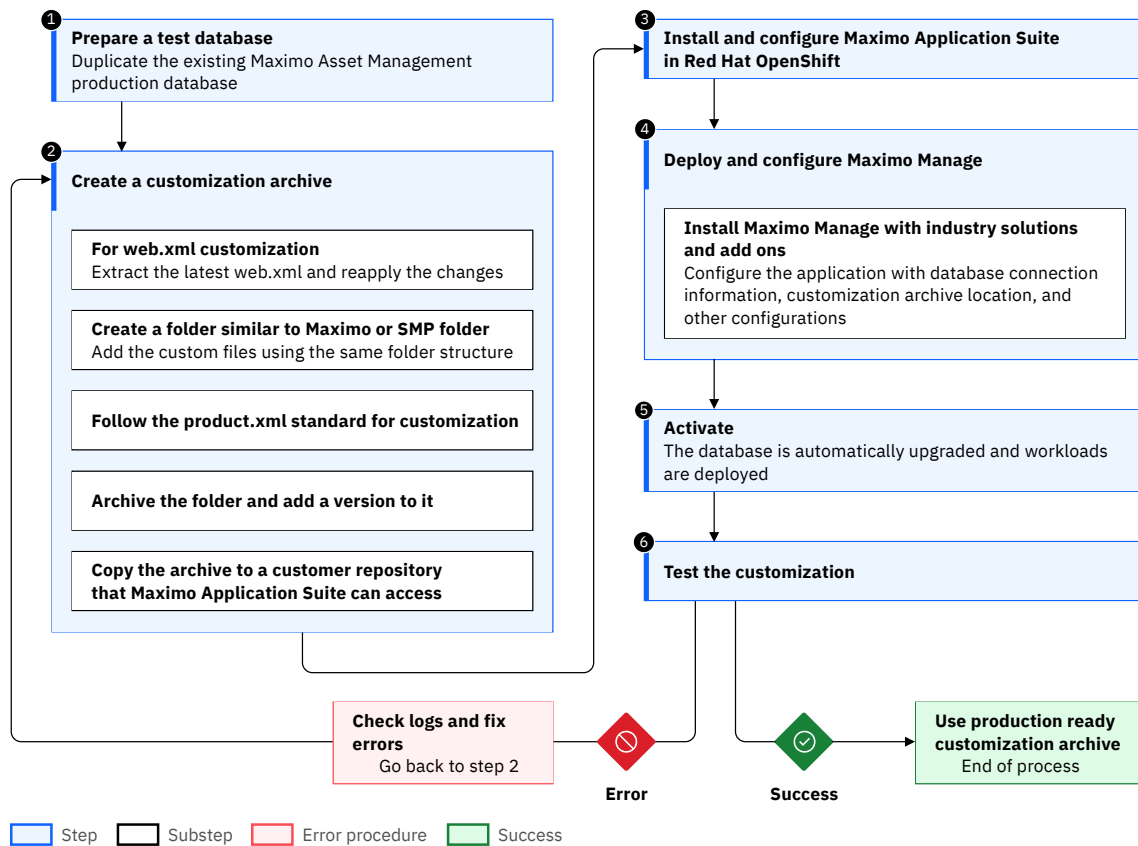
All customizations are preserved for Maximo Everyplace® during the upgrade process.

- Data model
- User interface and presentation layer
- Workflow processes
- Data validations and default values
- Escalations

When upgrading from Maximo Asset Management to Maximo Manage, the following items are not supported during the upgrade:

- The migration process does not support upgrading directly from Maximo Asset Management to Maximo Manage. You must install Maximo Application Suite first before you deploy Maximo Manage as an application within it.
- The migration process does not support migration of integration definitions specified in Maximo Asset Management. You must configure the integrations as part of the upgrade.
- Upgrading from one database platform to another. For example, you cannot upgrade from a Maximo Asset Management deployment that uses an Oracle database to a Maximo Manage deployment that uses a Db2 database.

The following diagram shows the migration process flow.



1 Prepare a test database

Duplicate the existing Maximo Asset Management production database to use as a test database.

2 Create a customization archive

For more information, see [Customization archive guidelines](#).

3 Install and configure Maximo Application Suite

Install and configure a Maximo Application Suite instance in your Red Hat OpenShift environment. For more information, see [Installing Maximo Application Suite](#).

Update all required properties and integrations after you install Maximo Application Suite. For more information, see [“Updating system settings path”](#) on page 563 and [“Integrating with external systems”](#) on page 553.

4 Deploy and configure Maximo Manage

Use the Maximo Application Suite user interface to configure Maximo Manage, industry solutions, and add-ons to use the new upgraded database, and other configurations. Specify the location of the customization archive. Maximo Manage application.

5 Activate Maximo Manage

Activate the Maximo Manage application. Activation updates the database and deploys workloads to the containers. For more information, see [Activating Maximo Manage](#).

6 Test customization

Log on to Maximo Application Suite and launch Maximo Manage. Execute any business workflow or go to any page to check if the customization was migrated from Maximo Asset Management to Maximo Manage. For more information on a specific customization scenario, see [Validating customization archive](#).

Planning to upgrade

All Maximo Manage implementations are unique, and the migration process is different for every deployment. However, some considerations in the process are common to every migration.

Related concepts

[“Planning” on page 132](#)

Upgrade checklist

You can use the upgrade checklist to check for tasks that you must do for upgrading from Maximo Asset Management to Maximo Manage.

Before you upgrade

- Ensure that you have Maximo Asset Management 7.6.0.10, 7.6.1.2, or 7.6.1.3 installed. For more information, see [Installing Maximo Asset Management](#).

Note: For IBM Control Desk, ensure that you have IBM Control Desk 7.6.1.5 and Maximo Asset Management 7.6.1.3 installed, to upgrade to Maximo Manage and Maximo IT.

- Plan for Maximo Application Suite installation requirements and preferences. For more information, see [Planning](#).
- Install and configure Red Hat OpenShift cluster for non-production and production environments, according to your requirements. For more information, see [Installing Red Hat OpenShift Container Platform](#).

Note: To know more about installing on Amazon Web Services, IBM Cloud , Microsoft Azure, or using a command-line interface, see [Supported installation paths](#).

- Become familiar with authentication, encryption and security, and SMTP configuration methods. For more information, see [Authentication and security](#) and [SMTP configuration](#).
- Plan your upgrade schedule. For more information, see [Planning your upgrade schedule](#).
- Review current database settings. For more information, see [Database settings](#).
- If you are using Oracle Connector or SAP Connector, prepare it in Maximo Asset Management before you upgrade. For more information, see [Preparing Oracle and SAP Connector before upgrade](#).
- Check for industry solutions and add-ons compatibility. For more information, see [Deployment of industry solutions and add-ons](#).
- Disable custom triggers in any table of your database. For more information, see [Custom triggers](#).
- Commit any database configuration changes that are pending. For more information, see [Configuring the database](#).
- Backup your production database. For more information, see [Backups](#).
- Prepare a test database as a duplicate of the Maximo Asset Management production database. For more information, see [Backups](#).
- Complete any post-installation tasks for Maximo Asset Management before you upgrade. For more information, see [Post installation tasks](#).
- Create a customization archive to store specific changes, such as Java classes, XML files, and database scripts. For more information, see [Migrating customizations using customization archive](#).
 - Create deployment descriptors, as needed. A deployment descriptor describes how a component, application, or module is to be deployed with specific security settings, container options, and configuration requirements. For more information, see [Deployment descriptors](#).
- Run Integrity checker in Maximo Asset Management 7.6.1.0, 7.6.1.2, or 7.6.1.3 and fix all errors reported. For more information, see [Running Integrity checker](#).
- Test the upgrade in testing environment. For more information, see [Testing the upgrade](#)
- Use Maximo Manage logs to check and fix any errors you may run into while upgrading. For more information, see [Troubleshooting the upgrade using Maximo Manage logs](#).

Tip: Keep your testing environment with application servers started, so that you can better measure the actual downtime it will take considering the production environment.

- Stop the application server only when you are nearing completion of Maximo Manage deployment as part of Maximo Application Suite.

During upgrade

- Install and configure Maximo Application Suite. For more information, see [Installing Maximo Application Suite](#).
- Configure SMTP in Maximo Application Suite before you create the admin user. Otherwise, the email with the generated password cannot be sent. For more information, see [SMTP server](#).
- Create and log on as the admin user in Maximo Application Suite. For more information, see [Administering users and user access](#).
- Prepare Maximo Manage for deployment. For more information, see [Preparing to upgrade](#).
- Deploy Maximo Manage. For more information, see [Deploying in Maximo Application Suite](#).
- Set server bundle properties in Maximo Application Suite. For more information, see [Adding server bundle properties](#).
 - Set `mxe.oslc.webappurl` to point to the route address for each server bundle, by using bundle level properties.
- Activate Maximo Manage. For more information see [Activating Maximo Manage](#).
- Check Maximo Manage deployment status. For more information, see [Checking Maximo Manage deployment status](#).
- Update system settings like `mxe.doclink.path01` and others with the new path. For more information, see [“Updating system settings path” on page 563](#).
- If the database requires an SSL connection, you must obtain the certificate for the database. For more information, see [Obtaining SS certificate for database](#).
- Import additional certificates as needed in Maximo Application Suite. For more information, see [Importing additional certificates in Maximo Application Suite](#).
- Enable monitoring in Red Hat OpenShift cluster. For more information, see [Installing Logging](#).
- Check Maximo Manage logs for any errors and fix them. For more information, see [Troubleshooting using Maximo Manage logs](#).
- Check for user synchronization from Maximo Application Suite to Maximo Manage. For more information, see [Managing users post upgrade](#).
- Configure Oracle Connector if you are using it. For more information, see [Configuring Oracle Connector after upgrade](#).
- Configure SAP Connector, if you are using it. For more information, see [Configuring SAP Connector after upgrade](#).
- **Tip:** The following points are optional as per your specific requirements.
- Maintain attached documents for your applications in persistent storage or cloud object storage, that is, S3. For more information, see [Configuring attached documents](#).
- Use the External Systems application in Maximo Manage to initiate export and import of data, for example, integrate data by using files. For more information, see [Exporting and importing file-based data](#).
- Use XSL maps to transform messages for outbound transactions in the integration framework provided with Maximo Manage. For more information, see [XSL mapping](#).

Upgrade prerequisites

Check your Maximo Asset Management version and get your environment ready for upgrading by configuring your Red Hat OpenShift cluster.

Product version

Before you upgrade to Maximo Manage, you must have Maximo Asset Management 7.6.0.10 or 7.6.1.2 or later installed.

Note: For IBM Control Desk, ensure that you have IBM Control Desk 7.6.1.5 and Maximo Asset Management 7.6.1.3 installed, to upgrade to Maximo Manage and Maximo IT.

Red Hat OpenShift cluster requirements

Prepare the Red Hat OpenShift cluster based on the following requirements:

1. Determine the capacity needed to upgrade the product.

For more information, see [Planning](#).

2. Secure a Red Hat OpenShift cluster for development, testing, or production.

You must ensure that the prerequisites to install Maximo Application Suite are in place.

Related concepts

[“Prerequisite software” on page 5](#)

Requirements and capacity planning

Use the IBM Maximo Application Suite sizing calculator to estimate the required sizing for your planned deployment.

The Maximo Application Suite sizing calculator is used to estimate your Red Hat OpenShift worker node configuration requirements, storage requirements, and memory requirements.

1. Download the calculator.

- [Sizing calculator for Maximo Real Estate and Facilities 9.1](#)
- [Sizing calculator for 9.0.1](#)
- [Sizing calculator for 9.0](#)
- [Sizing calculator for 8.11 and earlier](#)

2. Select or enter values for the yellow fields to match your planned application deployment.

The calculator provides estimated total system requirements in VPCs and Memory (GB) for your configuration in the Resulting Complete Environments Requirements section of the Output table.

Important: The information in this document represents the minimum resources that you need to successfully install Maximo Application Suite. A minimum of 300GB of storage per worker node is recommended for Maximo Application Suite build process. You might need more resources to support your specific workload. If needed, work with your IBM Sales representative to generate more accurate calculations based on your expected workload.

For more information, see [Sizing guidance](#).

Validating the upgrade process

Before you attempt the actual upgrade in a production environment, you could validate the upgrade process on multiple testing environments.

Procedure

1. In a test environment, check the logs for errors. For more information, see [Troubleshooting the upgrade using Maximo Manage logs](#).
 - a) If errors are found, fix them. If the problem is related to customization, use the admin image container in the Red Hat OpenShift environment to copy the entire build directory to a local development computer with the customization and compile.

- b) Create the customization archive again with the updated code, redeploy, and reactivate.
2. After successful testing, complete the upgrade by deploying in a production environment.
 - a) Configure a Red Hat OpenShift cluster.
 - b) Install Maximo Application Suite and all prerequisites.
 - c) Get all production configuration, database configuration, server bundles configurations, and customization archive, ready if they exist.
 - d) Deploy and activate Maximo Manage.

Preparing to upgrade

Before you upgrade, you must understand how authentication and security are implemented and how SMTP is configured. You prepare connectors and migrate customizations by using customization archives. Run the Integrity Checker utility to check for database errors before you start the upgrade process.

Reviewing database settings and backups

Before you start the upgrade, review your current database settings, disable custom triggers, commit any pending configuration changes, and backup the database and other important system files.

Review database settings

To ensure a successful upgrade, compare the configuration settings of your existing database with the default configuration settings of Maximo Manage. If your current values do not sufficiently match the default settings in Maximo Manage, it might cause problems during the upgrade process. Set configuration parameters that are equal to or greater than those parameters that required for Maximo Manage. For more information, see [Preparing your database for deployment](#).

Disable custom triggers

Disable all custom triggers that exist for any table in your Maximo Asset Management database, for example, stored procedures, triggers, views, and synonyms. The upgrade process does not re-create or remove these objects. Reapply custom triggers after your database is upgraded as part of activating Maximo Manage.

Commit any database configuration changes

Commit any configuration changes to the Maximo Asset Management database before you upgrade to Maximo Application Suite. Configuration changes were part of the postinstallation tasks that were required when you installed Maximo Asset Management.

To confirm that all changes are committed, run the following SQL query against the Maximo Asset Management database:

```
SELECT count(*) from maxobjectcfg where changed != 'N'  
SELECT count(*) from maxsysindexes where changed != 'N'
```

'N' indicates that a change is committed. If any positive row count values are returned for the query, you must apply or discard the configuration changes. Alternatively, you can use the **configdb.bat** command to commit configuration changes. For more information, see [Configuring initial data](#).

Back up the existing database and other files

Back up the existing Maximo Asset Management database, the contents of the Maximo Asset Management installation folder, and the deployment engine. If a failure occurs during upgrade, you might be required to restore the Maximo Asset Management database. Some upgrade tasks cannot be rolled back after they are committed to the database. If you have a backup of the database, you can restore your environment. By default, files are found in the C:\ibm\smp directory. Backing up this directory can be useful if you must rebuild Maximo Asset Management EAR files. Back up the deployment engine registry as described in the Maximo Asset Management installation information.

For more information, see [Backup and restoration for WebSphere Application Server](#) and [Backup and restoration for Oracle WebLogic Server](#).

Related information

Authentication and security

Before you upgrade to Maximo Manage, you must be familiar with authentication, encryption and security, and SMTP configuration.

Authentication

- If you select local authentication, the username and password of users are stored on MongoDB and Maximo Application Suite directly authenticates the users.
- LDAP or SAML can be configured for authentication before or after the installation of Maximo Manage by using the **Configuration** page in Maximo Application Suite. For more information, see [“Authentication methods” on page 605](#).
 - Only LDAP registries that are supported by Liberty runtime can be used. For more information, see [LDAP User Registry 3.0](#)
 - If you are currently using LDAP or SAML, your existing configuration can be used in Maximo Application Suite.
 - From Maximo Application Suite 8.8 onwards, Maximo Manage supports only API key-based authentication for integration with external applications and REST API transactions. For integration, XML along with SOAP and HTTP protocols use API keys. The existing REST APIs, for example, **maxrest** or **rest** support API keys as well as new REST APIs added from Maximo Application Suite 8.8 like **oslc**.

API key-based authentication is primarily used for machine to machine interactions and authentication. If you are using **maxauth** in Maximo Asset Management, after upgrading to Maximo Application Suite and deploying Maximo Manage, use API keys because **maxauth** is not supported in Maximo Manage. For more information, see [API keys application](#).

Encryption and security

If you are using custom encryption keys for CRYPTO and CRYTOX attributes, then the custom keys must be provided during the Maximo Manage application installation.

Custom encryption properties are specified in the `maximo.properties` file in Maximo Asset Management.

Note: The `maximo.properties` file is not used in Maximo Application Suite.

After the Maximo Manage application deployment, if encryption keys are not specified when you activate Maximo Manage with a fresh database, new encryption keys are automatically generated. Set the **autoGenerateEncryptionKeys** property to false if you do not want to generate the keys automatically. For more information, see [Disabling automatic generation of encryption keys and Database encryption](#).

SMTP configuration

SMTP configuration is retained in the database during the upgrade.

It is used by Maximo Manage for sending emails.

Maximo Application Suite also provides SMTP configuration. This configuration is used for sending welcome emails and user password emails. For more information, see [Configuring SMTP](#).

Preparing enterprise adapters before upgrade

If you are using IBM Maximo Enterprise Adapter for Oracle Applications or IBM Maximo Enterprise Adapter for SAP Applications, you must prepare the connector before you upgrade.

Procedure

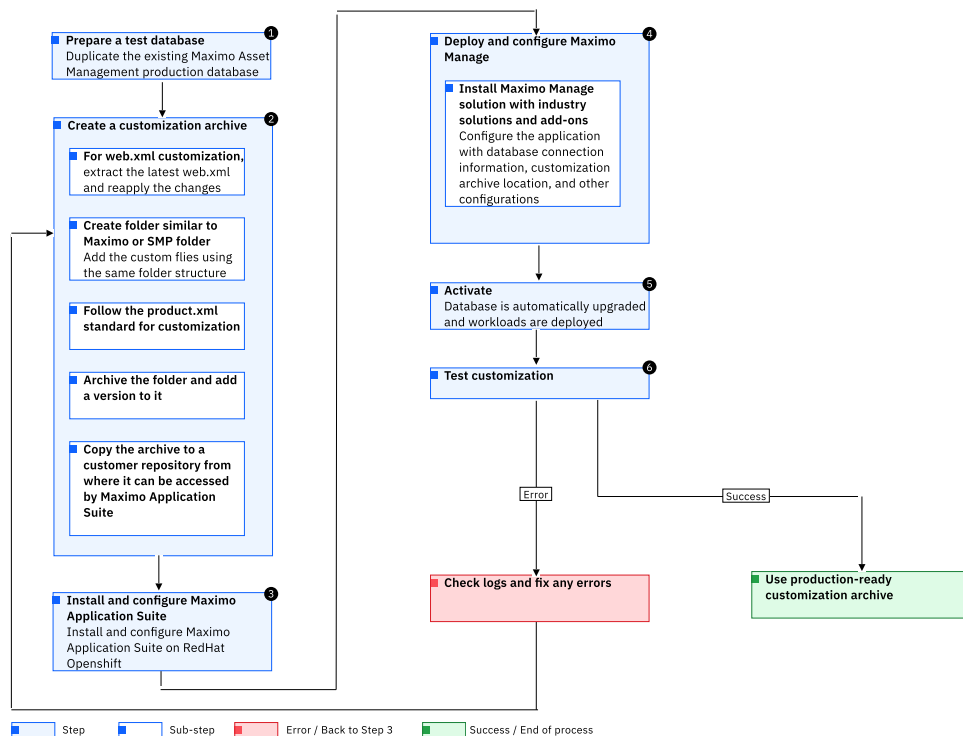
1. Stop all transactions between Maximo Asset Management and IBM Maximo Enterprise Adapter for Oracle Applications or IBM Maximo Enterprise Adapter for SAP Applications.
2. Process all transactions in the interface tables and integration queues.
3. Resolve any Oracle integration errors in Maximo Asset Management. For more information, see [Error management](#).
4. In the External Systems application, disable the external system for the enterprise adapter.
 - The OA12 external system for Oracle.
 - The SAP2005 external system for SAP Connector
5. If you are preparing an Oracle connector, back up PL/SQL user exits and any customization on the Oracle e-business suite.
6. Stop the Maximo server.
7. Back up the Maximo database.

Migrating customizations using customization archives

All customer-specific changes, such as Java classes, XML files, and database scripts, must be included in a customization archive. You create the customization archive in a location accessible to IBM Maximo Application Suite during deployment. The structure of the customization archive is the same as the Maximo Asset Management folder structure. Test the customization archive in a development or test environment before you apply it to the production environment.

About this task

The following diagram shows the customization process:



Procedure

1. Prepare the database.

Prepare a test database as a duplicate of the existing Maximo Asset Management production database.

2. Create a customization archive.

For more information, see [Creating customization archives](#) section. If your customization includes a web.xml file, such as customer servlet, filters, changes in the order of the servlet startup, context parameters, or session timeout:

- a) Install Maximo Manage with industry solutions and add-ons without customization on an empty database.
- b) Extract the web.xml file.
Use the **oc rsync** command to retrieve the web.xml file.
- c) Apply your changes.
- d) Copy the web.xml file to the customization archive in the appropriate directory.

Note: Ensure that the location of the customization archive is accessible by the Red Hat OpenShift.

3. Deploy the application.

Use Maximo Application Suite to configure Maximo Manage, industry solutions and add-ons to point to the database to upgrade and other configurations. Specify the location of the customization archive. Deploy.

For more information, see [Setting up a local Maximo Manage development environment](#).

4. Activate the application.

Maximo Manage updates the database and deploys workloads to the liberty containers.

5. Test the application.

6. Using the admin image pod that contains maxinst, in the Red Hat OpenShift console, copy the entire build directory with the customization and compile. You can use IDE to build the Maximo Manage project with customizations.

What to do next

After you fix any errors, create a customization archive again with the updated code. Deploy and activate the Maximo Manage application. After successful testing, use the customization archive in the production environment.

Creating customization archives

You can create a customization archive, which is a set of files that contain changes and customizations for the application.

Procedure

1. Follow the existing SMP or Maximo folder structure to create a customization folder.
2. Copy class files (Java customization) in the appropriate directory and extracted with the existing Maximo binary files in the correct package or module hierarchy.

The file might include the following information.

- Maximo Business Object (MBO) customization
- Field validation
- Maximo integration framework
- Others like customer-specific code

What to do next

- Version your customization archive.

- Customization archive location is specified during Maximo Manage application configuration from IBM Maximo Application Suite admin dashboard.
- Customization archive location must be accessible from Maximo Application Suite. It is a zip file that you can access through HTTP, HTTPS, FTP, or FTPS.
- Any database scripts can be added to customer directory and an `a_customer.xml` file must be used. For more information, see [Customization archive guidelines](#). All those files need to be part of the customization archive. Update db will run after deployment to apply all customer scripts.

Deployment descriptors

All deployment descriptors such as `web.xml`, `ejb-jar.xml`, and `webservices.xml` files that are customized by users can be put into the customization archive and overlaid on the files that are supplied by IBM.

- The following files need to be created:
 - For deployment of full Maximo Asset Management in one deployment (Maximo **-all**)
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maximouiweb\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maximo-x\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maximo-x\webmodule\WEB-INF\web-guest.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maximo-x\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maxrestweb\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\maximo-x\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\mboweb\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\meaweb\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-all\mboejb\ejbmodule\META-INF\ejb-jar.xml`
 - UI deployment
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-ui\webmodule\WEB-INF\web.xml`
 - Maximo Enterprise Adapter
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-mea\meaweb\webmodule\WEB-INF\web.xml`
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-mea\mboejb\ejbmodule\META-INF\ejb-jar.xml`
 - Report deployment
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-report\webmodule\WEB-INF\web.xml`
 - Cron deployment
 - `deployment\was-liberty-default\config-deployment-descriptors\maximo-cron\webmodule\WEB-INF\web.xml`
- IBM-specific file `ibm-ejb-jar-bnd.xml` is not used in Liberty deployment.

XSL customization

XSL customizations that are part of the Maximo EAR or WAR files are also part of the customization archive and are copied to final images. For more information, see [Rule-based customization](#).

Sample customization archive

A customization archive contains Java classes, XML files, scripts, servlet, and deployment descriptors.

A sample [customization archive](#) has the following elements.

- classes
 - applications\maximo\businessobjects\classes\cust\app\asset\Asset.class
 - applications\maximo\businessobjects\classes\cust\app\asset\AssetSet.class
 - applications\maximo\businessobjects\classes\cust\app\asset\FldAssetNewField.class
- product xml
 - applications\maximo\properties\product\a_customer.xml
- script
 - tools\maximo\en\cust\V7612_01.dbc
- servlets
 - applications\maximo\commonweb\classes\com\ibm\tivoli\maximo\oslc\provider\MYPingServlet.class
- deployment descriptors (web.xml)
 - deployment\was-liberty-default\config-deployment-descriptors\maximo-mea\meaweb\webmodule\WEB-INF\web.xml
 - deployment\was-liberty-default\config-deployment-descriptors\maximo-ui\meauwebmodule\WEB-INF\web.xml

Adding third-party JAR files

You can add third-party Java Archive files to the Manage lib folder in the customization archive to migrate any extended third-party functionality added in Maximo Asset Management to Maximo Manage.

About this task

Use the following steps to add a third-party JAR file name to the Manifest file, if the deployment has an all or mea bundle server type.

Procedure

1. Go to the admin pod, which is the name of the pod that contains **maxinst**, terminal and get the file.
2. Go to /opt/IBM/SMP/maximo/deployment/was-liberty-default deployment folder.
 - a) Get the maximo-all.xml file if the deployment is all bundle server type. In maximo-all.xml, go to the maximo.businessobjectclasspath property name and add the JAR file name in the path. When deployed it will update the classpath in the Manifest file.
 - b) Get the buildmaximomea-ear.xml file if the deployment has a mea bundle server type and update the classpath.
3. Copy the updated file in the Customization Archive folder in the same path and archive or compress the folder.

The file path for the **maximo-all.xml** file is <localdrive>\custasset_bin\deployment\was-liberty-default.
4. Deploy the customization archive and activate the changes.

For more information, see [Customizing the application](#).

Running Integrity Checker before upgrade

The Integrity Checker is a database configuration utility that you can use to assess the health of the base layer data dictionary.

About this task

You can use the Integrity Checker utility to ensure that the Maximo Asset Management database is ready for upgrade. When run in Report mode, the Integrity Checker utility checks the current database and reports errors. If the Integrity Checker reports an error, you must resolve it by running the Integrity Checker in Repair mode.

Procedure

1. On the system where Maximo Asset Management is installed, open a command prompt and change directory to the tools directory.
For example, `install_home\maximo\tools\maximo`.
2. Start the Integrity Checker utility by issuing the **integrityui.bat** command.
3. Select the **Check Integrity** tab.
4. Run the Integrity Checker in Report mode.
 - a) Ensure that the **Repair Mode** check box is cleared and then click **Run Integrity Checker**.
 - b) When the report dialog box appears, click **OK**.
Results are found in the `install_home\maximo\tools\maximo\log` directory in the file that is defined in the **Log File Name** field of the **Check Integrity** pane.
5. Optional: If any errors are reported, run the Integrity Checker in Repair mode.
 - a) Select the **Repair Mode** check box and then click **Run Integrity Checker**.
 - b) When the report dialog box appears, click **OK**.
 - c) Change directory to `install_home\maximo\tools\maximo\` and then run the **configdb** command.
For more information on Integrity checker warning and error messages, see [Integrity checker messages](#).



Attention: Although the Integrity Checker can repair many issues, you might need to resolve some errors manually by consulting the log files or opening a case with the IBM support team.

What to do next

Check the log file to ensure that all reported items are repaired. If further manual intervention is required, you must resolve the errors and then rerun the Integrity Checker in Report mode. Repeat the process until no more errors are reported.

Global property values

Set the **mxe.int.globaldir** and other properties when you upgrade from Maximo Asset Management to Maximo Manage.

If you want directories to exist for the doc-link table and for integration, check that **mxe.int.globaldir** and other URL properties are set. Check whether you need to mount additional persistent volumes or change directories for transient data.

For more information, see [Troubleshooting global property values](#).

Installing Maximo Application Suite

You must install Maximo Application Suite when you upgrade from Maximo Asset Management. Maximo Application Suite has multiple supported installation paths.

For more information, see [Installing Maximo Application Suite](#).

Setting up your database

Before you can deploy IBM Maximo Manage, you must configure your database and determine how your database is encrypted.

Database deployment best practices

The best practices to create and configure your database guides and helps you to set up your database before deploying Maximo Manage

Oracle Database

For more information, see [Best practices for system performance](#).

Configuring database instances

Maximo Manage supports multiple databases. You must configure your database and gather the information when you deploy the application.

Before you begin

If you intend to use the database instance for production, you can select **Deploy database on dedicated nodes** and adjust the number of nodes as needed on the **Configure** page.

When you are considering the size the database, remember the following guidelines:

- The storage class that you select can affect the size of your database. For example, if you plan to use Red Hat OpenShift Container Storage, it increases the storage space that you need for your database.
- Be aware of space requirements for services, such as IBM Cloud Pak for Data. If you deploy in a clustered environment, IBM Cloud Pak for Data space requirements change.
- Managed versus on-premises services, such as MongoDB and Kafka, impact your disk space needs.
- All-in-one pod bundles have different disk space requirements when compared to configuring separate resources for UI, report, cron, Maximo integration framework, and Maximo Mobile.
- The number of concurrent users affect resource needs. For example, for the UI resource, configure a Java virtual machine (JVM) that has two cores to support 50–75 concurrent users.

About this task

Depending on the type of database, you must configure one of the following databases:

- [Configuring IBM Db2](#)
- [Configuring IBM Db2 Warehouse](#)
- [Configuring Oracle Database](#)
- [Configuring Microsoft SQL Server](#)

For information about database version compatibility, see [Maximo Application Suite detailed system requirements](#).

Configuring IBM Db2

Before you deploy Maximo Manage in Maximo Application Suite, configure IBM Db2 for use by Maximo Manage.

Before you begin

Before you configure the database, install and deploy it. For more information, see [“Configuring database instances” on page 316](#).

For information about Db2 version compatibility, see [Maximo Application Suite detailed system requirements](#).

See the following guidelines on how to configure a database instance.

- Configure separate system, data, and backup storage when you create a Db2 instance.
- Increase the maxsequence cache to 50.
- Run **REORG INDEXERS/TABLES** and **RUNSTATS** daily.
- Separate system storage, user storage, backup storage, transaction logs storage, and temporary table space storage on different disks.
- Maximo Manage requires row-organized tables. By default, the IBM Db2 Warehouse database setting uses column-based table organization. Update the setting as needed.
- Maximo Manage does not support Massively Parallel Processing (MPP) or table partitioning. Archive records that are over a year old. InfoSphere Optim Data Growth Solution can be used for archiving. For more information, see [IBM Maximo Archiving 7.5.1 for IBM Maximo Asset Management](#) .
- An issue can occur when you load a large amount of data by using the Maximo Integration Framework. Increase the concurrently running statements that are allowed for a Db2 application. For more information, see [How many concurrently running statements allowed for a Db2 Java application and how to increase it?](#).
- If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:
 - Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
 - Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
 - Secure operator versions and disable auto-update for all components.
 - Use the IAM and LDAP service instead of the default authentication methods.
 - For more information, see [DB2 Performance Insight](#).

Note:

- Starting in 9.0.5 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x architecture, Db2 is not configured.
- Starting in 9.0.12 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM Power (ppc64le) architecture, internalDb2 is not configured. Db2 can be used as an external service.

For information about supported database versions, generate a Software Product Compatibility Report. For more information, see [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. On the **Supported Software** tab of the report, check for the supported database versions.

Configure your database with the following operating systems:

- Linux or UNIX
- Microsoft Windows

About this task

The commands in this task can be used to configure a Db2 database outside of the Red Hat OpenShift cluster, by using different operating systems such as Microsoft Windows, Linux , or UNIX .

Note: The commands in this task are not applicable for the configuration of Db2 Warehouse instance.

The commands in this task are examples of the commands that you must run. For example, maxdb80 is the name of the database. If maxdb80 is not your database name, ensure that you replace all instances with the correct database name.

Procedure

1. Log in to the system as a user that has administrative permissions.

2. If system users do not exist on the system, create the system users.

- Windows
 - db2admin
 - maximo
- Linux or UNIX
 - maximo for the Maximo database user
 - ctgfenc1 for the Db2 fenced user
 - ctginst1 for the Db2 instance owner

The most used administrative user is assigned the primary group of the instance owner to complete some of the following steps.

3. At the Db2 installation directory, set up the command-line environment.

- For Windows , run the following command:**db2cmd**
- For Linux or UNIX , ensure that the /opt/ibm/db2/V11.5/bin, /opt/ibm/db2/v11.5/instance, and /opt/ibm/db2/V11.5/adm directories are added to your PATH.

4. Run the following commands to create the database instance.

- Windows

Where *<administrator_password>* with the Db2 administrator password.

```
db2icrt -s ese -u db2admin,<administrator_password> -r 50005,50005 ctginst1
set db2instance=ctginst1
db2start
db2 update dbm config using SVCENAME 50005 DEFERRED
db2stop
db2set DB2COMM=tcPIP
db2start
```

- Linux or UNIX

```
db2icrt -s ese -u ctgfenc1 -p 50005 ctginst1
./home/ctginst1/sqllib/db2profile
db2start
db2 update dbm config using SVCENAME 50005 DEFERRED
db2stop
db2set DB2COMM=tcPIP
db2start
```

5. Run the following commands to create the database:

```
db2 create db 'maxdb80' ALIAS 'maxdb80' using codeset UTF-8 territory US pagesize 32 K
db2 connect to 'maxdb80'
db2 GRANT DBADM ON DATABASE TO USER db2admin (windows only)
db2 GRANT SECADM ON DATABASE TO USER db2admin (windows only)
db2 connect reset
```

6. Run the following command according to your operating system and bit size:

Operating system	Command
32-bit Microsoft Windows	db2 update db cfg for maxdb80 using MAXFILOP 32768 DEFERRED #32-bit Windows
64-bit Windows	db2 update db cfg for maxdb80 using MAXFILOP 65335 DEFERRED #64-bit Windows
32-bit UNIX	db2 update db cfg for maxdb80 using MAXFILOP 30720 DEFERRED #32-bit UNIX
64-bit UNIX	db2 update db cfg for maxdb80 using MAXFILOP 61440 DEFERRED #64-bit UNIX

7. Run the following commands to configure the database:

```
db2 update db cfg for maxdb80 using SELF_TUNING_MEM ON
db2 update db cfg for maxdb80 using APPGROUP_MEM_SZ 16384 DEFERRED
db2 update db cfg for maxdb80 using APPLHEAPSZ 2048 AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using AUTO_MAINT ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_TBL_MAINT ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_RUNSTATS ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_REORG ON DEFERRED
db2 update db cfg for maxdb80 using AUTO_DB_BACKUP ON DEFERRED
db2 update db cfg for maxdb80 using CATALOGCACHE_SZ 800 DEFERRED
db2 update db cfg for maxdb80 using CHNGPGS_THRESH 40 DEFERRED
db2 update db cfg for maxdb80 using DBHEAP AUTOMATIC
db2 update db cfg for maxdb80 using LOCKLIST AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using LOGBUFSZ 1024 DEFERRED
db2 update db cfg for maxdb80 using LOCKTIMEOUT 300 DEFERRED
db2 update db cfg for maxdb80 using LOGPRIMARY 20 DEFERRED
db2 update db cfg for maxdb80 using LOGSECOND 100 DEFERRED
db2 update db cfg for maxdb80 using LOGFILSIZ 8192 DEFERRED
db2 update db cfg for maxdb80 using SOFTMAX 1000 DEFERRED

db2 update db cfg for maxdb80 using PCKCACHESZ AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using STAT_HEAP_SZ AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using STMTHEAP AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using UTIL_HEAP_SZ 10000 DEFERRED
db2 update db cfg for maxdb80 using DATABASE_MEMORY AUTOMATIC DEFERRED
db2 update db cfg for maxdb80 using AUTO_STMT_STATS OFF DEFERRED
db2 update db cfg for maxdb80 using STMT_CONC LITERALS DEFERRED
db2 update alert cfg for database on maxdb80 using db.db_backup_req SET THRESHOLDSCHECKED
YES
db2 update alert cfg for database on maxdb80 using db.tb_reorg_req SET THRESHOLDSCHECKED YES
db2 update alert cfg for database on maxdb80 using db.tb_runstats_req SET THRESHOLDSCHECKED
YES
db2 update dbm cfg using PRIV_MEM_THRESH 32767 DEFERRED
db2 update dbm cfg using KEEPFENCED NO DEFERRED
db2 update dbm cfg using NUMDB 2 DEFERRED
db2 update dbm cfg using RQIOBLK 65535 DEFERRED
db2 update dbm cfg using HEALTH_MON OFF DEFERRED
db2 update dbm cfg using AGENT_STACK_SZ 1000 DEFERRED
db2 update dbm cfg using MON_HEAP_SZ AUTOMATIC DEFERRED
db2set DB2_SKIPINSERTED=ON
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_MINIMIZE_LISTPREFETCH=Y
db2set DB2_EVALUNCOMMITTED=YES
db2set DB2_FMP_COMM_HEAPSZ=65536
db2set DB2_SKIPDELETED=ON
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
```

8. For Linux or UNIX , log in to the system.

For example, log in as the ctginst1 user and then restart the Db2 command-line environment:

```
su - ctginst1
db2
```

9. Run the following command to stop the database:

```
db2stop force
```

10. Run the following command to start the database:

```
db2start
```

11. Run the following command to reconnect to the database:

```
db2 connect to 'maxdb80'
```

12. Run the following commands to create a buffer pool:

```
db2 CREATE BUFFERPOOL MAXBUFPOOL IMMEDIATE SIZE 4096 AUTOMATIC PAGESIZE 32 K
db2 CREATE REGULAR TABLESPACE MAXDATA PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
db2 CREATE TEMPORARY TABLESPACE MAXTEMP PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
BUFFERPOOL MAXBUFPOOL
db2 CREATE REGULAR TABLESPACE MAXINDEX PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
db2 GRANT USE OF TABLESPACE MAXDATA TO USER MAXIMO
```

13. Run the following command to create the schema:

```
db2 create schema maximo authorization maximo
```

14. Run the following commands to grant authority to the Maximo user:

```
db2 GRANT
DBADM, CREATETAB, BINDADD, CONNECT, CREATE_NOT_FENCED_ROUTINE, IMPLICIT_SCHEMA, LOAD, CREATE_EXTERN
AL_ROUTINE, QUIESCE_CONNECT, SECADM ON DATABASE TO USER MAXIMO
db2 GRANT USE OF TABLESPACE MAXDATA TO USER MAXIMO
db2 GRANT CREATEIN, DROPIN, ALTERIN ON SCHEMA MAXIMO TO USER MAXIMO
```

15. Run the following command to break the database connection:

```
db2 connect reset
```

Example

For example, you can configure the database by using an Amazon Web Services EC2 instance.

Configuring IBM Cloud Pak for Data

If you select Db2 as your database, you must configure Cloud Pak for Data.

About this task

If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:

- Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
- Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
- Secure operator versions and disable auto-update for all components.
- Use the IAM and LDAP service instead of the default authentication methods.
- For more information, see [DB2 Performance Insight](#).

Procedure

1. Open the Red Hat OpenShift console to log in to IBM Cloud Pak for Data.

If you do not have the Cloud Pak for Data admin password the first time you log in, you can get it through the Red Hat OpenShift console.

a) In the Red Hat OpenShift console, from the side navigation menu, click **Workloads > Secrets**.

b) In the **Project** field, select a Cloud Pak for Data namespace.

If you installed Maximo Application Suite on Amazon Web Services, the Cloud Pak for Data namespace is **cpd-services-uniqueid**.

c) Filter for admin-user-details.

d) Click the name of the administrator account.

e) Copy the value of the **initial_admin_password** field.

2. From the side navigation menu, click **Networking > Routes**.

3. On the **Routes** page, from the **Project** field, select a Cloud Pak for Data namespace.

4. Click the **Location** link to open the Cloud Pak for Data login page in a new browser tab.

5. Log in to Cloud Pak for Data as an administrator.

6. Select **Databases** and then click **Create Database**.

7. On the **Select a database** page, click **Next**.

8. On the **Configure** page, select **Single location for all data** and then click **Next**.

9. On the **Advanced** page, click **Next**.

10. On the **Storage** page, select the storage class and then click **Next**.

For Amazon Web Services customer-managed Red Hat OpenShift clusters that are provisioned through an automated deployment offering, select **oc-storagecluster-cephs**.

11. On the **Finalize** page, update the display name if needed, and then click **Create**.

Results

After the instance is created, a green icon appears on the database tile.

Configuring IBM Db2 Warehouse

Create a IBM Db2 Warehouse database for exclusive use by Maximo Manage. You cannot reuse a Db2 Warehouse database on Cloud Pak for Data that is already used by another deployed Maximo Application Suite application.

Before you begin

Create a Db2 Warehouse database on IBM Cloud Pak for Data. For more information, see [Creating Db2 instance by using IBM Cloud Pak for Data console](#).

After Db2 Warehouse databases are provisioned, a parameter is configured to specify user table organization on creation set as a column-organized table. However, to successfully deploy Maximo Manage on Db2 Warehouse, you must change this parameter so that the tables are created as row-organized tables.

Db2 Warehouse can have more than one provisioned database instance. Ensure that you choose the one that is used by Maximo Manage to set this configuration and use when you configure the database. If your instance of Db2 Warehouse was installed through Cloud Pak for Data, you can complete the following steps to find the Db2 administrator pod:

1. Log in to the Cloud Pak for Data interface as an administrator.
2. Select **Databases**.
3. On the tile of the Db2 Warehouse database instance that you provisioned for Maximo Manage, select **Details** from the menu.
4. Search for the deployment ID value, for example, db2wh-1652220906500619.
5. Copy the suffix of the deployment ID, for example, 1652220906500619.
6. Search for the database name. The name of the database is the value that you use in place of the `$DB_NAME` variable in command-line examples.
7. Open the Red Hat OpenShift console and from the side navigation menu, click **Workloads > Pods**.
8. From the **Project** menu, select your Cloud Pak for Data namespace.

If you installed Maximo Application Suite on Amazon Web Services, the Cloud Pak for Data namespace is **cpd-services-uniqueid**.

9. In the **Filter** field, enter the suffix of the deployment ID that you copied and append it to the administrator pod value, for example, 1652220906500619-db2u-0. The pod that is displayed is the db2u administrator pod where you run the **db2inst1** command.

See the following guidelines on how to configure a database instance.

- Configure separate system, data, and backup storage when you create a Db2 instance.
- Increase the maxsequence cache to 50.
- Run **REORG INDEXERS/TABLES** and **RUNSTATS** daily.
- Separate system storage, user storage, backup storage, transaction logs storage, and temporary table space storage on different disks.
- Maximo Manage requires row-organized tables. By default, the IBM Db2 Warehouse database setting uses column-based table organization. Update the setting as needed.

- Maximo Manage does not support Massively Parallel Processing (MPP) or table partitioning. Archive records that are over a year old. InfoSphere Optim Data Growth Solution can be used for archiving. For more information, see [IBM Maximo Archiving 7.5.1 for IBM Maximo Asset Management](#).
- An issue can occur when you load a large amount of data by using the Maximo Integration Framework. Increase the concurrently running statements that are allowed for a Db2 application. For more information, see [How many concurrently running statements allowed for a Db2 Java application and how to increase it?](#).
- If you use IBM Cloud Pak for Data, consider the following configuration and deployment options:
 - Avoid the IBM operator catalog to install IBM Cloud Pak for Data.
 - Use storage services, such as Portworx or Red Hat OpenShift Container Storage.
 - Secure operator versions and disable auto-update for all components.
 - Use the IAM and LDAP service instead of the default authentication methods.
 - For more information, see [DB2 Performance Insight](#).

Note: The following commands use the variable `$DB_NAME` to indicate the name of the database that you defined. Before you run the commands, ensure that you replace the `$DB_NAME` with the name of your database in the **export** command, which is available in the list of the commands.

```
su - db2inst1
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
db2 update db cfg for $DB_NAME using dft_table_org row
db2 terminate
db2 deactivate db $DB_NAME
db2stop
db2start
db2 activate db $DB_NAME
db2 connect to $DB_NAME
db2 get db cfg | grep DFT_TABLE_ORG
```

About this task

The following information is an example of the steps that you must complete and applies only to Cloud Pak for Data. If you complete different steps, ensure that a database schema and table spaces are configured and that the **ddl_constraint_def** parameter is set to yes.

The following steps and commands also use the variable `$DB_NAME` to indicate the name of the database that you defined. Before you run the commands, ensure that you replace the `$DB_NAME` with the actual name of your database.

Procedure

1. Confirm that the database is created.
 - a) As an administrator, log in to Red Hat OpenShift and in the Red Hat OpenShift cluster, from the navigation menu, click **Workloads > Pods**.
 - b) Locate and open the `db2wh-string-db2u-0` pod. *string* is a randomly generated set of numbers.
 - c) On the **Terminal** tab, run the following commands:

```
su - db2inst1
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
```

If the database is configured, the database connection information is returned. The following text is an example of this information, where BLUDB is the name of the database. You need these values later in the configuration.

```
sh-4.2$ su - db2inst1
Last login: Tue May 26 14:29:55 UTC 2020
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ db2 connect to BLUDB
```


Database Connection Information

```
Database server = DB2/LINUX8664 11.5.2.0
SQL authorization ID = DB2INST1
Local database alias = BLUDB
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

2. Create an administrative Cloud Pak for Data user.

To set Maximo Manage to connect to the database, you can use the db2inst1 user that comes by default with Db2 Warehouse, or you can create a different administrative user.

If you installed Db2 Warehouse without Cloud Pak for Data, you can create a different administrative user. For more information, see [Authentication options for Db2U](#).

If you installed Db2 Warehouse through Cloud Pak for Data, you can create a new administrative user through Cloud Pak for Data.

- a) As an administrator, log in to Cloud Pak for Data and from the side navigation menu, click **Administration > User management**.
- b) Create a user by specifying the following information:
 - Specify a user and username.
 - Specify an email address and password.
 - Select the **Administrator** role.
- c) From the side navigation menu, click **Data > Databases**.
- d) Click the three-dot icon for your Maximo Manage database and then click **Details**.
- e) From the drop-down menu, click **Manage access**.
- f) In the Maximo user row, click the edit icon.
- g) In the **Role** field, select **Admin** and then click **Save**.

Note: Take note of this username because it is used as the value of the `$DB_USERNAME` variable in commands in this procedure.

3. Prepare the database for the **maxinst** program. Run the following commands from a database browser or command-line:

- a) Run the following command to connect to the database as the db2inst1 user:

```
su - db2inst1
```

- b) Run the following command to connect to the database:

```
export DB_NAME=$DB_NAME
db2 connect to $DB_NAME
```

- c) Optional: For administrators who connect from outside Red Hat OpenShift, run the following command to connect to the Db2 pod. Replace the variable with the string value for your db2wh-*string*-db2u-0 pod that was accessed during step 1.

```
oc rsh -n <db2wh namespace> c-db2wh-*<string>*-db2u-0 /bin/bash
```

Note: The name of the c-db2wh-*<string>*-db2u-0 pod can be different if you installed Db2 Warehouse without Cloud Pak for Data. For example, if you installed the DB2u operator through [Db2](#), the name of the pod is similar to c-*\$DB2_INSTANCE_NAME*-db2u-0.

4. Configure the database.

- a) Run the following commands to configure the database:

Note: Choose the APPHEAPSZ db2 value to use in the following commands and replace the value of the `$HEAPVALUE` variable by at least 2048 in its specific **export** command line. If you plan

to install many Manage extensions at the same time, set the `$HEAPVALUE` variable to 16384 to prevent failure of the **maxinst** or **updated** processes.

```
su - db2inst1
export DB_NAME=$DB_NAME
export HEAPVALUE=$HEAPVALUE
db2 connect to $DB_NAME
db2 update db cfg for $DB_NAME using "APPHEAPSZ $HEAPVALUE AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_MAINT ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_TBL_MAINT ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_RUNSTATS ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_REORG ON DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_DB_BACKUP ON DEFERRED"
db2 update db cfg for $DB_NAME using "CATALOGCACHE_SZ 800 DEFERRED"
db2 update db cfg for $DB_NAME using "CHNGPGS_THRESH 40 DEFERRED"
db2 update db cfg for $DB_NAME using "DBHEAP AUTOMATIC"
db2 update db cfg for $DB_NAME using "LOCKLIST AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "LOGBUFSZ 1024 DEFERRED"
db2 update db cfg for $DB_NAME using "LOCKTIMEOUT 300 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGPRIMARY 20 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGSECOND 100 DEFERRED"
db2 update db cfg for $DB_NAME using "LOGFILSIZ 8192 DEFERRED"
db2 update db cfg for $DB_NAME using "SOFTMAX 1000 DEFERRED"
db2 update db cfg for $DB_NAME using "MAXFILOP 61440 DEFERRED"
db2 update db cfg for $DB_NAME using "PCKCACHESZ AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "STAT_HEAP_SZ AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "STMTHEAP AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "UTIL_HEAP_SZ 10000 DEFERRED"
db2 update db cfg for $DB_NAME using "DATABASE_MEMORY AUTOMATIC DEFERRED"
db2 update db cfg for $DB_NAME using "AUTO_STMT_STATS OFF DEFERRED"
db2 update db cfg for $DB_NAME using "STMT_CONC LITERALS DEFERRED"
db2 update alert cfg for database on $DB_NAME using "db.db_backup_req SET THRESHOLDSCHECKED YES"
db2 update alert cfg for database on $DB_NAME using "db.tb_reorg_req SET THRESHOLDSCHECKED YES"
db2 update alert cfg for database on $DB_NAME using "db.tb_runstats_req SET THRESHOLDSCHECKED YES"
db2 update dbm cfg using "PRIV_MEM_THRESH 32767 DEFERRED"
db2 update dbm cfg using "KEEPFENCED NO DEFERRED"
db2 update dbm cfg using "NUMDB 2 DEFERRED"
db2 update dbm cfg using "RQRIOBLK 65535 DEFERRED"
db2 update dbm cfg using "HEALTH_MON OFF DEFERRED"
db2 update dbm cfg using "AGENT_STACK_SZ 1000 DEFERRED"
db2 update dbm cfg using "MON_HEAP_SZ AUTOMATIC DEFERRED"
db2 update db cfg using "DDL_CONSTRAINT_DEF YES"
db2set DB2_SKIPINSERTED=ON
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_MINIMIZE_LISTPREFETCH=Y
db2set DB2_EVALUNCOMMITTED=YES
db2set DB2_FMP_COMM_HEAPSZ=65536
db2set DB2_SKIPDELETED=ON
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
```

b) Run the following command to create the buffer pool:

```
db2 CREATE BUFFERPOOL MAXBUFPOOL IMMEDIATE SIZE 4096 AUTOMATIC PAGESIZE 32 K
```

c) Run the following commands to create the table spaces:

```
db2 CREATE REGULAR TABLESPACE MAXDATA PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
db2 CREATE TEMPORARY TABLESPACE MAXTEMP PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
BUFFERPOOL MAXBUFPOOL
db2 CREATE REGULAR TABLESPACE MAXINDEX PAGESIZE 32 K MANAGED BY AUTOMATIC STORAGE
INITIALSIZE 5000 M BUFFERPOOL MAXBUFPOOL
```

d) Run the following command to create the schema:

Note:

- The `$DB_SCHEMA` variable is the name you give to the schema.
- The `$DB_USERNAME` variable is the user with administrative rights, which was created in step 2.

```
export DB_SCHEMA=$DB_SCHEMA
export DB_NAME=$DB_NAME
```

```
export DB_USERNAME=$DB_USERNAME
db2 CREATE SCHEMA $DB_SCHEMA AUTHORIZATION $DB_USERNAME
```

- e) Run the following commands to grant authority to the database user:

```
db2 GRANT
DBADM, CREATETAB, BINDADD, CONNECT, CREATE_NOT_FENCED_ROUTINE, IMPLICIT_SCHEMA, LOAD, CREATE_EXTE
RNAL_ROUTINE, QUIESCE_CONNECT, SECADM ON DATABASE TO USER $DB_USERNAME
db2 GRANT USE OF TABLESPACE MAXDATA TO USER $DB_USERNAME
db2 GRANT CREATEIN, DROPIN, ALTERIN ON SCHEMA $DB_SCHEMA TO USER $DB_USERNAME
```

- f) Run the following command to break the database connection:

```
db2 connect reset
```

5. Verify that the configuration is successful.

- In the Red Hat OpenShift console, from the side navigation menu, click **Workload > Pods**.
- In your Db2 Warehouse database instance, click the **Db2u** pod.
- Select the **Terminal** tab and run the following commands:

```
su - db2inst1
Export DB_NAME=<yourdbname>
Export DB_USERNAME=<yourdbusername> db2 connect to $DB_NAME user $DB_USERNAME using
<password>
```

Tip:

- Replace *<yourdbname>* with the database name that you obtained in a previous step.
- Replace *<yourdbusername>* with the username that you obtained in a previous step.
- Replace *<password>* with the username password that you set when you created and configured the user in a previous step.

The details for the database connection are displayed after the database is connected. The following text is an example of this information, where BLUDB is the database name and MAXIMO is the username.

```
sh-4.2$ su - db2inst1
Last login: Tue May 26 14:29:55 UTC 2020
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ export DB_NAME=BLUDB
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ echo $DB_NAME BLUDB
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ export DB_USERNAME=MAXIMO
[db2inst1@db2wh-1589293563350-db2u-0- Db2U db2inst1]$ db2 connect to $DB_NAME user
$DB_USERNAME using MAXIMO
```

Database Connection Information

```
Database server      = DB2/LINUX8664 11.5.2.0
SQL authorization ID = MAXIMO
Local database alias = BLUDB
```

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

- d) Run the following command to list the table spaces:

```
db2 list tablespaces
```

Confirm the Db2 MAXINDEX, MAXTEMP, and MAXDATA table spaces are listed. The following text is an example of the table spaces details:

```
NAME = MAXDATA
TYPE = Database Managed Space
CONTENTS = All permanent data. Regular table space.
STATE = 0x0000
Detailed explanation = Normal
```

```
TABLESPACE ID = 5
NAME = MAXTEMP
TYPE = System Managed Space
CONTENTS = System Temporary data
STATE = 0x0000
Detailed explanation = Normal
```

```
TABLESPACE ID = 6
NAME = MAXINDEX
TYPE = Database Managed Space
CONTENTS = All permanent data. Regular table space.
STATE = 0x0000
Detailed explanation = Normal
```

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

- e) Run the following command to view the default table organization:

```
db2 get db cfg | grep DFT_TABLE_ORG
```

Confirm that the configuration `db2 get db cfg | grep DFT_TABLE_ORG` is set to ROW. The following text is an example of the configuration:

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$ db2 get db cfg | grep DFT_TABLE_ORG
Default table organization (DFT_TABLE_ORG) = ROW
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

- f) Run the following command to view the default application heap:

```
db2 get db cfg | grep APPLHEAPSZ
```

Confirm that the value for APPLHEAPSZ in `db2 get db cfg | grep APPLHEAPSZ` is the same that you set. The following text is an example of the configuration:

```
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$ db2 get db cfg | grep APPLHEAPSZ
Default application heap (4KB) (APPLHEAPSZ) = AUTOMATIC(16384)
[db2inst1@db2wh-1589293563350-db2u-0 - Db2U db2inst1]$
```

6. Disable the non-SSL Db2 database instance port for security.

- a) Run the following command to edit the Db2 service configuration:

```
oc edit svc -n project c-service_name-db2u-engn-svc
```

Where

project

The name of the Red Hat OpenShift project where Db2 is deployed.

service_name

The identifier for the Db2 service instance.

For example, `c-db2oltp-1605022957148004-db2u-engn-svc`.

- b) Remove the following text from the `spec . ports` section:

```
- name: legacy-server
  nodePort: 30279
  port: 50000
  protocol: TCP
  targetPort: 50000
```

- c) Save the service.

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

Configuring Oracle Database

To configure Oracle Database for use with Maximo Manage, you create table spaces, create a database user, and configure database settings.

Before you begin

For information about installing and deploying Oracle Database, review the Oracle Database product documentation.

For information about supported database versions, generate a Software Product Compatibility Report. For more information, see [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. Check for the supported database versions on the **Supported Software** tab of the report.

Configure your database with the following operating system:

- Linux or UNIX
- Microsoft Windows

For more information about system performance, see [Best practices for system performance](#).

Procedure

1. Log in as the Oracle software user. Typically, this user is named `oracle`.
2. To manage requests to connect to the database, create the database listener.
3. Create a database for use by Maximo Manage.

For the database initialization parameters, change the values of the following parameters:

nls_length_semantics

Change this value to CHAR.

open_cursors

Change this value to 1000.

cursor_sharing

Set this value to FORCE.

Note: Ensure that the database character setting is set to the AL32UTF8 character set, which is required for Oracle Database.

4. In SQL*Plus, create a table space by running the following command. Replace the directory with the path to the database location.

```
Create tablespace maxdata datafile
'C:\oracle\product\12.1.0.1\db_1\dfs\maxdata.dbf'
size 1000M autoextend on;
```

To create table spaces for indexes, repeat the command and use a similar syntax.

5. Create a temporary table space by running the following command. Replace the directory with the path to the database location.

```
create temporary tablespace maxtemp tempfile
'C:\oracle\product\12.1.0.1\db_1\dfs\maxtemp.dbf'
size 1000M autoextend on maxsize unlimited;
```

6. To create the Maximo user and grant permissions, run the following command:

```
create user maximo identified by maximo default tablespace maxdata temporary
tablespace maxtemp;
grant connect to maximo;
grant create job to maximo;
grant create trigger to maximo;
grant create session to maximo;
grant create sequence to maximo;
grant create synonym to maximo;
grant create table to maximo;
```

```
grant create view to maximo;
grant create procedure to maximo;
grant alter session to maximo;
grant execute on ctxsys.ctx_ddl to maximo;
alter user maximo quota unlimited on maxdata;
```

If you created a separate table space for indexing, you must also grant access to that index table space to the Maximo user.

For example, if you created a separate table space for indexing that is called **TSI_MAM_OWN**, then run the following command:

```
alter user maximo quota unlimited on TSI_MAM_OWN
```

7. Create an Oracle preference arbitrary that is called **MAXIMO_STORAGE** and store the Oracle text indexes in dedicated table spaces.

- Store implicit indexes in the MAXINDX table space.
- Store implicit tables in the MAXDATA table space.
- Store implicit LOB tables in the LOB table space.

For example, run the following preference definition to split the implicit objects in a text index across three table spaces: MAXDATA, MAXINDX, and MAXLOBS.

```
begin
ctx_ddl.create_preference('MAXIMO_STORAGE', 'BASIC_STORAGE');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'I_TABLE_CLAUSE',
'tablespace MAXDATA LOB(token_info) store as (tablespace MAXLOBS
enable storage in row)');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'I_INDEX_CLAUSE',
'tablespace MAXINDX compress 2');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'K_TABLE_CLAUSE',
'tablespace MAXINDX');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'R_TABLE_CLAUSE',
'tablespace MAXDATA LOB(data) store as (tablespace MAXLOBS
cache)');
ctx_ddl.set_attribute('MAXIMO_STORAGE', 'N_TABLE_CLAUSE',
'tablespace MAXINDX');
end;
```

To create an Oracle text index, the preference definition must be specified in the **CREATE INDEX** clause as shown in the following example.

```
create index pm_ndx6 on pm (description) indextype is
ctxsys.context parameters ('lexer global_lexer language column
LANGCODE storage MAXIMO_STORAGE');
```

8. Set up Oracle text preferences and sublexer definitions.

- a) Use an SQL query tool to log on to the database as the maximo user, which is the schema owner, and run the following set of calls.

```
call ctx_ddl.drop_preference('global_lexer');
call ctx_ddl.drop_preference('default_lexer');
call ctx_ddl.drop_preference('english_lexer');
call ctx_ddl.drop_preference('chinese_lexer');
call ctx_ddl.drop_preference('japanese_lexer');
call ctx_ddl.drop_preference('korean_lexer');
call ctx_ddl.drop_preference('german_lexer');
call ctx_ddl.drop_preference('dutch_lexer');
call ctx_ddl.drop_preference('swedish_lexer');
call ctx_ddl.drop_preference('french_lexer');
call ctx_ddl.drop_preference('italian_lexer');
call ctx_ddl.drop_preference('spanish_lexer');
call ctx_ddl.drop_preference('portu_lexer');
call ctx_ddl.create_preference('default_lexer', 'basic_lexer');
call ctx_ddl.create_preference('english_lexer', 'basic_lexer');
call ctx_ddl.create_preference('chinese_lexer', 'chinese_lexer');
call ctx_ddl.create_preference('japanese_lexer', 'japanese_lexer');
call ctx_ddl.create_preference('korean_lexer', 'korean_morph_lexer');
call ctx_ddl.create_preference('german_lexer', 'basic_lexer');
call ctx_ddl.create_preference('dutch_lexer', 'basic_lexer');
call ctx_ddl.create_preference('swedish_lexer', 'basic_lexer');
```

```

call ctx_ddl.create_preference('french_lexer','basic_lexer');
call ctx_ddl.create_preference('italian_lexer','basic_lexer');
call ctx_ddl.create_preference('spanish_lexer','basic_lexer');
call ctx_ddl.create_preference('portu_lexer','basic_lexer');
call ctx_ddl.create_preference('global_lexer','multi_lexer');
call ctx_ddl.add_sub_lexer('global_lexer','default','default_lexer');
call ctx_ddl.add_sub_lexer('global_lexer','english','english_lexer','en');
call ctx_ddl.add_sub_lexer('global_lexer','simplified chinese','chinese_lexer','zh');
call ctx_ddl.add_sub_lexer('global_lexer','japanese','japanese_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','korean','korean_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','german','german_lexer','de');
call ctx_ddl.add_sub_lexer('global_lexer','dutch','dutch_lexer',null);
call ctx_ddl.add_sub_lexer('global_lexer','swedish','swedish_lexer','sv');
call ctx_ddl.add_sub_lexer('global_lexer','french','french_lexer','fr');
call ctx_ddl.add_sub_lexer('global_lexer','italian','italian_lexer','it');
call ctx_ddl.add_sub_lexer('global_lexer','spanish','spanish_lexer','es');
call ctx_ddl.add_sub_lexer('global_lexer','portuguese','portu_lexer',null);

commit;

```

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.
- The table space, index space, and schema values. These values are created as part of configuring the database.

If you modified the default role sets assigned to the user ID used to connect to the database, then you must explicitly grant the role sets to the Maximo user. If you restricted the default privileges that are granted to user IDs, you must also explicitly grant the role sets to the Maximo user. For example, if you do not grant a role such as the **select_catalog_role** role, you must explicitly grant that role to the Maximo user. Make the assignment by running the following SQL*Plus command:

```
grant select_catalog_role to maximo
```

Configuring Microsoft SQL Server

To configure Microsoft SQL Server for Maximo Manage, you create table spaces, create a database user, and configure database settings.

Before you begin

For information about configuring Microsoft SQL Server, review the Microsoft SQL Server product documentation.

See the following guidelines for configuring the database:

- If the original database was created in a version earlier than Microsoft SQL Server 2019, set the compatibility level to the older version to maintain the execution plan.
- Set the transactions isolation level by using the following commands:

```

ALTER DATABASE MyDatabase
SET ALLOW_SNAPSHOT_ISOLATION ON

ALTER DATABASE MyDatabase
SET READ_COMMITTED_SNAPSHOT ON

```

For information about supported database versions, you can generate a [Software Product Compatibility Report](#). Search for IBM Maximo Application Suite and select the suite version to generate the report. Check for the supported database versions in the **Supported Software** tab of the report.

Procedure

1. Configure the listener port.

The default instance of the Microsoft SQL Server Database Engine listens on TCP port 1433. Named instances of the Microsoft SQL Server Database Engine and Microsoft SQL Server Compact Edition are configured for dynamic ports, which means they select any available port when the service starts. When you connect to a named instance across a firewall, configure the Database Engine to listen on a specific port to open this port in the firewall.

2. Verify that you enabled the Full-text Search setting during the installation of Microsoft SQL Server.
3. Create a Microsoft SQL Server database.

- a) In Microsoft SQL Server Management Studio, select **New Database** from the databases folder.
- b) Specify a unique database name.
For example, enter maxdb80
- c) For the maxdb80 Logical Name, change the **Initial Size (MB)** field to 500 and also set the value of the **Autogrowth / Maxsize** field to **By 1 MB, Unlimited**.
- d) Optional: Modify the log settings to accommodate your production environment.
- e) To deploy Maximo Manage in a specific language, choose the default collation for the database.
For example, to deploy the application in English, select **Latin1_General_100_CI_AS_KS_SC_UTF8**.

Starting from Maximo Application Suite 9.0, Maximo Manage supports a Microsoft SQL Server database that uses Unicode. You must select a collation that has UTF8 in its name. Microsoft SQL Server supports multiple languages in the same database in instances where the chosen languages support the same Microsoft SQL Server collation. For example, English and French can be installed because both languages support the same Microsoft SQL Server collation. However, English and Japanese cannot be installed together because they have different Microsoft SQL Server collations.

For more information, review the Microsoft SQL Server Collation and Unicode support documentation.

4. Create the Maximo user for Microsoft SQL Server.

- a) In Microsoft SQL Server Management Studio, from the SQL Server Configuration Manager navigation, click **Databases**.
- b) Right-click the **maxdb80** database and select **New Query**
- c) Enter the following command to create the Maximo database user MAXIMO with a password that adheres to the password policy of the system.

```
sp_addlogin MAXIMO,password
go
```

This value is case-sensitive.

- d) Enter the following command to change the database owner to MAXIMO.

```
sp_changedbowner MAXIMO
go
```

What to do next

Gather the following information for your database:

- The host and hostname.
- The port.
- The database name.
- The username and password for the database user. These values are created as part of configuring the database.

- The table space, index space, and schema values. These values are created as part of configuring the database.

Database encryption

When you deploy Maximo Manage, fields that require security, such as passwords and API keys, are encrypted or reencrypted to provide security.

Database encryption overview

When you configure Maximo Manage, specify encryption keys and encryption algorithms to determine how the fields that require security are encrypted.

Important: Save the encryption secret, which contains the encryption keys, after the Maximo Manage deployment is completed. The same keys are used for configuration when you reinstall Maximo Manage with the same database.

The following table describes the Crypto and CryptoX encryption keys for Maximo Manage:

<i>Table 45. Encryption keys</i>	
Key	Description
MXE_SECURITY_CRYPTOKEY	Use it to encrypt Crypto fields, such as passwords. For Crypto encryption, if you specify a MXE_SECURITY_CRYPTOKEY value that matches the MXE_SECURITY_OLD_CRYPTOKEY value that was used in the previous deployment, no reencryption occurs. If you specify a key value during deployment that does not match the MXE_SECURITY_OLD_CRYPTOKEY value, the database is reencrypted.
MXE_SECURITY_OLD_CRYPTOKEY	Specifies the value for the previous Crypto encryption key that was used for the database.
MXE_SECURITY_CRYPTOXKEY	Used to encrypt CryptoX fields, including API keys, such as the electronic signature key. For CryptoX encryption, if you specify a MXE_SECURITY_CRYPTOXKEY value that matches the MXE_SECURITY_OLD_CRYPTOXKEY value that was used in the previous deployment, no encryption changes occur. CryptoX values cannot be decrypted, and the original value cannot be determined. If you specify a key value in a deployment that does not match the MXE_SECURITY_OLD_CRYPTOXKEY value, CryptoX values are set to null when encryption is run.
MXE_SECURITY_OLD_CRYPTOXKEY	Specifies the value for the previous CryptoX encryption key that was used for the database.

The following encryption properties are also supported:

<i>Table 46. Encryption properties</i>	
Encryption property	Description
MXE_SECURITY_CRYPTOKEY_ALGORITHM	The default value is AES.
MXE_SECURITY_CRYPTOKEY_MODE	The default value is CBC.

Encryption property	Description
MXE_SECURITY_CRYPTOMODULUS	
MXE_SECURITY_CRYPTOPADDING	The default value is PKCS5Padding.
MXE_SECURITY_CRYPTOSPEC	The length must be a multiple of 8.
MXE_SECURITY_CRYPTOX_ALGORITHM	The default value is AES.
MXE_SECURITY_CRYPTOX_MODE	The default value is CBC.
MXE_SECURITY_CRYPTOXMODULUS	
MXE_SECURITY_CRYPTOX_PADDING	The default value is PKCS5Padding.
MXE_SECURITY_CRYPTOX_SPEC	The length must be a multiple of 8.

Note: After the database is installed, only the Crypto and CryptoX encryption keys can be changed.

When you configure the database settings for deployment before you activate the application, you can add a value for the **MXE_SECURITY_CRYPTO_KEY** or **MXE_SECURITY_CRYPTOX_KEY** encryption keys. If you do not specify an encryption key secret in the Maximo Manage configuration when you activate, the system automatically generates keys. The system names the secret in the keys by using the following naming convention:

```
<workspaceId>-<appId>-encryptionsecret
```

For more information about how to specify the encryption secret in the Maximo Manage configuration, see [“Adding encryption key secrets” on page 322](#).

Because your database functions only with valid encryption keys, implement the following practices:

- Maintain your encryption keys in a vault or other secure management system for secrets.
- Specify your own values for encryption keys instead of using system-generated values. If you use system-generated values and do not create a backup, you cannot retrieve the keys. Without the keys, you cannot use your database.

Database encryption scenarios

Your deployment scenario determines your encryption and reencryption options. Scenarios include deploying a new database, deploying a previously encrypted database, or changing the encryption keys for deployment.

Deploying a new database

If you deploy a new database, you have two options for database encryption:

Option	Action	Result
Provide your own values for the encryption keys.	Specify the values that you want to use for the Crypto and CryptoX encryption keys when you configure your database during deployment.	The database is then encrypted by using the key values that you provide.

Table 47. Encryption for deploying a new database (continued)

Option	Action	Result
Use system-generated keys.	Do not specify key values when you configure your database.	If you do not specify the keys, a secret is automatically generated that contains the new MXE_SECURITY_CRYPTOKEY and MXE_SECURITY_CRYPTOXKEY encryption keys. Later, if you need the keys, you can view the keys in the secret.

Deploying a previously encrypted database

The following scenarios can occur if your deployment includes a database that was previously encrypted. For example, you might be upgrading from IBM Maximo Asset Management 7.6.0.10 or 7.6.1.2, which are not installed on Red Hat OpenShift but can be upgraded to Maximo Manage. You might also be upgrading from a previous Maximo Manage version. The scenarios apply both to a database that previously implemented the default encryption for Maximo Manage or a database that was encrypted by using a different algorithm or keys.

An existing database that used the default encryption for Maximo Manage and no values were provided for the Crypto or CryptoX key

The database for Maximo Manage no longer uses a default set of keys for encryption. If you have an existing database that used the default encryption, provide the **MXE_SECURITY_CRYPTOKEY** and **MXE_SECURITY_CRYPTOXKEY** values so that the database can be reencrypted. If your database was previously encrypted by using keys other than the default encryption for Maximo Manage, provide the old **MXE_SECURITY_OLD_CRYPTOKEY** and **MXE_SECURITY_OLD_CRYPTOXKEY** encryption keys. As a result, the database can be decrypted and then reencrypted.

When you configure the database, select one of the following options for reencryption:

Note: Reencryption always occurs in this scenario.

Table 48. Encryption for existing databases where no values were provided for the encryption keys.

Option	Action	Result
Provide your own values for the encryption keys.	<ol style="list-style-type: none"> 1. Enter the MXE_SECURITY_CRYPTOKEY and MXE_SECURITY_CRYPTOXKEY values. 2. Do not specify any key values when you configure your database. 	The database is reencrypted by using the values that you specified.
Use system-generated keys.	<ol style="list-style-type: none"> 1. Do not specify any key values when you configure your database. 	The system generates new keys, and the database is reencrypted with the system-generated keys.

An existing database that used values from your own Crypto or CryptoX keys

When you configure the database, select one of the following options for encryption:

Note: The first option is less likely to result in errors. You can change the keys later.

Table 49. Encryption for existing databases where you provided your own values for the encryption keys.

Option	Action	Result
Do not reencrypt the database.	<ol style="list-style-type: none"> 1. Specify the Maximo Manage security properties, including the MXE_SECURITY_OLD_CRYPTO_KEY and MXE_SECURITY_OLD_CRYPTOX_KEY encryption keys. 2. Specify the same values for the MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. 	Because you specified the same values for the old and new keys, the database is not reencrypted.
Reencrypt the database.	<ol style="list-style-type: none"> 1. Specify the Maximo Manage security properties, including the MXE_SECURITY_OLD_CRYPTO_KEY and MXE_SECURITY_OLD_CRYPTOX_KEY encryption keys. 2. Select one of the following options: <ul style="list-style-type: none"> • Specify values for the new MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. • To use system-generated keys, do not specify encryption key values. 	The database is reencrypted.

Changing the encryption keys for deployment

The following table describes the tasks to complete reencrypting the database when you want to change the **MXE_SECURITY_CRYPTO_KEY** and **MXE_SECURITY_CRYPTOX_KEY** encryption keys.

Table 50. Encryption by using new encryption keys

Option	Action	Result
Reencrypt the database by using new encryption keys.	<ol style="list-style-type: none"> 1. Set the MXE_SECURITY_OLD_CRYPTO_KEY and MXE_SECURITY_OLD_CRYPTOX_KEY encryption keys to the values that the database currently uses for the MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. 2. Select one of the following options: <ul style="list-style-type: none"> • Specify values for the new MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. • To use system-generated keys, do not specify values for the MXE_SECURITY_CRYPTO_KEY and MXE_SECURITY_CRYPTOX_KEY encryption keys. 	The database is reencrypted by using the new encryption keys.

Viewing database encryption history

You can view the history for database encryption. If you encrypt or reencrypt the database, a new entry is logged in the syschangetracker table in the Maximo database.

The syschangetracker table contains the following columns:

- PROCESSNAME
- MESSAGE
- CHANGEBY
- CHANGEDATE
- WORKSPACEID
- APPID
- INSTANCENAME
- SYSCHANGETRACKERID
- ROWSTAMP

If you have issues with encryption keys, the record in the syschangetracker table helps to identify whether reencryption occurred and when it occurred. The MESSAGE column includes the first 2 bytes and last 2 bytes of the encryption key. The MESSAGE column can help you identify the encryption keys that are being used.

Disabling automatic generation of encryption keys

By default, if encryption keys are not specified when you activate Maximo Manage with a fresh database, new encryption keys are automatically generated. If you do not want to automatically generate encryption keys, set the **autoGenerateEncryptionKeys** property to false.

About this task

The **autoGenerateEncryptionKeys** property controls whether encryption keys are automatically generated when you activate Maximo Manage with a fresh database. By default, this property is set to `true`, and encryption keys are automatically generated if no value is specified for them. Set the property to `false` if you do not want to generate keys automatically. If you set the **autoGenerateEncryptionKeys** property to `false` and you do not provide encryption keys, deployment fails. You also receive an error message that the property is set to `false` and encryption keys are missing.

If encryption keys are automatically generated, you can easily lose track of them, especially in development and test environments where databases are reused. If you set the **autoGenerateEncryptionKeys** property to `false`, users are forced to enter the key. You are less likely to lose keys that you generate and maintain. Securely store the encrypted keys, for example, by using a password keeper.

This property takes effect when you activate the application by using an API call. Updating the property for an already activated instance does not produce an effect until you make a change that is related to the encryption keys, such as when you delete the keys to trigger reencryption.

Procedure

1. In the Red Hat OpenShift console, from the side navigation menu, click **Administration > Custom Resource Definitions**.
2. On the **CustomResourcesDefinitions** page, select the ManageWorkspace custom resource definition record.
3. On the **CustomResourceDefinition details** page, on the **Instances** tab, select the instance for which you want to disable automatic generation of encryption keys.
4. On the **YAML** tab for the instance, set `spec.settings.deployment.autoGenerateEncryptionKeys` to `false`.
5. Save the custom resource.

Adding encryption key secrets

For a new installation, the encryption secret is used as the key to encrypt the database. For already encrypted databases, the encryption secret is used to restore the encryption keys in the future.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. In the side navigation menu, from the **Suite** application, select **Administration > Workspace**.
3. On the **Manage workspace details** page, click **Actions**, and select **Update configuration**.
4. In the Database connection row, click the edit icon.
5. Click **Show advanced settings** and set the **System managed** toggle to off.
6. Click **Add property** to add a row for an encryption key property.
7. In the Key column, enter the encryption key property.
For example, `MXE_SECURITY_CRYPT0_KEY`.
8. In the Value column, enter the value for the added encryption key property.
9. Click **Apply changes** to save and apply all the configuration changes.

Resetting the Crypto and CryptoX fields in the database

If you lose your Crypto and CryptoX encryption keys, you can run a reset script that clears the Crypto and CryptoX fields in the Maximo Manage database. Then, you can restart the server and reset your data.

About this task

When you run the `resetcryptocryptox.sh` script, the Crypto and CryptoX fields are cleared. After you restart the server, you can review the data that was cleared and reset the values that were cleared, such as properties or API keys.

Procedure

1. In the Red Hat OpenShift Container Platform, from the side navigation menu, click **Workloads** > **Pods**.
2. On the **Pods** page, open the maxinst pod.
3. On the **Terminal** tab, enter `resetcryptocryptox.sh`.
4. Restart the Maximo Manage server.
5. Add any variables, such as API keys or properties that the `resetcryptocryptox.sh` script cleared.
6. Redeploy or reactivate the application.

Troubleshooting database deployment

When you deploy the Maximo Manage application, issues that are related to the database deployment might cause the deployment process to fail.

To prevent or solve some database configuration issues that might result in issues during the Maximo Manage deployment, follow these practices:

- Review the database settings that you used to configure Maximo Manage. If something is wrong with the values, edit your database configuration on the **Manage Status** page and activate it again.
- Ensure that the JDBC URL format and details are correct and are using the correct JDBC protocol, database URL, name, and port. Test your connection to the database by using a database tool.
- If you are using the SSL-enabled database and port, verify that you selected the SSL option.
- To ensure that your database is ready, verify that you are using the correct username, password, schema name, tablespace name, and index tablespace name.
- Do not skip preparation steps. For example, the deployment process for Maximo Manage databases does not work if you skip the step to set its table organization to be row-based instead of column-based.
- If you plan to deploy languages, ensure that your database is properly configured. Set the configuration to support Maximo Manage or SQL for the language that is used for the databases.
- Confirm that you used the correct database certificate for the database in the Maximo Manage configuration and that the certificate is not expired.
- If you are deploying Maximo Manage and reusing a database that was previously used by Maximo Manage, pass the encryption keys that are saved from the previous deployment. Otherwise, the process generates new keys, which do not match the ones that are used to encrypt the database, and the deployment process does not proceed.
 - For example, you deactivated Maximo Manage and deleted the application. Now, you are redeploying the Maximo Manage application and activating it again, but you are still using the previously deployed Maximo Manage database.
- To deploy multiple Maximo Manage components at the same time, if you are using IBM Db2 Warehouse databases or other Db2 versions, increase the **APPHEAPSZ** value to at least 16384 to avoid failures during the database deployment. Alternatively, you can deploy fewer components simultaneously by deploying Maximo Manage components one at a time.
- Save your database in the same location as Maximo Manage. If the location of your database is different, the chances of database failures during the Maximo Manage deployment due to connection issues and latency are higher.

During the Maximo Manage deployment, database update interruptions might occur due to a situation, such as lost connection. The Maximo Manage operator reconciliation process usually resolves the problem.

If the failure persists, make sure that your database connection is working correctly. Before you remove, re-create the database, and retry the Maximo Manage deployment, you can try the following procedures:

1. Select the **Bypass upgrade version** checkbox and then reactivate the Maximo Manage application. If you select this option, the deployment can run, regardless of the current version of the installed database. It skips version validations that might be preventing the deployment to progress.
2. If the failure persists, review the log files to assess the situation and note the script number for the failed deployment. To review the log files, complete the following steps:
 - a. In the Red Hat OpenShift console, from the side navigation menu, click **Workloads > Pods**.
 - b. From the **Project** list, select the Manage project.
 - c. On the **Pods** page, select the administrative pod. The administrative pod includes manage-maxinst in the pod name, as shown in the following example:

```
env-managedev-manage-maxinst-7fd3c77492-kmj6z
```

- d. To review the logs, on the **Pod details** page, click the **Logs** tab. If you want to view all the logs, click the **Terminal** tab and view the logs in the `/opt/IBM/SMP/maximo/tools/maximo/log` directory.
- e. After you review the log files, manually correct the problems and run the database installation again.

Demo data

When you deploy Maximo Manage, you can add sample data to a demo database.

The sample data in the demo database is useful for development or test environments.

To set up a test or development environment with demo data, install an instance of Maximo Application Suite specifically for testing or development. Then, when you configure the database settings for your Maximo Manage deployment, select the option to install demo data.

Preparing the system for sample data

You can add sample data to a demo database when you deploy Maximo Manage. The sample data is useful for development or test environments. If you want to add sample data after Maximo Manage deployment is complete, you must re-create or clean your database and reconfigure Maximo Manage.

Before you begin

Ensure that Maximo Manage deployment is complete before you follow the steps to add sample data.

Logs of previous installations might be deleted when you delete the maxinst pod. If you want to retain these logs, save backup copies outside the pod or container before you delete the pod.

Procedure

1. In Maximo Application Suite, update the configuration to install demo data.
 - a) Log in to Maximo Application Suite by using administrator credentials.
 - b) Click the **Administration** icon.
Starting in Maximo Application Suite 9.1, select the **Suite > Administration** page.
 - c) Click the **Workspaces** tile.
 - d) In the Applications section, click the **Manage** tile.
 - e) From the **Actions** list, select **Update configurations**.
 - f) In the **Update Manage configuration** window, in the Database row, click the **Edit** icon.
 - g) In the Advanced settings section, select the **Install demo data** checkbox.
 - h) Click **Apply changes**.
 - i) Click **Confirm** in the confirmation dialog.

2. In the Red Hat OpenShift web console, check if demodata is set to true.
 - a) Log in to the Red Hat OpenShift web console by using your administrator credentials.
 - b) From the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
 - c) On the **CustomResourceDefinitions** page, search for manageworkspace.
 - d) Click **ManageWorkspace** and then click the **Instances** tab.
 - e) Select the custom resource of your instance and click the **YAML** tab.
 - f) In the `spec.settings.db.maxinst` section, check if `demodata: true` is set.

Tip: If you do not see the demodata section, add it.
 - g) Click **Save** to save your changes.
 - h) Click **Reload** to check whether your changes were applied.

Remember: If you do not see your changes, you might need to repeat this step because the automatic reload process can update the YAML file before you make your manual updates.

3. If you want to clean the database to force the maxinst process to restart, use a query tool by consulting your database provider. In your Maximo Manage database, delete the following database tables.
 - maxvars
 - apikeytoken
 - serversession
 - maxsession

Note: If you want to drop your entire database, back up your database and then complete the steps that are provided by your database provider.

4. If you deployed Maximo Utilities with Maximo Manage, delete the plusdcuchngstatseq sequence.

5. Restart the maxinst pod.

- a) In the Red Hat OpenShift web console, from the side navigation menu, click **Workloads** > **Pods**.
- b) On the **Pods** page, search for the maxinst pod.
For example, `<mas_instance_id>-<mas-workspace-id>-manage-maxinst-<randomstring>`.
- c) In the row for the maxinst pod, click the **Overflow** menu and select **Delete Pod**.
- d) In the **Delete Pod** dialog, click **Delete** to confirm deleting the maxinst pod.

Important: Deleting the pod removes any files that are copied to the pod container directly. The deleted files are different from the ones in the Maximo Manage image that is built during the Maximo Manage deployment process. The files that are part of the Maximo Manage image include the contents of the Maximo Manage components and customizations that were added through customization archives.

When the pod is deleted, another pod with a similar name is created.

- e) Select the new pod that was created and click the **Terminal** tab.
- f) Run the following commands.

```
cd logs
ls -ltr
```

Tip: If you do not see the maxinst log, wait for five minutes and run the `ls -ltr` command again. You can also click the **Logs** tab to check for the maxinst log.

What to do next

Wait for the maxinst restart process to complete and for the Maximo Manage server bundles to restart. You can monitor the progress in the Red Hat OpenShift web console or in the Maximo Application Suite user interface. For more information, see [Monitoring Manage application deployment by using the Maximo Application Suite interface](#).

When the maxinst restart process is completed, you can configure Maximo Manage. For more information, see [Configuring Maximo Manage after deployment](#).

Deploying Maximo Manage

You deploy Maximo Manage or Maximo Manage with Maximo Health operator, which is responsible to control the deployment process and to maintain and update the application as necessary.

Before you begin

- You also can check the actual supported matrix of your current Maximo Manage application that is deployed and the matrix of compatibility between the components, through the Red Hat OpenShift console:
 1. Go to the Red Hat OpenShift console, in the Workloads/Pods section, select the namespace of your Maximo Manage instance in Projects, for example, `mas-<yourmasinstancename>-manage`, and then click the `ibm-mas-manage-operator-somestring` pod.
 2. On the **Terminal** tab, select to connect to **webhook** container.
 3. Open the supported matrix JSON file that displays the components version that is supported by each Manage application version and look for the version that you installed. If you are not sure what the version is, click **Details** and click **Manager** in the **Containers** section. The version is displayed in the **Image version** tab or field. When you go back to the terminal, you can use this command to open the file:

```
cat /manage-admission/metadata/supported-versions-matrix.json
```

The file has a format like this:

```
"8.3.1": [
  {
    "name": "anywhere",
    "versions": [
      "8.0.1",
      "8.0.0"
    ]
  }
]
```

Note: In this example, the Manage application version is 8.3.1. The following list shows the Anywhere versions that are supported by this version.

When you select **latest**, the page displays a section that states the current Manage application version:

```
"latest": {
  "base": "8.3.1",
  "anywhere": "8.0.1"
}
```

4. Open the dependency matrix JSON file that displays the compatibility between each component by using the following command: **cat dependency-matrix.json** The file displays the components with the following attributes:

```
"aviation": {
  "description": "Maximo Aviation",
  "includesForbidden": [
    "hse",
    "serviceprovider"
  ],
  "includesCoexist": [
    "acm",
    "transportation"
  ],
  "conflict": [
    "civil",
    "health",
    "oilandgas",
    "oracleadapter",
    "nuclear",
    "sapadapter",
  ]
}
```

"spatial",
"utilities"

Note: includesForbidden The component cannot be co-deployed with the listed components.
includesCoexist The component can be co-deployed with the listed components.
conflict The component cannot be co-deployed with the listed components.

Procedure

1. From the Suite catalog, on the **Applications** tab, select the **Manage** with tile, and review the information on the **Setup** page.
 - a) Optional: If you plan to co-deploy industry solutions or add-ons with Manage, select the **Show all** option and then click the plus icon on the respective tile of each industry solution or add-on that you want to co-deploy to reserve the necessary AppPoints for them.
 - b) Click **Continue**.
2. In the **Administer application upgrades** window, in the **Upgrade strategy** field, select one of the following upgrade strategies.

Option	Description
Channel subscription	<ol style="list-style-type: none"> a. To make sure that updates are automatically offered when an updated version is added to a channel, select Channel subscription. You can choose whether this update occurs automatically with a new version or requires manual approval. If you choose an automatic upgrade strategy through Channel subscription, required downtime might occur before you have a chance to review changes in the new version or fix pack, back up the database, or take other preparatory action. Therefore, for production Maximo Manage deployments, set the approval mode to Manual. You can then be alerted about the availability of an update directly in the IBM Maximo Application Suite application catalog, in the Maximo Manage tile. However, the update is not started until you trigger it yourself. In this case, you can be better prepared by reviewing the changes, running backups of the Maximo Manage configuration and custom resource definitions, scheduling the update, and communicating the scheduled downtime to users. b. In the Channel field, select a channel. c. In the Custom source field, specify a source. d. In the Channel details section, select whether to approve the update automatically or to require manual approval. e. Click Subscribe to channel.
Manual	<ol style="list-style-type: none"> a. To manually update the application when you are notified that an updated version is available, select Manual. b. In the Version field, select a version. c. Click Deploy version.

3. Configure the database connection information for Maximo Manage.
 - a) In the Integrations and dependencies section, on the **Database connection** tile, select **Configure**.
 - b) On the **Database connection** page, click **Configure**.
 - c) In the **JDBC connection information** section, specify the following fields:

Note:

If you are upgrading from Maximo Asset Management to Maximo Application Suite, refer to the `maximo.properties` file of your Maximo Asset Management folder to get the values for hostname, port, database name to use in the jdbc url as required by Maximo Manage.

- i) If you want to use an SSL-enabled connection, specify this field by using one of the following JDBC URL formats. Ensure that the **Port** that is used contains the SSL-enabled port of the database.

Oracle Database

TCPS is the protocol to use for Oracle SSL connections. For Oracle SSL database connections in Maximo Manage, you must specify `SID=<TNS Service ID>`. You can use the following URL as an example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCPS)(Host=mymaximodb.com)
(Port=2484))(CONNECT_DATA=(SID=MAXDB)))
```

Microsoft SQL Server Database

For SQL Server SSL database connections in Maximo Manage, you must specify `encrypt=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:sqlserver://mymaximodb.com:1433;databaseName=MAXDB;encrypt=false;
```

IBM Db2 database

For Db2 SSL database connections in Maximo Manage, you must specify `sslConnection=true`. Ensure that you use a semicolon to end the JDBC connection string. You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB:sslConnection=true;
```

- a) In the **User name** field, specify the database username.
 - b) In the **Password** field, specify the database user password.
 - c) Ensure that you select the **SSL Enabled** option.
- ii) If you want to use non-SSL enabled connection, specify the **Connection string** field by using one of the following JDBC URL formats, depending on the database you are using.

Oracle Database

You can use the following URL as an example:

```
jdbc:oracle:thin:@mymaximodb.com:1521:MAXDB
```

Microsoft SQL Server database

You can use the following URL as an example:

```
jdbc:sqlserver://
mymaximodb.com:1433;databaseName=MAXDB;integratedSecurity=false;encrypt=false;
```

IBM Db2 database

You can use the following URL as an example:

```
jdbc:db2://mymaximodb.com:50001/MAXDB
```

- a) In the **User name** field, specify the database username.
 - b) In the **Password** field, specify the database user password.
 - c) Ensure that you do not select the **SSL Enabled** option.
- d) In the Additional driver options section, in the **Driver options** field, add more driver options, which are separated by a semicolon.
- Typically, you can specify JDBC options as part of the URL for the database. However, in some cases you might want to specify JDBC options in the **Driver options** field.

For example, your URL might exceed the maximum length that is allowed, or you might want to configure a JDBC option that cannot be included in the connection URL. You cannot specify the same JDBC option in both the URL and the **Driver options** field. If you do, JDBC driver errors might cause the connection to fail.

- If you specify an extra JDBC option for your database, the CustomProxyDriver acts as a proxy driver that routes the database requests to the actual driver for your type of database.
- e) If you chose to use an SSL-enabled database connection, in the Security > Certificates (optional) section, click **Add+** to display the fields to include in your database certificate.
 - i) In the **Alias** field, specify an alias name to identify the certificate, for example, DB2WHcert.
 - ii) In the **Certificate content** field, copy and paste your certificate in the format that is mentioned in the field content. You can retrieve a PEM certificate for your database. The file must be a Base-64 encoded X.509 file. You do not need to retrieve a private key. For more information, see the documentation for your database. After you copy and paste the text into the field, including the BEGIN CERTIFICATE, and END CERTIFICATE text, click **Confirm**.
 - f) Click **Save**. The **Database connection** page is closed.
 - g) Click in the **Database connection** tile to verify your database connection. Expand the **Status** icon that is loading in the Configuration-scope-Workspace-application section to display some tiles. Click **Select** once the **Status** icon is ready to close the page.

Note: Click **Save and Select** in the previous step if you do not want to wait for the database connection verification after you complete the fields.
 - h) If you plan to deploy Maximo Optimizer, see [Deploying Maximo Optimizer](#).
4. In the **Components** section, select the industry solutions or add-ons and the version that you are also activating with Maximo Manage.
 - If you select **latest** from the **New version** list for the selected industry solution or add-on, the version of the component that is supported by the current Maximo Manage application version is co-deployed with Maximo Manage base component.

Note: If your configurations settings were set to the latest for the components before you upgrade the application, the components will be automatically updated after the upgrade. If you select a specific version instead, then the components are not updated until you change their versions to the current, supported version of the new application version that was now updated, both by selecting the exact version or by selecting **latest**.
 - If you select one component that is going to be deployed by selecting the **latest** option, you must select **latest** for any other component you co-deploy. If you select one component that is going to be deployed with an exact version number, you must select the exact version number for any other extra component that you co-deploy. You cannot mix exact version numbers and the latest version in the components you want to co-deploy.
 - You also can click in the **New version** column for a component and select **Select version**. In the **Select unlisted version** dialog, you can specify a valid version that is supported by the current application version that is deployed in the **Version** field and click **Save**.

For more information, see [Deployment of industry solutions and add-ons](#) section, including access to the compatibility matrix that shows the compatibility between application version and components versions. Some components might not be compatible with each other.

5. Click **Show advanced settings** to view and specify the configuration settings, such as database, server bundle, language settings, and others.

- a) In the Database section, clear the **System managed** checkbox and manually configure the database.

Schema

Enter the name of the schema that is configured in your database. For more database configuration information for Maximo Manage, see [“Setting up your database”](#) on page 301.

Encryption secret (optional)

This value is optional if you are deploying Maximo Manage in your database for the first time, and your database is not encrypted. Enter your encryption keys for this parameter. For more encryption settings information, see [Database encryption](#).

In the Key/Value table, click **Add property +**. In the Key column, enter `MXE_SECURITY_CRYPTOX_KEY` and in the Value column, enter your encryption key value.

In the Key/Value table, click **Add property +**. In the Key column, enter `MXE_SECURITY_CRYPTO_KEY` and in the Value column, enter your encryption key value.

Table space

If the default value does not match your database configuration, enter the name of the table space that was configured in your database. For more database configuration information for Maximo Manage, see [“Setting up your database” on page 301](#).

Index space

If the default value does not match your database configuration, enter the name of the index table space that was configured in your database. For more database configuration information for Maximo Manage, see [“Setting up your database” on page 301](#).

Install demo data

If you are deploying a test or demonstration environment for Maximo Manage, you can install sample data.

The sample data in the demo database is useful for development or test environments.

To set up a test or development environment with demo data, install an instance of Maximo Application Suite specifically for testing or development. Then, when you configure the database settings for your Maximo Manage deployment, select the option to install demo data.

Note: You cannot add sample data after Maximo Manage is deployed because the database is updated without sample data. To add the sample data after deploying Maximo Manage, you must re-create or clean your database, and reconfigure Maximo Manage.

Db2 Vargraphic

If you use Db2 and you plan to install a language other than English for your base language or as an extra language, select this option. If you intend to add more languages later, select this option during your initial deployment. This option does not affect the Maximo Manage deployment if it is selected and you are using a database other than Db2.

Bypass upgrade version check

Select this option to skip validation of the IBM Maximo Asset Management version you are upgrading to Maximo Manage. Select this option to continue a failed upgrade that failed during the `maxinst` or `updatedb` process.

6. If you want to set Maximo Manage with a language different from English or include other languages in your Maximo Manage deployment, clear the **System managed** checkbox in the Languages section. Then, in the **Base** field, select your preferred language to be the base language and in the **Additional** field, order the list of other languages. For more information, see [Language support](#).

Note: If you are selecting other languages, make sure that you do not select a language in the **Additional** field that was selected in the **Base** field. For example, if you set the base as **EN**, do not select **EN** in the **Additional** field.

7. Configure server bundles for your deployment.
 - a) If you want to deploy Maximo Manage with more than one server bundle or with customized configurations for it, under the Server Bundles section, clear the **System managed** option. A table with **Name**, **Pod count**, **Type**, and **Additional Properties** is displayed. This table has the **Default**, **User synchronization**, and **Mobile** optional fields available. When you clear the **System managed** option, a line with the server bundle named as *all* is set with one Pod with the Type *all* and Route subdomain as *all*.
 - b) Click **Add bundle** to add more server bundles.
 - c) Select which server bundle to set as respectively the Default, the server bundle to be used for user synchronization, and the server bundles for Mobile.
 - d) Optional: Change the name, pod count, type, route subdomain name and other customized configurations according to your preference.
 - e) Click the **View** label under **Additional Properties** column to view the **Route subdomain** value.

The **Additional server bundle properties** page for your selected server bundle is displayed with the **Route subdomain** and **Additional server config** fields available and a list of properties you can define by selecting **Add property +** option in the **Bundle level properties** table. For more information, see [Server bundle overview](#).

8. If you want to include specific customizations through a customization archive, clear the **System managed** checkbox in the Customizations section. In the **File address** field, specify the location of the customization archive and if you must enter credentials to access the file, specify them in the **Credentials (optional)** field. For more information, see the related sections in [Customizing the application](#).

- a) In the Customization section of the configuration window, specify the URL for the customization archive file.

The following URL protocols are supported:

- HTTP
- HTTPS
- FTP
- FTPS

To include more customization archive files, click **Add customization archive**.

- b) Optional: If you applied password security to the file, in the **Credentials** field, specify the user ID and password in the following format:

```
user=your user name password=your password
```

9. If you do not want the server bundles to start after the database operations of the Maximo Manage deployment are completed, clear the **System managed** checkbox in the **Server mode** section.

Then, set **Mode** to **Off** to prevent access to the Maximo Manage application after deployment. You can restart the server bundle or bundles when you change the configuration to **On** and activate Maximo Manage again.

10. To connect to PVCs under the **Persistent volume claims** section, clear the **System managed** to **Off**, and click **Add PVC**. A table with the following columns is displayed.

Option	Description
PVC name	User-defined name of the persistent volume claim, maximum of 63 characters
Volume name	Leave blank as it is provisioned dynamically.
Size	Amount of storage that is required for this persistent claim, for example, 60G
Mount path	Mount path for the volume within the Maximo Manage pod.

When you configure the PVCs in OpenShift Container Platform cluster on deployments, use the default storage class name `StorageClasses ocs-storagecluster-cephfs` to create a **ReadWriteMany** (rwx) PVC. The storage in the volume that you provisioned are available to all server bundles in the workspace. You can also configure a PVC for specific server bundles in a deployment. If you configure a PVC for a server bundle, the mount path that you specify for the server bundle PVC overrides the path that you specify for the deployment.

Tip: To configure a PVC in OpenShift Container Platform cluster on IBM Cloud platform, use the default `StorageClasses ibmc-file-gold-gid` (instead of `StorageClasses ocs-storagecluster-cephfs`) to create a **ReadWriteMany** PVC.

11. If you select **Asset Configuration Manager** or **Aviation** in the list of components, the Build data interpreter section is displayed. If you want to customize the configuration for the build data interpreter (BDI), clear the **System managed** checkbox. You can then specify a BDI version instead of latest. In the **BDI version** field, click **Add instance+**. You can customize each instance by selecting **View** in the Configuration column. The **BDI Configuration** page is displayed. You can change and save the configuration. Then, you can return to the **BDI configuration** page, select **Reset to Defaults**, and click **Save** to return to the default settings.

12. If you want to use an earlier build for deployment, in the Build section, set **System managed** to **Off**. Then, in the **Build tag** field, specify the build tag. Build images are tagged with a timestamp, for example buildtag: 202011092887843.
13. To connect to any external systems that Maximo Manage with is integrated with, import the certificate for the system.
14. Optional: You might want to specify the time zone that your database server is configured to use. In the Server time zone section, clear the **System Managed** checkbox. In the **Time Zone** field, select the time zone of your database server.
15. If you are deploying Maximo Health as part of Maximo Manage, you can set following configurations.
 - You can enable asset investment optimization. In the Asset investment optimization section, clear the **System managed** checkbox. Then, select **Asset investment optimization**. When asset investment optimization is enabled, the **Asset investment optimizer** page is available in Maximo Health. Ensure that you deploy and configure Maximo Scheduler Optimization before you enable asset investment optimization.
 - If you have IBM Watson Studio and want to use existing models of Maximo Health from Health and Predict - Utilities, in the IBM Watson Studio section, clear the **System Managed** checkbox and then specify the **Watson Studio Project ID**.
 - If you want the Maximo Health to deploy the out of box data loader configuration files and create the integration server, in the IBM App Connect section, clear the **System Managed** checkbox and then specify the **Dashboard URL**.

Activating Maximo Manage

Before you can grant users access and start working with Maximo Manage, you must activate the application. Activating Maximo Manage triggers the second phase of the application's deployment.

What to do next

- Complete the post-deployment configurations, such as giving permission to the Maximo Manage admin user to connect to Maximo Manage. After the Maximo Manage admin user is synchronized, log out and log in again to Maximo Application Suite as the admin user and access Maximo Manage.
- When new versions are available, system administrators can update the deployed application. To update to a new version, in Maximo Application Suite, in **Suite administration**, select **Applications** from the side navigation menu, and click **Update available for Manage**.

When updates are required, system administrators can also reconfigure and update initial implementations on the **Manage workspace details** page. To reconfigure and update changes, select **Workspaces** from the side navigation menu, and click **Manage**. On the **Manage workspace details** page, click **Actions**, and select **Update configuration**.

You can perform a rolling upgrade from a previous version of Maximo Manage to a later version of Maximo Manage. With a rolling upgrade, the interruption to Maximo Manage is minimal when you upgrade both operator and operands. During a rolling upgrade from one version to another version, the user session becomes invalid. Therefore, you must authenticate to continue your work. For more information, see [Upgrade and Rollback](#).

Customer-managed **Migrating**

You must upgrade to Maximo Application Suite and then deploy Maximo Manage when you upgrade from Maximo Asset Management. You must perform some key tasks and understand important information after you deploy and activate Maximo Manage.

Related concepts

[“Installing Maximo Application Suite” on page 218](#)

Report migration

Report components are automatically upgraded as a part of the database migration process. However, you must back up existing report data from Maximo Asset Management and import the data into the newly installed Maximo Application Suite.

Business Intelligence and Reporting Tools (BIRT)

You can use BIRT reporting in Maximo Application Suite. Maximo Application Suite supports BIRT version 4.8.

- When you migrate from Maximo Asset Management 7.6.1.2 to Maximo Application Suite, BIRT reports are automatically upgraded from version 4.3.1 to version 4.8.
- Because Maximo Asset Management 7.6.1.3 already supports BIRT version 4.8, migrating to Maximo Application Suite does not require an upgrade for BIRT.

To migrate existing report properties and designs, export the report properties and design from Maximo Asset Management and import them in Maximo Application Suite. For more information, see [Administering reports](#).

Cognos Analytics

You can use Cognos Analytics with Maximo Application Suite. Migrating from Maximo Asset Management 7.6.1 to Maximo Application Suite upgrades Cognos Analytics to version 11.2.4. To migrate Cognos Analytics reports, back up your current data from the Cognos content store and import the data in Cognos Analytics version 11.2.4. For more information, see [Backing up and restoring Cognos Analytics](#).

External report integration

Any external report integrations are migrated with the customization archive. For more information, see [Migrating customizations using customization archive](#).

Integrating with external systems

Maximo Manage has an in-built Java Messaging Service (JMS) provider and it is possible to use an external JMS provider as well.

About this task

Maximo Application Suite 8.8 onwards supports an in-built JMS provider that you can use for integration purposes in Maximo Manage.

Procedure

1. Procure, install, and configure the external JMS provider. For example, IBM MQ or Liberty messaging engine.
2. Register the JMS provider queues in Maximo Manage.
3. Associate the registered queues by using the **External Systems** application in Maximo Manage. For more information, see [Adding JMS queues to an external system](#).

Note: If you are using the Liberty messaging engine, you must register the queues and connection factories in Maximo Manage by using the server.xml file for Liberty. For more information, see [Enabling JMS messaging for Liberty](#).

Adding server bundle properties

Use Maximo Application Suite to add or update server bundle properties.

Procedure

1. In Maximo Application Suite, click the **Administration** icon.
Starting in Maximo Application Suite 9.1, select the **Suite > Administration** page.
2. From the side navigation menu, select **Workspace**.

3. On the **Workspace** page, click the **Manage** tile.
4. On the **Manage workspace details** page, click **Actions** and then click **Update configuration**.
5. In the **Update Manage configuration** dialog, in the **Activation configuration** section, click the **Edit** icon for **Server bundles**.
6. In the **Server bundles** section, click **View** for the server to which you want to add properties.
7. In the **Additional server bundle properties** dialog, in the **Bundle level properties** section, add the bundle-specific properties .
8. Click **Save**.
9. Click **Activate** to activate your changes.
10. Validate the bundle-level properties by using the following API call.

```
{UI_SERVER_URL}/api/service/system?
action=wsmethod:getProperty&propName=mxe.int.webappurl&apikey={API_KEY}
```

Sample custom resource:

```
{
  "spec": {
    "bindings": {
      "jdbc": "workspace-application"
    },
    "components": {
      "base": {
        "version": "latest"
      }
    },
    "settings": {
      "db": {
        "maxinst": {
          "bypassUpgradeVersionCheck": false,
          "db2Vargraphic": true,
          "demodata": false,
          "indexSpace": "maximo",
          "tableSpace": "maximo"
        },
        "dbSchema": "MAXIMO"
      },
      "deployment": {
        "serverBundles": [
          {
            "bundleType": "ui",
            "isDefault": true,
            "name": "default",
            "replica": 1,
            "routeSubDomain": "ui",
            "bundleLevelProperties": "mxe.int.webappurl=http://localhost/
ui\nmxe.webclient.activitydashboard=TRUE"
          },
          {
            "name": "cron",
            "replica": 1,
            "bundleType": "cron",
            "isDefault": false,
            "routeSubDomain": "cron",
            "bundleLevelProperties": "mxe.int.webappurl=http://localhost/
cron\nmxe.cronTaskMonitorInterval=30"
          },
          {
            "name": "report",
            "replica": 1,
            "bundleType": "report",
            "isDefault": false,
            "routeSubDomain": "report",
            "bundleLevelProperties":
            "mxe.report.reportsInAPage=10\nmxe.int.webappurl=http://localhost/report"
          },
          {
            "name": "mea",
            "replica": 1,
            "bundleType": "mea",
            "isDefault": false,
            "routeSubDomain": "mea",
            "bundleLevelProperties":
            "mxe.adminEmail=email@ibm.com\nmxe.int.webappurl=http://localhost/meaweb"
          }
        ]
      }
    }
  }
}
```

```
}
  }
    }
      }
        ]
          }
```

Server bundles

You can add or update server bundle properties when you deploy IBM Maximo Manage. For each server bundle, a service and a route that points to the service are created.

When the IBM Maximo Manage application is deployed in IBM Maximo Application Suite, a Maximo Manage deployment custom resource (CR) is created.

- The Maximo Manage deployment CR contains the configuration that you entered, which includes database URL, server bundle types, deployment size, and other information.
- The database username or password can be specified in Maximo Application Suite and persisted as a secret in the cluster.
- The Maximo Manage **crypto** or **cryptox** properties can be updated in Maximo Application Suite and persisted as a secret in the cluster.
- Both secrets, the username or password and the **crypt** or **cryptox** values can also be updated in Red Hat OpenShift. Any change to the values automatically redeploys the application completely or partially.

Server bundle

A server bundle or workload is a logical abstraction for a deployed group of point of deployment (pods) in a cluster. These server bundles do the same function and provide an access point as a service. These server bundles can be accessed as a service internally and through a route externally. A route is a way to display a service by giving it an externally reachable hostname. By using a route or service, Red Hat OpenShift balances the load to the pods that are included in a server bundle. Each server bundle defines replica size, subdomain, and other properties.

- For each server bundle, a service is created with name that is appended by **-<serverbundlename>**.
- A route is created with name that is appended by **-<serverbundlename>**.
- A default route is created to point to the service that ends with **-<serverbundlename>**.
 - Maximo Application Suite uses the default route to establish the default URL link to Maximo Manage.
- You can add or update server bundle properties. For more information, see [Adding server bundle properties](#).

The following diagram illustrates how Red Hat OpenShift Container Platform routers provide external hostname mapping and balancing of load for service end points over protocols. The router uses the hostname to determine where to send the external client request.

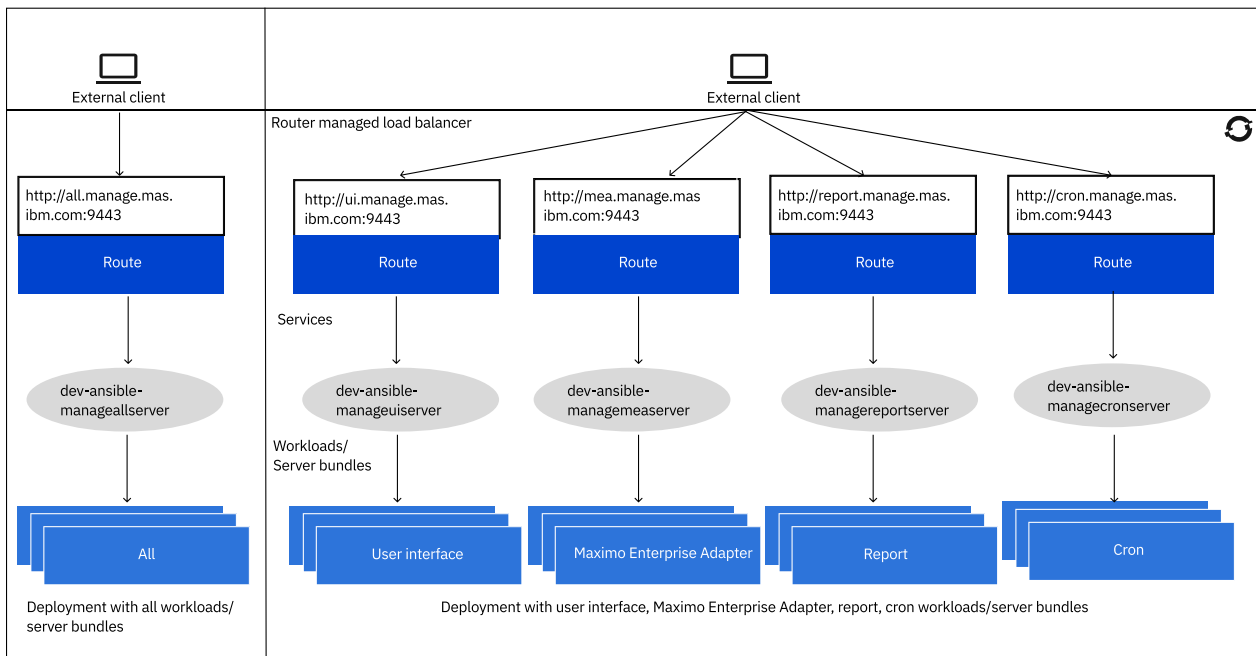


Figure 6. External host mapping and load balancing

- The Maximo Manage application can be deployed with different server bundles or workloads for the processing and isolation needs.
- The deployment can be All bundle server type or a combination of the four bundle server types: UI, cron, report, and Maximo Enterprise Adapter.
 - If All bundle server type is not deployed, and you used a combination of the four bundle server types, you must use the UI bundle server type.
- Each server bundle can have its own server properties.

The following table shows the five different server bundle types.

Bundle server type	Description
All	Contains all the code.
UI	Contains UI code and supporting code. It is the interface for accessing Maximo Manage application.
Maximo Enterprise Adapter	Displays the enterprise web services API.
report	Contains the code that is needed to enable BIRT Report Only Server (BROS). It is used to separate the work load that is related to the execution of reports that are submitted in Maximo Manage.
cron	This bundle contains the code that is needed to run Maximo Manage cron tasks.

Server bundle properties

The server bundles have the following properties:

- The server bundle properties can be set in Maximo Application Suite UI or in the CR.
- A configmap <workspaceid>-<serverbundlename>-bundleproperty is created for the server properties during deployment or operator reconciliation. It is mounted to /config/manage/

properties on the pod. The **bundleLevelProperties** file on the pod must not be updated manually.

- Maximo Manage server process automatically detects the change and updates the Maximo Manage property cache.
- If the property value needs to be modified, update the CR directly or in Maximo Application Suite. The Maximo Manage Operator reconciles the changes. The Maximo Manage server process updates the Maximo Manage property cache with the updated value.

Liberty server XML

You might need to customize the Liberty `server.xml` file, for example, when creating queues.

- The custom server XML can be set in Maximo Application Suite UI. For example, to create queues in JMS server.
 - If you create it in the CR directly, then you must create the secret manually. Set this secret to **additionalServerConfig.secretname** for the bundle server in the CR.
- On deployment, a secret is created for the custom `server.XML` file. It is mounted to **/config/manage/serverxml**. This location is included in Liberty `server.xml` file. The configuration is applied to the Liberty server.

Sample custom `server.xml` file:

```
<?xml version="1.0" encoding="UTF-8"?>
<server description="new server">
  <featureManager>
    <feature>wmqJmsClient-2.0</feature>
    <feature>jmsMdb-3.2</feature>
  </featureManager>
  <logging traceSpecification="JMSApi=all:WAS.j2c=all"/>
  <variable name="wmqJmsClient.rar.location" value="/wmq/wmq.jmsra.rar"/>
  <jmsConnectionFactory jndiName="jms/maximo/int/cf/intcf" connectionManagerRef="MIFJMS">
    <properties.wmqJms
      transportType="CLIENT"
      hostName="mifjmsmanager-afd7.qm.us-south.mq.appdomain.cloud"
      port="31440"
      channel="CLOUD.APP.SVRCONN"
      applicationName="maxliberty"
      userName="{{username}}"
      password="{{yourpassword}}"
      queueManager="MIFJMSMANAGER"/>
  </jmsConnectionFactory>
  <connectionManager id="MIFJMS" maxPoolSize="20"/>
  <jmsQueue id="sqout" jndiName="jms/maximo/int/queues/sqout">
    <properties.wmqJms baseQueueName="sqout" baseQueueManagerName="MIFJMSMANAGER"/>
  </jmsQueue>
  <jmsQueue id="sqin" jndiName="jms/maximo/int/queues/sqin">
    <properties.wmqJms baseQueueName="sqin" baseQueueManagerName="MIFJMSMANAGER"/>
  </jmsQueue>
  <jmsQueue id="jms/maximo/int/queues/cqin" jndiName="jms/maximo/int/queues/cqin">
    <properties.wmqJms baseQueueName="cqin" baseQueueManagerName="MIFJMSMANAGER"/>
  </jmsQueue>
  <jmsQueue id="jms/maximo/int/queues/cqinerr" jndiName="jms/maximo/int/queues/cqinerr">
    <properties.wmqJms baseQueueName="cqinerr" baseQueueManagerName="MIFJMSMANAGER"/>
  </jmsQueue>
</server>
```

User migration

You can create new users in IBM Maximo Application Suite users or migrate from Maximo Asset Management.

Migrated users

Users are migrated to Maximo Application Suite from IBM Maximo Asset Management during the upgrade.

Users are created with the **Set in Manage** access type. The user access is managed in the IBM Maximo Manage application by using security groups. For more information, see [Configuring security groups](#).

If SMTP is configured, the migrated users receive the following emails:

- Welcome to IBM Maximo Application Suite.
- Your IBM Maximo Application Suite password.

On the **Users** page, you can replace passwords in edit mode by using the **Replace forgotten password** link. A MAXADMIN user is created by default. User authorization or application access is done in the Maximo Manage application by using security groups. .

Customer-managed

Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 8.11

In Maximo Application Suite 8.11 and earlier versions, the person ID is equal to the user ID in Maximo Manage. If you are migrating users from Maximo Asset Management to Maximo Application Suite from an LDAP server and the person ID is different than the user ID, you can configure the user data to maintain the person ID data.

Before you begin

Before you can configure the user data to maintain the person ID, install Maximo Application Suite 8.11.7 or 8.10.10 fix packs or later. You can install fix packs by using a channel subscription. For more information, see [“Upgrading IBM Maximo Application Suite by using the channel subscription method” on page 477](#).

If you are using Maximo Application Suite 9.0, see [“Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 9.0” on page 559](#)

About this task

When a user is added in Maximo Application Suite, the user ID is created as the primary identifier for that user. When user synchronization occurs with Maximo Manage, the user ID is added in Maximo Manage as the user ID and the person ID.

If you need to maintain separate person ID data to meet your business needs, create an automation script in Maximo Manage that copies the value from the LDAP `employeeNumber` field to the `personid` field during the user synchronization process. You can also use this automation script to map to other fields in Maximo Manage.

If the field you are using in the LDAP server is not `employeeNumber`, you must also map the field to `employeeNumber` in the `ScimCfg` LDAP configuration record.

Procedure

1. Create an automation script that copies the value from the `employeeNumber` field to the `personid` field,
 - a) In Maximo Manage, open the Automation Scripts application.
 - b) From the **More Actions** menu, select **Create > Script for Integration**.
 - c) In the Integration Details section, select **Enterprise Services** and in the Enterprise Service field, select `MASPERUSER`.
 - d) Select **Request**, **User Exit**, and **After External Exit**.
 - e) In the Script Details section, enter the following code as Jython code:

```
from com.ibm.tivoli.maximo.oslc import OslcUtils
from com.ibm.tivoli.maximo.oslc.provider import OslcJSONStructureData

data = irData.getDataAsBytes()
dataEr = erData.getDataAsBytes()

jo = OslcUtils.bytesToJSONObject(data)
erJo = OslcUtils.bytesToJSONObject(dataEr)
```

```

if erJo.get("owner").upper() == 'SCIM':
    extjo = jo.get("extension")
    empno = extjo.get("employeeNumber")
    jo.put("personid", empno)
    irData = OslcJSONStructureData(jo, "MASPERUSER", "PERSON", userInfo, "Sync", True)

```

2. If the field in the LDAP server is not `employeeNumber`, update the `ScimCfg` LDAP configuration record to map the field to `employeeNumber`.
 - a) In Red Hat OpenShift, from the side navigation menu, click **Administration > CustomResourceDefinitions**.
 - b) On the **CustomResourceDefinitions** page, search for and open the `ScimCFG` custom resource definition.
 - c) On the **Instances** tab, search for and open the CR that starts with the Maximo Application Suite instance ID.
For example, `<your_mas_instance>-scim-default-system`.
 - d) On the **YAML** tab, in the `spec:` section under `usersync`, map the field to `employeeNumber`.

For example, if the field in LDAP is named `employeeID`, add the following configuration in `spec.usersync` to map to `employeeNumber`.

```

userSync:
  mappings:
    extensions:
      employeeNumber: employeeID

```

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Customer-managed

Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 9.0

In Maximo Application Suite 9.0, the person ID is equivalent to the user ID in Maximo Manage. If you are migrating users from Maximo Asset Management to Maximo Application Suite from an LDAP server and the person ID is different than the user ID, you can configure the user data to maintain the person ID data.

Before you begin

If you are using Maximo Application Suite 8.11 or earlier, see [“Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 8.11”](#) on page 558

About this task

When a user is added in Maximo Application Suite, the user ID is created as the primary identifier for that user. When user synchronization occurs with Maximo Manage, the user ID is added in Maximo Manage as the user ID and the person ID.

If you need to maintain separate person ID data, you can customize the mapping for the person ID to synchronize with the value from the LDAP server during the user synchronization process. For example, you can customize the mapping for the person ID to map to the `employeeNumber` field in LDAP.

Procedure

1. On the **Suite administration** page, from the side navigation menu, select **Configurations**, click **User registry synchronization** and then **Edit**.
2. Specify the following LDAP domain attributes for the LDAP server.

LDAP URL

The URL for the LDAP server.

Base DN

The path in the object hierarchy of the directory server.

Bind DN

The user and location that is used to bind to an LDAP server.

Maximum user synchronization

The maximum number of users that are synchronized between the LDAP server and the user registry. If the search results of the LDAP database exceed this limit, the synchronization process is canceled. This property is the `customMaxSearchResults` property in the `ScimCfg` custom resource.

LDAP type

The type of LDAP server that you are using. For example, select Microsoft Active Directory if that is the LDAP server that you are using.

3. Add or retrieve the [CA certificate](#) for the LDAP server.
4. Specify the user synchronization, such as User Base DN, ID map, and filter.
5. In the User mapping section, switch **Use default mapping** to off and enter the custom mapping for the person ID user data to synchronize with the employee number from the LDAP server.
 - a) Click **Add custom mapping**.
 - b) In the Maximo Application Suite field column, enter `person.personid` in the field after extension.

By specifying the `person` value, the `personid` maps to the `person` table in Maximo Manage
 - c) In the LDAP field column, enter `employeeNumber`.
6. Optional: Specify group synchronization information, such as Group Base DN, filter, ID map, and member ID map.
7. Optional: Enter the custom mapping for the group data to synchronize with the LDAP server.
8. Optional: Assign the default application entitlement and application access to apply to all synced users.

You can also modify the entitlement and access for individual synced users on the **Users** page.

Results

When you save the mapping changes, the configuration is processed, and the user synchronization changes are applied in the next scheduled synchronization.

Related concepts**[LDAP user registry synchronization](#)**

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Managing users post upgrade

Migrated users are synchronized by a cron task, and you can view the synchronization information. If your migrated users have different user IDs and login IDs and you are upgrading to Maximo Application Suite 8.7 or earlier, you must re-create your existing users. User entitlement to applications and AppPoints, and license consumption reports are viewable in the **Suite administration** dashboard. You can also change user passwords by editing the user profile on the **Users** page in Maximo Application Suite.

Synchronizing migrated users

The `MASUSERSYNC` cron task synchronizes users from Maximo Manage to Maximo Application Suite.

Before you begin

You must have administrator rights to set up the `MASUSERSYNC` cron task instance. If the users are not migrated to Maximo Application Suite, check the database log and fix any errors.

Procedure

1. On the **Users** page in Maximo Application Suite, create an administrative user . For more information, see [Administering users and user access](#).
2. Log in to Maximo Application Suite by using the administrative user's credentials.
3. Click the **Administration** icon to open the **Suite navigator** page and then click the Manage tile.
4. From the side navigation menu, click **System Configuration > Platform Configuration > Cron Task Setup**.

Note: Starting in Maximo Application Suite 9.1, the Manage application is available in the side navigation menu. You can access the **Cron Task Setup** page directly from the side navigation menu in **System configuration > Platform configuration** under the Manage application.

5. Search for and select the MASUSERSYNC cron task.
6. On the **Cron Task** tab, click **New Row**.
7. Specify a name and a schedule for the cron task instance.
The date is shown as a string in the **Schedule** field. Do not change the string in the **Schedule** field. Click **Set Schedule** to change the schedule.
8. Specify a user who has the necessary privileges. The user must have access for the actions that the cron task performs.
9. Optional: Select the **Active** check box if you want to activate the cron task.
10. Save your changes.
11. Select the **Reload Request** action.
12. Select the instance and click **OK** to run the cron task.

User synchronization information

You can view information about user synchronization in IBM Maximo Manage.

- When a user is created and is assigned the Manage role, the user record is synchronized into Maximo Manage. The user sync process pulls the data from the user registry in IBM Maximo Application Suite and pushes it to Maximo Manage.
- If a user is assigned to the NO_ACCESS role or is deleted in Maximo Application Suite, the user becomes inactive in Maximo Manage.
- The user synchronization process from Maximo Application Suite to Maximo Manage uses the Maximo integration framework. It uses an enterprise service to post data to Maximo Manage. If a IBM Maximo Application Suite user is assigned the Manage role and the user record has a sync status of PENDING or DELETE_PENDING, the user synchronization process processes the user record. After successful synchronization, the status is changed to SYNCED.

Sample of inbound message from Maximo Application Suite

```
{ "_id": "joesmith5", "username": "joesmith5", "permissions": { "systemAdmin": true,
"userAdmin": true }, "issuer": "local", "displayName": "Joe Smith", "familyName": "Smith",
"givenName": "Joe", "title": "Supervisor", "preferredLanguage": "EN", "locale": "en_US",
"phoneNumbers": [ { "value": "555-555-5555", "type": "work" }, { "value": "555-555-4444",
"type": "work" } ], "addresses": [ { "streetAddress": "100 Universal City Plaza", "locality":
"Hollywood", "region": "CA", "postalCode": "91608", "country": "USA", "formatted": "100
Universal City Plaza\nHollywood, CA 91608 USA", "type": "work", "primary": true } ], "emails":
[ { "value": "joesmith5@us.ibm.com", "type": "work", "primary": true }, { "value":
"joesmith51@gmail.com", "type": "work" } ], "extension": { "employeeNumber": "701984",
"costCenter": "4130" }, "entitlement": { "application": "PREMIUM", "admin": "ADMIN_PREMIUM" },
"workspaces": { "space1": { "permissions": { "workspaceAdmin": true }, "applications":
{ "manage": { "role": "ADMIN" }, "iot": { "role": "ADMIN" }, "health": { "role": "USER" },
"monitor": { "role": "ADMIN" } } } }, "added": { "id": "admin", "timestamp":
"2020-08-10T18:01:36.694331" }, "updated": { "id": "admin", "timestamp":
"2020-08-10T18:08:17.455782" }, "sync": { "status": "SUCCESS", "timestamp":
"2020-08-10T18:10:16.731047" }, "applications": { "manage": { "sync": { "state": "PENDING",
"reason": "", "timestamp": "2020-08-10T18:08:49.604430" } }, "monitor": { "sync": { "state":
"SUCCESS", "reason": "", "timestamp": "2020-08-10T18:08:17.471947" } }, "health": { "sync":
{ "state": "SUCCESS", "reason": "", "timestamp": "2020-08-10T18:09:18.131171" } }, "predict":
{ "sync": { "state": "SUCCESS", "reason": "", "timestamp": "2020-08-10T18:08:17.471947" } } },
"token":
```

```
"1000:8744077b1411c0601e4912d556d93ff859089bfd16863f16:591949ad7b4d4f7017de846a3f3b2609ac4caef4e
be09448" }
```

Sample of output from default user exit

```
{ "addressline1": "100 Universal City Plaza", "city": "Hollywood", "country": "USA",
"displayname": "Joe Smith", "email": [ { "_action": "AddChange", "emailaddress":
"joesmith5@us.ibm.com", "isprimary": 1, "type": "work" }, { "_action": "AddChange",
"emailaddress": "joesmith51@gmail.com", "isprimary": 0, "type": "work" } ], "extension":
{ "costCenter": "4130", "employeeNumber": "701984" }, "firstname": "Joe", "language":
"EN", "lastname": "Smith", "locale": "en_US", "maxuser": [ { "groupuser": [ { "_action":
"AddChange", "groupname": "MAXADMIN" }, { "_action": "AddChange", "groupname": "TOOLMGR" } ] },
"inactivesites": 1, "loginid": "joesmith5", "statusdate": "2020-10-30T15:43:44-04:00",
"userid": "joesmith5" } ], "personid": "joesmith5", "phone": [ { "_action": "AddChange",
"isprimary": 1, "phonenumber": "555-555-5555", "type": "work" }, { "_action": "AddChange",
"isprimary": 0, "phonenumber": "555-555-4444", "type": "home" } ], "postalcode": "91608",
"stateprovince": "CA", "statusdate": "2020-10-30T15:43:44-04:00", "title": "Supervisor" }
```

Troubleshooting user migration

The migration process uses an internal API to migrate user IDs from Maximo Asset Management to Maximo Application Suite. If you have users with a different user ID and login ID and you are upgrading to Maximo Application Suite 8.7 or earlier, you must re-create your existing users.

About this task

Maximo Application Suite has two fields for user management, username and user ID. If any of your Maximo Asset Management users have a user ID different from their login ID, you must delete and re-create the users after you upgrade to Maximo Application Suite.

Procedure

1. If the user has only Maximo Manage license entitlement:
 - a) In the Maximo Application Suite user interface, login as an administrator and go to **Suite Administration > Users**.
 - b) Search for the user by using the Maximo Asset Management login ID as the search criterion.
 - c) Create a new user based on the information of the existing user and set the user ID as a unique key and login ID.
For more information, see [Administering users and user access](#).
 - d) Delete the old user.
 - e) Run the **MASUSERSYNC** cron task to synchronize the users from Maximo Application Suite to Maximo Manage
For more information, see [Synchronizing migrated users](#).
2. If the user has entitlement to multiple products in Maximo Application Suite, for example, Maximo Manage and IoT:
 - a) In the Maximo Application Suite user interface, login as an administrator and go to **Suite Administration > Users**.
 - b) Create a user based on the information of the existing user and set the user ID as a unique key and login ID.
 - c) Change the login ID of the existing user by adding **_bk** to it.
 - d) Delete the email address of the existing user to avoid duplicate email addresses.
 - e) Do not synchronize the user to any products in Maximo Application Suite.

Tip: In the case of a user in multiple products, preserve the original user information from Maximo Asset Management even if the user cannot log in to the system by using the Maximo Asset Management login ID. If you want to keep an old user, you can change the **Username** and use the user. If you decide to delete it, you can delete it yourself and add the entitlement to the new user. To find these users easily, search for **Usernames** with "**_bk**".

User entitlement

License entitlement for Maximo Manage can be determined by using a cron task, `MasUserAnalyzer`.

- In Maximo Manage, use a cron task to calculate user entitlement.
 - The `MasUserAnalyzer` cron task calculates the user entitlement, which can be Limited, Base, Premium, or Self-Service.
 - The `MasUserReporting` cron task sends the entitlement change to Maximo Application Suite.

Note: If the entitlement is greater in Maximo Manage, a message is displayed on the **Suite administration** dashboard.

- For more information on how the cron task works, see [License type information](#).

To generate License consumption reports, open the **Suite administration** dashboard and click **License consumption**.

Changing user passwords

User passwords are changed on the **Users** page in Maximo Application Suite.

Procedure

1. Maximo Application Suite, from the side navigation menu, select **Users**.
2. Click on the user record to go to **View user** page.
3. Click the **Edit** icon to edit the user details.
4. In the **Authentication** section, click **Replace forgotten password**.
5. Select **Send password by email** to receive an email that contains the password.
6. Select **Autogenerated** to automatically generate the password or select **Custom** to manually specify the new password.
7. Click **Save changes**.

Changing current user password

On the **Suite administration** dashboard, you can change the password for the currently logged in user.

Procedure

1. In Maximo Application Suite, in the main menu bar, click the **Profile** icon.
2. Click **Manage profile**.
3. On the **Change password** tab, specify the current password and the new password and then click **Save**.

Updating system settings path

Update system settings with the new values.

About this task

The system settings depend on your current Maximo Application Suite setup. For example,

- `mxe.doclink.path01`
- `mxe.doclink.doctypes.defpath`
- `mxe.doclink.doctypes.topLevelPaths`

Secured attachments must always be enabled for Maximo Application Suite or Maximo Manage.

Procedure

1. Update the path in the doc-link table:

```
update docinfo set urlname = replace(urlname, 'C:', '');
```

```
update docinfo set urlname = replace(urlname, '\\', '/');
```

```
update docinfo set urlname = replace(urlname, 'DOCLINKS', 'doclinks');
```

2. Verify the doc information table because the directories are case-sensitive in Maximo Manage.

Using certificates

You can use route or TLS certificates, SSL certificates, external certificates.

Route or TLS certificates

Maximo Application Suite uses cert-manager for automatic management and issuance of TLS certificates for application routes. During installation, you can provide a cluster issuer that is based on a trusted certificate authority (CA) for signing the certificates that are generated for your Maximo Application Suite domains.

By default, Maximo Application Suite provides a cluster issuer that generates Maximo Application Suite certificates that are signed by a self-signed CA. To use your own cluster issuer, include the following parameter when you run the Maximo Application Suite installer: `-c myClusterIssuerName`. For more information, see [System requirements](#).

Obtaining SSL certificate for database

If the database requires an SSL connection, you must obtain the certificate for the database. You can use the **openssl** command to connect to the database host and port that is specified in the database URL

Procedure

1. Run the **openssl** command: `openssl s_client -showcerts -connect databasehost:databaseport`
2. While activating Maximo Manage, go to the **Database** section and open it.
3. Give the certificate an alias that does not conflict with other certificates in the trust store.
4. Make sure that **SSL Enabled** is set to Yes in the UI.
5. Specify the certificate on the `jdbccfg` CR in the certificates section as shown in the following example.

Sample Custom Resource (CR)

```
apiVersion: config.mas.ibm.com/v1
kind: JdbcCfg
metadata:
  name: "mng-jdbc-system"
  namespace: "mas-mng-core"
  labels:
    mas.ibm.com/configScope: system
    mas.ibm.com/instanceId: "mng"
spec:
  displayName: IBM Cloud Databases for Db2
  config:
    url: "jdbc:db2://dashdb-txn-sbox-yp-lon02-02.services.eu-gb.ibm.com:50001/
    BLUDB;sslConnection=true"
    sslEnabled: true
    credentials:
      secretName: db2-masdev-lite-credentials
  certificates:
    - alias: part1
      crt: |
        -----BEGIN CERTIFICATE-----
        MIIG7zCCBdegAwIBAgIQBMX5yCOP3RiCSrij07HJjANBgkqhkiG9w0BAQsFADBN
        MQswCQYDVQQGEwJVUzEVMBMGA1UEChMMRG1naUNlcnQgSW5jMSQwCQYDVQQDEw5E
        aWdpQ2VydCBTSEEyIFNlY3VyZSB0Z2ZlZ2ZlZ2ZlZ2ZlZ2ZlZ2ZlZ2ZlZ2ZlZ2Zl
        MjIwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEwMjEw
        MBQGA1UEBxMNT3ZlcmxhbmQgUGFyYzE0MDIwMjEwMjEwMjEwMjEwMjEwMjEwMjEw
        VVNTkVUyYybnQUNISU5FUyBDT1J0T1JBVElPTjE1MCMGA1UEAwwcKi5zZXJ2aWN1
        cy5ldS1nYi5ibHVlbw14Lm5ldDCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoC
```

```

ggEBAJzzmjnE+tJk1ZkrEEHP4RkDdJ+Tmho0diJdK8soPNVSvs5SuWNWKnBDDohE
B/jEiSgmClWKPFzAAAn+PlTU2djntLma6LWATyc0zx0E2r+221a1yg+UhZr0BfQ9e
fAy3wQcu7Aylsq80sMgM33+U0aM254uurrkg3x0RV1Do1Y1sAzW2/wJmSilM0vyf
hwZUad9hNwi/1bEt6z4WpN/231bvDzeTlr6jznHFArQ8e6/AV98orFd82NxZcM6K
ByGtQpyXkBNf0wy/8kbyYQqzisoWLnanyxAHJABnUESjS3WmvRQ8H1QW7RT63LwL
9s13dErNg7D814L9NVWsiJMaTPcCAwEAAa0CA4IwggN+MB8GA1UdIwQYMBaFAFA+A
YRyCMWHVLYjnYUY4tCzhxtniMBOGA1UdDgQWBSuP2nUw2cmgLQqvH6vI70k1zQ/
jjBDBgNVHREEPDA6ghwqLnN1cnZpY2VzLmV1LWdiLmJsdWvtaXgubmV0ghpzZXJ2
awNlcy51dS1nYi5ibHV1bW14Lm51dDAOBgNVHQ8BAf8EBAMCBaAwHQYDVR01BBYw
FAYIKwYBBQUHAWEGCCsGAQUFBwMCMGSA1UdHwRkMGiW16AtoCuGKWh0dHA6Ly9j
cmwzLmRpZ21jZjXJ0LmNvbS9zc2NhLXNoYTIITZzYuY3J5MC+gLaArlodHRwOi8v
Y3J3NC5kaWdpY2VydC5jb20vc3NjYS1zaGEyLWc2LmNybDBMBG9VHSAERTBDMDCG
CWCSSAGG/WwBATAqMCgGCCsGAQUFBwIBFhxdHRwczovL3d3dy5kaWdpY2VydC5j
b20vQ1BTMAgGBmeBDAECAjB8BgggrBgEFBQcBAQRwMG4wJAYIKwYBBQUHMAggGgH0
dHA6Ly9vY3NwLmRpZ21jZjXJ0LmNvbTBGbggrBgEFBQcwoAoY6aHR0cDovL2NhY2Vy
dHMuzG1naWw1cnQuY29tL0RzZ21DZXJ0U0hBMlN1Y3VyZVZV1cnZlckNBLmNydDAM
BgNVHRMBAf8EAjAAMIIBfwYKKwYBBAHWeQIEAgSCAw8EggFrAwkAdwCkuQmQtBhY
Fie7E6LMZ3AKPDWYBPkb37jdd800yA3cEAAAAXAh/gncAAAEAwBIMEYCIQC75oq7
nsysXvTj9uwoH+4p3/LZD4mEGzC27BQPaFFsvwIhAKinjPXsguG3Q/7EZJey/orx
hTjJDio6yKkvGQR8SIysAHYAikVFB11VJFawP6Ev8fdthuAjJmOtwEt/XcaDXG7i
DwIAAAFWIf4KMgAABAMARzBFAiEAzt1/4xYHslgBIauTztvYEKmxVFijFKvYAXF3
v8FzDzMCIA59GDHJfiqC18angaQzKyDIuYuRPTDMNixgo0ht/mjUAHYAqcjKsd8i
RkoQxqE6CUKHxk4xixsD6+tLx2jwkGKWbVYAAAFWIf4JrgAABAMARzBFAiEASB5k
q51RvXRXk0wyTWX92Q0c+IrL5+0rUwfBXeY1jPgCIH944k+IpKAFN5vM5YnGk/X1
ryONIoMwKDOtqdpvyqARMA0GCSqGSIb3DQEBcwUAA4IBAQBShftEwr1tIjh4a1of
Sc+BScv7NdRXHIOJDg2lQZ3mhBq7Mttw/cAwP1EKvEw/31KB2iQLJN90Q8grTwr
NMBSeNu4b1CTJY+vBYRKKfYEqJH74oHURu4d+9wZl0ZUcHJvXj1vgBR/80+7YV2Y
y02u/4sJJj3yVNa/RzroI6oS+01w0znzc5Io+vst50hveVmiwaHH4fNux00BqHE
Asy2nFSpvzNS/dlMGgM6XoEU46CMS00RIoxoMEWRbDk20PdCtKsg+ySkIYS/ylyN
vdCl1LW0hHqLrG5ZCVQoVgr92vLtxys+rHAeqJISdq3o16QV3iGpBXjv9hww9hpi
XhsQ
-----END CERTIFICATE-----

```

Importing certificates in Maximo Application Suite

You can import certificates into Maximo Application Suite. For example, Maximo Manage might integrate with external systems that contain self-signed certificates, such as Kafka or an SMTP server. You must give an alias to the certificate so that this name is used as the alias for this certificate when you add it to the truststore of Liberty.

About this task

The actual certificate must contain the -----BEGIN----- and -----END CERTIFICATE----- lines when you copy and paste the certificate itself. It must be in the PEM format. In a custom resource, the certificate content must be added in the `crt` property, and because it is multiline, you must add `|`. You can specify multiple certificates that each have a different alias.

Procedure

1. Log in to Maximo Application Suite as a system administrator.
2. On the side navigation menu, in **Suite > Administration** page, click **Configurations**.
3. In the **Other** section, click the row where the external system that Maximo Manage is integrating with is displayed, for example, Apache Kafka.
4. In the dialog, click the Edit icon.
5. Click **Add** to add the certificate.
6. Specify a certificate alias and paste the certificate content in PEM format.
7. Click **Confirm**.
8. Click **Save**.

The following example shows a custom resource (CR).

```

deployment:
  ...
  importedCerts:
    - alias: kafka
      crt: |
        -----BEGIN CERTIFICATE-----
        MIIDLTCcAhWgAwIBAgIJAIyuocAUfASaMA0GCSqGSIb3DQEBcwUAMC0xEzARBGNV
        BAoMcm1vLnN0cm1temkxFjAUBGNVBAMMDWnsdXN0ZXItY2EgdjAwHhcNMjEwNDEx

```

```
MjIzMjUyWhcNMjIwNDAMjIzMjUyWjAtMRMwEQYDVQKDApby5zdHJpbXppMRYw
FAYDVQDDA1jbHVzdGVyLWNhIHVwMIIIBjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
CgKCAQEAsGgoHx3zN5qJs06L/f0r7az9eY3sHH1Ne9cZpe8awMttQD35thL9sT8g
ahStB0uhE/KbhFuZGcKT1Q8w3vJxaqSLeepTgfk9YWMF0LZbr0XbEg8v+apuvLLN
a91EI46yuiZcKUMlMA7WwqF1CHAaca68z5nkdPDf2BvB2Tmy1UkayjiDm9sPukUE
qRxdWTw7Z0j8PSBt2KZP9xyCmA6F7M7KuPr700wH+0291mAvzMmp1f/1bg2jw9e
rz38jmcXrXVv6I2otJHjTY+wRGEVWafP5vWet4vNPXHtvi+e3w+HXAvGEstUJQo
tzPf66+sFfXUI4sT80jJbPfmXrfPgwIDAQABo1AwTjAdBgNVHQ4EFgQU63t8rSXJ
sJMoJAY0wNEHJ/CkjaYwHwYDVR0jBBgwFoAU63t8rSXJsJMoJAY0wNEHJ/CkjaYw
DAYDVR0TBAAUwAwEB/zANBgkqhkiG9w0BAQsFAA0CAQEAFKEVmjYp1FbM9pfYHpwk
5+01NnD3JaCCXDdpP3ccH5ABpxLsD0jf/c12sUxEoxe+h1LnsxUBVnr6QuU4Iipe
50BT82SNzoJU19w9Qp5IYeb2KF4oCbb80bejz6RXdJsuMk4pxKo1E1tIDYmuXZi7
Wk8Np588ZHUdzka7dZklr9CtDLywuJGxHxc0t8R2wccFvGAANG8vMMzUU3DTWk+d
eMumY6m0Q/BnvPrIrrL1/45Gv0v23G5oDLGLkSNmM3UIH+6q18z/vyN5D9K07xUn
ThpsmIQVwX0i079qcZxBaceMK1CTxwdMUTg2MfImTqc66/SAj/cIzEN+320gxEnI
ZA==
```

```
-----END CERTIFICATE-----
```

```
- alias: smtp
  crt: |
-----BEGIN CERTIFICATE-----
...
-----END CERTIFICATE-----
```

Verifying the migration

After you have migrated to Maximo Manage, you must check some aspects so that it works as it is supposed to.

Running Integrity Checker after upgrade

The Integrity Checker is a database configuration utility that you can use to check the health of the database that you migrated from Maximo Asset Management to Maximo Manage. To run Integrity Checker, use the ToolsAPI API.

Before you begin

To work with APIs, you must use a browser or an API platform, such as Postman.

About this task

Authorized users must have an API key and must belong to the security group that enables the signature options for the APIs for script commands. The object structure that specifies the necessary signature options is available to the default administrator group, such as MAXADMIN. The default administrator group is specified by the value of the *ADMINGROUP* variable name in the MAXVARS table.

Procedure

1. Log in as an administrator in Maximo Manage and in the API Keys application, create API keys for the users. This action registers the keys to the users. Copy the keys and provide them to your users. An API key is displayed only one time and cannot be retrieved later.
2. In the Security Groups application, select the administrator group that can work with the APIs.
3. Select the **Object Structures** tab and search for TOOLSAPI in the **Object Structures** section.
4. From the **Actions** menu, click **Grant Listed Object Structures** and select an access level from the list.
5. In the **Options for Tool APIs** section, select the row for the signature option and click **Grant Listed Options for This Object Structure** to grant access to the signature options that are associated with the APIs.

Important: Stop the Maximo Manage pods before you run the Integrity Checker utility.

For more information on Integrity checker warning and error messages, see [Integrity checker messages](#).

6. **Tip:** You can use the Swagger interface to authorize APIs. For more information see, [Swagger Interface for APIs](#).

Use any of the APIs that are applicable to your upgrade.

API	Action	HTTP method	Signature option	API request type
icheckerreport	Generate the integrity checker log.	POST	ICREPORT	Asynchronous
toolslog	Get a specified tools log or get a list of all tools logs.	GET	GETLOG	Synchronous
submitUploadLogRequest	Upload logs from Maximo Manage UI, Cron, Maximo Enterprise Adapter, or Report pods to S3 Cloud Object Storage.	POST	GETLOG	Asynchronous
icheckerrepair	Start integrity checker repair.	POST	ICREPAIR	Asynchronous
managestart	Start all Maximo Manage pods.	POST	MANAGESTART	Asynchronous
managestop	Stop all Maximo Manage pods.	POST	MANAGESTOP	Asynchronous
installexternalcert	Import an external certificate to a server bundle truststore.	POST	INSTALLEXTERNALCERT	Synchronous
validatedbformas	Validate the database for the migration to Maximo Manage.	POST	VALIDATEDBFORMAS	Synchronous

Note: Do not use the **SetAdminMode** command from Manage tools to set the admin mode. The **SetAdminMode** command is not available in Maximo Application Suite. You must use the Maximo Application Suite user interface or use the REST APIs to manage the admin mode.

7. Use any of the following commands in Postman to check the health of your database.

Action	API request
Generate the integrity checker log.	POST <code>https://host:port/toolsapi/toolservice/icheckerreport</code>
Get an integrity checker log.	GET <code>https://host:port/toolsapi/toolservice/toolslog?logfile=name of report from icheckerreport request</code>
Get a list of all tools logs.	GET <code>https://host:port/toolsapi/toolservice/toolslog</code>
Run the integrity checker utility.	POST <code>https://host:port/toolsapi/toolservice/icheckerrepair</code>
Upload logs from Maximo Manage pods to S3 Cloud Object Storage.	POST <code>https://host:port/maximo/api/service/logging?action=wsmethod:submitUploadLogRequest</code>
Stop the Maximo Manage pods.	POST <code>http://host:port/toolsapi/toolservice/managestop</code>
Start the Maximo Manage pods.	POST <code>http://host:port/toolsapi/toolservice/managestart</code>

Action	API request
Import an external certificate to a server bundle truststore.	POST http://host:port/toolsapi/toolservice/installexternalcert
Validate the database for the migration to Maximo Manage.	POST http://host:port/toolsapi/toolservice/validatedbformas

Configuring Oracle connector after upgrade

You must define a message provider, messaging queues, and assign queues if you are using Maximo Connector for Oracle Applications for integration in Maximo Manage.

Procedure

1. Define the message provider and queues, which are Kafka or JMS.
2. Assign queues to External System OA12.
3. Configure Kafka crontask according to the [MIF guide](#) and to update settings for JMS cron task, see [Configuring JMS servers](#).
4. If you are using Kafka, disable **JMSQSEQCONSUMER** crontask.
5. Re-create interface tables and PL/SQL objects in Oracle E-Business Suite12.x MAXORA schema.
6. Reapply customizations on PL/SQL side.
7. Enable the external system.

Configuring SAP Connector after upgrade

You must define a message provider, messaging queues, and assign queues if you are using Maximo Connector for SAP Applications for integration in Maximo Manage after you upgrade.

Procedure

1. Define the message provider and queues as either Kafka or JMS.
2. Assign queues to the external system SAP2005.
3. Configure Kafka cron task. For more information, see [Integration by using Apache Kafka](#).
4. Update settings for the JMS cron task. For more information, see [Configuring JMS servers](#).
5. If you are using Kafka, disable **MSQSEQCONSUMER** cron task.
6. Generate **APIKey** for integration user **sapadmin**. For more information, see [Generating API keys](#).
7. Change connection parameters in SAP PO or Maximo HTTP channel.
8. Enable the external system SAP2005.

Checking Maximo Manage deployment status

Verify that Maximo Manage is deployed successfully by checking the **Custom Resource Definitions** page on the Red Hat OpenShift console.

Before you begin

Maximo Manage must be activated first.

Procedure

1. In the Red Hat OpenShift console, select **Administration > Custom Resource Definition**.
2. In the **Name** field, search for `manageworkspace` and then click the Manage Workspace custom resource definition to open it.
3. To select your project instance, select the **Instances** tab.

4. Check the project status on the **Details** tab of the **Custom Resource Definitions** page.

The project status can be either true or false based on various conditions.

For example, the project is build ready when the build is completed, or it is deployment ready if all the server bundles are running.

Accessing Maximo Manage

You can access Maximo Manage in different ways after you deploy it in Maximo Application Suite.

Procedure

- In Maximo Application Suite, click the **Administration** icon to open the **Suite navigator** page and on the **Applications** tab, click the **Manage** tile.
- Click the **AppSwitcher** icon and then click **Manage**.
- In Red Hat OpenShift web console, use the `ui` bundle or `all` bundle location link. Append `/maximo` to the hyperlink text and use the complete link to access Maximo Manage in a supported browser.
- Starting in Maximo Application Suite 9.1, access Maximo Manage on the side navigation menu.
 - Select Manage from **Suite > Administration > Suite > Applications** tab.
 - Select Manage from **Suite > Administration > Catalog** page.

What to do next

To go to Maximo Application Suite from Maximo Manage, click **IBM Maximo Application Suite** from the main menu.

Accessing database after upgrade

You might want to access the migrated database after you finish the upgrade, for example, to connect the database to a database utility tool, such as DBeaver.

Procedure

1. Log on to Red Hat OpenShift web console and from the side navigation menu, click **Networking > Routes**.
2. Click **Create Route** to access the Db2 host outside the cluster.
3. On the **Create Route** page, select the **YAML view** radio button and provide the YAML that includes the external hostname from the route and the external port from Services.

Tip: Search for a service name, for example, `c-db2w-manage-db2u-engn-svc` and a sample URL might be `https://ukiot-mas2-0026a8a1020f89b5eb8fa6780c129be5-0000.eu-gb.containers.appdomain.cloud`.

You can modify the following example YAML to use the values for your route:

```
kind: Route
apiVersion: route.openshift.io/v1
metadata:
  name: db2u-tls-route
  namespace: db2u
  uid: a9b76ae3-76f5-49e8-868f-ce060174936d
  resourceVersion: '114046'
  creationTimestamp: '2022-09-13T20:24:44Z'
  labels:
    formation_id: db2u-shared
  annotations:
    kubernetes.io/last-applied-configuration: >
      {"apiVersion":"route.openshift.io/v1","kind":"Route","metadata":
{"labels":{"formation_id":"db2u-shared"},"name":"db2u-tls-route","namespace":"db2u"},"spec":
{"host":"db2u-shared-db2u.monitordemo1-822c5cdfc486f5db3c3145c89ca6409d-0000.us-
south.containers.appdomain.cloud","port":{"targetPort":"ssl-server"},"tls":
{"nsecureEdgeTerminationPolicy":"None","termination":"passthrough"},"to":
{"kind":"Service","name":"c-db2u-shared-db2u-engn-
svc","weight":100},"wildcardPolicy":"None"}}

```

```

managedFields:
- manager: OpenAPI-Generator
  operation: Update
  apiVersion: route.openshift.io/v1
  time: '2022-09-13T20:24:44Z'
  fieldsType: FieldsV1
  fieldsV1:
    'f:metadata':
      'f:annotations':
        .: {}
        'f:kubect1.kubernetes.io/last-applied-configuration': {}
      'f:labels':
        .: {}
        'f:formation_id': {}
    'f:spec':
      'f:host': {}
      'f:port':
        .: {}
        'f:targetPort': {}
      'f:tls':
        .: {}
        'f:termination': {}
      'f:to':
        'f:kind': {}
        'f:name': {}
        'f:weight': {}
      'f:wildcardPolicy': {}
- manager: openshift-router
  operation: Update
  apiVersion: route.openshift.io/v1
  time: '2022-09-13T20:24:44Z'
  fieldsType: FieldsV1
  fieldsV1:
    'f:status':
      'f:ingress': {}
    subresource: status
spec:
  host: >-
    db2u-shared-db2u.monitordemo1-822c5cdfc486f5db3c3145c89ca6409d-0000.us-
    south.containers.appdomain.cloud
  to:
    kind: Service
    name: c-db2u-shared-db2u-engn-svc
    weight: 100
  port:
    targetPort: ssl-server
  tls:
    termination: passthrough
    wildcardPolicy: None
status:
  ingress:
    - host: >-
        db2u-shared-db2u.monitordemo1-822c5cdfc486f5db3c3145c89ca6409d-0000.us-
        south.containers.appdomain.cloud
        routerName: default
        conditions:
          - type: Admitted
            status: 'True'
            lastTransitionTime: '2022-09-13T20:24:44Z'
        wildcardPolicy: None
        routerCanonicalHostname: >-
            router-default.monitordemo1-822c5cdfc486f5db3c3145c89ca6409d-0000.us-
            south.containers.appdomain.cloud

```

4. In Maximo Application Suite, from the side navigation menu, click **Configurations**.
5. In the **Storage** section, open the **Database connection** row and view and note the database connection string, username, and password.

Troubleshooting global property values

When migrating a Maximo SaaS Flex 7.6.1.2 database, you might see an error in the **mx.int.globaldir** property value.

Symptoms

If you have a database in Maximo SaaS Flex 7.6.1.2 and want to refresh this database to Maximo Application Suite, you might see the following error in the **mxe.int.globaldir** property value.

```
[Default Executor-thread-7] ERROR maximo.graphite - [DatabaseResourceLoader] Expand App manage-shell failed.
Message: /opt/IBM/WebSphere/AppServer/profiles/ctgAppSrv01/MAXIMO_b1dk/maximo/maf/manage-shell/8.0.0.0-0/app-source.zip
java.nio.file.NoSuchFileException: /opt/IBM/WebSphere/AppServer/profiles/ctgAppSrv01/MAXIMO_b1dk/maximo/maf/manage-shell/8.0.0.0-0/app-source.zip
\tat sun.nio.fs.UnixException.translateToIOException(UnixException.java:92) ~[?:?]
\tat sun.nio.fs.UnixException.rethrowAsIOException(UnixException.java:111) ~[?:?]
\tat sun.nio.fs.UnixException.rethrowAsIOException(UnixException.java:116) ~[?:?]
```

Causes

In your database, the **mxe.int.globaldir** property value is set to something that does not exist in the image.

Resolving the problem

Make sure that the **globaldir** value is either a mounted persistent volume, or it is a directory that already exists in the image because it is used as a default system directory. If you want to temporarily use a non persistent volume by default, set the value for **mxe.int.globaldir** as blank or null.

You can also set the environment variable through the **MAF_APP_ROOT** encryption property. For example, set **MAF_APP_ROOT** to `/tmp`.



Attention: Any Maximo properties that refer to a directory or **doclink** root must be validated. If you need persistence, a Persistent Volume must be mounted to host that directory, or, point to any existing writable directory, for example, any sub-folder under `/tmp`.

Accessing Maximo Manage logs by using Red Hat OpenShift web console

View the Maximo Manage logs for Server or System.Out, MAXINST, Update database, User sync, Workspace operator, and Build from the Red Hat OpenShift console.

Procedure

1. Log into the Red Hat OpenShift console.
2. From the side navigation menu, click **Workload > Pods** menu.
3. Select your Maximo Manage project name from the Project drop down.
4. Select the pod from the list for which you want to view logs.
5. Click the **Logs** tab to view the streaming log for that pod.
6. To push the server logs to any S3 compatible object storage:
 - a) Set up S3 credentials, by creating the following four environment variables in your Maximo Manage deployment, `LOG_BUCKETNAME`, `LOG_S3ACCESSKEY`, and `LOG_S3ENDPOINTURL`, and `LOG_S3SECRETKEY`.
 - b) Create a log request, that is a POST request by using any HTTP REST Client, for example, download and use the Postman tool to create the log request.
Request URL: `http://manageserver:7001/maximo/oslc/service/logging?action=wsmethd:submitUploadRequest`

Note: Replace the Maximo Manage server in the URL with your Maximo Manage server host or IP address. A record for Maximo Manage is created in the LOGREQUESTDET table. The Maximo Manage records are removed by the LOGREQUESTCLEANUP cron task after the logs are posted to the object storage. The cron task must be activated manually, and then it runs once daily,

- c) Use the S3 browser to view the logs by using your S3 credentials.

Each Maximo Manage server handles the latest log request. It compresses the log files in the log directory and uploads it to cloud object storage.

What to do next

You can install Red Hat OpenShift logging by deploying Red Hat OpenShift Elasticsearch and Red Hat OpenShift logging operators. For more information, see [Installing Logging](#).

Verifying system settings after upgrade

After the upgrade process is complete, verify and update specific settings in Maximo Application Suite, Maximo Manage, and the Red Hat OpenShift web console.

- In the System Properties application in Maximo Manage, check that the `mxe.int.dfltuser` property is set to a user with administrative privileges. By default, this admin user is used to synchronize the Maximo Application Suite users to Maximo Manage.
- Check that each user has only one primary email address. Maximo Manage business logic does not allow multiple primary email addresses.
- If multiple users have the same login IDs that differ only in the usage of uppercase or lowercase letters, correct the login IDs so that they are each unique. Two users cannot have the same login IDs.
- Check that `ui`, `report`, `mea`, and `cron` routes are present in the Red Hat OpenShift cluster.
- Check that `ui`, `report`, `mea`, and `cron` services are running in the Red Hat OpenShift cluster.
- Customize the Liberty `server.xml` file, if necessary in Maximo Application Suite.
- Set the `mxe.hostname` property to the maximo-all or maximo-ui route host.

Updating statistics

As a system administrator, you can analyze tables in IBM Maximo Manage to ensure that the Oracle Database cost-based optimizer has up-to-date statistics.

Procedure

1. In Maximo Manage, , click **System Configuration > Platform Configuration > Database Configuration**.
2. In the **More Actions** menu, click **Update Statistics**.
3. Click **OK**.

Overview of migrating TRIRIGA to Maximo Real Estate and Facilities

A high-level overview of the steps that are needed to migrate IBM TRIRIGA or IBM TRIRIGA Application Suite to IBM Maximo Real Estate and Facilities Platform 9.1, including the TRIRIGA users and licenses.

Before you begin

- The source environment must be IBM TRIRIGA Application Platform 4.5 or 5.0 and IBM TRIRIGA Application 11.5 or 11.6.
- The target environment must be Maximo Application Suite 9.1 or later.
- During the migration, you must deploy and activate Maximo Real Estate and Facilities Platform 9.1.2 or later.
- The TRIRIGA AES keystore credentials must be known and working.



Warning: For data encrypted in TRIRIGA to be available and readable in Maximo Real Estate and Facilities, you must create a secret in Red Hat OpenShift Container Platform before you activate Maximo Real Estate and Facilities.

1. Obtain the appropriate Maximo Application Suite licenses and install Maximo Application Suite, see [“Installing Maximo Application Suite” on page 218](#).

2. Review the changes between TRIRIGA and Maximo Real Estate and Facilities, see [“What's changed in Maximo Real Estate and Facilities”](#) on page 573.
3. Prepare your database, see [“Migrating the application database”](#) on page 575.
4. Configure your AES encryption, see [“Migrating AES reversible encryption”](#) on page 575.
5. Deploy and activate Maximo Real Estate and Facilities, which includes configuring the Maximo Real Estate and Facilities database, see [“Deployment and activation settings for Maximo Real Estate and Facilities”](#) on page 576.
6. Create the mandatory initial Maximo Real Estate and Facilities FACILITIESADMIN administrator user in IBM Maximo Application Suite, see [“Administering Maximo Real Estate and Facilities users”](#) on page 385.
7. Migrate your TRIRIGA user files, see [“Migrating TRIRIGA user files”](#) on page 576.
8. Use the User Migration Tool to migrate your TRIRIGA or IBM TRIRIGA Application Suite users and licenses to Maximo Application Suite users and licenses.

All Maximo Real Estate and Facilities users and licenses are managed in IBM Maximo Application Suite.

Note: The User Migration Tool is covered under existing licenses so no additional license is needed.

Note: For security purposes, all users are migrated as concurrent rather than authorized users. Also certain Enterprise license users need more access than other users. After the migration, log in to Maximo Application Suite administration to update the appropriate users to authorized user and to assign licenses to advanced users.

For more information, see [“Migrating users and licenses”](#) on page 576.

9. You are now on IBM TRIRIGA Application 11.5 or 11.6 on Maximo Real Estate and Facilities 9.1.2 or later. Complete any post-migration tasks, see [“Post-migration tasks”](#) on page 579.

To benefit from the most recent updates, plan to upgrade to Maximo Real Estate and Facilities Application 9.1 or later, see [Upgrading the Maximo Real Estate and Facilities application](#).

What's changed in Maximo Real Estate and Facilities

IBM TRIRIGA has been rebranded to Maximo Real Estate and Facilities and is now part of Maximo Application Suite. Some processes, such as licensing and user management, that were managed at the product-level are now managed at the suite level.

Changes in Maximo Real Estate and Facilities

Changes in licensing model

Maximo Application Suite uses a different licensing model and uses AppPoints to track application usage, runtime, and user access. AppPoints are allocated in your organization as defined by your license entitlement. You can configure your environment to enforce the AppPoint entitlement. IBM TRIRIGA Application Suite users are familiar with the concept of AppPoints. For information about Maximo Application Suite licensing and AppPoints, see [“Licensing in Maximo Application Suite”](#) on page 78.

Changes in authentication and user management

In IBM TRIRIGA and IBM TRIRIGA Application Suite, user authentication is configured in the application.

In Maximo Real Estate and Facilities, users, and their access and entitlement to Maximo Real Estate and Facilities are now managed in IBM Maximo Application Suite. Security groups in the application are unchanged. You must migrate existing users to the **Users** application in Maximo Application Suite where you can view and edit their details.

Users are given Maximo Real Estate and Facilities access with Self-Service, Base, Limited, or Premium entitlement, see [Administering users and user access in Maximo Application Suite in 9.1](#). Users that are given access to Maximo Real Estate and Facilities are automatically synchronized with Maximo Real Estate and Facilities.

Important: For the first login to Maximo Real Estate and Facilities, you must create a mandatory initial administrator user with a user ID of FACILITIESADMIN. Regardless of your selection, the FACILITIESADMIN user always defaults to Base entitlement in Maximo Real Estate and Facilities.

Configuring more granular access to specific functional areas and capabilities in Maximo Real Estate and Facilities is unchanged. Administrators still configure application security groups as before in Maximo Real Estate and Facilities administration. Maximo Real Estate and Facilities administrators can assign access to security groups and other permissions either individually or in bulk. For more information about administering user access within Maximo Real Estate and Facilities, see [Administering user access and permissions](#).

Middleware implementation

Maximo Application Suite is deployed on Red Hat OpenShift Container Platform. For more information, see [IBM Maximo Application Suite technical overview](#).

Deployment architecture

Workloads are deployed on Red Hat OpenShift Container Platform pods and the server deployments have changed. TRIRIGA was deployed on JVMs with the WebSphere Liberty application server running directly on hardware or VM and in Maximo Real Estate and Facilities the workloads are running on JVMs in pods. You no longer have TRIRIGA process servers and the workload is distributed automatically. For more information see [“Workload sizes and deployments”](#) on page 374.

Installation and configuration

Installation, configuration, and deployment are done by the Red Hat OpenShift operator in Maximo Application Suite. The entire deployment is based on Red Hat OpenShift.

Maximo Application Suite installation and configuration process includes customer-managed or IBM managed installations. You can choose from multiple platforms and environments. For more information, see [Installing Maximo Application Suite](#).

Upgrade process

You can upgrade Maximo Application Suite automatically or manually. For more information see, [Upgrading IBM Maximo Application Suite](#).

- The Maximo Real Estate and Facilities Platform is upgraded or patched by updating Maximo Real Estate and Facilities Operator in Red Hat OpenShift Container Platform, see [Upgrading](#).
- The Maximo Real Estate and Facilities Application is upgraded or patched as before, see [Upgrading the IBM® Maximo Real Estate and Facilities application](#).

Changes due to containerization

Containerization is the packaging of software code with just the operating system (OS) libraries and dependencies that are required to run the code to create a single lightweight executable package that is infrastructure-agnostic. System properties, file paths, and storage change when you upgrade from IBM TRIRIGA to Maximo Real Estate and Facilities because of containerization.

System properties

- The properties files, such as TRIRIGAWEB.properties, are still used. Certain properties are now automatically set by the operator and must not be changed in the properties files, see [TRIRIGAWEB.properties](#).

User files and persistent storage

The file paths for user files change in IBM Maximo Real Estate and Facilities and now only the userfiles, config, and log directories are persistent.

Context path

The context path is no longer used in Maximo Real Estate and Facilities.

Changes in application server

You do not need to install or migrate IBM WebSphere Application Server Liberty, it is embedded in the Maximo image and deployed automatically by Maximo Real Estate and Facilities.

Pod-specific URLs

You must use pod-specific URLs to access the Administrator Console and to configure Cloud Login. You can find the specific URLs for a pod by logging in to the Red Hat OpenShift Container Platform.

Migrating the application database

You can use your existing TRIRIGA application database at its current location or prepare a new database and move the TRIRIGA application database. You will need the database credentials for when you deploy and activate Maximo Real Estate and Facilities.

For more information about the database options in Maximo Real Estate and Facilities, see [“Preparing your Maximo Real Estate and Facilities database”](#) on page 375.

Prepare the application database.

1. Optional: Move the TRIRIGA database.
2. Clean up the database by purging the data from the AGENT_REGISTRY and AGENT_STARTUP tables.
3. Note the database credentials for when you deploy and activate Maximo Real Estate and Facilities.

Optional: Move the TRIRIGA database

1. Prepare a new database for Maximo Real Estate and Facilities
2. Do a full backup of the TRIRIGA application database from the source environment with reference to your database vendor documentation.
3. Restore the TRIRIGA application database into the prepared Maximo Real Estate and Facilities database with reference to your database vendor documentation.

Application database cleanup

Ensure that the application database is no longer used by TRIRIGA and is not yet connected to Maximo Real Estate and Facilities.

For both Db2 and Oracle Database, enter the following commands to purge the data from AGENT_REGISTRY and AGENT_STARTUP tables.

```
TRUNCATE TABLE <schema>.AGENT_STARTUP;
TRUNCATE TABLE <schema>.AGENT_REGISTRY;
DELETE FROM <schema>.ENVIRONMENT_PROPERTIES where ENVIRONMENT= 'platform.class.upgrade' and
PROPERTY = 'lastRun';
DELETE FROM <schema>.ENVIRONMENT_PROPERTIES where
PROPERTY='com.tririga.util.startup.upgrade.AgentManagementUpgrade' and
ENVIRONMENT='ibs.startup.upgrade';
```

For Oracle Database, the schema name is typically the same as the username.

Migrating AES reversible encryption

TRIRIGA and Maximo Real Estate and Facilities use AES reversible encryption to securely store some of the data you add, such as passwords to external systems used by Integration Objects. The AES Encryption key is stored in a secured vault, which can be unlocked only by its password. This password is typically set while installing TRIRIGA.



Warning: For data encrypted in TRIRIGA to be available and readable in Maximo Real Estate and Facilities, you must create a secret in Red Hat Openshift Container Platform **before** you activate Maximo Real Estate and Facilities.

For more information about AES encryption for database passwords, see [AES encryption](#).

The secret must be in the mas-`<instanceId>`-facilities namespace, as specified below.

```
kind: Secret
apiVersion: v1
metadata:
  name: <workspaceId>-facilities-vs--sn
  namespace: mas-<instanceId>-facilities
data:
  pwd: <password you entered while installing TRIRIGA>
type: Opaque
```

For more information, see [Configuring and activating Maximo Real Estate and Facilities](#).

Deployment and activation settings for Maximo Real Estate and Facilities

When deploying and activating IBM Maximo Real Estate and Facilities, configure the application database and any custom application server settings. You must have the database credentials and have configured the AES secret.

For information about deploying and activating Maximo Real Estate and Facilities, see [“Deploying Maximo Real Estate and Facilities in Maximo Application Suite”](#) on page 373.

Migrating changes to server.xml

1. If you have customized the `server.xml` file, you can include those changes in **Liberty Server XML extensions**, see [Configuring and activating Maximo Real Estate and Facilities](#).
2. Adding custom JAR, WAR, or EAR files is not supported.

Migrating changes in jvm.options

As of Maximo Real Estate and Facilities 9.1.3, changes to `jvm.options` are not supported.

Migrating TRIRIGA user files

Some file paths have changed in IBM Maximo Real Estate and Facilities and now only the `userfiles`, `config`, and `log` directories are persistent.

Back up the `config`, `userfiles`, and `logs` directories from TRIRIGA, and any other directories that you might have configured Integration Objects to write to. If you have Integration Objects that write to other directories, update them to write to the `userfiles` directory.

Any files you need must be copied from the source TRIRIGA environment to the corresponding directories in Maximo Real Estate and Facilities. For example, any Exported Scheduled Reports that you want to have accessible in Maximo Real Estate and Facilities.

- `/home/default/userfiles`
- `/home/default/config`
- `/home/default/log`

You can use the Red Hat Openshift CLI to copy the files to those directories. Ensure that you copy files on the `appserver-0` pod.

Migrating users and licenses

You can migrate your IBM TRIRIGA or IBM TRIRIGA Application Suite users to IBM Maximo Real Estate and Facilities by using the User Migration Tool. You can review the available users, update their people records before migrating, and choose which licenses to assign to selected users on migration. The migration status indicates the status for each user and displays their new license.

Before you begin

- The source environment must be IBM TRIRIGA Application Platform 4.5 or 5.0 and IBM TRIRIGA Application 11.5 or 11.6.
- The target environment must be Maximo Real Estate and Facilities Platform 9.1.2 or later.

Check the pod logs to ensure that the agents are running correctly.

Check that Workflow is running.

About this task

On migration, each user is registered in Maximo Application Suite and assigned access to Maximo Real Estate and Facilities with the appropriate Self-Service, Base, Limited, or Premium license and as concurrent users. Where users exist in Maximo Application Suite, their details are merged with the values from Maximo Application Suite taking precedence.

Note: The system user is no longer used in Maximo Real Estate and Facilities and is not migrated.

Users with IBM TRIRIGA or IBM TRIRIGA Application Suite Enterprise licenses are also migrated as concurrent users, even if they were authorized users in IBM TRIRIGA. Enterprise licenses are calculated as follows:

TRIRIGA	MAS equivalent	Calculation
IBM TRIRIGA Reservation Manager (3023)	MAS Base (5051)	Users who have only this enterprise license are given MAS Self-Service (5053). Most users with this license need only self service, but users who need to access Reserve metadata must be given MAS Base (5051) in Maximo Application Suite administration.
IBM TRIRIGA Businesses Connect (3021)	MAS Base (5051)	Users who have only this enterprise license are given MAS Self-Service (5053). Most users with this license need only Self-Service, but integration users must be given MAS Base (5051) in Maximo Application Suite administration.
IBM TRIRIGA Workplace Services (3056)	MAS Self-Service (5053)	Users who have only this enterprise license are given MAS Self-Service (5053).
IBM Facilities and Real Estate Management on Cloud Workplace Services User (3057)	MAS Self-Service (5053)	Users who have only this enterprise license are given MAS Self-Service (5053).

Note: IBM TRIRIGA Reporting install and IBM TRIRIGA Reporting nonproduction install are not currently available in Maximo Real Estate and Facilities.

The following options are available on the **Migration Options** tab.

- Users are migrated asynchronously in batches and, if needed, you can configure the batch size for performance reasons. By default the batch size is 499, which is the maximum value that is supported. If needed, expand **Batch Size** and enter a new batch size value.
- You can recalculate the projected licenses to reflect any manual changes to people records, any changes to ignored licenses, or changes during processing. Expand **Recalculate Projected Licenses**. Select the users whose licenses you want to recalculate and click **Recalculate**.
- You can specify licenses to ignore when calculating projected licenses. Select the **Migration Options** tab, and expand **Ignore These Licenses**. By default, the following obsolete licenses are ignored: 10000, 10001, 3046, 3051, 3054, and 3055. You can ignore more licenses by adding license IDs to the default

comma-separated list. If you add license IDs, you must recalculate the projected licenses. For more information about TRIRIGA license IDs, see [License IDs and names](#).

Procedure

1. Download and install the User Migration Tool from [IBM Fix Central](#).

Note: The User Migration Tool is covered under existing licenses so no additional license is needed.

2. Log in to Maximo Real Estate and Facilities as the FACILITIESADMIN user or an administrator with equivalent privileges. The **Mas Core - Admin Portal** opens.
3. Select **Maximo Application Suite - User Migration**.

On the **User Migration** tab, you can see the list of users that are available for migration and their details:

- **Actual License:** Maximo Application Suite licenses that were manually added to their people record. You can delete IBM TRIRIGA licenses but not Maximo Application Suite licenses on a people record.
- **Projected License:** A Maximo Application Suite license that is calculated from their existing TRIRIGA licenses.

You can switch between a details or a people view in user lists. For example:

- **Display Details:** A link to the details view is displayed.
- **Display People:** A link from their user name to their people record is displayed.

You can click a user name to open their people record and change their details prior to migration.

4. From the **Maximo Application Suite Licences** menu, select the license that you want to apply to the selected users during the migration:
 - Select **Actual or Projected License** to apply the actual license. If an actual license is not set, the projected license is applied. This option is applied automatically if you choose to migrate all users.
 - Select **Self-Service, Limited, Base, or Premium** to apply a specific license.
5. You can select the users that you want to migrate or migrate all users. Again, you can click a username to update a people record.
 - a) To migrate selected users, review your selections and click **Review Selected Users**. Your selected licenses are applied.
 - b) To migrate all users, click **Review All Users**. All users are migrated with their actual or projected licenses.
6. Review your selections.
 - a) To go back and change your selections, click **Back**.
 - b) To migrate the users to Maximo Application Suite with their new licenses, click **Migrate Users**.
7. The migration process starts and the users are displayed in the **Users in Process of Migration** section. Batches are run asynchronously, so click **Refresh** to see the latest status. If needed, you can click **Clear the User Migration Queue** to stop the migration process for the current batch of users.
8. Select the **Status** tab to review the migration status for each user. Batches are run asynchronously, so click **Refresh** to see the latest status.

The status for each migrated user is shown in the following sections:

- **Successful** - Users that were successfully migrated. You can review the new licenses that are assigned to each user.
- **Errors** - Users that were not migrated are listed with an error message.
- **Warnings** - User that were successfully migrated but with warnings.
 - A warning is displayed for each successfully migrated user that existed in Maximo Application Suite so that you can review their merged details.

- A warning is displayed for users who were assigned a Base or Premium license and who might need a Maximo Application Suite administrator role. Authorized access for administrators can be assigned only in Maximo Application Suite administration.
9. Log in to Maximo Application Suite as an administrator to update the following users.

For security purposes, all users are migrated as concurrent rather than authorized users. Also certain Enterprise license users need more access than other users. After the migration, update the following users:

- Update users with a Base or Premium license from concurrent to authorized users.
- For Enterprise licenses, update Reservation Manager users who need to access Reserve metadata, and Business Connect integration users, from a Self-Service license to a Base license and from concurrent to authorized users.

For more information, see [Administering users and user access in Maximo Application Suite in 9.1](#).

Post-migration tasks

After you migrate to IBM Maximo Real Estate and Facilities Platform 9.1.2 or later, complete the following post-migration tasks.

For security purposes, all users are migrated as concurrent rather than authorized users. After the migration, log in to Maximo Application Suite administration and update the following users:

- Update users with a Base or Premium license from concurrent to authorized users.
- For Enterprise licenses, update Reservation Manager users who need to access Reserve metadata, and Business Connect integration users, from a Self-Service license to a Base license and from concurrent to authorized users.

For more information, see [“User access and entitlements in Maximo Application Suite 9.1” on page 782](#).

Remove TRIRIGA license artefacts

After you migrate to IBM Maximo Real Estate and Facilities, some reports, documents, and queries become obsolete because they apply only to IBM TRIRIGA license data that existed before the upgrade. After you have completed any needed license audits on TRIRIGA licenses, you can manually delete the following TRIRIGA external reports, documents and queries.

During a migration, TRIRIGA licenses are converted to Maximo Application Suite licenses for use in Maximo Real Estate and Facilities.

External report names

1. triLicense - BIRT - Daily License Usage Report for Selected Month
2. triLicense - BIRT - Daily TAS License Usage Report for Selected Month
3. triLicense - BIRT - Daily TAS User Load Report for Selected Month
4. triLicense - BIRT - Daily User Load Report for Selected Month
5. triLicense - BIRT - Hourly License Usage Report for Selected Day
6. triLicense - BIRT - Hourly TAS License Usage Report for Selected Day
7. triLicense - BIRT - Hourly TAS User Load Report for Selected Day
8. triLicense - BIRT - Hourly User Load Report for Selected Day
9. triLicense - BIRT - Monthly License Usage Report for Selected Year
10. triLicense - BIRT - Monthly TAS License Usage Report for Selected Year
11. triLicense - BIRT - Monthly TAS User Load Report for Selected Year
12. triLicense - BIRT - Monthly User Load Report for Selected Year

Document filenames and file paths

Filename	File path
DailyLicenseUsageReportforSelectedMonth.zip	\\ROOT\TRIRIGA\System Reports\Daily License Usage Report for Selected Month (BIRT)
DailyTASLicenseUsageReportforSelectedMonth.zip	\\ROOT\TRIRIGA\System Reports\Daily TAS License Usage Report for Selected Month (BIRT)
DailyTASUserLoadReportforSelectedMonth.zip	\\ROOT\TRIRIGA\System Reports\Daily TAS User Load Report for Selected Month (BIRT)
DailyUserLoadReportforSelectedMonth.zip	\\ROOT\TRIRIGA\System Reports\Daily User Load Report for Selected Month (BIRT)
HourlyLicenseUsageReportforSelectedDay.zip	\\ROOT\TRIRIGA\System Reports\Hourly License Usage Report for Selected Day (BIRT)
HourlyTASLicenseUsageReportforSelectedDay.zip	\\ROOT\TRIRIGA\System Reports\Hourly TAS License Usage Report for Selected Day (BIRT)
HourlyTASUserLoadReportforSelectedDay.zip	\\ROOT\TRIRIGA\System Reports\Hourly TAS User Load Report for Selected Day (BIRT)
HourlyUserLoadReportforSelectedDay.zip	\\ROOT\TRIRIGA\System Reports\Hourly User Load Report for Selected Day (BIRT)
MonthlyLicenseUsageReportforSelectedYear.zip	\\ROOT\TRIRIGA\System Reports\Monthly License Usage Report for Selected Year (BIRT)
MonthlyTASLicenseUsageReportforSelectedYear.zip	\\ROOT\TRIRIGA\System Reports\Monthly TAS License Usage Report for Selected Year (BIRT)
MonthlyTASUserLoadReportforSelectedYear.zip	\\ROOT\TRIRIGA\System Reports\Monthly TAS User Load Report for Selected Year (BIRT)
MonthlyUserLoadReportforSelectedYear.zip	\\ROOT\TRIRIGA\System Reports\Monthly User Load Report for Selected Year (BIRT)

Query names

1. triLicense - REPORT - All Non TAS Licenses
2. triLicense - REPORT - TAS And Application Builder Licenses

Customer-managed

Configuring Maximo Application Suite

Use the Suite administration Configurations page to access and manage the IBM Maximo Application Suite configuration settings for your environment.

When you deploy and activate a Maximo Application Suite application, you must provide the required settings for the application and for any supporting tools. You can preconfigure these settings to require no further input during deployment.

The following sections describe the Maximo Application Suite configurations and the required parameters. Any configurations that require pod downtime can result in system or application outages for users.

Maximo Application Suite configurations are set at a defined scope, such as system or application.

System scope

The configuration is set for and can be used across the whole suite. Example: A JDBC configuration that can be used by all applications in the suite.

Workspace scope

The configuration is set for and can be used in the default workspace. Example: The JDBC connection that can be used by all applications in the suite.

Application scope

The configuration is set for and can be used by a single application. Example: The JDBC connection that is used by the application.

Workspace-application scope

The configuration is set for and can be used by a single application in the default workspace. Example: The JDBC connection that is used by the application in the default workspace.

Setting up IBM Maximo Application Suite

After you install IBM Maximo Application Suite, the setup program guides you through the initial configuration.

Before you begin

1. Complete the installation.

Obtain the link to the Maximo Application Suite setup program and the login credentials that you need to complete the setup process.

For more information about obtaining the login credentials, see [how to locate the default username and password](#).

2. Enable login for Maximo Application Suite self-signed certificates.

If you are using self-signed certificates in a development or test environment, you must manually enable login by using either of the following methods.

- Download the certificates from the cluster and add them to your local certificate manager.
- In your browser, go to the Maximo Application Suite API URL `https://api.<mas_domain>/` and accept the certificate security risks. After you accept the risks, an AIUC01999E error is displayed. This message is expected. You can now continue with the setup process.

If the Maximo Application Suite dashboard does not load after you login for the first time and instead see a spinning wheel, see [how to troubleshoot the issue](#).

About this task

The Maximo Application Suite setup configurations are set at the System scope. For more information about configuration scopes, see [Configure Maximo Application Suite](#).

Procedure

1. Log in to the Maximo Application Suite setup program by using the superuser credentials that were created during installation.
`https://admin.<mas_domain>/initialsetup`

Important: Treat the superuser account the same way that you treat the root account on your servers. Use it only for the initial setup. As part of the setup, you create a default administrator user account that has access to the Maximo Application Suite administrative interface. Use this administrative account to add and manage users, deploy applications, and more.

For more information about obtaining the superuser credentials, see [how to locate the default username and password](#).

2. Configure MongoDB.

MongoDB is used as the data dictionary for Maximo Application Suite and its applications. It is also used as the default user registry.

Specify the following MongoDB information:

Hostname and port

You can configure one or more MongoDB hostname and port combinations.

Authentication mechanism

Specify the mechanism that is used to authenticate Maximo Application Suite when it connects to MongoDB. Select the closest match to the mechanism that is configured for your MongoDB cluster. For example, if your cluster uses the SCRAM-SHA-256 mechanism, select **DEFAULT (SCRAM)**.

To authenticate by using LDAP, specify **PLAIN** as the authentication mechanism.

Auth db

Provide the name of the authentication database. If you are authenticating with LDAP, the value must be `$external`.

MongoDB login credentials

At a minimum, the MongoDB administrator needs table creation privileges.

Note: The MongoDB verification might take up to a minute. The configuration cannot be modified after the MongoDB verification is complete. MongoDB is a prerequisite for Maximo Application Suite. Changing the configuration requires careful coordination and possible data migration to avoid service outages. System administrators can change the configuration in the Red Hat OpenShift console. For assistance with changing the MongoDB configuration, contact your IBM representative.

For more information, see [Installing MongoDB](#).

3. Upload a CA certificate.

If the service uses the transport layer security (TLS) communication protocol and is not secured with a certificate that is issued by a well-known certificate authority (CA), then provide the certificate of the CA that issued the service's certificate. Because the CA might use intermediate CAs, you can provide more than one certificate.

For each certificate that you provide, the following details are displayed:

- The name of the certificate issuer.
- The name of the subject, such as the organization, that the certificate is issued to.
- The start and end dates of the certificate's validity period. If the validity of any certificate that you provide expires soon, a warning message appears.

You can automatically retrieve or manually add certificates.

Important: If your MongoDB cluster uses self-signed CA certificates that you must retrieve or add a certificate.

- Automatically retrieving certificates

In the certificates section, click **Retrieve**. If the connection credentials that you specify are correct, all CA certificates that are configured on the server are automatically retrieved and displayed.

These certificates are not validated. Verify that only the correct certificates are retrieved and remove any unexpected certificates.

After you retrieve the certificates, you can manually add more certificates.

- Manually adding certificates

In the certificates section, click **Add manually** and specify the following values for each certificate that you want to add:

Alias

An alphanumeric identifier that is in the range 3 - 50 characters long.

Certificate content

The content of a certificate file in either the X.509 or PEM formats.

For more information, see [“Configuring certificate authority certificates”](#) on page 590.

4. Configure a Simple Mail Transfer Protocol (SMTP) server connection to enable email notifications for system events, such as new user welcome emails and password reset communication. For more information, see [“Setting up email notifications”](#) on page 634.
5. Configure analytics data.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator”](#) on page 7.

- If you are using Maximo Application Suite 8.11, 8.10 or earlier versions, you must migrate your User Data Services to Data Reporter Operator. For more information, see [“Migrating Maximo Application Suite from User Data Services to Data Reporter Operator”](#) on page 8.
- If you are using Maximo Application Suite 9.0, 8.11.7, 8.10.10 or later versions, configure IBM Data Reporter Operator.

The IBM Data Reporter Operator accepts events and transforms them into reports that are submitted to the Data Service of the IBM Metrics Operator.

- a. Enter the following information to configure Data Reporter Operator for Maximo Application Suite:

- **URL** - This URL is the DRO URL endpoint. To find it, go to your Red Hat OpenShift console, switch to `ibm-common-services` project, then **Networking** > **Routes**. Copy the URL displayed under the Location column for the `dro-endpoint` route.

For example, `https://dro-endpoint-ibm-common-services.<your-cluster-domain>`

- **API Key** - This API key is the DRO API Key credential. To find it, go to your Red Hat OpenShift console, switch to `ibm-common-services` project, then **Workloads** > **Secrets** > **Search and select the secret named dro-api-key**. Under the **Data** section, copy the `apikey` value.

For example, `k2wnQY...`

- **Email** - Enter a contact email address to use for DRO communication. The email address does not have to match an existing Maximo Application Suite user.
- **Given Name** - Enter the given name of the owner of the provided contact email address that is used for DRO communication.
- **Surname** - Enter the surname of the owner of the provided contact email address that is used for User Data Services communication.
- **Certificates** - Enter the chain of SSL certificates for your DRO. To retrieve the certificates, you can click the **Retrieve button (under Certificates section)** while configuring DRO into Maximo Application Suite. The DRO certificates to configure in Maximo Application Suite will vary according to the cloud service provider's cluster that is hosting your DRO installation.

- b. Click **Add** to add the **intermediate of the certificate chain**.

- c. Enter an **alias**. **Example:** `drocertpart1`.

- d. Enter the **Certificate content**. Include **the Let's Encrypt R3 intermediate certificate**, issued to **US, Let's Encrypt, R3**. For more information, see [certificate content](#). **Example:**

```
-----BEGIN CERTIFICATE-----
MIIF5jCCBM6gAwIBAgISA0Y...
-----END CERTIFICATE-----
```

- e. Click **Confirm**. The first part of this certificate should include valid dates and look like the following example:

```
Issued to: US, Let's Encrypt, R3
Issued by: US, Internet Security Research Group, ISRG Root X1
Valid from: Thu Aug 01 2024
```

```
Valid to: Mon Sep 15 2025
```

```
This is the intermediate certificate which is required for the SSL connection to DRO endpoint.
```

- f. Click **Add** to add the **root of the certificate chain**.
- g. Enter an **alias**. **Example:** drocertpart2.
- h. Enter the **Certificate content**. Include the ISRG Root X1 self-signed certificate. For more information, see [certificate content](#). Example:

```
-----BEGIN CERTIFICATE-----  
MIIFazCCA10gAw...  
-----END CERTIFICATE-----
```

- i. Click **Confirm**. The **second part of this certificate** should have valid dates and look like the following example:

```
Issued to: US, Internet Security Research Group, ISRG Root X1  
Issued by: US, Internet Security Research Group, ISRG Root X1  
Valid from: Thu Jun 04 2015  
Valid to: Mon Jun 04 2035
```

```
This is the root certificate which is required for the SSL connection to DRO endpoint.
```

- j. **Save** the DRO configuration.
- k. Now, wait for the DRO configuration to reconcile, this process might take up to 10 minutes. The configuration will be successfully completed when the configuration status is set to Ready. Example:
Configuration Ready - DRO configuration was successfully verified

6. Configure the Suite License Service.

The Suite License Service (SLS) stores and manages the Maximo Application Suite license.

Each Maximo Application Suite instance can be connected to a unique SLS instance. Two or more Maximo Application Suite instances can also share an SLS and the corresponding license file.

Enter the following SLS information to configure Maximo Application Suite:

- URL - The URL for the SLS server.
- Registration key - Enter the SLS registration key.

Depending on your environment, the SLS configuration might take 10 minutes or more to complete.

7. Optional: Upload your license key file.

If the IBM Suite License Service that you configured for use with Maximo Application Suite includes a valid license file, you do not need to upload a license file. You can continue with the next configuration step.

To activate Maximo Application Suite, you must provide your license key from the [IBM License Key Center](#). The login information is provided in the license Key Center welcome letter. For more help on licensing, see the [IBM Support - Licensing page](#).

- a) Log in to the license Key Center.
- b) Select your company name.
- c) Select the **IBM AppPoints** product line.
- d) Select the IBM Maximo Application Suite... license key name.
- e) Select the product or sales order for which to create the license key.
- f) Enter the number of keys to generate. These correspond to the AppPoints that are allocated to the license key.
- g) Provide the Maximo Application Suite license server parameters.

Use the parameters that are displayed in the **Advanced settings > license key** section of the Maximo Application Suite setup program, or provide the following parameters:

- For Configuration, specify a Single License Server.
- For Host ID type, specify the Ethernet address.
- For Host ID, specify the host ID that was generated when you installed the Suite License Service (SLS). To display this ID, connect to your Red Hat OpenShift cluster and run the following command:

```
oc -n <sls_project_namespace> get licenseservice sls
```

For example, if the namespace of the SLS project is mas-sls-dev5, run the following command:

```
oc -n mas-sls-dev5 get licenseservice sls
```

In the command output, the host ID is displayed in the LICENSEID column.

- For Hostname, specify a hostname of your choice, for example: sls-mas
- For Port, specify 27000.

h) Download the key and then upload it to the Maximo Application Suite setup program.

8. Create the workspace.

The Maximo Application Suite workspace is a unique collection of configuration settings for your instance of Maximo Application Suite. Enter the following information to create your Maximo Application Suite workspace:

- Workspace ID

The workspace ID forms part of the Maximo Application Suite URL, for example:

```
https://<workspace_id>.home.<mas_domain>
```

Note: The workspace ID must be 3 - 12 characters in length, and can contain only lowercase letters and numbers. The first character must be a letter.

- Workspace display name

The display name is shown in your Maximo Application Suite user interface.

9. Review the setup configuration.

Your Maximo Application Suite setup is now complete. Verify that all configuration settings are done and then click **Finish** to complete the setup.

What to do next

After the Maximo Application Suite setup is complete, you can start to use your environment by going to the Maximo Application Suite administration or the Maximo Application Suite navigator page:

```
https://admin.<mas_domain>  
https://<workspace_id>.home.<mas_domain>
```

As the Maximo Application Suite superuser, you can now continue configuring your environment to suite your enterprise needs:

- [Configure authentication](#)

Maximo Application Suite supports local user authentication by MongoDB and authentication by using LDAP or SAML.

- [Configure LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

- [Create administrator user accounts](#)

The initial superuser account is used to complete the Maximo Application Suite setup. You can add application administrator users or system administrator users for day-to-day administrative tasks.

- [Getting started](#)

With the setup completed, your users can log in and start to use Maximo Application Suite.

Related concepts

Simple Mail Transfer Protocol

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

Related tasks

Migrating Maximo Application Suite from User Data Services to Data Reporter Operator


As an IBM Maximo Application Suite administrator, configure the IBM Data Reporter Operator (DRO) to collect and process metrics for licensing compliance. New and existing Maximo Application Suite users can install or migrate to DRO by using the IBM Maximo Application Suite command line interface (CLI), ansible role, or manually.

Customer-managed

Storing configuration values as secrets

Maximo Application Suite uses Red Hat OpenShift for managing digital authentication credentials (secrets).

Input parameters that are stored as secrets are identified by a capital S.

After you submit a parameter as a secret, the secret value cannot generally be viewed by using the Maximo Application Suite user interface. Only secret parameter fields that contain an eye icon  can be viewed.

Red Hat OpenShift administrators can view Maximo Application Suite secret values at: Red Hat OpenShift dashboard > Projects > mas-`<instance_name>-core` > Workloads > Secrets.

Important: When you edit a Maximo Application Suite configuration that uses one or more secrets, such as a database password, you must reenter the secret values before you can save the updated configuration.

Procedure

- Updating secrets from the user interface.
 1. Open the configuration that includes the secret that you want to update, such as a database password or similar.
 2. Edit the configuration and enter the updated parameter value.
 3. Click **Save**.

The secret is updated.

- Updating the super user password.

The super user is a special Maximo Application Suite user that is not editable from the user interface. The super user password is stored as a secret.

To change the password for the super user:

1. Log in to the Red Hat OpenShift web console as an administrator.
2. Provide a new secret for: Red Hat OpenShift dashboard > Projects > mas-`<instance_name>-core` > Workloads > Secrets > `<instance_name>-credentials-superuser`
3. Delete the following pods, and then wait until they become ready again: `admin-dashboard` and `coreidp`.

Configure the global image pull secret so that you can pull image sources from the IBM Entitled Registry.

Procedure

1. Configure the pull secret to pull images from the IBM Entitled Registry, the global image pull secret must contain your [IBM entitlement API key](#).

```
export IBM_ENTITLEMENT_SERVER=cp.icr.io
export IBM_ENTITLEMENT_USER=cp
export IBM_ENTITLEMENT_KEY=xxxxx
```

2. Create an environment variable that points to a temporary directory on your workstation. For example:

```
export WORK_ROOT=$HOME/temp/work
```

Note: If the temporary directory you choose does not exist, you must create it before proceeding with the next step.

3. Perform the `oc login` command.

You can access your Red Hat OpenShift cluster by using `oc` command directly from a Terminal in the client machine `oc` was installed to.

To do so:

- a) Go to the Red Hat OpenShift web console.
- b) Click your login name and select the option: Copy login command.
- c) Click View token.
- d) Copy the entire command line under the Log in section with this token.
- e) Paste it in the Terminal of the client machine where `oc` was installed to and run the command. You should see a message saying that you can access a number of projects and a default project selected.

From now on, next time you run `oc` commands in this opened Terminal in the client machine, it will be running the `oc` commands in your OpenShift cluster.

4. Download the pull secret to the temporary directory:

```
oc get secret/pull-secret \
  -n openshift-config \
  --template='{{index .data ".dockerconfigjson" | base64decode}}' > ${WORK_ROOT}/
  global_pull_secret.cfg
```

5. Add the new pull secret to the local copy of the `global_pull_secret.cfg` file:

```
oc registry login \
  --registry="${IBM_ENTITLEMENT_SERVER}" \
  --auth-basic="${IBM_ENTITLEMENT_USER}:${IBM_ENTITLEMENT_KEY}" \
  --to=${WORK_ROOT}/global_pull_secret.cfg
```

6. Update the global pull secret on your cluster:

```
oc set data secret/pull-secret \
  -n openshift-config \
  --from-file=.dockerconfigjson=${WORK_ROOT}/global_pull_secret.cfg
```

Certificate management

Maximo Application Suite uses IBM Certificate Manager service to automatically manage SSL/TLS certificates for your apps and services and ensure that certificates are valid and up to date. Alternatively, you can also enable manual certificate management to upload your public transport layer security (TLS) certificates.

Customer-managed **Creating a ClusterIssuer**

Certificate Manager arranges to have certificates issued when applications first need them and also to renew them when they expire.

There are multiple ways to issue certificates that is described in the [cert-manager documentation](#).

One approach is to create a new certificate authority (CA). This might be either a new root CA, or an Intermediate CA signed by your enterprise CA (if you have one).

If you are planning to do this for a production instance, review it with your Enterprise IT security department and ensure that you are doing it on a secure machine.

For Maximo Application Suite instances installed in Cloud environments, we recommend you to use custom cluster issuers signed by the Let's Encrypt certificate authority.

Procedure

1. Install openssl on your local secure machine. Make sure that the following sections are included in your openssl config file (you can take a copy of the file and add the lines if they aren't there).

```
[ v3_ca ]
basicConstraints = critical,CA:TRUE
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer:always

[ v3_req ]
# Extensions to add to a certificate request
basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
```

2. Use openssl to generate a public/private key pair for the CA.

```
openssl genrsa -out ca.key 4096
```

3. Use openssl to generate a CA certificate. You need so set the -days parameter to the lifetime that you want for this certificate and -config to the path to the openssl configuration file (you can omit this if your default configuration file contains the lines mentioned in point 1).

```
openssl req -x509 -new -nodes -key ca.key -days 3650 -reqexts v3_req -extensions v3_ca -out ca.crt -config ~/openssl.conf
```

4. During this process, you may be asked to enter information to be included in the certificate. As this is a CA certificate you can fill any values you like for these fields, though you should use meaningful ones since you are going to distribute the certificate to end users. You do not need to use a host name for the Common Name.

```
Country Name (2 letter code) []:
State or Province Name (full name) []:
Locality Name (eg, city) []:
Organization Name (eg, company) []:
Organizational Unit Name (eg, section) []:
Common Name (eg, fully qualified host name) []:
Email Address []:
```

The ca.key file contains the private key for the CA, so make sure to keep it secure. The ca.crt file contains the CA's certificate and public key.

5. Create Red Hat OpenShift Secret containing the private key and the certificate.

```
oc create secret tls mas-ca-key-pair --cert=ca.crt --key=ca.key -n cert-manager
```

6. Create a ClusterIssuer resource that can issue certificates as this CA. You can do this from the command line as follows:
Copy this code into a file called `clusterIssuer.yaml`:

```
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: mas-ca-issuer
spec:
  ca:
    secretName: mas-ca-key-pair
```

7. Apply the file as follows:

```
oc apply -f clusterIssuer.yaml
```

Alternatively you can create the ClusterIssuer using the Red Hat OpenShift dashboard. Do this by clicking the + button and pasting the clusterIssuer YAML in directly (it's worth keeping a copy in the file though so that you have a record of what you have done). After you paste the YAML, check that the ClusterIssuer has been created.

Related concepts

[Red Hat OpenShift certificate manager](#)

The Red Hat OpenShift certificate manager service helps you manage and deploy SSL/TLS certificates for your apps and services. It provides you with a security-rich repository for your certificates and their associated private keys, and helps prevent outages by sending you notifications when your certificates are about to expire.

Customer-managed

Disabling default certificate authorities

Starting in Maximo Application Suite 8.11, you can disable default certificates. Maximo Application Suite provides a built-in set of certificate authority (CA) certificates and by default automatically trusts a certificate if that certificate is issued by one of these CAs. To disable the default CAs that are provided, you can update the custom resource (CR) file for Maximo Application Suite. If you disable the default trust then you need to specifically configure certificates and CAs for all external systems that Maximo Application Suite connects to.

About this task

The `trustDefaultCAs` variable in the custom resource (CR) is set to `True` to trust certificates that are issued by one of the default certificate authorities. If the `trustDefaultCAs` variable is not included in the CR, the use of the default certificate authorities is assumed as `True`.

To disable and prevent Maximo Application Suite from automatically trusting these default certificate authorities, you must set the `trustDefaultCAs` variable in the CR is set to `False`.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **CustomResourceDefinitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR.
3. In the **CustomResourcesDefinitions** window, on the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the **spec.settings** section, add or change the `trustDefaultCAs` variable with the setting of `False`.
5. Save the CR changes.

This update might take a few minutes to process. To validate that the update is applied, check that the setting for the **trustDefaultCAs** variable in the **status.settings** section is set to `False`.

What to do next

You will have to provide the required CA certificates when configuring all TLS connections from Maximo Application Suite to external services.

Related tasks

[Configuring certificate authority certificates](#)

Customer-managed

Configuring certificate authority certificates

If the service that you are connecting to uses the transport layer security (TLS) communication protocol and is not secured with a certificate that is issued by a certificate authority (CA) that is trusted by default by Maximo Application Suite, you must provide the certificate of the CA that issued the service's certificate. Since the CA might use intermediate CAs, you can provide more than one certificate.

If you configured **trustDefaultCAs** to `false` in the custom resource (CR), you must provide the CA for each service that you are connecting to.

About this task

Services that Maximo Application Suite might connect to include, MongoDB, Kafka, Db2, Watson Studio, Object Storage and SMTP. For more information, see [“Prerequisite software” on page 5](#).

For each certificate that you provide, the following details are displayed:

- The name of the certificate issuer.
- The name of the subject, such as the organization, that the certificate is issued to.
- The start and end dates of the certificate's validity period. If the validity of any certificate that you provide expires soon, a warning message appears.

Procedure

You can automatically retrieve or manually add certificates.

1. Automatically retrieving certificates

In the certificates section, click **Retrieve**. If the connection credentials that you specify are correct, all CA certificates that are configured on the server are automatically retrieved and displayed.

These certificates are not validated. Verify that only the correct certificates are retrieved and remove any unexpected certificates.

After you retrieve certificates, you can manually add more certificates.

2. Manually adding certificates

In the certificates section, click **Add manually** and specify the following values for each certificate that you want to add:

- Alias
 - An alphanumeric identifier that is between 3 and 50 characters long.
- Certificate content
 - The content of a certificate file in either the X.509 or PEM formats.

Related tasks

[Disabling default certificate authorities](#)

Configuring the size of public certificate resources

For certificates that are managed by IBM Cloud Certificate Manager, you can change the private key size of public certificates that are provided by Maximo Application Suite.

About this task

By default, certificates that are issued by Certificate Manager use a private key size of 2048 bits. You can change the default private key size for Maximo Application Suite public certificates that are issued by Certificate Manager to one of the following values:

- 2048
- 4096
- 8192

Changing the private key size does not apply to manual certificates.

Procedure

1. In the Red Hat OpenShift Container Platform console, from the side navigation menu, click **Administration > CustomResourceDefinitions**.
2. On the **CustomResourcesDefinitions** page, click **Suite**.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the **settings** section, update the size for `certManager.certificates.privateKey` to one of the following values:
 - 2048
 - 4096
 - 8192

The following example shows changing the size to 8192.

```
settings:
  certManager:
    certificates:
      privateKey:
        size: 8192
```

5. Save your changes.

Results

When you add or change the private key size setting, Maximo Application Suite issues new public certificates with the specified private key size.

Manual certificate management

When you configure the suite, you can enable manual certificate management to upload your public transport layer security (TLS) certificates in Maximo Application Suite.

Related tasks

[Enabling manual certificate management](#)

By default, Maximo Application Suite uses IBM Certificate Manager to automatically control certificate management. To upload your own public transport layer security (TLS) certificates, you must first enable manual certificate management by updating the custom resource (CR) file for Maximo Application Suite.

[Uploading public certificates in Red Hat OpenShift](#)

After you enable manual certificate management, you can add your public certificates in Red Hat OpenShift for the applications that are deployed in your cluster by directly applying the secrets to your cluster.

Certificate configuration

After you enable certificate management, you can add certificates by adding secrets to your cluster in Red Hat OpenShift or by uploading certificates for your instance in the Maximo Application Suite user interface.

The following certificates must be available for each application in Maximo Application Suite:

tls.crt

The server certificate to access Maximo Application Suite and the suite applications.

tls.key

The server certificate key for use of the server certificate in Maximo Application Suite.

ca.crt

The public certificate of the certificate authority (CA) that authorizes your server certificate.

You can add a different certificate for each application in Maximo Application Suite. However, the certificate authority must be the same across all applications. For example, if you have one certificate for Maximo Manage that is signed by your internal CA, other application certificates must also be signed by the same CA. Alternatively, you can have a single generated certificate that includes all Subject Alternative Names (SAN).

The following tables provide the Subject Alternative Names (SAN) that you include in your certificate for Maximo Application Suite applications. If you cannot use wildcards in your certificate, you must include all endpoints individually.

Use the following environment variables that were defined during installation:

masdomain

The domain for your Maximo Application Suite instance is set during installation, with the option for default values or custom configuration. You can locate the domain in your Red Hat OpenShift cluster in the suite custom resource (CR), specifically in the spec field.

workspaceid

The identifier for the unique workspace in which applications are deployed.

Subject Alternative names

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

If you use a single certificate for Maximo Application Suite, use the following Subject Alternative Names:

<i>Table 53. Subject Alternative Names (SAN) with and without wildcard for a single certificate</i>		
Application	SAN with wildcards	SAN without wildcards
Maximo Application Suite core	<ul style="list-style-type: none"> • '*.<masdomain>' • '*.home.<masdomain>' 	<ul style="list-style-type: none"> • admin.<masdomain> • api.<masdomain> • auth.<masdomain> • home.<masdomain> • <workspaceid>.home.<masdomain>

Table 53. Subject Alternative Names (SAN) with and without wildcard for a single certificate (continued)

Application	SAN with wildcards	SAN without wildcards
IoT	<ul style="list-style-type: none"> • '*.iot.<masdomain>' • '*.messaging.iot.<masdomain>' In 8.11, the following SAN with wildcards are applicable: <ul style="list-style-type: none"> • *.edgeconfig.iot.<masdomain> • *.edgeconfigapi.iot.<masdomain> 	<ul style="list-style-type: none"> • <workspaceid>.iot.<masdomain> • <workspaceid>.messaging.iot.<masdomain> • messaging.iot.<masdomain> In 8.11, the following SAN without wildcards are applicable: <ul style="list-style-type: none"> • edgeconfig.iot.<masdomain> • edgeconfigapi.iot.<masdomain> • <workspaceid>.edgeconfig.iot.<masdomain> • <workspaceid>.edgeconfigapi.iot.<masdomain>
Maximo Monitor	<ul style="list-style-type: none"> • '*.monitor.<masdomain>' • '*.api.monitor.<masdomain>' 	<ul style="list-style-type: none"> • admin.monitor.<masdomain> • api.monitor.<masdomain> • <workspaceid>.monitor.<masdomain> • <workspaceid>.api.monitor.<masdomain>
Maximo Optimizer	<ul style="list-style-type: none"> • '*.optimizer.<masdomain>' • '*.api.optimizer.<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid>.optimizer.<masdomain> • <workspaceid>.api.optimizer.<masdomain>

Table 53. Subject Alternative Names (SAN) with and without wildcard for a single certificate (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Manage	<ul style="list-style-type: none"> • '* .manage .<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid> .manage .<masdomain> • <workspaceid>-all .manage .<masdomain> • <workspaceid>-cron .manage .<masdomain> • <workspaceid>-mea .manage .<masdomain> • <workspaceid>-report .manage .<masdomain> • <workspaceid>-ui .manage .<masdomain> • maxinst .manage .<masdomain> <p>Starting in 9.1, the following SAN without wildcards is applicable:</p> <ul style="list-style-type: none"> • <workspaceid>-foundation .manage .<masdomain>
Maximo Health	<ul style="list-style-type: none"> • '* .health .<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid> .health .<masdomain> • <workspaceid>-all .health .<masdomain> • <workspaceid>-cron .health .<masdomain> • <workspaceid>-mea .health .<masdomain> • <workspaceid>-rpt .health .<masdomain> • <workspaceid>-ui .health .<masdomain> • maxinst .health .<masdomain>
Maximo Predict	<ul style="list-style-type: none"> • '* .predict .<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid> .predict .<masdomain> • predict .<masdomain>
Maximo Assist or Maximo Collaborate	<ul style="list-style-type: none"> • '* .assist .<masdomain>' or • '* .collaborate .<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid> .assist .<masdomain> or • <workspaceid> .assist .<masdomain>
Maximo Health and Predict - Utilities	<ul style="list-style-type: none"> • '* .hputilities .<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid> .hputilities .<masdomain>

Table 53. Subject Alternative Names (SAN) with and without wildcard for a single certificate (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Visual Inspection	• '*.visualinspection.<masdomain>'	• <workspaceid>.visualinspection.<masdomain>

Table 53. Subject Alternative Names (SAN) with and without wildcard for a single certificate (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Real Estate and Facilities	<ul style="list-style-type: none"> • *.facilities.<masdomain> • *.<workspaceid>.facilities.<masdomain> 	<ul style="list-style-type: none"> • facilities.<masdomain> • <workspaceid>.facilities.<masdomain> • multiagents.<workspaceid>.facilities.<masdomain> • dataconnectagent.<workspaceid>.facilities.<masdomain> • extendedformulaagent.<workspaceid>.facilities.<masdomain> • formularecalcagent.<workspaceid>.facilities.<masdomain> • incomingmailagent.<workspaceid>.facilities.<masdomain> • objectmigrationagent.<workspaceid>.facilities.<masdomain> • objectpublishagent.<workspaceid>.facilities.<masdomain> • pmscheduleragent.<workspaceid>.facilities.<masdomain> • reportqueueagent.<workspaceid>.facilities.<masdomain> • reservesmtpagent.<workspaceid>.facilities.<masdomain> • wfagent.<workspaceid>.facilities.<masdomain> • wffutureagent.<workspaceid>.facilities.<masdomain> • wfnotificationagent.<workspaceid>.facilities.<masdomain> • appserver.<workspaceid>.facilities.<masdomain> • pod-0.<workspaceid>.facilities.<masdomain> • pod-1.<workspaceid>.facilities.<masdomain> • ... • pod-n.<workspaceid>.facilities.<masdomain> • dwfagent-

If you use a certificate for each application in Maximo Application Suite, use the following Subject Alternative Names:

<i>Table 54. Subject Alternative Names (SAN) with and without wildcard for multiple application certificates</i>		
Application	SAN with wildcards	SAN without wildcards
Maximo Application Suite core	<ul style="list-style-type: none"> '*.<masdomain>' '*.home.<masdomain>' 	<ul style="list-style-type: none"> admin.<masdomain> api.<masdomain> auth.<masdomain> home.<masdomain> <workspaceid>.home.<masdomain>
IoT	<ul style="list-style-type: none"> *.<masdomain> '*.iot.<masdomain>' '*.messaging.iot.<masdomain>' <p>In 8.11, the following SAN with wildcards are applicable:</p> <ul style="list-style-type: none"> *.edgeconfig.iot.<masdomain> *.edgeconfigapi.iot.<masdomain> 	<ul style="list-style-type: none"> iot.<masdomain> <workspaceid>.iot.<masdomain> <workspaceid>.messaging.iot.<masdomain> messaging.iot.<masdomain> <p>In 8.11, the following SAN without wildcards are applicable:</p> <ul style="list-style-type: none"> edgeconfig.iot.<masdomain> edgeconfigapi.iot.<masdomain> <workspaceid>.edgeconfig.iot.<masdomain> <workspaceid>.edgeconfigapi.iot.<masdomain>
Maximo Monitor	<ul style="list-style-type: none"> *.<masdomain> '*.monitor.<masdomain>' '*.api.monitor.<masdomain>' 	<ul style="list-style-type: none"> monitor.<masdomain> admin.monitor.<masdomain> api.monitor.<masdomain> <workspaceid>.monitor.<masdomain> <workspaceid>.api.monitor.<masdomain>
Maximo Optimizer	<ul style="list-style-type: none"> *.<masdomain> '*.optimizer.<masdomain>' '*.api.optimizer.<masdomain>' 	<ul style="list-style-type: none"> optimizer.<masdomain> <workspaceid>.optimizer.<masdomain> <workspaceid>.api.optimizer.<masdomain>

Table 54. Subject Alternative Names (SAN) with and without wildcard for multiple application certificates (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Manage	<ul style="list-style-type: none"> • *.<masdomain> • '*.manage.<masdomain>' 	<ul style="list-style-type: none"> • manage.<masdomain> • <workspaceid>.manage.<masdomain> • <workspaceid>-all.manage.<masdomain> • <workspaceid>-cron.manage.<masdomain> • <workspaceid>-mea.manage.<masdomain> • <workspaceid>-report.manage.<masdomain> • <workspaceid>-ui.manage.<masdomain> • maxinst.manage.<masdomain> <p>Starting in 9.1, the following SAN without wildcards is applicable:</p> <ul style="list-style-type: none"> • <workspaceid>-foundation.manage.<masdomain>
Maximo Health	<ul style="list-style-type: none"> • *.<masdomain> • '*.health.<masdomain>' 	<ul style="list-style-type: none"> • health.<masdomain> • <workspaceid>.health.<masdomain> • <workspaceid>-all.health.<masdomain> • <workspaceid>-cron.health.<masdomain> • <workspaceid>-mea.health.<masdomain> • <workspaceid>-rpt.health.<masdomain> • <workspaceid>-ui.health.<masdomain> • maxinst.health.<masdomain>
Maximo Predict	<ul style="list-style-type: none"> • '*.predict.<masdomain>' 	<ul style="list-style-type: none"> • <workspaceid>.predict.<masdomain> • predict.<masdomain>

Table 54. Subject Alternative Names (SAN) with and without wildcard for multiple application certificates (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Assist or Maximo Collaborate	<ul style="list-style-type: none"> • *.<masdomain> • '*.assist.<masdomain>' or '*.collaborate.<masdomain>' 	<ul style="list-style-type: none"> • assist.<masdomain> or collaborate.<masdomain> • <workspaceid>.assist.<masdomain> or <workspaceid>.collaborate.<masdomain>
Maximo Health and Predict - Utilities	<ul style="list-style-type: none"> • *.<masdomain> • '*.hputilities.<masdomain>' 	<ul style="list-style-type: none"> • hputilities.<masdomain> • <workspaceid>.hputilities.<masdomain>
Maximo Visual Inspection	<ul style="list-style-type: none"> • *.<masdomain> • '*.visualinspection.<masdomain>' 	<ul style="list-style-type: none"> • visualinspection.<masdomain> • <workspaceid>.visualinspection.<masdomain>

Table 54. Subject Alternative Names (SAN) with and without wildcard for multiple application certificates (continued)

Application	SAN with wildcards	SAN without wildcards
Maximo Real Estate and Facilities	<ul style="list-style-type: none"> • *.<masdomain> • *.facilities.<masdomain> • *.<workspaceId>.facilities.<masdomain> 	<ul style="list-style-type: none"> • facilities.<masdomain> • <workspaceid>.facilities.<masdomain> • multiagents.<workspaceid>.facilities.<masdomain> • dataconnectagent.<workspaceid>.facilities.<masdomain> • extendedformulaagent.<workspaceid>.facilities.<masdomain> • formularecalcagent.<workspaceid>.facilities.<masdomain> • incomingmailagent.<workspaceid>.facilities.<masdomain> • objectmigrationagent.<workspaceid>.facilities.<masdomain> • objectpublishagent.<workspaceid>.facilities.<masdomain> • pmscheduleragent.<workspaceid>.facilities.<masdomain> • reportqueueagent.<workspaceid>.facilities.<masdomain> • reservesmtpagent.<workspaceid>.facilities.<masdomain> • wfagent.<workspaceid>.facilities.<masdomain> • wffutureagent.<workspaceid>.facilities.<masdomain> • wfnotificationagent.<workspaceid>.facilities.<masdomain> • appserver.<workspaceid>.facilities.<masdomain> • pod-0.<workspaceid>.facilities.<masdomain> • pod-1.<workspaceid>.facilities.<masdomain> • ... • pod-n.<workspaceid>.facilities.<masdomain>

Enabling manual certificate management

By default, Maximo Application Suite uses IBM Certificate Manager to automatically control certificate management. To upload your own public transport layer security (TLS) certificates, you must first enable manual certificate management by updating the custom resource (CR) file for Maximo Application Suite.

About this task

Enabling certificate management manually is available to upload public TLS certificates. Internal certificates are still created and managed by using IBM Certificate Manager.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. In the **CustomResourcesDefinitions** window, on the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the **spec.settings** section, change the **manualCertMgmt** variable from `False` to `True`.
5. Validate that the **spec.settings.manualCertMgmt** setting is `True`.
6. Save the CR changes.

Results

You can add certificates by updating the secret names in the namespace of each cluster or uploading the certificates in the Maximo Application Suite user interface.

To upload in the user interface, log in to Maximo Application Suite and select the suite administration page on the side navigation. You can upload the certificates on the **Certificates** page or when you initially deploy and activate applications.

To revert to controlling certificates automatically by using IBM Certificate Manager, you can change the **manualCertMgmt** to `False`. By switching to automatic certificate management, the certificates that you uploaded manually are replaced and cannot be retrieved from the system. If you change back to manual certificate management, you must upload your certificates again.

Related concepts

[Manual certificate management](#)

When you configure the suite, you can enable manual certificate management to upload your public transport layer security (TLS) certificates in Maximo Application Suite.

Related tasks

[Uploading public certificates in Red Hat OpenShift](#)

After you enable manual certificate management, you can add your public certificates in Red Hat OpenShift for the applications that are deployed in your cluster by directly applying the secrets to your cluster.

Uploading public certificates in Red Hat OpenShift

After you enable manual certificate management, you can add your public certificates in Red Hat OpenShift for the applications that are deployed in your cluster by directly applying the secrets to your cluster.

Procedure

1. Log in to the Red Hat OpenShift web console as an administrator.
2. Go to the application namespace that you want to add your certificates to.

- a) To access the Maximo Application Suite namespace, select **Red Hat OpenShift dashboard > Projects > mas-`<instance_name>-core`**.
 - b) To access the namespace for applications in Maximo Application Suite, select **Red Hat OpenShift dashboard > Projects > mas-`<instance_name>-<appName>`**.
3. Update the secret name in each application that is deployed for the following certificates.
- `tls.crt`
 - `tls.key`
 - `ca.crt` - The `ca.crt` certificate is optional.

If the secret does not exist for the application, create a secret by using the secret name that is provided for the application, which is shown in the following table.

<i>Table 55. Secret names for each application</i>	
Application	Secret name
Maximo Application Suite core	INSTANCE_ID-cert-public
IoT	INSTANCE_ID-public-tls
Maximo Monitor	INSTANCE_ID-public-tls
Maximo Health	INSTANCE_ID-WORKSPACE_ID-cert-public-81
Maximo Manage	INSTANCE_ID-WORKSPACE_ID-cert-public-81
Maximo Real Estate and Facilities	INSTANCE_ID-WORKSPACE_ID-public-facilities-tls
Maximo Predict	INSTANCE_ID-public-predict-tls
Maximo Visual Inspection	public-visualinspection-tls
Maximo Assist or Maximo Collaborate*	public-assist-tls or public-collaborate-tls
Maximo Optimizer	INSTANCE_ID-cert-optimizer-public

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Related concepts

[Manual certificate management](#)

When you configure the suite, you can enable manual certificate management to upload your public transport layer security (TLS) certificates in Maximo Application Suite.

Related tasks

[Enabling manual certificate management](#)

By default, Maximo Application Suite uses IBM Certificate Manager to automatically control certificate management. To upload your own public transport layer security (TLS) certificates, you must first enable manual certificate management by updating the custom resource (CR) file for Maximo Application Suite.

Customer-managed Authentication options for Db2U

When using the standalone Db2U operator to provision a Db2 instance there are two options for user authentication.

1. Local OS user authentication. For example, using the `db2inst1` user.
2. LDAP authentication. For example, using the local LDAP registry that is provisioned with the Db2 instance.

1. Get the pod name of the local LDAP service:

```
oc get pod -n db2u | grep ldap
```

Sample output

```
c-db2u-manage-ldap-8469ff9f7b-g4qlp 1/1 Running 0 2m27s
```

2. Create the user in the local LDAP registry:

```
oc exec -it c-db2u-manage-ldap-8469ff9f7b-g4qlp -n db2u -- /opt/ibm/ldap_scripts/addLdapUser.py -u user1 -r admin -p password
```

Sample output

```
Next UID will be 5003
Adding user1 to LDAP server
Updating LDAP password for user user1
Added user to LDAP server
```

3. Test that the newly created user can connect to Db2:

```
oc exec -it c-db2u-manage-db2u-0 -n db2u -c db2u -- su -lc "db2 connect to bludb user user1 using password" db2inst1
```

Sample output

```
Database Connection Information

Database server          = DB2/LINUX8664 11.5.7.0
SQL authorization ID    = USER1
Local database alias    = BLUDB
```

Note:

User authentication with an LDAP server will incur extra latency at connection time that is not required when using OS user authentication.

Customer-managed **Storage**

The storage configurations control Maximo Application Suite storage options at the system and application scope.

Database connection

A database is used for Maximo Application Suite application data and data analytics storage.

Configuration parameters

The following parameters are configurable:

- URL

The URL is of the following form:

```
protocol://<hostname>:<port>/<database_name>
```

- Protocol
Example: jdbc:db2
- Hostname
- Port

- Database name
- Maximo Manage only: SSL Enabled
 - Important:** SSL enabled must be set if you are adding certificates for the Maximo Manage connection.
- Retrieve or add a CA certificate.
- Username
- Password
- Maximo Manage only: Optional driver options

Tip: If your Db2 Warehouse is running under Cloud Pak for Data you can obtain the database name, SSL port number and certificate from the Cloud Pak for Data administration console.

Maximo Monitor and IoT tool requirements

The System-scoped Db2 Warehouse configuration for the Maximo Monitor application and the IoT tool must fulfill the following requirements:

- URL:
 - Append `;sslConnection=true` to the URL so that it has the form:
`jdbc:db2://hostname:port/database_name;sslConnection=true`
 - Use the secure (SSL) port number.
- Set the SSL Enabled slider to **Yes**.
- Retrieve or add the [CA certificate](#) that is used by the Db2 Warehouse service.

Required by

- Maximo Monitor at the System scope.
- Maximo Predict

Object Storage

Object Storage is a required component for Maximo Collaborate. For more information, see [Install Ceph Object Storage](#).

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Configuration parameters

The following parameters are configurable:

- URL
 - S3-compatible object storage URL.
- Username
- Password

Retrieve or add a [CA certificate](#).

Required by Maximo Collaborate at the System scope.

Related information

[Sizing guidance](#)

User authentication

Users authentication is managed by using local user authentication, Lightweight Directory Access Protocol (LDAP) authentication, Security Assertion Markup Language (SAML) authentication, and OpenID Connect (OIDC) authentication. After you set up identity providers, you can configure them to provide multiple

login options that users can authenticate to when they log in. You can also specify a default identity provider to be the primary login option and enable seamless login for SAML.

Customer-managed **Authentication methods**

Maximo Application Suite supports MongoDB, Lightweight Directory Access Protocol (LDAP) authentication, and Security Assertion Markup Language (SAML) authentication methods for local user authentication.

Regardless of where authentication is managed, access management and user privileges authorization is managed by Maximo Application Suite.

Starting in Maximo Application Suite 8.11, you can configure local, SAML, or LDAP authentication to provide multiple login options that users can authenticate to when they log in. You can also specify a default identity provider to be the primary login option for users on the suite login page. For more information, see [configuring default identity providers](#).

Local authentication by MongoDB

With local authentication, Maximo Application Suite provides single sign-on (SSO) for all fully integrated applications.



Attention: Starting in Maximo Application Suite 8.11, when you create a user, if an authentication type such as **Local**, **LDAP**, or **SAML** is not selected, the user cannot access Maximo Application Suite.



Attention: Starting in Maximo Application Suite 9.0, when you upgrade a user from Maximo Asset Management 7.6, you can select ownership as follows:

- **local** ownership for **Local** authentication method.
- **scim** ownership for **LDAP** authentication method.
- **local** or **scim** ownership for **SAML** authentication method.

LDAP authentication

With LDAP, the user authentication is managed by your LDAP server. You can configure your Maximo Application Suite environment to use your own corporate LDAP server. Maximo Application Suite provides SSO for all fully integrated applications, and you can also configure external applications to use the same LDAP server.

With LDAP enabled, you can:

- Select to use LDAP authentication when you [create new users](#). LDAP uses its own username to link to Maximo Application Suite users.
- [Synchronize your LDAP user registry](#) with Maximo Application Suite, immediately setting up your suite users from your existing user registry.

Important: For synchronization, secure LDAP (LDAPS) is the only allowed protocol. Non-TLS connections are not supported.

You can configure Maximo Application Suite to use LDAP at setup or later. For more information about configuring Maximo Application Suite for LDAP, see [“Configuring LDAP authentication” on page 608](#).

SAML authentication

With SAML, the user authentication is managed by your SAML server. When SAML is enabled, you can complete the following tasks:

- Select to use SAML authentication when you [create new users](#). SAML uses its own ID to link to Maximo Application Suite users.

- Set up SSO for Maximo Application Suite and for any external application that supports SAML and that is accessed from the same browser.

You can configure Maximo Application Suite to use SAML at setup or later. For more information about configuring Maximo Application Suite for SAML, see [“Configuring SAML authentication” on page 607](#).

Starting in Maximo Application Suite 8.11, SAML authentication supports the following types of global configuration:

Non-Default (either local or ldap is default)

When SAML is configured but not the default identity provider (IdP), the option to log in by using SAML is available on the Maximo Application Suite login page as an alternative option. If users select this option, they are directed to use the SAML authentication.

Set as the default identity provider with seamless login enabled

When SAML is set as the default IdP with seamless login enabled, the authentication occurs directly in the SAML IdP. With seamless login, users are directed to the SAML IdP to authenticate instead of the Maximo Application Suite login page.



Attention: If you enable seamless login, then the login page is not shown. If you need to display a security message to comply with federal regulations, make sure that seamless login is disabled. Otherwise, users do not see the system notification that might be enabled on the login page.

Set as the default identity provider with seamless login disabled

When SAML is set as the default IdP but with seamless login disabled, the option to log in by using SAML is available on the Maximo Application Suite login page as the primary login option.

For more information, see [configuring default identity providers](#).

SAML standard also provides the following authentication methods that Maximo Application Suite supports.

Service Provider (SP) initiated

If users access any Maximo Application Suite endpoint, such as Manage or Monitor applications, an internal OIDC process is triggered in Maximo Application Suite where one of the following scenarios occurs:

- If SAML is set as the default IdP but with seamless login disabled, the Maximo Application Suite login page is shown where users can either click **Continue** to redirect them to the SAML IdP login page or choose an alternative login option.
- If SAML is set as the default IdP with seamless login enabled, the user is directed to the SAML IdP login page instead of the Maximo Application Suite login page.
- If the user already logged in, the Maximo Application Suite application opens immediately.

IdP initiated

For IdP initiated, users access Maximo Application Suite from the SAML IdP portal. If SAML is set as the default IdP with seamless login enabled, the user goes directly to the Maximo Application Suite page that is setup in the IdP **relayState** parameter. IdP administrators can configure any Maximo Application Suite application, such as the Manage application, in the **relayState** so that the application page opens directly.

If you need to keep another IDP as the default while seamless IDP initiated login is still required, you can use a different endpoint for the **relayState** parameter: `https://auth.<masdomain>/idplogin/idpinitiated`. To access a specific application in Maximo Application Suite, use the following parameter values:

- **appid** - The ID of the application that you want to access.
- **wsid** - The ID of the workspace in your Maximo Application Suite environment.
- **apppath** - An additional path that can be added as part of the URL.

For example, if you want all your users to have direct access to the Manage application in the workspace *main*, set the **relayState** parameter as `https://auth.<masdomain>/idplogin/idpinitiated?appid=manage&wsid=main`

Related concepts

[Identity providers](#)

[Identity provider prerequisites for Maximo Application Suite.](#)

Related tasks

[Configuring LDAP authentication](#)

By configuring LDAP authentication with Maximo Application Suite, user authentication is managed by your LDAP server.

[Configuring SAML authentication](#)

By configuring SAML user authentication with Maximo Application Suite, you integrate Maximo Application Suite as a service provider (SP) with your organization's SAML Identity Provider (IdP).

Related reference

[“Streamlined login” on page 612](#)

As an IBM Maximo Application Suite administrator, configure the properties that control the streamlined login experience for your users.

Configuring SAML authentication

By configuring SAML user authentication with Maximo Application Suite, you integrate Maximo Application Suite as a service provider (SP) with your organization's SAML Identity Provider (IdP).

About this task

Your Maximo Application Suite server acts as a service provider for the SAML identity provider (IdP). You need to provide a preferred service provider name and select a name identifier format, or you can use the default values. The information is written to a service provider metadata file that you use to configure your SAML identity provider (IdP).

Starting in Maximo Application Suite 9.0, you can enable initiated logout for the SAML service provider so that current user sessions are logged out before another user logs in with the same credentials.

Procedure

1. From the **Suite administration** menu, select **Configurations** from the side navigation menu and then click **SAML authentication**.
2. Create SAML service provider information.
 - a) Specify the display name.
 - b) Specify your preferred service provider name that is used to register the Maximo Application Suite service provider.
 - c) Select a name identifier format that is used with the SAML server.
 - d) To enforce the logout of a user from the SAML service provider before another user logs in with the same credentials, select **Enable initiated logout**.

Note: When you configure your SAML IdP, you must also select the option to initiate single logout that is described in step 3b.
 - e) Generate the metadata file and then download the file that you use to configure the data with SAML identity provider.

If you change the service provider name or the user identifier format after the initial configuration, you must save and generate the Service provider metadata file to register the changes with your SAML provider.

3. Register with the SAML identity provider.
 - a) In your IdP, configure your SAML IdP to recognize Maximo Application Suite.

You can upload the metadata file to your SAML IdP or use the contents to configure the SAML IdP to recognize Maximo Application Suite requests.

- b) Enable the single logout to initiate single logout with Maximo Application Suite and add the single logout URL.

The single logout URL is the single logout service in the metadata file. For example, `https/auth.<domain>/ibm/saml20/<saml_name>/slo`.

- c) After you configured your SAML IdP, download the SAML IdP metadata XML file to import it into Maximo Application Suite.

4. To complete the SAML configuration with Maximo Application Suite, in the **SAML authentication** page, upload the SAML IdP metadata XML file.

Related concepts

Authentication methods

Maximo Application Suite supports MongoDB, Lightweight Directory Access Protocol (LDAP) authentication, and Security Assertion Markup Language (SAML) authentication methods for local user authentication.

Configuring LDAP authentication

By configuring LDAP authentication with Maximo Application Suite, user authentication is managed by your LDAP server.

Procedure

1. From the **Suite administration** menu, select **Configurations** from the side navigation menu and then click **LDAP authentication**.
2. Specify the LDAP configuration parameters.

Parameter	Details	Example
URL	The URL for the LDAP server is in the format: <code>protocol://<hostname>:<port></code> Note: Secure LDAP (LDAPS) is the only allowed protocol. Non-TLS connections are not allowed.	Example: <code>ldaps://MSAD2021.fyre.ibm.com:636</code>
Base DN	The path in the object hierarchy of the directory server.	Example: <code>OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com</code>
Bind DN	Bind DN is used to bind to an LDAP server. Administrators must have sufficient privileges to search for users under user search DN or groups under group search DN.	Example: <code>CN=wilson,OU=users,OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com</code>
Bind PW	LDAP admin password	
userIdMap	The field that is used for user IDs	Example: <code><sAMAccountName></code>

3. Add or retrieve [CA certificate](#).

What to do next

You can synchronize your LDAP user registry with Maximo Application Suite to create users from your existing user registry.

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

[Authentication methods](#)

Maximo Application Suite supports MongoDB, Lightweight Directory Access Protocol (LDAP) authentication, and Security Assertion Markup Language (SAML) authentication methods for local user authentication.

Configuring OIDC authentication

Starting in Maximo Application Suite 9.1, you can configure OpenID Connect (OIDC) authentication with Maximo Application Suite for user authentication.

Before you begin

Register Maximo Application Suite with the OpenID Connect provider (OP) that you are using. Use the following URL to register:

```
https://auth.<masdomain>/oidcclient/redirect/default-oidc/
```

After you register, you receive a client ID and a client secret. You use this information when you are configuring the OIDC with Maximo Application Suite.

About this task

When you configure OIDC authentication, you must specify the discovery endpoint to retrieve information about the OIDC provider. If you don't provide the discovery endpoint, then you must provide the information for the following fields:

- Issue identifier
- Authorization endpoint
- Token endpoint
- JSON Web Key sets URI (JWKS) URI

Procedure

1. On the side navigation menu, select **Suite > Administration > Configurations** and then click **OIDC authentication**.
2. Specify the OIDC provider information.

Field	Description
Display name	The name of the provider.

Table 57. OIDC configuration attributes (continued)	
Field	Description
User identifier	<p>The JSON Web Token that connects to the OIDC provider. Select one of the following token properties that you want to use as the user identifier:</p> <ul style="list-style-type: none"> • sub • preferred_username • email • name
Client ID registered with the OIDC provider	<p>The Client ID that you registered with the OIDC provider. You receive this information after you register the client with the OpenID Connect provider (OP) that you are using.</p>
Client secret	<p>The secret is to authenticate to the OIDC provider. You receive this information after you register the client with the OpenID Connect provider (OP) that you are using.</p>
Signature algorithm	<p>The algorithm your OIDC provider supports. The default is typically R256.</p> <p>To find your token endpoint by using a browser, enter <code>https://<your_domain>/.well-known/openid-configuration</code>.</p>
Discovery endpoint	<p>The URL for the OIDC discovery endpoint to retrieve information about the OIDC provider. The URL format is <code>https://<your_domain>/.well-known/openid-configuration</code></p> <p>If set, other endpoints are automatically discovered.</p> <p>Note: If you don't specify the discovery endpoint, then you must specify the following fields:</p> <ul style="list-style-type: none"> • Issue identifier • Authorization endpoint • Token endpoint • JSON Web Key sets URI (JWKS) URI
Token endpoint authentication method	<p>You can select one of the following methods:</p> <ul style="list-style-type: none"> • Post • Basic • Private key JWT <p>Select Basic if the method is not known.</p>

Table 57. OIDC configuration attributes (continued)	
Field	Description
Token endpoint signature algorithm	The algorithm that your OIDC provider supports. The default is typically R256. To find your token endpoint by using a browser, enter <code>https://<your_domain>/.well-known/openid-configuration</code> .

3. If the **Discover endpoint** field is not defined, enter the following information.

Table 58. OIDC configuration attributes	
Field	Description
Issue identifier	The URL for the issuer identifier.
Authorization endpoint	The URL for the authorization endpoint.
Token endpoint	The URL for the token endpoint.
JSON Web Key sets (JWKS) URI	The JWKS URI.

4. Add or retrieve [CA certificate](#).

This certificate is the same as the one used for HTTPS at `https://<your_domain>/.well-known/openid-configuration` and can be download from a browser.

5. Save your changes.

Results

When you configure the authentication in the user interface, the name `default-oidc` is used to identify the configuration. For any additional configurations that you can create by using APIs, the identity name is the name that you can provide. For more information, see [“Configuring multiple identity providers for same authentication type” on page 617](#)

Configuring default identity providers

If you configure more than one identity provider, such as LDAP or SAML, you can specify which identity provider is the primary login option for users by updating authentication options on the **Suite administration** page. Alternatively, you can update the custom resource file in Red Hat OpenShift Container Platform.

About this task

Starting in Maximo Application Suite 8.11, you can select one of the following values to set a default identity provider:

- **Local**
- **default-saml**
- **default-ldap**

Starting in Maximo Application Suite 9.1, you can also set **default-oidc** default identity provider.

If you specify SAML as the default identity provider, you can enable seamless login so that users authenticate to Maximo Application Suite by using the login page that uses the SAML identity provider.



Attention: If you enable seamless login, then the Maximo Application Suite login page is not shown. If you need to display a security message to comply with federal regulations, ensure that seamless login is disabled. Otherwise, users do not see any system notification that might be shown on the Maximo Application Suite login page. For more information, see [“Enabling login notification” on page 713](#).

Procedure

1. To configure a default identity provider in the user interface, specify the identity provider on the **Authentication** page.
 - In Maximo Application Suite 9.1, from the side navigation menu, select **Suite > Administration > Authentication**.
 - In Maximo Application Suite 9.0 and earlier, on the **Suite administration** page, select **Users** from the side navigation menu and then select the **Authentication** tab.
 - a) In the Default login section, select the default identity provider from the list.
 - b) To enable seamless login for users to authenticate to Maximo Application Suite by using login page that uses the SAML identity provider, select the **Enable** check box.
 - c) Save your changes.
2. **Customer-managed**
To configure a default identity provider in the custom resource file, specify the identity provider for the defaultIDP.
 - a) In the Red Hat OpenShift Container Platform console, from the side navigation menu, in the Administration section, select **CustomResourceDefinitions**.
 - b) In the **CustomResourcesDefinitions** window, select the Suite CR file.
 - c) On the **Instances** tab, select the instance that you want to update.
 - d) On the **YAML** tab, change the value for `spec.settings.sso.defaultIDP` to the one of the following values.
 - `local`
 - `default-saml`
 - `default-ldap`
 - `default-oidc`
 - e) To enable seamless login for users to authenticate to Maximo Application Suite by using the login page that uses the SAML identity provider, change the `seamlessLogin` to `true`.

Results

When multiple identity providers are configured, users have multiple options to choose from when they log in to Maximo Application Suite where the default identity provider is set as the primary login option.

Related tasks

[Enabling login notification](#)

Starting in Maximo Application Suite 8.11, you can create and display a system message on the login page to provide security and privacy information to users.

Streamlined login

As an IBM Maximo Application Suite administrator, configure the properties that control the streamlined login experience for your users.

In the Suite custom resource (CR), the `spec.sso.defaultIDP` and `spec.sso.seamlessLogin` properties control the streamlined login experience in Maximo Application Suite.

You can enable the seamless login experience in the Maximo Application Suite. In the **Suite administration** page, click **Users** and then click the **Authentication** tab. In the **Default login behavior** section, enable seamless login for SAML authentication.

Dedicated login pages for identity providers

When a user accesses a Maximo Application Suite protected page, and if the user is not authenticated, the user is redirected to the default login page. The default login page requests user credentials and also has other login options, such as LDAP and SAML if these options are configured in Maximo Application Suite. However, a user can access specific endpoints that take them directly to a version of the login page that is configured with an identity provider option without other login options. The main format of the dedicated login pages for identity providers is `https://auth.<masdomain>/idplogin/loginpage<query parameters>`. For SAML dedicated login pages, no credentials are required for Maximo Application Suite, so the user is automatically redirected to the dedicated login page for the identity provider.

Important: A dedicated login page must not be used by default. It is an alternative mechanism to access Maximo Application Suite without going to the default login page. For example, for SAML seamless integration, you must set SAML as default and seamless instead of using a SAML dedicated login page.

Required query parameters

You can pass only the `&idp=<idp type>:<idp id>` query parameter to the dedicated login page endpoint. For local authentication, the `<idp id>` cannot be included, and the query parameter is `&idp=local`. For LDAP and SAML identity providers, the type and ID are needed. For example, for LDAP, the query parameter is `&idp=ldap:default-ldap`, where the IDP ID is `default-ldap`. For SAML, the query parameter is `&idp=saml:default-saml`, where the identity provider ID is `default-saml`. After the user logs in, they are redirected to the Maximo Application Suite home page by default, unless the user appends optional query parameters to the dedicated login endpoint.

Optional query parameters

Optional query parameters can be added to the dedicated login endpoint to control which application the user is redirected to after login. To redirect to a specific application page, both workspace ID `&wsid=<workspace id>` and application ID `&appid=<application id>` need to be provided. For example, if a user logs in using the SAML integration and lands on the IBM Maximo Manage application page that uses workspace `masdev`, which assumes that the user has access to the page, then the dedicated login page URL is `https://auth.<masdomain>/idplogin/loginpage?idp=saml:default-saml&wsid=masdev&appid=manage`. A user can also use the `appath=<app path>` query parameter so that a user can land on a specific application page path after login.

Self-registration for users

Starting from Maximo Application Suite 9.0, users can self-register to create their own login accounts and use the applications that they have access to. Before users can self-register, an administrator must enable and configure access options that are associated with each identity provider that is configured.

Configuration

As an administrator, when you enable self-registration in Maximo Application Suite, you can either set automatic approval of self-registered users or manually approve each user. To manually approve users, you must activate the user before the user logs in to Maximo Application Suite. You can change the initial application entitlement for users from self-service to a concurrent entitlement. You can also specify the email address of administrators who are available to support users during the registration process. This email address is provided in the users' email notifications. The email addresses of administrators who can approve requests can also be added.

To enable self-registration in Suite administration, from the side navigation menu, select **Users** and then click the **Authentication** tab. In the **User self-registration** section, select the identity provider that you want to enable self-registration for. Users can then register to create their accounts by selecting **Register** on the login page.

If you configured self-registration to automatically approve and activate users, the user can log in to Maximo Application Suite. Otherwise, you need to manually activate the user to approve their request. An email is sent to notify any administrator who has permission to approve the request. An email is also sent to the user to inform them that their request to access was submitted.

To approve a self-registered account in Suite administration, from the side navigation menu, select **Users** and click the **Users** tab. Select the self-registered user and click **More actions** > **Activate**. The user can then log in to Maximo Application Suite.

Note: Self-registered users who need approval are indicated as Pending on the Self-registered column of the user record list. To show this table column in the user record list, select **Column selection** and then **Self-registered**.

After a user is approved, an application administrator might need to set further detailed application privileges for each individual application. For example, a user who needs access to Maximo Manage must wait until the application administration for Maximo Manage approves access. For more information, see [Configuration of self-registered users in Maximo Manage](#).

Self-registration for local users

After a user registers on the login page by providing information such as username, email address, and password, the new user account is created in the local database. The users credentials, including username and password, are securely stored.

A verification email is sent to the user with a code that the users entered on the login page. After the user enters the correct code, the user can log in immediately if automatic approval is enabled. Otherwise, the user must wait for administrator approval.

Self-registration for SAML users

With self-registered accounts for Security Assertion Markup Language (SAML), when a user attempts to log in for the first time, the user accounts are automatically created in Maximo Application Suite.

When a user requests access for the first time, a SAML authentication request is sent to the identity provider. The user is authenticated in the identity provider where their credentials are validated in its database. A SAML assertion is generated that contains the user's information, such as username, attributes, and roles.

If the user does not exist in Maximo Application Suite, the SAML response from the identity provider can include attributes that trigger just-in-time (JIT) account creation. The Maximo Application Suite automatically creates a user account based on the information that is provided in the SAML assertion.

With the user now authenticated and an account that is created, the user is allowed to access the requested resource or application.

Although the user is created during this authentication process, if manual approval is enabled, users cannot log in until the administrator manually approves and activates the user.

Self-registration for LDAP users

With LDAP self-registration, users directly communicate with the LDAP server for registration. If manual approval is enabled, users cannot log in until the administrator manually approves and activates the user.

Related concepts

Simple Mail Transfer Protocol configuration

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

Enabling self-registration for users

Starting in Maximo Application Suite 9.0, you can enable self-registration so that users can create their own login accounts and use the applications that they have access to.

Before you begin

Before you enable self-registration, validate that the Simple Mail Transfer Protocol (SMTP) is configured. To provide email notifications, such as new user welcome emails and password communication, an SMTP server connection is required. Otherwise, users cannot self-register and an error message is shown on the registration page.

About this task

You can enable self-registration and access options that are associated with each identity provider that is configured.

Procedure

1. In Suite administration, select **Users** from the side navigation menu and click the **Authentication** tab.
2. In the User self-registration section, select the identity provider that you want to enable self-registration for.
3. To automatically approve self-registered users, select **Automatically approve and activate self-registered users**.

If you choose to manually approve users, you must activate the user before the user can log in.

4. Change the initial application entitlement for users from self-service to a concurrent entitlement.
5. Specify the contact information of administrators who are available for support and to approve requests.

The administrators email address for support is provided in the users' email notifications.

What to do next

If you chose to manually approve users, you must activate the user.

To approve a self-registered account in Suite administration, from the side navigation menu, select **Users** and click the **Users** tab. Select the self-registered user and click **More actions** > **Activate**. The user can then log in to Maximo Application Suite.

After a user is approved, an application administrator might need to set further detailed application privileges for each individual application. For example, a user who needs access to Maximo Manage must wait until the application administration for Maximo Manage approves access. For more information, see [Configuration of self-registered users in Maximo Manage](#).

Customer-managed

Configuring user authentication sessions

To increase application security and ensure that AppPoints are promptly returned when users close their browsers or tabs without first logging out, you can configure user authentication session behavior. You configure user authentication session behavior by changing the expiration time for the access and refresh token in the custom resource (CR) file for Maximo Application Suite.

About this task

In Maximo Application Suite 8.10, the default expiration time is changed to 30 minutes for the access token and 12 hours for refresh token.

The expiration time change works only if your applications are running from the following versions:

- Maximo Assist 8.6 and later

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

- Maximo Health 8.5 and later
- Maximo Manage 8.5 and later
- Maximo Monitor 8.9 and later
- Maximo Optimizer 8.3 and later
- Maximo Visual Inspection 8.7 and later
- IoT 8.6 and later

Note: You cannot change the expiration time in earlier versions of these applications.

In Maximo Application Suite 8.9, the default expiration time was 12 hours for the access token and 1 week for refresh token. If you want to revert to these expiration time values, complete the following steps.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the `spec.settings` section, change the length of time for `refreshTokenTimeout` and `accessTokenTimeout`.
For example, to revert to the default behavior that was set in version 8.9, update `accessTokenTimeout` from 30 minutes to 12 hours and update `refreshTokenTimeout` from 12 hours to 168 hours. By changing these values, when a user closes a browser without logging out, the AppPoints are returned within 12 hours. The access token is valid and can be refreshed for up to 7 days, which is 168 hours, before users must log in to their session.

```
spec:
  settings:
    sso:
      refreshTokenTimeout: 168h
      accessTokenTimeout: 12h
```

Note: To avoid unexpected behavior, the minimum that you can set the expiration time is 15 minutes for the access token, and 12 hours for refresh token.

5. Save the CR changes.

Customer-managed **Configuring single sign-on properties**

Starting in Maximo Application Suite 9.0, you can change the default single sign-on (SSO) token name LTPAtoken2 to avoid conflicts with the same cookie name that is generated by other software.

About this task

If you are using earlier versions of Maximo Application Suite, you can change the default single sign-on (SSO) token name from Maximo Application Suite 8.10.15 and 8.11.12.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the `spec.settings.sso` section, customize the SSO token properties.
 - In Maximo Application Suite 8.10.15 and 8.11.12 and later versions, you can customize the following properties.

Properties	Default values
<code>disableLtpaCookie</code>	True
<code>allowDefaultSsoCookieName</code>	True
<code>useOnlyCustomCookieName</code>	False
<code>allowCustomCacheKey</code>	True

Table 59. Single sign-on properties for 8.10.15 and 8.11.12 (continued)	
Properties	Default values
ssoCookieName	LtpaToken2

- In Maximo Application Suite 9.0 and later versions, you can customize the following properties in the Suite CR. The default behavior uses the cookie name **ltpatoken2_[instance name]**.

Table 60. Single sign-on properties for 9.0	
Properties	Default
disableLtpaCookie	False
allowDefaultSsoCookieName	False
useOnlyCustomCookieName	True
allowCustomCacheKey	False
ssoCookieName	ltpatoken2_instanceid

5. Save the CR changes.

Configuring multiple identity providers for same authentication type

Starting in IBM Maximo Application Suite 9.0, you can configure multiple identity providers for the same authentication type, such as OIDC, SAML or LDAP, for user authentication.

Before you begin

Generate an API key and ensure that the API key that you generate has system admin permission. For more information about generating an API key, see [Maximo Application Suite Admin APIs](#).

Procedure

- Configure the SAML authentication type
 - a) Run the following command:

```
curl --location \
  --request PUT https://api.<masDomain>/config/saml/<idpId> \
  --header "x-access-token: <access token>" \
  --header "Content-Type: application/json" \
  --insecure \
  --data '{
    "spInitiatedLogout": true,
    "displayName": "SAML",
    "issuer": "saml",
    "serviceName": "<service provider>",
    "nameIDFormat": "<name ID format>"
  }'
```

Where

<masDomain>

The domain name for the Maximo Application Suite instance.

<idpID>

The unique name for the identity provider.

<access token>

The authentication token that was obtained from the API key generation.

<service provider>

The unique name of the service provider.

<name ID format>

The format that is used for the name ID. Use one of the following values:

- customize
- email
- encrypted
- entity
- kerberos
- persistent
- unspecified
- windowsDomainQualifiedName
- x509SubjectName

b) Download the SAML configuration .xml file.

c) Access the following URL:

`https://auth.<masDomain>/ibm/saml20/<serviceProviderName>/samlmetadata`

Tip: Wait for few minutes for the URL to become available.

d) Upload the SAML configuration .xml file to your identity provider.

e) Download the SAML metadata from your identity provider.

f) Run the following command to upload the SAML metadata from your identity provider to Maximo Application Suite

```
curl --location \
  --request PUT https://api.<masDomain>/config/saml/<idpId>/metadata \
  --header "x-access-token: <access token>" \
  --insecure \
  --form 'file=@"<file path>/<metadata filename>.xml"'
```

Where

<masDomain>

The domain name for the Maximo Application Suite instance.

<idpID>

The unique name for the identity provider.

<access token>

The authentication token that was obtained from the API key generation.

<file path>

The location to which the metadata file is downloaded.

<metadata filename>

The name of the metadata file that is downloaded.

The command generates the required resources.

Tip: Wait for the operator to reconcile and the user interface to refresh and reflect the updates.

- Configure the LDAP authentication type

a) Run the following command:

```
curl --location \
  --request PUT https://api.<masDomain>/config/ldap/<idpId> \
  --header "x-access-token: <access token>" \
  --header "Content-Type: application/json" \
  --insecure \
  --data '{
    "displayName": "LDAP",
    "url": "ldaps://<ldap url>",
    "baseDN": "dc=example,dc=com",
    "bindDN": "uid=<user id>,dc=example,dc=com",
```

```

"bindPassword": "<password>",
"userIdMap": "cn",
"certificates": [
  {
    "alias": "intermediate",
    "cert": "-----BEGIN CERTIFICATE-----<cert value><snip>-----END CERTIFICATE-----"
  }
]
}

```

Where

<masDomain>

The domain name for the Maximo Application Suite instance.

<idpID>

The unique name for the identity provider.

<access token>

The authentication token that was obtained from the API key generation.

<user id>

The name of the user for authentication.

<password>

The password that is set for the user.

<cert value>

The values or content available in the certificate.

The command generates the required resources.

Tip: Wait for the operator to reconcile, and the user interface to refresh and reflect the updates.

User synchronization

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups with an LDAP server. By using user registry synchronization, you can automate your user management by synchronizing users and groups between an LDAP server and your local user registry.

LDAP user registry synchronization

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

User synchronization

Synchronization is one way, from the LDAP server to Maximo Application Suite, and all updates on the LDAP side are merged over at synchronization time.

If your environment is configured for [LDAP authentication](#), you can sync the complete user and group registry, or you can sync a filtered subset of the LDAP user registry.

If your environment is configured for local or SAML authentication, you can specify an external LDAP server to sync the users and groups with.

During synchronization, the mapped LDAP users are automatically added as Maximo Application Suite users. When the user is synchronized from LDAP, the field name owner is set to scim in MongoDB, which means it is an externally managed user.

The synced users are initially added with only application entitlement but can be assigned administration entitlement if needed. Users can also be granted specific access to applications during the initial synchronization.

Generally, synced personal information cannot be updated in Maximo Application Suite. Only user entitlement and application access can be managed in Maximo Application Suite. If the user authentication is local, passwords can also be managed in Maximo Application Suite.

Important: Because the synchronization runs on a schedule, discrepancies might be temporarily introduced between synchronizations. For example, if a previously synced user ID is removed from LDAP, that user ID can still be used to log in to Maximo Application Suite until the user removal is synced. If the synced LDAP users are using local authentication, the user ID still has access. If LDAP or SAML authentication is used, the login fails because the user is no longer active in LDAP.

User and group synchronization with SCIM 2.0

Starting in Maximo Application Suite 9.0, you can synchronize users and groups from an external identity provider (IdP) by using the System for Cross-domain Identity Management (SCIM) 2.0 protocol. For more information, see [“User synchronization with SCIM 2.0”](#) on page 625

User and group synchronization operations

The following user and group synchronization operations are supported.

User operations

<i>Table 61. User synchronization operations</i>	
Operation	Description
Insert	Adds a user if it does not exist in the Maximo Application Suite user registry. If the user was previously deactivated as a result of being removed from the LDAP server, but is now added back into the LDAP server then they are reactivated. The user is initially set up with an identity provider (IDP) issuer, entitlement, and application access. If SMTP is configured , newly created users also receive a welcome email.
Update	Updates a user if it exists in the Maximo Application Suite user registry. User entitlement and application access are not updated.
Skip	Skips user update if no changes in the LDAP server occurred since the last synchronization. The verified field is <code>ldap.meta.lastModified</code> .
Delete	Deactivates the user in Maximo Application Suite if the user is removed from the LDAP server.

Group operations

<i>Table 62. Group synchronization operations</i>	
Operation	Description
Insert	Adds a group if it does not exist in the Maximo Application Suite user registry.
Update	Updates a user if it exists in the Maximo Application Suite user registry. User entitlement and application access are not updated.

Groups are always updated, and Maximo Application Suite does not delete them as part of the synchronization process.

LDAP configuration attributes

Maximo Application Suite LDAP filter configuration is based on IBM Liberty.

The following configuration examples are based on Microsoft Active Directory. Refer to the IBM Liberty documentation for other types of user registries.

Parameter	Details	Example
URL	The URL for the LDAP server is in the format: protocol://<hostname>:<port> Note: Secure LDAP (LDAPS) is the only allowed protocol. Non-TLS connections are not allowed.	Example: ldaps://MSAD2021.fyre.ibm.com:636
Base DN	The path in the object hierarchy of the directory server.	Example: OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com
Bind DN	Bind DN is used to bind to an LDAP server. Administrators must have sufficient privileges to search for users under user search DN or groups under group search DN.	Example: CN=wilson,OU=users,OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com
Bind PW	LDAP admin password	
Certificate	The TLS certificate for your LDAP Server. You can also add multiple certificates if you are using a certificate chain.	
User Base DN	The user-level path in the object hierarchy of the directory server. If not provided, the baseDN is used by default.	Example: OU=users,OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com
userFilter	The query that is used to search the users in the directory.	Example 1: (&(sAMAccountName=%v)(objectcategory=user)) Example 2: (&(sAMAccountName=%v)(objectclass=user))
userIdMap	The field that is used for user IDs	Example: user:<sAMAccountName>
Group Base DN	The group-level path in the object hierarchy of the directory server. If not provided, the baseDN is used by default.	Example: OU=groups,OU=FYRE,DC=MSAD2021,DC=fyre,DC=ibm,DC=com
groupFilter	The query that is used to search the groups in the directory.	Example 1: (&(cn=%v)(objectcategory=group)) Example 2: (&(cn=%v)(objectclass=group))

Table 63. LDAP configuration attributes (continued)

Parameter	Details	Example
groupIdMap	The field that is used for group ID.	Example: *:cn
groupMemberIdMap	An LDAP filter that identifies the group memberships for users.	Example: memberOf:member

User and group registry mapping

The Maximo Application Suite user data model is based on the custom mapping of the user data that synchronizes with the LDAP server.

Starting in Maximo Application Suite 9.0, you can configure user registry synchronization to map data for users from LDAP with Maximo Application Suite in the user interface.

To map user or group data, on the **Suite administration** page, select **Configurations** from the side navigation menu and then click **User registry synchronization**. You can also use a default value that is set by the system. If you don't specify custom field values, then default values are used.

User mapping

You can use the following user properties to map between Maximo Application Suite and LDAP by specifying the field in LDAP that maps to the property field in Maximo Application Suite.

Standard user properties

- **id**
- **username**
- **displayName**
- **title**
- **familyName**
- **givenName**
- **email**
- **phoneNumber**

Extensions

- **employeeNumber**
- **costCenter**
- **organization**
- **division**
- **department**
- **manager**

Extensions support custom property names that you can map to LDAP attributes.

Starting in Maximo Application Suite 9.0.2, to ensure that no standard values of LDAP users are used when you map user property values, set the **forceMappedValue** parameter value to **True** in **spec.config** property of the ScimCFG Red Hat OpenShift custom resource.

```

---
spec:
  config:
    forceMappedValue: true
    ...
    userSync:
    ...
    mappings:
      standard:

```

```
displayName: myDisplayName
phoneNumber: myPhoneNumber
...
```

- By setting the **forceMappedValue** value to True, user synchronization can be changed to use the custom mapping when pulling data from LDAP users, and if the property name is not found in the LDAP user, then the NOT_PROVIDED string is used as value for the corresponding Maximo Application Suite property.
- For the **email** and **phoneNumber** mappings, if the custom mapped property names are not found in the LDAP user, then empty arrays are set for Maximo Application Suite user's **emails** and **phoneNumbers** properties instead of reverting to the default **emails** and **phoneNumber** properties of LDAP users.

Group mapping

You can use the following group properties to map between Maximo Application Suite and LDAP by specifying the field in LDAP that maps to the property field in Maximo Application Suite.

Standard user properties

- **id**
- **displayName**

Extensions properties support custom property names that you can map to LDAP attributes.

If you are mapping the following properties, review the following considerations:

Note:

When you use custom mapping, field values must be unique in LDAP. Allowed characters for **id** and **username** are letters and numbers, -, @, ., _ without spacing.

id

After the initial synchronization, if you configure the property mapping for **id**, you must manually delete the existing SCIM owned users who have the SCIM value in the **owner** property from the Maximo Application Suite database.

You can delete the existing SCIM owned users by running the following command in MongoDB

```
db.getCollection("User").deleteMany({"owner": "scim"})
```

Deleting existing SCIM owned users is required because users who are included in the next scheduled synchronization cron job are considered new users with different IDs. Otherwise, the synchronization of those users will fail since existing users in the database have the same username and email.

username

Configuring the property mapping for **username** might affect the user's ability to authenticate.

email and phoneNumber

When **email** or **phoneNumber** properties are not mapped, one or more LDAP emails and one or more LDAP phone numbers are copied into Maximo Application Suite user records. This synchronization is based on the definition of phone numbers and emails that is available for each user in the LDAP server.

If you specify the mapping for the **email** or the **phoneNumber** property, then only one email and only one phone number that is specified by the mapped value is copied from the LDAP user into the Maximo Application Suite.

If users have multiple emails or phone numbers that might be duplicated, then the synchronization might fail. To make sure that the data synchronizes correctly, use **email** and **phoneNumber** information that is unique to each user.

For information and examples about mapping user and group data, see [“Mapping LDAP users from Microsoft Active Directory”](#) on page 632 and [“Mapping groups from LDAP to display group descriptions”](#) on page 633.

Extensions

In Maximo Application Suite 8.11 and earlier versions, the extension fields that are listed in the following table are not included in the Maximo Application Suite user interface. They are part of the Maximo Application Suite user object and might be used by applications.

The following table shows how the LDAP fields map to the Maximo Application Suite user object.

LDAP field	Details
employeeNumber	A string identifier, typically numeric or alphanumeric, that is assigned to a person, typically based on order of hire or association with an organization.
costCenter	Identifies the name of a cost center.
organization	Identifies the name of an organization.
division	Identifies the name of a division.
department	Identifies the name of a department.
manager	The user's manager. A complex type that optionally allows service providers to represent an organizational hierarchy by referencing the "ID" attribute of another user.

Customizations in ScimCfg Custom Resource

Starting in Maximo Application Suite 8.9, you can change some configurations in the ScimCfg Custom Resource as the configurations are not available in the Maximo Application Suite user interface.

Add the following properties in `spec.config` of the ScimCfg Custom Resource:

customMaxSearchResults

Use the `customMaxSearchResults` to configure the maximum number of entries that can be returned in a search.

ldapType

Use the `ldapType` property to override the default Custom setting for `ldapType` in the server's `<ldapRegistry>` configuration. The following LDAP servers are supported::

- Custom, which is the default value.
- IBM Lotus Domino
- IBM SecureWay Directory Server
- IBM Tivoli® Directory Server
- Microsoft Active Directory
- Netscape Directory Server
- Novell eDirectory
- Sun Java System Directory Server

customLdapRegistryExtensions

Use the `customLdapRegistryExtensions` property to override the default settings for `<ldapRegistry>` configuration.

For more information about supported settings, see [LDAP User Registry](#).

The exceptions are `ldapType` because this property is configurable in a separate stand-alone property, the properties that are exposed through the Scim Sync CRUD APIs, User Interface

properties such as `host`, `port`, `bindDN`, `bindPassword`, `baseDN`, and some properties that cannot be changed such as `id` and `realm`.

The remaining properties can be added under `customLdapRegistryExtensions`. For more information about adding the properties, see the format that is specified by [Websphere Liberty](#).

Note: The properties are not exposed through LDAP sync or User Interface. However, updating LDAP sync configuration through the API or UI does not override these settings.

Limitations

- In Maximo Application Suite 9.0 and earlier, user synchronization is supported for a single LDAP server.
- User and group sync are done in the same job.
- Synchronization by using external SCIM APIs are not supported.

Related concepts

[Identity providers](#)

Identity provider prerequisites for Maximo Application Suite.

Related tasks

[Configuring LDAP authentication](#)

By configuring LDAP authentication with Maximo Application Suite, user authentication is managed by your LDAP server.

[Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 8.11](#)

[Mapping LDAP users from Microsoft Active Directory](#)

[Mapping groups from LDAP to display group descriptions](#)

Starting in Maximo Application Suite 9.0, you can configure the properties from LDAP to map to Maximo Application Suite by configuring group mapping in user registry synchronization. When synchronizing LDAP groups from an LDAP server to the Maximo Application Suite database, group description is not part of the group data that is copied from LDAP.

[Mapping LDAP fields as person ID for Maximo Manage in Maximo Application Suite 9.0](#)

Related information

[LDAP user mapping from Microsoft Active Directory](#)

User synchronization with SCIM 2.0

Starting in Maximo Application Suite 9.0, you can synchronize users and groups from an external identity provider by using the System for Cross-domain Identity Management (SCIM) 2.0 protocol.

SCIM 2.0 support simplifies user management by enabling external identity providers, such as Okta or Microsoft Entra, to synchronize users and groups through a push mechanism. The SCIM specification defines a common schema for user and group resources with a protocol that defines how operations are performed on those resources. SCIM facilitates the exchange of these resources across different services.

By using SCIM capabilities in Maximo Application Suite, existing users and groups that are defined in an identity provider are automatically synchronized with Maximo Application Suite when those resources are assigned access. Therefore, you do not need to manually add users or groups into the Suite. Also, updates or removal of access is automatically synchronized when the appropriate configuration is applied in the identity provider.

SCIM APIs

Maximo Application Suite can use SCIM 2.0 to integrate with identity providers by using a set of API endpoints from Maximo Application Suite. These APIs implement the SCIM 2.0 protocol and are initiated by the identity provider, which acts as a client, to manage user and group lifecycle operations.

Note: These APIs are to be used only with integration between Maximo Application Suite and an identity provider for user and group synchronization. This information is provided for reference usage by experienced administrators.

The Maximo Application Suite APIs are accessed by using the base URL: `https://api.{mas-instance-id}.{domain}/scim/v2/{profileId}`.

You can use the following endpoints to apply create, retrieve, update, and delete operations on user resources,

Endpoint	Method	Description
<code>/scim/v2/{profileId}/Users</code>	POST	Pushes a new user into Maximo Application Suite.
<code>/scim/v2/{profileId}/Users</code>	GET	Retrieves users with filtering and paging.
<code>/scim/v2/{profileId}/Users/{userId}</code>	GET	Retrieves a specific user.
<code>/scim/v2/{profileId}/Users/{userId}</code>	PUT	Updates a specific user.
<code>/scim/v2/{profileId}/Users/{userId}</code>	PATCH	Updates any changed fields for a specific user.
<code>/scim/v2/{profileId}/Users/{userId}</code>	DELETE	Deletes a specific user.

You can use the following endpoints to apply create, retrieve, update, and delete operations on group resources.

Endpoint	Method	Description
<code>/scim/v2/{profileId}/Groups</code>	POST	Pushes a new group into Maximo Application Suite.
<code>/scim/v2/{profileId}/Groups</code>	GET	Retrieves groups with filtering and paging.
<code>/scim/v2/{profileId}/Groups/{groupId}</code>	GET	Retrieves a specific group.
<code>/scim/v2/{profileId}/Groups/{groupId}</code>	PUT	Updates a specific group
<code>/scim/v2/{profileId}/Groups/{groupId}</code>	PATCH	Updates any changed fields for a specific group
<code>/scim/v2/{profileId}/Groups/{groupId}</code>	DELETE	Deletes a specific group

Maximo Application Suite SCIM profiles

The `{profileId}` in the base URL is a user-selected ID for a Maximo Application Suite resource that is called the Maximo Application Suite SCIM profile.

A SCIM profile is a collection of Maximo Application Suite configurations that acts as a template to configure users and groups as they are added into Maximo Application Suite. You use a SCIM profile because the SCIM schema defines the general users and groups attributes and does not include specific Maximo Application Suite elements that are necessary for a functional user, such as the ability to log in, entitlements, and application permissions. This Maximo Application Suite SCIM profile is applied to the following configurations:

- Workspace access

- Application access
- Entitlement
- Identity configuration

The following fields are also included in the configuration:

id

The *{profileId}* that is used in the base URL and is how Maximo Application Suite identifies which profile to apply to incoming API requests.

version

The version of the profile model resource. The only valid value in Maximo Application Suite 9.0 is 1.

By including the *{profileId}* in the base URL that is used to configure an external identity provider SCIM integration, the provision requests that the identity provider are linked to a Maximo Application Suite SCIM profile. The configuration can be applied to the new users or groups at the point of synchronization. Synchronized users can then log in and access applications without the need for further manual administrative action.

The Maximo Application Suite SCIM profile is used to configure the user or group once, at the point that the user or group is created. The following rules are applied when you use the Maximo Application Suite SCIM profile:

- If the profile does not exist when a user or group is added, the configuration is not applied. Administrators must manually configure the users.
- If the profile is updated after a user or group is imported, those changes apply only to new users or groups that are created after the updates. Any changes are not retroactively applied to any existing users or groups.
- If an administrator updates a user or group in Maximo Application Suite, those changes are not modified or undone by the application of a Maximo Application Suite SCIM profile.

In Maximo Application Suite 9.0, you can configure the SCIM profile resource only by using APIs.

Endpoint	Method	Description
/scim/v2/Profiles	POST	Creates a Maximo Application Suite SCIM profile.
/scim/v2/Profiles	GET	Retrieves Maximo Application Suite SCIM profiles with paging.
/scim/v2/Profiles{profileId}	GET	Retrieves a specific Maximo Application Suite SCIM profile.
/scim/v2/Profiles/{profileId}	PUT	Updates a specific Maximo Application Suite SCIM profile.
/scim/v2/Profiles/{profileId}	DELETE	Deletes a specific Maximo Application Suite SCIM profile.

A Maximo Application Suite SCIM profile has the following structure:

```

{
  "id": "myprofilename",
  "version": 1,
  "identities": [
    {
      "type": "local"
    },
    {
      "id": "default-saml",
      "type": "saml",
      "samlId": "userName"
    }
  ]
}

```

```

    {
      "id": "default-ldap",
      "type": "ldap",
      "ldapUsername": "userName"
    }
  ],
  "entitlement": {
    "application": "BASE"
  },
  "workspaces": [
    {
      "id": "workspace1",
      "applications": [
        "manage"
      ]
    }
  ]
}

```

You can provide the following main elements to the configuration:

- [Identities](#)
- [Entitlement](#)
- [Workspaces](#)

Identities

Identities specify the authentication mechanisms that are configured when a new user is added to Maximo Application Suite. This mechanism is an array where multiple configurations can be added when necessary to enable multiple authentication options for users.

Starting in Maximo Application Suite 9.0, the following authentication type options are available for user identities:

local

User authentication is managed by a local Maximo Application Suite username and password.

When you specify `local` as the type of the identity configuration, additional configuration is not required. The SCIM `userName` field of the imported user is used as the Maximo Application Suite username and a random password is generated and emailed to the user.

Saml

User authentication is managed by a SAML identity provider.

When you specify `saml` as the type of identity configuration, the following fields are also required.

id

The Maximo Application Suite ID of the identity provider.

samlId

The name of a SCIM user attribute that is extracted from the imported user that represents the SAML ID.

When you specify the `samlId` field, you are not specifying a specific value for any particular user. You are indicating which field of the user record is used as the SAML ID, for example, whatever the value of the `userName` field is.

The value of this field supports a limited form of expressions to target the appropriate field. The expressions are based on the SCIM patch path expressions. For more information, see [Modifying with Patch in the System for Cross-domain Identity Management \(SCIM\) specification](#).

The PATH attribute follows the ABNF syntax rule `PATH = attrPath / valuePath [substrata]`. Where `attrPath` and `valuePath` are defined by the ABNF filtering rules.

For more information, see [Filtering for system for Cross-domain Identity Management 2.0 \(SCIM\) protocol](#).

Use a limited expression when you need to specify that the primary email field is to be used as the `samlId` but is not the same as the `userName`. When you use a limited expression, the `samlId` can be specified as `emails[primary eq true].value`

Ldap

User authentication that is managed by an LDAP identity provider.

When you specify `ldap` as the type of configuration, the following fields are also required.

id

The Maximo Application Suite ID of the identity provider.

ldapUsername

The name of a SCIM user attribute that is extracted from the imported user to represent the LDAP username.

When you specify the `ldapUsername` field, you are indicating which field of the user record is used as the LDAP username, for example, whatever the value of the `userName` field is.

The value of this field supports a limited form of expressions to target the appropriate field. The expressions are based on the SCIM patch path.

Entitlement

Entitlement specifies the level of entitlement that is granted to the user. Starting in Maximo Application Suite 9.0, you can automate the assignment of application- level entitlement by using only this configuration.

The following entitlements can be assigned:

- NONE
- SELF_SERVICE
- LIMITED
- BASE
- PREMIUM

For more information, see [“User entitlement and access in Maximo Application Suite in 9.0 and earlier” on page 797.](#)

Workspaces

You configure workspaces to specify access to Maximo Application Suite workspaces and applications. This configuration is applied to both users and groups when they are pushed into Maximo Application Suite.

You specify the following fields in the workspace configuration:

id

The identity of the Maximo Application Suite workspace that the user or group is added to.

applications

A list of applications, by ID, that the user or group is added to.

The applications are listed by Maximo Application Suite application IDs that a user is granted a USER role to or that a group is synchronized with.

You can query the application IDs by using a GET request with the following API: `https://api.{mas-instance-id}.{mas-domain}/applications`. For more information, see [API reference get information about all applications](#).

Related information

[SCIM 2.0 overview](#)

[SCIM 2.0 schema](#)

Configuring Maximo Application Suite to synchronize user and groups with SCIM 2.0

Starting in Maximo Application Suite 9.0, you set up the user and group synchronization in the identity provider by using the SCIM 2.0 API endpoints from Maximo Application Suite. You initially create an API Key in Maximo Application Suite to generate a JSON Web Token.

About this task

To invoke the new SCIM 2.0 APIs with the identity provider, an authentication token, which is called a JSON Web Token, is required. This token is obtained by using a Maximo Application Suite API key. The API key requires `user admin` permissions.

By default the JSON Web Token has a short expiry time and requires frequent regeneration and reconfiguration of the connection details in the identity provider. You can specify an expiry duration for the token as part of the authenticated request by specifying an HTTP Header as part of the request, with the name `mas-jwt-expiry-duration` and a value that is an ISO8601 duration. For example, you can specify `P90D` for 90-day expiry.

Procedure

1. In Maximo Application Suite administration, create an API key.
 - a) From the side navigation menu, click **API keys** and click **Create API key**.
 - b) Enter the description and specify the authentication token expiry.
 - c) For suite administrative access that is applicable to the API key, select **User management**.
 - d) Click **Submit**.
 - e) Copy the API key and authentication token details.

If authentication token details are lost, you cannot recover the details. To create a token, you must create an API key.
2. To generate a JSON Web Token, issue a GET request to the `/v1/authenticate` API with Basic Auth.

For more information about generating a GET request, see [Get all API keys](#) and [Get specific API key](#).

 - a) Specify `Basic Auth` as the authentication type.
 - b) Enter the API key ID and authentication token as username and password.
 - c) Specify a custom expiry duration for the token as part of the authenticated request.

For example, enter `P90D` for a 90-day duration. The JSON Web Token is generated, which you can use in the API calls that you want to make. The response to the authenticated request contains a token field.
 - d) Copy the JSON Web Token details.

For more information, see [Obtain a JWT token using API key credentials](#).
3. Create a Maximo Application Suite SCIM profile to specify the Maximo Application Suite configuration that is applied to users and groups when they are synchronized from the identity provider to Maximo Application Suite.

For more information, see [Create a new MAS SCIM profile](#).
4. Configure the identity provider.
 - a) In the identity provider, create an application to represent Maximo Application Suite.
 - b) Enable the SCIM 2.0 provisioning in the application.
 - i) Specify the base URL for integration by using `https://api.{mas-instance-id}.{domain}/scim/v2/{profileId}`.

- ii) Provide the JSON Web Token for the header-based authentication that you created from the API key.
 - iii) Validate that the identity provider can connect to Maximo Application Suite and issue SCIM requests.
5. Assign users and groups to the application in the identity provider to initiate the synchronization of users and groups with Maximo Application Suite.

Related concepts

User synchronization with SCIM 2.0

Starting in Maximo Application Suite 9.0, you can synchronize users and groups from an external identity provider by using the System for Cross-domain Identity Management (SCIM) 2.0 protocol.

Configuring multiple LDAP user registry synchronizations

Starting in IBM Maximo Application Suite 9.1, you can configure multiple LDAP user registry synchronizations to synchronize users and groups from more than one LDAP server to your local Maximo Application Suite user registry. If you already configured LDAP user registry synchronization for one LDAP server and you need to synchronize users from a different server, you can set up LDAP user registry synchronization for that LDAP server.

Before you begin

Generate an API key and ensure that the API key that you generate has system admin permission. For more information about generating an API key, see [Maximo Application Suite Admin APIs](#).

About this task

In Maximo Application Suite 9.0 and earlier, the LDAP user registry synchronization supported only one configuration by using a default configId. You can now set up LDAP user registry synchronization with multiple LDAP servers by using the PUT `https://api.<masDomain>/config/scim/<configId>` API and a custom configId. Even if an LDAP user synchronization is already in place for a default configId, you can run other registry synchronizations in parallel for different LDAP servers.

Procedure

To configure the LDAP user registry synchronization for a **configId** that is not default, run the following command.

```
curl --location \
  --request PUT https://api.<masDomain>/config/scim/<configId> \
  --header "x-access-token: <access token>" \
  --insecure \
  --data '{
    "entitlement": {
      "application": "<application entitlement>"
    },
    "groupSync": {
      "groupBaseDN": "<Group Base DN>",
      "groupFilter": "<groupFilter>",
      "groupIdMap": "<groupIdMap>",
      "groupMemberIdMap": "<groupMemberIdMap>",
      "mappings": <Group mappings>,
      "syncGroups": true
    },
    "issuer": "<identity provider ID>",
    "ldapAuth": {
      "baseDN": "<Base DN>",
      "bindDN": "<Bind DN>",
      "bindPassword": "<Bind PW>",
      "certificates": [
        {
          "alias": "intermediate",
          "crt": "-----BEGIN CERTIFICATE-----<cert value><snip>-----END CERTIFICATE-----"
        }
      ]
    }
  },
  "url": "<URL>"
```

```

    },
    "podTemplates": [],
    "syncFrequency": "<synchronization frequency>",
    "userSync": {
      "mappings": <User mappings>,
      "userBaseDN": "<User Base DN>",
      "userFilter": "<userFilter>",
      "userIdMap": "<userIdMap>"
    },
    "workspaces": [
      {
        "applications": <applications>,
        "id": "<workspace id>"
      }
    ]
  },
]
}

```

Where

<configId>

A string with no spaces or no special characters is the identifier of this LDAP user registry synchronization.

<application entitlement>

The application entitlement that all synchronized users have in the Maximo Application Suite user registry. Administrator entitlement is always set to NONE for synchronized users.

<access token>

The authentication token that was obtained from the API key generation.

<identity provider ID>

The identity provider ID, such as, default-ldap, default-saml, or local.

<synchronization frequency>

The frequency that synchronization jobs run in cron syntax.

<workspace id>

The workspace ID.

<application>

Applications for which the synchronized users are assigned the USER role. For the Manage application, the MANAGEUSER role is assigned.

For all other properties in userSync, groupSync, and ldapAuth objects, refer to LDAP User Registry Synchronization for a detailed explanation of each property. For more information, see [“LDAP user registry synchronization”](#) on page 619.

Results

The command generates the resources that are necessary to synchronize the LDAP user and group for the provided <configId>. The **owner** property of all users and groups that are synchronized by this configuration is set to scim-<configId>. The **owner** of users and groups that are synchronized by default configId is set to scim.

Mapping LDAP users from Microsoft Active Directory

Starting in Maximo Application Suite 9.0, you can set up user mapping to map users from LDAP with Maximo Application Suite by using user registry synchronization in the user interface.

When synchronizing LDAP users from Microsoft Active Directory into the Maximo Application Suite database, some LDAP user properties might not match to the corresponding Maximo Application Suite user properties.

For example, an LDAP user who is called *John Doe* might have the **givenName** property set to John Doe and the **displayName** property set to John Doe Doe.

About this task

With user mapping, you can map the user properties between Maximo Application Suite and LDAP by specifying the **LDAP field** to map to the **Maximo Application Suite field**. Alternatively, you can use a default value that is set by the system. For more information, see [User mapping](#)

Procedure

1. On the **Suite administration** page, select **Configurations** from the side navigation menu and then click **User registry synchronization**.
2. In the User mapping section, map the user data to synchronize with the LDAP server.
For example, to map the LDAP user John Doe with the **givenName** set to John and the **displayName** set to John Doe, configure the following data:
 - For the **givenName** property in the column Maximo Application Suite field, enter **givenName** as the property in the column LDAP field.
 - For the **displayName** property in the column Maximo Application Suite field, enter **displayName** as the property in the column LDAP field.
3. Optional: Select **Use default mapping** to use values that are set by the system for LDAP fields.
If you don't specify custom field values, then default values are used.
4. In the LDAP domain attributes section, enter **Bind DN** and **Bind Password**.
Every time User Registry Synchronization configuration changes, you must update these security fields with the Bind DN and Bind password.
5. Save your changes.

Results

After you save the user mapping updates, the configuration is processed and in the next scheduled synchronization cron job, the user synchronization changes are applied.

When Maximo Application Suite connects to LDAP systems, the SCIM specification is processed internally. This specification uses a set of standard properties from the LDAP registry to form the **givenName** and **displayName** for users who are created in Maximo Application Suite. By using user mapping, an administrator can synchronize the **givenName** and **displayName** to the same attributes within the LDAP directory, which prevents the need for complex naming formats.

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Mapping groups from LDAP to display group descriptions

Starting in Maximo Application Suite 9.0, you can configure the properties from LDAP to map to Maximo Application Suite by configuring group mapping in user registry synchronization. When synchronizing LDAP groups from an LDAP server to the Maximo Application Suite database, group description is not part of the group data that is copied from LDAP.

About this task

Some applications, such as Maximo Manage, are unable to display group descriptions since the description is not available in the Maximo Application Suite group records after synchronizing groups from LDAP.

With group mapping, you can map the group properties between Maximo Application Suite and LDAP by specifying the **LDAP field** to map to the **Maximo Application Suite field**. Alternatively, you can use a default value that is set by the system. For more information, see [Group mapping](#)

Procedure

1. On the **Suite administration** page, select **Configurations** from the side navigation menu and then click **User registry synchronization**.
2. In the Group mapping section, specify the custom mapping for the group data to synchronize the LDAP server with Maximo Application Suite so that the groups contain a description field.
 - a) Select **Add custom mapping**.
 - b) For the Maximo Application Suite field, enter `extension.description`.
 - c) In the column for LDAP field, enter the attribute that is the attribute for LDAP server groups. For example, enter `description` if `description` is the valid attribute for the LDAP server groups.
3. Optional: Select **Use default mapping** to use values that are set by the system for LDAP fields.
If you don't specify custom field values, then default values are used.
4. In the LDAP domain attributes section, enter **Bind DN** and **Bind Password**.
Every time User Registry Synchronization configuration changes, you must update these security fields with the Bind DN and Bind password.
5. Save your changes.

Results

When you save the mapping changes, the configuration is processed. In the next scheduled synchronization cron job, the user synchronization changes are applied.

For example, if you added **description** as the attribute for the LDAP server groups to **extensions.description**, then it is mapped in the groups and is saved in the Maximo Application Suite database.

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Setting up email notifications

You can enable email notifications for Maximo Application Suite. Starting in Maximo Application Suite 9.0.3, 8.11.15 or 8.10.18, you can configure how emails are used by creating custom templates, changing the language of emails, and disabling emails.

Customer-managed

Simple Mail Transfer Protocol configuration

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

To configure SMTP, in Suite administration, select **Configurations** from the side navigation and then select **SMTP**.

As part of SMTP configuration you define two suite email addresses:

- [“Sender email address” on page 635](#)

The sender email address for all Maximo Application Suite generated emails, such as new user information and password reset confirmation.

- [“Suite system email address” on page 635](#)

The suite system email address receives Maximo Application Suite system announcements, including password emails for users that do not have an associated email address.

- Welcome email

- Self-registration (external IDP)
- Administrator password reset
- User linked to local authentication
- Password expiration
- Account lockout

SMTP setup

The following configuration information is required when setting up SMTP for Maximo Application Suite:

Connection settings

- SMTP host and port (Default: 587)

The connection information for your SMTP server.

- Security protocol for the server

Maximo Application Suite supports: None, STARTTLS, and SSL/TLS

Note: If you choose STARTTLS for IBM Maximo Real Estate and Facilities, ensure that your SMTP server supports TLS 1.3.

- Authentication credentials for the SMTP server, if required by your SMTP server.

Sender email address

The sender email address and optional name is used for all emails that are sent by the suite.

Tip: If the email address is not monitored, consider including a do not reply statement in the optional name field.

Example

IBM Maximo Application Suite (unmonitored)

Suite system email address

The suite system email address is a catchall address for Maximo Application Suite system notifications.

The Suite system email address must be a monitored email address.

Custom certificate

If you are using a secure socket layer (SSL) communication protocol and the service is secured with a certificate that is not provided by a trusted certificate authority, you can optionally retrieve or add a custom certificate that is valid for your organization.

Email notifications

Emails are sent by Maximo Application Suite when the following events occur:

- A user is created

Two emails are sent, a welcome email with the user information, and one with the initial password.

- A user requests a password reset

Users can request password resets at the Maximo Application Suite login prompt. Two emails are sent, a password reset confirmation that contains a reset code, and a password reset complete.

- An administrator resets a user password.

A password reset email is sent.

- Starting in Maximo Application Suite 9.0, a user self-registers to create their own account

A verification email is sent to the user with a code that the users entered on the login page. After the user enters the correct code, the user can log in immediately if automatic approval is enabled. Otherwise, the user must wait for administrator approval.

If self-registration approval is required, an email is sent to notify any administrator who has permission to approve the request. An email is also sent to the user to inform them that their request to access was submitted.

Email preferences

By default, password emails are sent. If you prefer to not send passwords by email, you can deselect the option when you create a user. Users must then contact their Maximo Application Suite administrator for their password.

Starting in Maximo Application Suite 9.0.3, 8.11.15 or 8.10.18, you can configure how emails are used by creating custom templates, changing the language of emails, and disabling emails.

Related concepts

[Simple Mail Transfer Protocol](#)

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

[Self-registration for users](#)

Starting from Maximo Application Suite 9.0, users can self-register to create their own login accounts and use the applications that they have access to. Before users can self-register, an administrator must enable and configure access options that are associated with each identity provider that is configured.

Related tasks

[Configuring emails as optional in Maximo Application Suite 9.0.14](#)

In Maximo Application Suite 9.0, a user's email address is required by default. Starting in Maximo Application Suite 9.0.14, by updating the Suite custom resource in Red Hat OpenShift Container Platform, you can change this setting so that the email address is optional.

Configuring emails as optional in Maximo Application Suite 9.0.14

In Maximo Application Suite 9.0, a user's email address is required by default. Starting in Maximo Application Suite 9.0.14, by updating the Suite custom resource in Red Hat OpenShift Container Platform, you can change this setting so that the email address is optional.

About this task

Note: In Maximo Application Suite 9.1, a user's email address is already optional.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, change the value for `settings.userDataValidation.emailRequired` to `false`.
5. Save the CR changes.

Results

When email addresses are set to optional, you are not required to specify an email address when you create a user record.

If the user does not have a specified email address, system-generated emails, such as welcome or password emails, are sent to the email address for the default suite system that is specified in the Simple Mail Transfer Protocol (SMTP) configuration. For more information, see [Suite system email address](#).

Related concepts

[Simple Mail Transfer Protocol configuration](#)

A Simple Mail Transfer Protocol (SMTP) server connection is required to enable email notifications for Maximo Application Suite system events such as new user welcome emails and password reset communication.

Creating custom email templates

Starting in Maximo Application Suite 9.0.3, you can customize the emails templates that are used in Maximo Application Suite by creating a ConfigMap in Red Hat OpenShift Container Platform.

About this task

If you are using earlier versions of Maximo Application Suite, you can create custom emails in Maximo Application Suite 8.11.15 or 8.10.18.

Only emails that are sent by Maximo Application Suite system events can be customized. Emails that are sent by suite applications, such as Maximo Manage, cannot be customized.

The following email templates, including the template arguments, can be used as a baseline to create a custom template. The template arguments that are associated with each template are variables that represent specific values, such as the users display name and password.

Template name	Description	Template argument
AdminPasswordReset	Sends a request to the administrator to reset a users password if the user clicks Forgot password on the login page.	<ul style="list-style-type: none"> <code>user.displayName</code> <code>password</code> String. The user's new password is in plain text.
UsageAlertCritical	Informs the administrator if AppPoints usage exceeded the capacity amount.	<ul style="list-style-type: none"> <code>user.displayName</code> The following arguments are generated within sendmail. <ul style="list-style-type: none"> <code>links.login</code> <code>links.saas_usage</code>
UsageAlertWarning	Informs the administrator if AppPoints usage is close to capacity.	<ul style="list-style-type: none"> <code>user.displayName</code> <code>alert.threshold</code> The following arguments are generated within sendmail. <ul style="list-style-type: none"> <code>links.login</code> <code>links.saas_usage</code>
UserAccountLocked	Informs a user that their account is locked.	<ul style="list-style-type: none"> <code>user.displayName</code> <code>minutes</code> If included, <code>minutes</code> is the time after which the account is unlocked. If not included, only an administrator can unlock the account.

Template name	Description	Template argument
UserPasswordReset	Sends a temporary password to the user if a user needs to change their password or forgot their password.	<ul style="list-style-type: none"> • <code>user.displayName</code> • <code>code</code> String. The password reset code. • <code>expirationMinutes</code> Number. The code expiration time in minutes.
UserPasswordResetConfirmation	Sends a confirmation email after a user resets their password.	<ul style="list-style-type: none"> • <code>user.displayName</code> • <code>password</code> String. The user's new password is in plain text.
UserPasswordResetFailed	Sends a confirmation email if a users password did not reset.	<ul style="list-style-type: none"> • <code>user.displayName</code>
UserSelfRegistrationAdminNotification	Informs the administrator to approve an account request from a user who self-registered.	<ul style="list-style-type: none"> • <code>primaryEmail</code> • <code>phoneNumber</code> <p>The following arguments are generated within sendmail.</p> <ul style="list-style-type: none"> • <code>links.admin_login</code>
UserSelfRegistrationConfirmed	Sends the account information to a user that self-registered after their account is ready to use.	<ul style="list-style-type: none"> • <code>user.displayName</code> • <code>user.username</code> • <code>adminSupportEmail</code> String. Support email. <p>The following arguments are generated within sendmail.</p> <ul style="list-style-type: none"> • <code>links.login</code>
UserSelfRegistrationInitiated	Sends a verification email with a code that the user can enter to complete the self-registration process if a user self-registers on the login page.	<ul style="list-style-type: none"> • <code>displayName</code> • <code>code</code> String. The self-registration code. • <code>expirationMinutes</code> Number. The code expiration time in minutes.
UserSelfRegistrationPendingAdminApproval	Informs a user who self-registered that the administrator needs to approve their request to create an account.	<ul style="list-style-type: none"> • <code>user.displayName</code> • <code>adminSupportEmail</code> String. Support email.
WelcomeUsername	Sends the username to the user after the administrator creates a user.	<ul style="list-style-type: none"> • <code>user.displayName</code> • <code>user.username</code> • <code>passwordEmailToFollow</code>

Template name	Description	Template argument
		Boolean. Set if the user is to be informed that their password is sent in a subsequent email. <ul style="list-style-type: none"> selfRegisteredUser Boolean. Set if the user self-registered.
WelcomePassword	Sends the initial password to the user after the administrator creates a user.	<ul style="list-style-type: none"> user.displayName password String. The user's initial password is in plain text.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Workloads section, select **ConfigMap**.
2. Click **Create ConfigMap** and select **YAML view**.
3. Enter the YAML syntax for the ConfigMap.

The following example is for a custom email template to confirm a password reset.

```

kind: ConfigMap
apiVersion: v1
metadata:
  name: custom-email-templates
  namespace: mas-myinstanceid-core
data:
  UserPasswordReset: |
    MIME-Version: 1.0
    Content-Type: multipart/alternative; boundary=edf79c96c6bc4320ae4106a24320edb0;
    charset="UTF-8"
    Subject: Password reset confirmation / Confirmação de reconfiguração de senha: IBM
    Maximo Application Suite

    This is a multipart message in MIME format.
    template_name='{{ meta.template_name }}'

    --edf79c96c6bc4320ae4106a24320edb0
    Content-Type: text/plain; charset="UTF-8"

    -----
    Password reset confirmation: IBM Maximo Application Suite
    -----

    Dear {{ user.displayName }},

    An IBM Maximo Application Suite password reset was requested for the account that is
    associated with this email address.
    If you did not make the request, you can ignore this email and no further action is
    taken.

    To complete the reset, enter the included confirmation code in the password reset dialog.

    Confirmation code: {{ code }}

    {% if expirationMinutes -%}
    This code expires in {{ expirationMinutes }} minutes.
    {% endif -%}

    -----
    Need help?
    For technical support, visit IBM Support page: {{ links.ibm_support }}
    For suite support, see IBM Documentation: {{ links.mas_kc }}
    -----
    (c) 2023 IBM All rights reserved
  
```

```
-----  
Confirmação de reconfiguração de senha: IBM Maximo Application Suite  
-----
```

```
Caro(a) {{ user.displayName }},
```

```
Uma reconfiguração de senha do IBM Maximo Application Suite foi solicitada para a conta  
que está associada a este endereço de e-mail.  
Caso não tenha feito a solicitação, ignore esse e-mail e nenhuma outra ação será tomada.
```

```
Para concluir a reconfiguração, insira o código de confirmação incluído na caixa de  
diálogo de reconfiguração de senha.
```

```
Código de confirmação: {{ code }}
```

```
{% if expirationMinutes -%}  
Esse código expira em {{ expirationMinutes }} minutos.  
{% endif -%}
```

```
-----  
Precisa de ajuda?
```

```
Para obter suporte técnico, visite a página de suporte da IBM: {{ links.ibm_support }}  
Para obter suporte a conjuntos, consulte a documentação da IBM: {{ links.mas_kc }}
```

```
-----  
(c) 2023 IBM Todos os direitos reservados
```

```
--edf79c96c6bc4320ae4106a24320edb0
```

If you use an existing email template as a baseline for a custom template, be aware of existing templates reference files, such as `CommonHeader.html.j2`, `CommonFooter.html.j2`, `CommonFooter.txt.j2`, and `CommonPreamble.txt.j2`. If those references are kept in the custom templates, make sure the content of these referenced files are also added to the custom-email-templates configmap, for example:

```
data:  
  ...  
  CommonPreamble.txt.j2: |  
    This is a multipart message in MIME format.  
    template_name='{{ meta.template_name }}'
```

4. To use the custom template, specify the template mode as `custom` in the `Smtpcfg` custom resource.
 - a) In the Administration section, select **Custom Resource Definitions**.
 - b) In the **CustomResourcesDefinitions** window, select the `Smtpcfg` file.
 - c) On the **Instances** tab, select the instance that you want to update.
 - d) On the **YAML** tab, in the `spec.config` section, enter `custom` for the `templateMode` property.

```
spec:  
  config:  
    templateMode: Custom
```

5. Save your changes.

What to do next

Validate that the custom template is working as expected by initiating the scenario where the email is sent. If the new template contains any errors, the email is not sent to the user.

You can also validate that the emails are sent as expected by checking the logs in the `<instance_name>-sendmailapi` and `<instance_name>-coreapi` pods.

Disabling email notifications

Starting in Maximo Application Suite 9.0.3, you can configure the `Smtpcfg` custom resource file to disable email templates that are used by Maximo Application Suite. When an email template is disabled, any email that uses that template is not sent to users.

About this task

If you are using earlier versions of Maximo Application Suite, you can disable email notifications in Maximo Application Suite 8.11.15 or 8.10.18.

Only emails that are sent by Maximo Application Suite system events can be disabled. Emails that are sent by suite applications, such as Maximo Manage, cannot be disabled.

The following email templates can be disabled when you specify them in the `disabledTemplates` property of the `Smtpcfg` custom resource:

Template name	Description
AdminPasswordReset	Sends a request to the administrator to reset a users password if the user clicks Forgot password on the login page.
UsageAlertCritical	Informs the administrator if AppPoints usage exceeded the capacity amount.
UsageAlertWarning	Informs the administrator if AppPoints usage is close to capacity.
UserAccountLocked	Informs a user that their account is locked.
UserPasswordReset	Sends a temporary password to the user if a user needs to change their password or forgot their password.
UserPasswordResetConfirmation	Sends a confirmation email after a user resets their password.
UserPasswordResetFailed	Sends a confirmation email if a users password did not reset.
UserSelfRegistrationAdminNotification	Informs the administrator to approve an account request from a user who self-registered.
UserSelfRegistrationConfirmed	Sends the account information to a user that self-registered after their account is ready to use.
UserSelfRegistrationInitiated	Sends a verification email with a code that the user can enter to complete the self-registration process if a user self-registers on the login page.
UserSelfRegistrationPendingAdminApproval	Informs a user who self-registered that the administrator needs to approve their request to create an account.
WelcomeUsername	Sends the username to the user after the administrator creates a user.
WelcomePassword	Sends the initial password to the user after the administrator creates a user.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the `Smtpcfg` file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab, in the `spec.config` section, enter the name of the email template that you want to disable in the `disabledTemplates` property.

For example, to disable the emails that are sent to new users about their username and password, specify the `WelcomeUsername` and `WelcomePassword` template name in the `disabledTemplates` property.

```
spec:
  config:
    disabledTemplates:
      - WelcomeUsername
      - WelcomePassword
```

5. Save the custom resource changes.

Changing the language of email notifications

Starting in Maximo Application Suite 9.0.3, you can configure the `SmtPCfg` custom resource file to change the language of system event emails from English to another language. In the `SmtPCfg` custom resource file, specify the template mode and the default language to determine which language is used.

About this task

If you are using earlier versions of Maximo Application Suite, you can change the language of system event emails in Maximo Application Suite 8.11.15 or 8.10.18.

You can change the language only for emails that are sent by Maximo Application Suite system events. You cannot change the language for emails that are sent by suite applications, such as Maximo Manage.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the `SmtPCfg` file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab in the `spec.config` section, enter the template mode in the `templateMode` property:

Template mode	Description
Default	The emails that are sent are provided in the language that is specified in the <code>defaultLanguage</code> property.
User	<p>If the user set their preferred language, the emails are provided in the user's preferred language. For more information, see “Setting language and time zone preferences for users” on page 807.</p> <p>If users have not set their preferred language, the emails are sent in the language that is defined in the <code>defaultLanguage</code> property.</p>

5. In the `spec.config` section, specify the language code in the `defaultLanguage` property. You can choose the following languages.

Language code	Language
cs	Czech
da	Danish
de	German

Language code	Language
es	Spanish
fi	Finnish
fr	French
hr	Croatian
hu	Hungarian
it	Italian
ja	Japanese
ko	Korean
nb	Norwegian
nl	Dutch
pl	Polish
pt-BR	Brazilian Portuguese
sk	Slovak
sl	Slovenian
sv	Swedish
tr	Turkish
zh-TW	Traditional Chinese
zh	Simplified Chinese

Note: If you enter a language code that is not supported or if you do not set the default language, the emails are sent in English.

6. Save your changes.

Results

Depending on your configuration of the template mode and the default language, the appropriate language is used.

However, if you created a custom email template and set the template mode as `custom`, then emails are sent in the language that was used in the custom template. For more information, see [“Creating custom email templates”](#) on page 637.

Example

In the following example, the administrator specifies the template mode as `User` and the default language as Brazilian Portuguese. When an email is sent to a user, the email is provided in the user's preferred language that is set in their profile account. If the user does not have a preferred language set, then the email is sent in Brazilian Portuguese.

```
spec:
  config:
    templateMode: User
    defaultLanguage: pt-BR
```

In the following example, the administrator specified the template mode as `Custom` and Spanish as the default language. When an email is sent, the email template that is defined in the custom-email-

templates ConfigMap is used, if available. If a custom version of the email template does not exist in the ConfigMap, then the email is sent in Spanish.

```
spec:  
  config:  
    templateMode: Custom  
    defaultLanguage: es
```

Customer-managed **Configuring external launchers**

External launchers are products that are managed outside of Maximo Application Suite. By configuring an external launcher, you enable product for use with Maximo Application Suite and add a tile in the Suite navigator.

Before you begin

Before you can add the product in Maximo Application Suite, install and configure the product.

About this task

To add an external launcher to the Suite navigator, you specify the solution portal URL in Maximo Application Suite.

Available external launchers	Type	Notes
Manage (external)	Application	In Maximo Application Suite 8.6.0, the Maximo Manage application can be used both integrated with Maximo Application Suite and externally as a stand-alone but linked application. You can access Maximo Manage from the Suite navigator. For more information, see the IBM Maximo Asset Management documentation .

Table 68. Available external launchers (continued)

Available external launchers	Type	Notes
MRO Inventory Optimization	Application	<p>Note:</p> <p>Starting in Maximo Application Suite 9.0, MRO Inventory Optimization is no longer available to be added as an externally configured application and must be accessed by the dedicated URL. The information that is provided is applicable to Maximo Application Suite 8.11 and earlier versions. If MRO Inventory Optimization is configured as an external launcher and you are upgrading to Maximo Application Suite 9.0, you must remove MRO Inventory Optimization before you can complete the upgrade.</p> <p>MRO Inventory Optimization application is not integrated with Maximo Application Suite. It is used as a stand-alone but linked product that requires an externally purchased license. For more information, see Configuring MRO Inventory Optimization.</p>
Maximo APM for E&U	Industry solution	Maximo APM for E&U is not integrated with Maximo Application Suite. It is used as a stand-alone but linked application that requires an externally purchased license.
Scheduler Optimization	Add-on	You use IBM Maximo Scheduler Optimization to optimize scheduling for your Maximo Manage organization. For more information, see Configuring Maximo Scheduler Optimization .

Procedure

1. On the **Suite administration** page, select **Workspace** from the side navigation menu and then the **External launchers** tab.
2. Click **Edit** and select the external launcher that you want to add.
3. Enter the solution portal URL of the external environment.

- Important:** The solution portal URL must be well-formed, including the initial `http://` or `https://`.
4. Save your changes.

Results

You can access the product login page from the Suite navigator.

To remove a product as an external launcher, you can delete the solution portal URL on the **External launcher** page. When the product URL is removed, users can no longer access from the suite navigator.

Related concepts

[Maximo Application Suite Industry solutions](#)

On the Industry solutions tab in the IBM Maximo Application Suite catalog, administrators who have workspace management access can add and remove industry solutions.

Customer-managed **Configuring the minimum password length**

You can change the default minimum password length in Maximo Application Suite by creating a secret.

Procedure

1. Create the yaml file and add the following information:

```
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: %instance ID%-password-policy
  namespace: "mas-%instance ID%-core"
stringData:
  minLength: "8"
```

2. Replace **%instance ID%** with your Maximo Application Suite instance name in the name and namespace fields.
3. Specify the minimum length for the password.
4. To create the secret, run the following command:
oc create -f <filename>.yaml
5. To apply the password policy update, delete the coreapi pods.

Changing privacy access for obtaining user data

Starting in Maximo Application Suite 8.11.15, you can configure the level of permissions to access user data. If you use APIs to retrieve user data, you can view that data. However, you can configure the Suite custom resource (CR) file to control whether that information is available to all users.

About this task

Requesting user information by using the GET /users and GET /users/userid APIs, enables users to view that data. By setting permissions of the userDataPrivacyAccess property in the Suite CR file, you can control who can access this information, whether it is full access, access to non-sensitive information, or no access. User administrators who are assigned user management privileges always have access to this data. Access for all other users, such as application users and suite administrators, depends on the permission setting.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the spec.settings section, change the permission for the userDataPrivacyAccess property.

Permission	Description
ALL	This setting is the default setting for user privacy access. All users have access to the user data.
NON_SENSITIVE_DATA	Application users and suite administrators can retrieve only the ID and username information of any user. User administrators, who are responsible for managing users, continue to have access to all user data.
NO_ACCESS	Application users and suite administrators cannot retrieve any user data from any user API. The user APIs return 403. User administrators, who are responsible for managing users, continue to have access to all user data.

For example, if you set `userDataPrivacyAccess` to `NO_ACCESS`, only user administrators have access to all user data. All other users do not have access to this information.

```
spec:
  settings:
    userDataPrivacyAccess: NO_ACCESS
```

5. Save the CR changes.

Enabling special characters for user ID and username

Starting in Maximo Application Suite 9.0.2, you can include special characters for user ID and username. User IDs and username can typically contain only alphanumeric characters and some special characters. However, you can enable Maximo Application Suite the use of all special characters by updating the custom resource file in Red Hat OpenShift Container Platform.

About this task

If you are using earlier versions of Maximo Application Suite, you can enable special characters for user ID and username in Maximo Application Suite 8.10.17 and 8.11.14.

By default, user IDs and usernames can contain only alphanumeric characters and the underscore (`_`), hyphen (`-`), at (`@`) and period (`.`) special characters.

As an alternative, if you install Maximo Application Suite by using a command-line interface (CLI), you can enable special characters by entering `Yes` to the prompt `Do you want to allow special characters for user IDs and usernames?`.

If you are upgrading to 8.10, 8.11 or 9.0, and users include special characters in their user IDs or usernames, enable the use of all special characters before you upgrade.

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, add the `userDataValidation` section, and enter the value for the `allowSpecialChars` property as `True`.

If you set `allowSpecialChars` to `True`, users can include special characters in their user IDs and usernames.

```
spec:
  settings:
    userDataValidation:
      allowSpecialChars: True
```

5. Save the custom resource changes.

What to do next



Attention: Maximo Application Suite

- If usernames include the opening bracket (`[`), users might not be able to log in. If you are upgrading from Maximo Asset Management to Maximo Application Suite, you can change the username to remove any opening brackets after the upgrade.
- In Maximo Application Suite 8.10.x and 8.11.x, if special characters are used in the user id or username, you cannot access IBM Watson IoT Platform.

If you import multiple users by using the `.csv` template and use double bytes characters, save the file as CSV UTF-8 format.

If you later change the `allowSpecialChars` property from `True` to `False`, unexpected behavior might occur.

Customer-managed Customizing workloads

As an administrator, you can manually configure workloads so that IBM Maximo Application Suite can scale them to match demand. You can modify pod specifications, such as replicas, container resources, affinity, anti-affinity, and tolerations.

podTemplates

Each supported pod is handled by the custom resource object. These custom resources support a list that is called **podTemplates**. You can configure values for **replicas**, **containers**, **affinity**, and **tolerations** for a pod in **podTemplates** of the **spec** property. During reconciliation, the operator overrides the default values with the configured specifications. Also, the operator automatically considers default values when it removes a manually configured pod or **podTemplates**.

Customer-managed Supported pods

To customize the workload in IBM Maximo Application Suite, you can modify the configurations only for supported pods that are handled by specific custom resource objects.

Note:

- The list of pods that are supported for the IBM Maximo Health and IBM Maximo Collaborate application are different than the supported pods for Maximo Application Suite and its other applications. For more information, see [Supported pods for Maximo Health](#) and [Supported pods for Maximo Collaborate](#).
- The list of pods that are supported for the Suite License Service are different than the following list of supported pods. For more information, see [Supported pods for workload customization in Suite License Service](#).

Pods that are handled by the **BasCfg** custom resource

The following deployment pods are handled by the **BasCfg** custom resource object:

Table 69. Deployment pods that are handled by the **BasCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
adoptionusageapi	Container	adoptionusageapi	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
adoptionusageapi	InitContainer	adoptionusageapi-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
milestonesapi	Container	milestonesapi	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
suds	Container	suds	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi

The following CronJob pods are handled by the **BasCfg** custom resource object:

Table 70. CronJob pods that are handled by the **BasCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
accappoints	Container	accappoints	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
accappoints	InitContainer	accappoints-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
adoptionusage-reporter	Container	adoptionusage-reporter	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
adoptionusage-reporter	InitContainer	adoptionusager-reporter-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi

Table 70. CronJob pods that are handled by the **BasCfg** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
usage-daily	Container	usage-daily	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
usage-daily	InitContainer	adoptionusagem etering-daily- init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
usage- historical	Container	usage- historical	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
usage- historical	InitContainer	adoptionusagem etering- historical- init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
usage-hourly	Container	usage-hourly	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
usage-hourly	InitContainer	adoptionusagem etering- hourly-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi

Pods handled by Suite custom resource

The following deployment and CoreIDP pods are handled by the **Suite** custom resource object:

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object.

Pod name	Container type	Container name	Default replicas	Default resources
entitymgr- suite	Container	manager	1	requests: cpu: 0.5 memory: 64Mi limits: cpu: 1.5 memory: 512Mi

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object. (continued)

Pod name	Container type	Container name	Default replicas	Default resources
entitymgr-addons	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.5 memory: 384Mi
entitymgr-bascfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-coreidp	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.5 memory: 512Mi
entitymgr-idpcfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-jdbccfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.5 memory: 384Mi
entitymgr-kafkacfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.4 memory: 256Mi
entitymgr-mongocfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-objectstorage	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.4 memory: 256Mi

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object. (continued)

Pod name	Container type	Container name	Default replicas	Default resources
entitymgr-pushnotificationcfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-scimcfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-slscfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.5 memory: 512Mi
entitymgr-smtpcfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
entitymgr-watsonstudiocfg	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.4 memory: 256Mi
entitymgr-ws	Container	manager	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
coreapi	Container	coreapi	3	requests: cpu: 0.3 memory: 500Mi limits: cpu: 1 memory: 2Gi
internalapi	Container	internalapi	1	requests: cpu: 0.1 memory: 300Mi limits: cpu: 1.2 memory: 2Gi

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object. (continued)

Pod name	Container type	Container name	Default replicas	Default resources
mobileapi	Container	mobileapi	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.8 memory: 2Gi
mobileapi	InitContainer	mobileapi-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
catalogapi	Container	catalogapi	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
catalogapi	InitContainer	catalogapi-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
catalogmgr	Container	catalogmgr	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi
catalogmgr	InitContainer	catalogmgr-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
groupsync-coordinator	Container	groupsync-coordinator	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
groupsync-coordinator	InitContainer	groupsync-coordinator-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object. (continued)

Pod name	Container type	Container name	Default replicas	Default resources
usersync-coordinator	Container	usersync-coordinator	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
usersync-coordinator	InitContainer	usersync-coordinator-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
workspace-coordinator	Container	workspace-coordinator		requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
workspace-coordinator	InitContainer	workspace-coordinator-init		requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi
monagent-mas	Container	monagent-mas		requests: cpu: 0.01 memory: 196Mi limits: cpu: 0.5 memory: 512Mi
admin-dashboard	Container	admin-dashboard		requests: cpu: 0.01 memory: 196Mi limits: cpu: 0.8 memory: 512Mi
homepage	Container	homepage		requests: cpu: 0.01 memory: 196Mi limits: cpu: 0.8 memory: 512Mi
navigator	Container	navigator		requests: cpu: 0.01 memory: 196Mi limits: cpu: 0.8 memory: 512Mi

Table 71. List of deployment and CoreIDP pods handled by the **Suite** custom resource object. (continued)

Pod name	Container type	Container name	Default replicas	Default resources
coreidp	Container	coreidp	1	requests: cpu: 0.02 memory: 500Mi limits: cpu: 1.5 memory: 1Gi
coreidp	InitContainer	coreidp-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 512Mi
coreidp-login	Container	coreidp-login	1	requests: cpu: 0.01 memory: 256Mi limits: cpu: 0.8 memory: 512Mi

The following Job pods are handled by the **Suite** custom resource object:

Table 72. List of Job pods that are handled by the **Suite** custom resource object.

Pod name	Container type	Container name	Default replicas	Default resources
ltpakeygenerator	Container	liberty	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi
oidcclientreg	Container	liberty	1	requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 256Mi

Pods that are handled by the **PushNotificationCfg** custom resource

The following deployment pods are handled by the **PushNotificationCfg** custom resource:

Table 73. Deployment pods that are handled by the **PushNotificationCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
push-notification-service	Container	push-notification-service	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.8 memory: 256Mi

Table 73. Deployment pods that are handled by the **PushNotificationCfg** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
push-notification-service	InitContainer	push-notification-service-init	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi

Pods that are handled by the ScimCfg custom resource

The following deployment pods are handled by the **ScimCfg** custom resource:

Table 74. Deployment pods that are handled by the **ScimCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
scimsync	Container	scimsync	1	requests: cpu: 0.2 memory: 256Mi limits: cpu: 0.5 memory: 500Mi

The following CronJob pods are handled by the **ScimCfg** custom resource:

Table 75. CronJob pods that are handled by the **ScimCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
scim-cronjob	Container	scimsync	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 256Mi

Pods that are handled by the SlsCfg custom resource

The following deployment pods are handled by the **SlsCfg** custom resource:

Table 76. Deployment pods that are handled by the **SlsCfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
licensing-mediator	Container	licensing-mediator	1	requests: cpu: 0.01 memory: 128Mi limits: cpu: 0.4 memory: 1Gi

Pods that are handled by the Smtpcfg custom resource

The following deployment pods are handled by the **Smtpcfg** custom resource:

Table 77. Deployment pods that are handled by the **Smtpcfg** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
sendmailapi	Container	sendmailapi	1	<pre>requests: cpu: 0.01 memory: 256Mi limits: cpu: 0.4 memory: 1Gi</pre>

Supported pods for Maximo Manage

To customize the workload in IBM Maximo Manage, modify the configuration for the pod that is handled by the custom resource object.

Maximo Manage supports the following **podTemplates** fields:

podTemplates field name	Example
replicas and resources	<pre>apiVersion: apps.mas.ibm.com/v1 kind: ManageWorkspace metadata: name: inst1-masdev namespace: mas-inst1-manage labels: mas.ibm.com/applicationId: manage mas.ibm.com/instanceId: inst1 mas.ibm.com/workspaceId: masdev spec: podTemplates: - name: monitoragent replicas: 2 containers: - name: monitoragent resources: limits: cpu: 0.25 memory: 350Mi requests: cpu: 0.1 memory: 256Mi</pre>
affinity	<pre>apiVersion: apps.mas.ibm.com/v1 kind: ManageWorkspace metadata: name: inst1-masdev namespace: mas-inst1-manage labels: mas.ibm.com/applicationId: manage mas.ibm.com/instanceId: inst1 spec: podTemplates: - name: monitoragent affinity: nodeAffinity: preferredDuringSchedulingIgnoredDuringExecution: - weight: 1 preference: matchExpressions: - key: runtimeType operator: In values: - frontend podAffinity: requiredDuringSchedulingIgnoredDuringExecution: - labelSelector: matchExpressions:</pre>

podTemplates field name	Example
	<pre> - key: security operator: In values: - S1 topologyKey: topology.kubernetes.io/zone podAntiAffinity: preferredDuringSchedulingIgnoredDuringExecution: - weight: 100 podAffinityTerm: labelSelector: matchExpressions: - key: security operator: In values: - S2 topologyKey: topology.kubernetes.io/zone </pre>
tolerations	<pre> - name: monitoragent tolerations: - key: "key1" operator: "Exists" effect: "NoSchedule" </pre>
securityContext - both at the pod and init or container level and nodeSelector	<pre> apiVersion: apps.mas.ibm.com/v1 kind: ManageWorkspace metadata: name: tfin-masdev namespace: mas-tfin-manage labels: mas.ibm.com/applicationId: manage mas.ibm.com/instanceId: tfin mas.ibm.com/workspaceId: masdev spec: podTemplates: - name: manage-maxinst nodeSelector: reservedFor: MAS securityContext: level: 's0:c30,c0' seLinuxOptions: null containers: - name: manage-maxinst-maxinst securityContext: fsGroup: 1000870000 seLinuxOptions: level: 's0:c30,c0' seccompProfile: type: RuntimeDefault </pre>
hostAliases and hostnames	<pre> spec: podTemplates: - name: manage-maxinst hostAliases: - ip: "10.10.1.1" hostnames: - "ldap1.com" - "ldap2.com" </pre>
topologySpreadConstraints	<pre> apiVersion: apps.mas.ibm.com/v1 kind: ManageServerBundle metadata: name: all namespace: mas-tfin-manage labels: app.kubernetes.io/instance: tfin app.kubernetes.io/managed-by: ibm-mas- manage </pre>

podTemplates field name	Example
	<pre> app.kubernetes.io/name: ibm-mas-manage mas.ibm.com/applicationId: manage mas.ibm.com/instanceId: tfin mas.ibm.com/workspaceId: masdev spec: podTemplates: - name: all topologySpreadConstraints: - labelSelector: matchLabels: mas.ibm.com/appType: serverBundle mas.ibm.com/appTypeName: all mas.ibm.com/applicationId: manage mas.ibm.com/instanceId: tfin mas.ibm.com/workspaceId: masdev maxSkew: 2 topologyKey: topology.kubernetes.io/zone whenUnsatisfiable: ScheduleAnyway </pre>

Remember: For the build-config pod, only resources and nodeSelector podTemplates are applicable.

Note: For the ManageServerBundle pod, securityContext, affinity, nodeSelector, hostAliases, hostname, and topologySpreadConstraints were previously handled through the passThroughDeployment spec only. Starting in Maximo Application Suite 9.0, podTemplates is the new approach and takes higher precedence when both podTemplates and the spec are applied.

ManageApp custom resource object

The following deployment pods are handled by the **ManageApp** custom resource object:

<i>Table 78. Deployment pods that are handled by the ManageApp custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
entitymgr- primary-entity	Container	manager	1	<pre> resources: requests: cpu: 0.01 memory: 64Mi limits: cpu: 0.2 memory: 512Mi </pre>
entitymgr- appstatus	Container	manager	1	<pre> resources: requests: cpu: 0.2 memory: 300Mi limits: cpu: 0.8 memory: 1024Mi </pre>

Table 78. Deployment pods that are handled by the **ManageApp** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
entitymgr-bdi	Container	manager	1	resources: requests: cpu: 0.03 memory: 128Mi limits: cpu: 0.8 memory: 1024Mi
entitymgr-ws	Container	manager	1	resources: requests: cpu: 0.2 memory: 500Mi limits: cpu: 0.8 memory: 2Gi
entitymgr-acc	Container	manager	1	resources: requests: cpu: 30m memory: 128Mi limits: cpu: 800m memory: 1Gi
usersyncagent	Container	manage- usersyncagent	1	resources: requests: cpu: 0.03 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
groupsyncagent	Container	manage- groupsyncagent	1	resources: requests: cpu: 0.03 memory: 128Mi limits: cpu: 0.25 memory: 256Mi

Table 78. Deployment pods that are handled by the **ManageApp** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
ibm-mas-imagestitching-operator	Container	imagestitching	1	<pre>resources: requests: cpu: 0.2 memory: 300Mi ephemeral-storage: 2Mi limits: cpu: 0.5 memory: 1024Mi ephemeral-storage: 2Gi</pre>
healthext-entitymgr-ws	Container	healthext	1	<pre>resources: requests: cpu: 0.1 memory: 128Mi ephemeral-storage: 2Mi limits: cpu: 0.5 memory: 512Mi ephemeral-storage: 2Gi</pre>
ibm-mas-slackproxy-operator	Container	slackproxy	1	<pre>resources: requests: cpu: 0.5 memory: 300Mi ephemeral-storage: 2Mi limits: cpu: 1Gi memory: 1Gi ephemeral-storage: 2Gi</pre>

Note: For all the deployments in Table 1, more than one replica is not advisable.

ManageWorkspace custom resource object

The following deployment pods are handled by the **ManageWorkspace** custom resource object:

Table 79. Deployment pods that are handled by the **ManageWorkspace** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
monitoragent	Container	monitoragent	1	<pre>resources: requests: cpu: 0.1 memory: 256Mi limits: cpu: 0.25 memory: 350Mi</pre>

Table 79. Deployment pods that are handled by the **Manageworkspace** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
manage-maxinst	Container	manage-maxinst-maxinst	1	resources: requests: cpu: 0.5 memory: 500Mi limits: cpu: 2 memory: 4Gi
Server bundle name which is dynamic while activating Manageworkspace , for example, mea	Container	Server bundle name which is dynamic while activating Manageworkspace , for example, mea	1	resources: requests: cpu: 0.5 memory: 2Gi limits: cpu: 6 memory: 10Gi
Server bundle name which is dynamic while activating Manageworkspace , for example, mea	Container	monitoragent	1	resources: requests: cpu: 0.1 memory: 256Mi limits: cpu: 1 memory: 512Mi
healthext-model-engine	Container	healthext-model-engine	1	resources: requests: cpu: 0.1 memory: 128Mi ephemeral-storage: 2Mi limits: cpu: 1 memory: 1Gi ephemeral-storage: 2Gi
imagestitching	Container	image-stitching	1	resources: requests: cpu: 1 memory: 4Gi limits: cpu: 3 memory: 16Gi
slackproxy	Container	slack-proxy	1	resources: requests: cpu: 1 memory: 4Gi limits: cpu: 3 memory: 16Gi

Table 79. Deployment pods that are handled by the **ManageWorkspace** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
Dynamic name that matches the spec.bdiConfiguration.name from ManageWorkspace custom resource	Container	bdiservice	1	resources: requests: cpu: 1 memory: 1Gi limits: cpu: 3 memory: 16Gi
build-config	Container	adminbuild	1	resources: limits: cpu: 2 memory: 512Mi ephemeral- storage: 100Gi requests: cpu: 1 memory: 256Mi ephemeral- storage: 30Gi
build-config	Container	bundlebuild	1	resources: limits: cpu: 2 memory: 512Mi ephemeral- storage: 100Gi requests: cpu: 1 memory: 256Mi ephemeral- storage: 30Gi

Note: For all the deployments in Table 2, more than one replica is not advisable.

BDI custom resource object

The following deployment pods are handled by the **BDI** custom resource object:

Table 80. Deployment pods that are handled by the **BDI** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
Dynamic name that matches the spec.bdiConfiguration.name from ManageWorkspace custom resource	Container	bdiservice	1	resources: requests: cpu: 1 memory: 1Gi limits: cpu: 3 memory: 16Gi

Imagestitching custom resource object

The following deployment pods are handled by the **Imagestitching** custom resource object:

Table 81. Deployment pods that are handled by the **Imagestitching** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
imagestitching	Container	image-stitching	1	<pre>resources: requests: cpu: 1 memory: 4Gi limits: cpu: 3 memory: 16Gi</pre>

SlackProxy custom resource object

The following deployment pods are handled by the **SlackProxy** custom resource object:

Table 82. Deployment pods that are handled by the **SlackProxy** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
slackproxy	Container	slack-proxy	1	<pre>resources: requests: cpu: 1 memory: 4Gi limits: cpu: 3 memory: 16Gi</pre>

HealthExtWorkspace custom resource object

The following deployment pods are handled by the **HealthExtWorkspace** custom resource object:

Table 83. Deployment pods that are handled by the **HealthExtWorkspace** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
healthext-model-engine	Container	healthext-model-engine	1	<pre>resources: requests: cpu: 0.1 memory: 128Mi ephemeral-storage: 2Mi limits: cpu: 1 memory: 1Gi ephemeral-storage: 2Gi</pre>

ManageAccelerators custom resource object

The following deployment pods are handled by the **ManageAccelerators** custom resource object:

Table 84. Deployment pods that are handled by the **ManageAccelerators** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
healthext-entitymgr-acc	Container	healthext-acc	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 500m memory: 512Mi ephemeral- storage: 2Gi</pre>

HealthExtAccelerators custom resource object

The following deployment pods are handled by the **HealthExtAccelerators** custom resource object:

Table 85. Deployment pods that are handled by the **HealthExtAccelerators** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
healthext-acc-job	Container	healthext-acc-job	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 1000m memory: 1Gi ephemeral- storage: 2Gi</pre>
uninstall-health-acc-job	Container	uninstall-health-acc-job	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 1000m memory: 1Gi ephemeral- storage: 2Gi</pre>

Supported pods for IBM Maximo Manage without custom resource objects

You cannot modify the configurations for supported pods for Maximo Manage unless they are handled by specific custom resource objects. The `ibm-mas-manage-operator` and the `ibm-truststore-mgr-controller-manager` pods cannot be modified.

The following deployment pods are not handled by any custom resource object because they are used as a starting point to install managed operators. After the managed operators are available, users can use custom resource objects. The **container type** for `ibm-mas-manage-operator` and `ibm-truststore-mgr-controller-manager` pods is **container** with a **default replica** value of 1.

Table 86. Deployment pods that are not handled by any custom resource object

Pod name	Container name	Default resources
ibm-mas-manage-operator	webhook	<pre>resources: requests: cpu: 0.2m memory: 50Mi limits: cpu: 250m memory: 1Gi</pre>
ibm-mas-manage-operator	manager	<pre>resources: requests: cpu: 200m memory: 300Mi limits: cpu: 1Gi memory: 1Gi</pre>
ibm-truststore-mgr-controller-manager	manager	<pre>resources: requests: cpu: 0.1m memory: 64Mi limits: cpu: 5Gi memory: 1Gi</pre>

Supported pods for Maximo Health

To customize the workload in IBM Maximo Health, modify the configuration for the pod that is handled by the custom resource object.

The custom resource object that is used to modify the deployment pod depends on whether the Maximo Health application is installed with Maximo Manage or as a stand-alone application.

stand-alone Maximo Health

The custom resource object is **HealthExtWorkspace**.

Maximo Health installed with Maximo Manage

The custom resource object is **ManageWorkspace**.

HealthExtWorkspace custom resource object

The following deployment pods are handled by the **HealthExtWorkspace** custom resource object:

Table 87. Deployment pods that are handled by the **HealthExtWorkspace** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
healthext-model-engine	Container	healthext-model-engine	1	<pre>resources: requests: cpu: 0.1 memory: 128Mi ephemeral-storage: 2Mi limits: cpu: 1 memory: 1Gi ephemeral-storage: 2Gi</pre>

HealthApp custom resource object

The following deployment pods are handled by the **HealthApp** custom resource object:

*Table 88. Deployment pods that are handled by the **HealthApp** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
healthext-entitymgr-ws	Container	healthext	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 500m memory: 512Mi ephemeral- storage: 2Gi</pre>

HealthWorkspace custom resource object

The following deployment pods are handled by the **HealthWorkspace** custom resource object:

*Table 89. Deployment pods that are handled by the **HealthWorkspace** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
healthext-model-engine	Container	healthext-model-engine	1	<pre>resources: requests: cpu: 0.1 memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 1 memory: 1Gi ephemeral- storage: 2Gi</pre>

HealthExtAccelerators custom resource object

The following deployment pods are handled by the **HealthExtAccelerators** custom resource object:

*Table 90. Deployment pods that are handled by the **HealthExtAccelerators** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
healthext-acc-job	Container	healthext-acc-job	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral- storage: 2Mi limits: cpu: 1000m memory: 1Gi ephemeral- storage: 2Gi</pre>

Table 90. Deployment pods that are handled by the **HealthExtAccelerators** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
uninstall-health-acc-job	Container	uninstall-health-acc-job	1	<pre>resources: requests: cpu: 100m memory: 128Mi ephemeral-storage: 2Mi limits: cpu: 1000m memory: 1Gi ephemeral-storage: 2Gi</pre>

Related tasks

“Customizing workload scale” on page 703

You can scale a pod horizontally and vertically by setting container or initContainer resources or pod replicas. By default, some pods have more than one replica. Depending on the requirements of your workloads, you can set the values for these replicas to less than or more than the default value.

“Customizing workload affinity” on page 705

As an administrator, you can customize and improve control of the pod scheduling process. You can constrain a pod so that it is restricted or preferred to run on particular nodes.

“Customizing workload tolerations” on page 706

The tolerations allow a pod to be scheduled on a node that has a matching taint. A taint is a key value pair that is assigned to a node. When a pod does not have a matching taint, it cannot be scheduled on that node.

Supported pods for Maximo Collaborate

To customize the workload in IBM Maximo Collaborate, modify the configuration for the pods that are handled by custom resource objects.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

During deployment of Maximo Collaborate, horizontal pod autoscale is enabled by default based on CPU and memory resource utilization. Maximo Collaborate automatically scales the internal components. To scale the deployment pods manually by using podTemplates, you must disable horizontal pod autoscale by setting the `spec.settings.common.podautoscale` to `false` in the **CollaborateApp** custom resource.

Deployment pods that are handled by the CollaborateApp custom resource

The following deployment pods are handled by the **CollaborateApp** custom resource object:

Table 91. Deployment pods that are handled by the **CollaborateApp** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api	Container	ema-api	1	<pre>requests: cpu: 500m memory: 576Mi limits: cpu: 1000m memory: 2Gi</pre>

Table 91. Deployment pods that are handled by the **CollaborateApp** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
adminconsole	Container	ema-admin-console	1	requests: cpu: 10m memory: 128Mi limits: cpu: 1 memory: 1Gi
antivirus	Container	ema-anti-virus	1	requests: cpu: 20m memory: 4Gi limits: cpu: 1 memory: 4Gi
crawler	Container	ema-crawler	1	requests: cpu: 10m memory: 256Mi limits: cpu: 1 memory: 2Gi
diagnosisengine	Container	<ul style="list-style-type: none"> • ema-diagnosis-engine • ema-diagnosis-engine-dt 	1	requests: cpu: 200m memory: 128Mi limits: cpu: 3 memory: 3Gi
diagnosisengine	Container	ema-diagnosis-engine-bn	1	requests: cpu: 500m memory: 160Mi limits: cpu: 3 memory: 3Gi
diagnosisengine	Container	ema-diagnosis-engine-dt	1	requests: cpu: 100m memory: 192Mi limits: cpu: 3 memory: 3Gi
haproxy	Container	haproxy	1	requests: memory: 128Mi cpu: 50m limits: memory: 1Gi cpu: 1
multitenant	Container	ema-multi-tenant	1	requests: cpu: 100m memory: 128Mi limits: cpu: 1 memory: 1Gi

Table 91. Deployment pods that are handled by the **CollaborateApp** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
studio	Container	ema-studio	1	requests: cpu: 200m memory: 160Mi limits: cpu: 1 memory: 1Gi
technician	Container	ema-technician	1	requests: cpu: 200m memory: 196Mi limits: cpu: 1 memory: 1Gi
voicegateway	Container	collaborate-voice-gateway	1	requests: cpu: 200m memory: 196Mi limits: cpu: 1 memory: 1Gi
voicemanagement	Container	collaborate-voice-management	1	requests: cpu: 200m memory: 196Mi limits: cpu: 1 memory: 1Gi
voicesessionmxinspect	Container	collaborate-voice-session-mxinspect	1	requests: cpu: 200m memory: 196Mi limits: cpu: 1 memory: 1Gi

Statefulset pods that are handled by the CollaborateApp custom resource object

The following Statefulset pods are handled by the **CollaborateApp** custom resource object:

Table 92. Statefulset pods that are handled by the **CollaborateApp** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
couch	Container	<ul style="list-style-type: none"> • db • mgmt 	3	db: requests: cpu: '0.5' memory: 576Mi limits: cpu: '2' memory: 2Gi mgmt: requests: cpu: '0.5' memory: 576Mi limits: cpu: '2' memory: 2Gi

Table 92. Statefulset pods that are handled by the **CollaborateApp** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
redis	Container	redis	3	<pre>requests: cpu: 50m memory: 700Mi limits: cpu: 1 memory: 700Mi</pre>
redis	Container	sentinel	3	<pre>requests: memory: 200Mi cpu: 100m limits: memory: 200Mi cpu: 1</pre>

Note:

- In couch pod, customization of the CouchDB size is not supported as it is exposed during the deployment of Maximo Collaborate.
- In redis pod, customization of the Scale out/in by is not supported.

Related tasks

[“Customizing workload scale” on page 703](#)

You can scale a pod horizontally and vertically by setting container or initContainer resources or pod replicas. By default, some pods have more than one replica. Depending on the requirements of your workloads, you can set the values for these replicas to less than or more than the default value.

[“Customizing workload affinity” on page 705](#)

As an administrator, you can customize and improve control of the pod scheduling process. You can constrain a pod so that it is restricted or preferred to run on particular nodes.

[“Customizing workload tolerations” on page 706](#)

The tolerations allow a pod to be scheduled on a node that has a matching taint. A taint is a key value pair that is assigned to a node. When a pod does not have a matching taint, it cannot be scheduled on that node.

Supported pods for Maximo Real Estate and Facilities

To customize the workload in IBM Maximo Real Estate and Facilities, modify the configuration for the pod that is handled by the custom resource object. The usersyncagent pod is handled by the FacilitiesApp custom resource and the rest of the pods are handled by the FacilitiesWorkspaces custom resource.

Maximo Real Estate and Facilities supports the following podTemplates fields:

- replicas
- resources
- affinity
- tolerations
- securityContext - both at the pod and init or container level
- nodeSelector
- hostAliases
- hostname
- topologySpreadConstraints

usersyncagent custom resource object

The following deployment pods are handled by the **usersyncagent** custom resource object:

<i>Table 93. Deployment pods that are handled by the usersyncagent custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
usersyncagent	Container	usersyncagent	1	<pre>resources: requests: cpu: 0.1 memory: 300Mi ephemeral- storage: 8Mi limits: cpu: 1 memory: 1Gi ephemeral- storage: 1Gi</pre>

datainit custom resource object

The following deployment pods are handled by the **datainit** custom resource object:

<i>Table 94. Deployment pods that are handled by the datainit custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
datainit	Container	datainit	1	<pre>resources: requests: cpu: 0.01 memory: 64Mi ephemeral- storage: 128Mi limits: cpu: 1 memory: 1Gi ephemeral- storage: 2Gi</pre>

appserver custom resource object

The following deployment pods are handled by the **appserver** custom resource object:

<i>Table 95. Deployment pods that are handled by the appserver custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
appserver	Container	appserver	1	<pre>resources: requests: memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 128Mi</pre>

Table 95. Deployment pods that are handled by the **appserver** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

multiagents custom resource object

The following deployment pods are handled by the **multiagents** custom resource object:

Table 96. Deployment pods that are handled by the **multiagents** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
multiagents	Container	multiagents	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 1Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

dataconnectagent custom resource object

The following deployment pods are handled by the **dataconnectagent** custom resource object:

Table 97. Deployment pods that are handled by the **dataconnectagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
dataconnectagent	Container	dataconnectagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

extendedformulaagent custom resource object

The following deployment pods are handled by the **extendedformulaagent** custom resource object:

Table 98. Deployment pods that are handled by the **extendedformulaagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
extendedformulaagent	Container	extendedformulaagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>

Table 98. Deployment pods that are handled by the **extendedformulaagent** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

formularecalcagent custom resource object

The following deployment pods are handled by the **formularecalcagent** custom resource object:

Table 99. Deployment pods that are handled by the **formularecalcagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
formularecalcagent	Container	formularecalcagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

incomingmailagent custom resource object

The following deployment pods are handled by the **incomingmailagent** custom resource object:

Table 100. Deployment pods that are handled by the **incomingmailagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
incomingmailagent	Container	incomingmailagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral-storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral-storage: 2Gi</pre>

objectmigrationagent custom resource object

The following deployment pods are handled by the **objectmigrationagent** custom resource object:

Table 101. Deployment pods that are handled by the **objectmigrationagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
objectmigrationagent	Container	objectmigrationagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral-storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral-storage: 2Gi</pre>

objectpublishagent custom resource object

The following deployment pods are handled by the **objectpublishagent** custom resource object:

*Table 102. Deployment pods that are handled by the **objectpublishagent** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
objectpublishagent	Container	objectpublishagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

maintenanceagent custom resource object

The following deployment pods are handled by the **maintenanceagent** custom resource object:

*Table 103. Deployment pods that are handled by the **maintenanceagent** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
maintenanceagent	Container	maintenanceagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

reportqueueagent custom resource object

The following deployment pods are handled by the **reportqueueagent** custom resource object:

*Table 104. Deployment pods that are handled by the **reportqueueagent** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
reportqueueagent	Container	reportqueueagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

reservesmtpagent custom resource object

The following deployment pods are handled by the **reservesmtpagent** custom resource object:

*Table 105. Deployment pods that are handled by the **reservesmtpagent** custom resource object*

Pod name	Container type	Container name	Default replicas	Default resources
reservesmtpagent	Container	reservesmtpagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>

Table 105. Deployment pods that are handled by the **reservesmtpagent** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

scheduleraagent custom resource object

The following deployment pods are handled by the **scheduleraagent** custom resource object:

Table 106. Deployment pods that are handled by the **scheduleraagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
scheduleraagent	Container	scheduleraagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

wfagent custom resource object

The following deployment pods are handled by the **wfagent** custom resource object:

Table 107. Deployment pods that are handled by the **wfagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
wfagent	Container	wfagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

wffutureagent custom resource object

The following deployment pods are handled by the **wffutureagent** custom resource object:

Table 108. Deployment pods that are handled by the **wffutureagent** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
wffutureagent	Container	wffutureagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral- storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

wfnotificationagent custom resource object

The following deployment pods are handled by the **wfnotificationagent** custom resource object:

<i>Table 109. Deployment pods that are handled by the wfnotificationagent custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
wfnotificationagent	Container	wfnotificationagent	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 2Gi ephemeral-storage: 2Gi</pre>
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral-storage: 2Gi</pre>

dwfagent-mycustomagent(1-n) custom resource object

The following deployment pods are handled by the **dwfagent-mycustomagent(1-n)** custom resource object. You can have multiple dedicated workflow agents. These custom agents are named by the users who create them for specific workflows.

<i>Table 110. Deployment pods that are handled by the dwfagent-mycustomagent(1-n) custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
dwfagent-mycustomagent(1-n)	Container	dwfagent-mycustomagent(1-n)	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral-storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral-storage: 2Gi</pre>

Table 110. Deployment pods that are handled by the **dwfagent-mycustomagent(1-n)** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
update	initContainer	update	1	<pre>resources: requests: cpu: 0.02 memory: 1.3Gi ephemeral- storage: 128Mi limits: cpu: 2 memory: 6Gi ephemeral- storage: 2Gi</pre>

Supported pods for IoT tool

To customize the workload in IoT tool, modify the configuration for the pods that are handled by custom resource objects.

IoT custom resource object

Note: Manager containers should not have replicas set to a value more than "1".

The following deployment pods are handled by the **IoT** custom resource object:

Table 111. Deployment pods that are handled by the **IoT** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
iot-ltpakeys-generator	Container	iot-ltpakeys-generator	1	<pre>requests: cpu: 0.5 memory: 1Gi limits: cpu: 2 memory: 4Gi</pre>
monagent-iot	Container	monagent-iot	1	<pre>requests: cpu: 0.01 memory: 196Mi limits: cpu: 0.5 memory: 512Mi</pre>
actions-operator	Container	manager	1	<pre>requests: cpu: 0.1 memory: 256Mi limits: cpu: 1 memory: 512Mi</pre>
auth-operator	Container	manager	1	<pre>requests: cpu: 0.1 memory: 256Mi limits: cpu: 1 memory: 512Mi</pre>

Table 111. Deployment pods that are handled by the **IoT** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
datapower-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
devops-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
dm-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
dsc-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
edgeconfig-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
fpl-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
guardian-operator	Container	manager	1	requests: cpu: 0.1 memory: 256Mi limits: cpu: 1 memory: 512Mi
mbgx-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi

Table 111. Deployment pods that are handled by the **IoT** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
mfgx-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
monitor-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
orgmgmt-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
provision-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
registry-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
state-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi
webui-operator	Container	manager	1	requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi

Table 111. Deployment pods that are handled by the **IoT** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
workspace-operator	Container	manager	1	resources: requests: cpu: 100m memory: 256Mi limits: cpu: 1000m memory: 512Mi

Actions custom resource object

The following deployment pods are handled by the **Actions** custom resource object:

Table 112. Deployment pods that are handled by the **Actions** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
action-mgr-invoker	Container	action-mgr	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 1 memory: 2Gi
action-mgr-invoker	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
action-mgr-resolver	Container	action-mgr	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 1 memory: 2Gi
action-mgr-resolver	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
api-actions	Container	api-actions	1	requests: cpu: 0.05 memory: 640Mi limits: cpu: 2 memory: 2Gi

Table 112. Deployment pods that are handled by the **Actions** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
api-actions	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-actions	Container	deprovagent-actions	1	requests: cpu: 0.05 memory: 64Mi limits: cpu: 0.25 memory: 256Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

Auth custom resource object

The following deployment pods are handled by the **Auth** custom resource object:

Table 113. Deployment pods that are handled by the **Auth** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-authorization	Container	api-authorization	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
api-authorization	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
api-authentication	Container	api-authentication	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
api-authentication	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi

Table 113. Deployment pods that are handled by the **Auth** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
auth-store	Container	auth-store	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
auth-store	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-auth	Container	deprovagent-auth	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
masuseragent	Container	masuseragent	1	requests: cpu: 0.1 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi

DataPower custom resource object

The following deployment pods are handled by the DataPower custom resource object:

Table 114. Deployment pods that are handled by the **DataPower** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
datapower	Container	datapower	1	requests: cpu: 4 memory: 4400Mi limits: cpu: 8 memory: 4400Mi

Table 114. Deployment pods that are handled by the **DataPower** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
datapower	Container	datapower-config	1	<pre>requests: cpu: 0.1 memory: 128Mi limits: cpu: 0.5 memory: 256Mi</pre>

Devops custom resource object

The following deployment pods are handled by the **Devops** custom resource object:

Table 115. Deployment pods that are handled by the **Devops** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-devops	Container	api-devops	1	<pre>requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi</pre>
api-status	Container	api-status	1	<pre>requests: cpu: 0.01 memory: 196Mi limits: cpu: 2 memory: 2Gi</pre>

Dm custom resource object

The following deployment pods are handled by the **Dm** custom resource object:

Table 116. Deployment pods that are handled by the **Dm** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-devicemgmt	Container	api-devicemgmt	1	<pre>requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi</pre>
api-devicemgmt	Container	statsd	1	<pre>requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi</pre>

Table 116. Deployment pods that are handled by the **Dm** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
deprovagent-dm	Container	deprovagent-dm	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
devicemgmt-server	Container	devicemgmt-server	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
devicemgmt-server	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

Dsc custom resource object

The following deployment pods are handled by the **Dsc** custom resource object:

Table 117. Deployment pods that are handled by the **Dsc** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-connectors	Container	api-connectors	1	requests: cpu: 0.1 memory: 640Mi limits: cpu: 2 memory: 2Gi
api-connectors	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi

Table 117. Deployment pods that are handled by the **Dsc** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
api-importconnectors	Container	api-importconnectors	1	requests: cpu: 0.1 memory: 640Mi limits: cpu: 2 memory: 2Gi
api-importconnectors	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
connectors-store	Container	connectors-store	1	requests: cpu: 0.1 memory: 640Mi limits: cpu: 2 memory: 2Gi
connectors-store	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
data-import-connector	Container	data-import-connector	1	requests: cpu: 1 memory: 2Gi limits: cpu: 0.5 memory: 640Mi
data-import-connector	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
historian-configuration-manager	Container	historian-configuration-manager	1	requests: cpu: 2 memory: 2Gi limits: cpu: 0.2 memory: 512Mi
historian-connector	Container	historian-connector	1	requests: cpu: 1 memory: 2Gi limits: cpu: 0.1 memory: 640Mi

Table 117. Deployment pods that are handled by the **Dsc** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
historian-connector	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 100m memory: 64Mi limits: cpu: 1 memory: 512Mi

Edgeconfig custom resource object

The following deployment pods are handled by the **Edgeconfig** custom resource object:

Table 118. Deployment pods that are handled by the **Edgeconfig** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
edgeconfig-configui	Container	edgeconfig-configui	1	requests: cpu: 0.3 memory: 0.5Gi limits: cpu: 2 memory: 2Gi
edgeconfig-server	Container	edgeconfig-server	1	requests: cpu: 0.3 memory: 0.5Gi limits: cpu: 2 memory: 2Gi

Fp1 custom resource object

The following deployment pods are handled by the **Fp1** custom resource object:

Table 119. Deployment pods that are handled by the **Fp1** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-pipeline	Container	api-pipeline	1	requests: cpu: 0.1 memory: 640Mi limits: cpu: 2 memory: 2Gi

Table 119. Deployment pods that are handled by the **Fpl** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
deprovagent-fpl	Container	deprovagent-fpl	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
functionsexecutor	Container	functionsexecutor	1	requests: cpu: 0.1 memory: 1Gi limits: cpu: 16 memory: 16Gi
functionsexecutor	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
pipelinerouter	Container	pipelinerouter	1	requests: cpu: 0.1 memory: 1Gi limits: cpu: 12 memory: 8Gi
pipelinerouter	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 100m memory: 64Mi limits: cpu: 1 memory: 512Mi

Guardian custom resource object

The following deployment pods are handled by the **Guardian** custom resource object:

Table 120. Deployment pods that are handled by the **Guardian** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-riskmgmt-secguardian	Container	api-riskmgmt-secguardian	1	requests: cpu: 0.01 memory: 512Mi limits: cpu: 2 memory: 2Gi
api-riskmgmt-secguardian	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-guardian	Container	deprovagent-guardian	1	requests: cpu: 0.1 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
provagent-guardian	Container	provagent-guardian	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

Mbgx custom resource object

The following deployment pods are handled by the **Mbgx** custom resource object:

Table 121. Deployment pods that are handled by the **Mbgx** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-mbgadmin	Container	api-mbgadmin	1	requests: cpu: 0.5 memory: 1Gi limits: cpu: 1 memory: 2Gi

Table 121. Deployment pods that are handled by the **Mbgx** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
api-mbgadmin	Container	statsd	1	<pre>requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi</pre>
monagent- mserver	Container	monagent- mserver	1	<pre>requests: cpu: 0.1 memory: 128Mi limits: cpu: 0.4 memory: 1Gi</pre>
monagent- mserver	Container	statsd	1	<pre>requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi</pre>
graphite- exporter	Container	graphite- exporter	1	<pre>requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi</pre>

Mfgx custom resource object

The following deployment pods are handled by the **Mfgx** custom resource object:

Table 122. Deployment pods that are handled by the **Mfgx** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
msproxy	Container	msproxy	1	<pre>requests: cpu: 1 memory: 2Gi limits: cpu: 0.2 memory: 768Mi</pre>
msproxy	Container	monitor	1	<pre>requests: cpu: 0.1 memory: 128Mi limits: cpu: 1 memory: 384Mi</pre>

Table 122. Deployment pods that are handled by the **Mfggx** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
msproxy	Container	statsd	1	<pre>requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi</pre>
graphite-exporter	Container	graphite-exporter	1	<pre>requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi</pre>

Monitor custom resource object

The following deployment pods are handled by the **Monitor** custom resource object:

Table 123. Deployment pods that are handled by the **Monitor** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
entity-connector	Container	entity-connector	1	<pre>requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi</pre>
entity-connector	Container	statsd	1	<pre>requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi</pre>
entity-manager	Container	entity-manager	1	<pre>requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi</pre>
entity-manager	Container	statsd	1	<pre>requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi</pre>

Orgmgmt custom resource object

The following deployment pods are handled by the **Orgmgmt** custom resource object:

Table 124. Deployment pods that are handled by the **Orgmgmt** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-organizations	Container	api-organizations	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
api-organizations	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
config-store	Container	config-store	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi
config-store	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-orgmgmt	Container	deprovagent-orgmgmt	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi
monagent-org	Container	monagent-org	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi
monagent-org	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
org-store	Container	org-store	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 2 memory: 2Gi

Table 124. Deployment pods that are handled by the **Orgmgmt** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
org-store	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
orgdeprovgr	Container	orgdeprovgr	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi
orgpoolmgr	Container	orgpoolmgr	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi
orgpoolmgr	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
orgprovgr	Container	orgprovgr	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

Provision custom resource object

The following deployment pods are handled by the **Provision** custom resource object:

Table 125. Deployment pods that are handled by the **Provision** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-resourcecontroller	Container	api-resourcecontroller	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 1 memory: 2Gi
api-resourcecontroller	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
api-s2s	Container	api-s2s	1	requests: cpu: 0.05 memory: 384Mi limits: cpu: 1 memory: 2Gi
api-s2s	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-provision	Container	deprovagent-provision	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
s2s-store	Container	s2s-store	1	requests: cpu: 0.1 memory: 384Mi limits: cpu: 1 memory: 2Gi
s2s-store	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 100m memory: 64Mi limits: cpu: 1 memory: 512Mi

Registry custom resource object

The following deployment pods are handled by the **Registry** custom resource object:

<i>Table 126. Deployment pods that are handled by the Registry custom resource object</i>				
Pod name	Container type	Container name	Default replicas	Default resources
api-messagesight	Container	api-messagesight	1	requests: cpu: 0.1 memory: 512Mi limits: cpu: 2 memory: 2Gi
api-org	Container	api-org	1	requests: cpu: 0.1 memory: 640Mi limits: cpu: 2 memory: 4Gi
api-org	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-registry	Container	deprovagent-registry	1	requests: cpu: 0.1 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
device-store	Container	device-store	1	requests: cpu: 0.1 memory: 512Mi limits: cpu: 2 memory: 2Gi
device-store	Container	statsd	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 100m memory: 64Mi limits: cpu: 1 memory: 512Mi

State custom resource object

The following deployment pods are handled by the **State** custom resource object:

Table 127. Deployment pods that are handled by the **State** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-schemas	Container	api-schemas	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
api-schemas	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
api-state	Container	api-state	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
api-state	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-state	Container	deprovagent-state	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
mqtt-connector	Container	mqtt-connector	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
mqtt-connector	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
state-updater-devices	Container	state-updater	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi

Table 127. Deployment pods that are handled by the **State** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
state-updater-devices	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
state-updater-things	Container	state-updater	1	requests: cpu: 0.05 memory: 512Mi limits: cpu: 2 memory: 2Gi
state-updater-things	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

WebUI custom resource

The following pods are handled by the **WebUI** custom resource object:

Table 128. Deployment pods that are handled by the **WebUI** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
api-dashboard	Container	api-dashboard	1	requests: cpu: 0.05 memory: 640Mi limits: cpu: 2 memory: 2Gi
api-dashboard	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi

Table 128. Deployment pods that are handled by the **WebUI** custom resource object (continued)

Pod name	Container type	Container name	Default replicas	Default resources
dashboard	Container	dashboard	1	requests: cpu: 0.05 memory: 192Mi limits: cpu: 0.8 memory: 512Mi
dashboard	Container	statsd	1	requests: cpu: 0.01 memory: 16Mi limits: cpu: 0.2 memory: 512Mi
deprovagent-webui	Container	deprovagent-webui	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 0.25 memory: 256Mi
graphite-exporter	Container	graphite-exporter	1	requests: cpu: 0.05 memory: 128Mi limits: cpu: 2 memory: 2Gi

Related tasks

[“Customizing workload scale” on page 703](#)

You can scale a pod horizontally and vertically by setting `container` or `initContainer` resources or pod `replicas`. By default, some pods have more than one replica. Depending on the requirements of your workloads, you can set the values for these replicas to less than or more than the default value.

[“Customizing workload affinity” on page 705](#)

As an administrator, you can customize and improve control of the pod scheduling process. You can constrain a pod so that it is restricted or preferred to run on particular nodes.

[“Customizing workload tolerations” on page 706](#)

The `tolerations` allow a pod to be scheduled on a node that has a matching taint. A taint is a key value pair that is assigned to a node. When a pod does not have a matching taint, it cannot be scheduled on that node.

Supported pods for IBM Data Dictionary

To customize the workload in IBM Data Dictionary, modify the configuration for the pods that are handled by custom resource objects.

Note: Currently no Deployments support more than 1 replica.

Deployment pods that are handled by the **AssetDataDictionary** custom resource

The following pods are handled by the **AssetDataDictionary** custom resource object:

Table 129. Deployment pods that are handled by the **AssetDataDictionary** custom resource object

Pod name	Container type	Container name	Default replicas	Default resources
graph-store	Container	graph-store	1	<pre>resources: limits: cpu: '4' memory: 10Gi requests: cpu: '2' memory: 5Gi</pre>
maximo-connector	Container	maximo-connector	1	<pre>resources: limits: cpu: '2' memory: 4Gi requests: cpu: '1' memory: 2Gi</pre>
router	Container	router	1	<pre>resources: limits: cpu: '2' memory: 4Gi requests: cpu: '1' memory: 2Gi</pre>
series-store	Container	series-store	1	<pre>resources: limits: cpu: '2' memory: 4Gi requests: cpu: '1' memory: 2Gi</pre>
user-store	Container	user-store	1	<pre>resources: limits: cpu: '2' memory: 4Gi requests: cpu: '1' memory: 2Gi</pre>

Customer-managed

Customizing workload scale

You can scale a pod horizontally and vertically by setting `container` or `initContainer` resources or pod `replicas`. By default, some pods have more than one replica. Depending on the requirements of your workloads, you can set the values for these replicas to less than or more than the default value.

About this task

In Kubernetes, a quality of service class is assigned to every pod based on the resource requests and the limits of its component containers. The resources can be set to the following quality of service classes:

- Guaranteed
- Burstable

- BestEffort

When a node runs out of resources, Kubernetes evicts nodes in the following order:

1. BestEffort pods
2. Burstable pods
3. Guaranteed pods

For more information, see [Configure Quality of Service for Pods](#).

For more information about supported pods and their default values, see [“Supported pods” on page 648](#).



Warning: Setting pod resource limits to less than their default values can result in throttling or slowdowns. The replicas for entity manager pods, Jobs pods, and CronJobs pods are set to 1. You cannot change the pod resource limits for these pods because of the nature of their workloads.

Note: `initContainers` do not need to be modified because they are designed for the basic task of initialization. The main load is within the `Containers`.

Procedure

1. In the Red Hat OpenShift web console, from the side navigation menu, click **Administration** > **CustomResourceDefinitions** and locate the custom resource that you want to customize.
2. Click **Instances** and open the instance of the custom resource that you want to work on.
3. Click the **YAML** tab to open the editor.
4. If the `podTemplates` field does not exist, add the `podTemplates` field inside the `spec` property.

Note:

- For IBM Maximo Health, inside the `spec` property, add the `components`, `health`, and then the `podTemplates` field.
- For the IoT tool, inside the `spec` property, add the `components` section. In the `components` section, add all the component names, and then the `podTemplates` field. For example,

```
spec:
  components:
    actions: # actions is the IoT component name
    podTemplates: {} # templates for the actions component operator
```

5. Add or change the pod entry according to your requirement.
For example in the following sample code, `admin-dashboard` is the pod entry that is edited.
6. Add the field `replicas` and `container` or `initContainer` according to your requirements.

```
kind: Suite
apiVersion: core.mas.ibm.com/v1
metadata:
  name: inst1
  namespace: mas-inst1-core
  labels:
    mas.ibm.com/instanceId: inst1
spec:
  podTemplates:
    - name: admin-dashboard
      replicas: 2
      containers:
        - name: admin-dashboard
          resources:
            requests:
              cpu: 0.01
              memory: 196Mi
            limits:
              cpu: 0.8
              memory: 512Mi
```

7. Click **Save**. Wait for the operator to reconcile and apply your changes.

Customizing workload affinity

As an administrator, you can customize and improve control of the pod scheduling process. You can constrain a pod so that it is restricted or preferred to run on particular nodes.

About this task

You can restrict or prefer the pods to run on specific nodes by using `nodeAffinity`, `podAffinity`, or `podAntiAffinity` methods. For more information, see [Assigning Pods to Nodes](#).

Procedure

1. In the Red Hat OpenShift web console, from the side navigation menu, click **Administration** > **CustomResourceDefinitions** and locate the custom resource that you want to customize.
2. Click **Instances** and open the instance of the custom resource that you want to work on.
3. Click the **YAML** tab to open the editor.
4. If the `podTemplates` field does not exist, add the `podTemplates` field inside the `spec` property.

Note:

- For IBM Maximo Health, inside the `spec` property, add the `components`, `health`, and then the `podTemplates` field.
- For the IoT tool, inside the `spec` property, add the `components` section. In the `components` section, add all the component names, and then the `podTemplates` field. For example,

```
spec:
  components:
    actions: # actions is the IoT component name
    podTemplates: {} # templates for the actions component operator
```

5. Add or change the pod entry according to your requirement.
For example in the following sample code, `admin-dashboard` is the pod entry that is edited.
6. Add the field `affinity` inside of which you can add `nodeAffinity`, `podAffinity`, and `podAntiAffinity` according to your requirements.

```
kind: Suite
apiVersion: core.mas.ibm.com/v1
metadata:
  name: inst1
  namespace: mas-inst1-core
  labels:
    mas.ibm.com/instanceId: inst1
spec:
  podTemplates:
    - name: admin-dashboard
      affinity:
        nodeAffinity:
          preferredDuringSchedulingIgnoredDuringExecution:
            - weight: 1
              preference:
                matchExpressions:
                  - key: runtimeType
                    operator: In
                    values:
                      - frontend
        podAffinity:
          requiredDuringSchedulingIgnoredDuringExecution:
            - labelSelector:
                matchExpressions:
                  - key: security
                    operator: In
                    values:
                      - S1
              topologyKey: topology.kubernetes.io/zone
        podAntiAffinity:
          preferredDuringSchedulingIgnoredDuringExecution:
            - weight: 100
              podAffinityTerm:
                labelSelector:
                  matchExpressions:
```

```
- key: security
  operator: In
  values:
  - S2
topologyKey: topology.kubernetes.io/zone
```

7. Click **Save**. Wait for the operator to reconcile and apply your changes.

Customer-managed **Customizing workload tolerations**

The tolerations allow a pod to be scheduled on a node that has a matching taint. A taint is a key value pair that is assigned to a node. When a pod does not have a matching taint, it cannot be scheduled on that node.

About this task

You can apply tolerations to the pods so that pods are scheduled with the respective taints. For more information, see [Taints and Tolerations](#).

Procedure

1. In the Red Hat OpenShift web console, from the side navigation menu, click **Administration > CustomResourceDefinitions** and locate the custom resource that you want to customize.
2. Click **Instances** and open the instance of the custom resource that you want to work on.
3. Click the **YAML** tab to open the editor.
4. If the `podTemplates` field does not exist, add the `podTemplates` field inside the `spec` property.

Note:

- For IBM Maximo Health, inside the `spec` property, add the `components`, `health`, and then the `podTemplates` field.
- For the IoT tool, inside the `spec` property, add the `components` section. In the `components` section, add all the component names, and then the `podTemplates` field. For example,

```
spec:
  components:
    actions: # actions is the IoT component name
    podTemplates: {} # templates for the actions component operator
```

5. Add or change the pod entry according to your requirement.
For example in the following sample code, `coreapi` is the pod entry that is edited.
6. Add the field `tolerations` according to your requirements.

```
- name: coreapi
  tolerations:
  - key: "onlycoreapi"
    operator: "Exists"
    effect: "NoSchedule"
```

7. Click **Save**. Wait for the operator to reconcile and apply your changes.

Customizing hostAliases in podTemplates

Map a hostname to an IP address, so you can bypass DNS resolution and create direct mapping between the hostname and IP address.

About this task

To map a hostname to an IP address, configure the `hostAliases` field in the `podTemplates` custom resource definition. You can complete the configuration for the following types of pods.

- Deployment pods
- Cron job pods

- Job pods

IBM Maximo Application Suite supports LDAPS, which means verification of the LDAP server's certificate is necessary. The certificates are issued to hostnames and not IP addresses. If the connection is made with an IP address instead of the hostname, the certificate check fails and IBM Maximo Application Suite rejects the LDAP connection.

With IBM Maximo Application Suite, in podTemplates use the hostAliases setting to map the LDAP server's hostname and IP address, for the supported custom resource definitions (CRD). The hostAliases mapping helps ensure that the hostname on the certificate matches the IP address that IBM Maximo Application Suite uses to connect.

LDAP is not the only scenario where hostAliases are needed. Many users require similar configurations for SMTP, peer-to-peer connectivity, or VPN setups where DNS resolution is unavailable or unreliable.

As the administrator, complete the following steps to customize hostAliases in podTemplates.

Procedure

1. In the Red Hat OpenShift web console, open your IBM Maximo Application Suite instance, which is also known as a IBM Maximo Application Suite project.
2. From the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
3. Open the custom resource where you want to customize the mapping a hostname to an IP address. For example, you want to add a hostname-IP map in the entitymgr-idpcfg pod. Lookup supported pods to identify the custom resource that handles podTemplates for the entitymgr-idpcfg pod. Since it is at the suite level, lookup the Suite custom resource.
4. Click **Instances** and open the instance of the custom resource that you want to work on.
5. Click the **YAML** tab to open the editor.
6. In the spec property, add the podTemplates field, if the podTemplates field does not exist.
7. In podTemplates field, add in the following details. Replace the values for ip and hostnames.

```
spec:
  podTemplates:
    - name: coreapi
      hostAliases:
        - ip: "10.10.1.1"
          hostnames:
            - "ldap1.com"
            - "ldap2.com"
```

8. Click **Save**. Wait for the operator to reconcile and apply your changes.
9. Verify your configuration changes.
 - a) From the side navigation menu, click **Administration**.
 - b) Click **Deployments, CronJobs, or Jobs**.
 - c) In the **Name** field, enter the pod name that you configured.
 - d) In the search results list, select your pod name.
 - e) Click the **YAML** tab and verify the configuration.

Enabling defaultJMS for Maximo Manage

Enable the property defaultJMS if you want to use the Java Messaging Service (JMS) provider that is out of the box in Maximo Manage.

Before you begin

If you complete the following procedure, the related server bundles for the Maximo Manage application stop and restart. Therefore, you must schedule downtime for your Maximo Manage environment before you complete the procedure.

About this task

In some situations, you might want to configure your own message service. However, if you want to use the JMS that is out of the box, you must enable the property `defaultJMS`. As the administrator, complete the steps in the following procedure.

Procedure

1. In the Red Hat OpenShift web console, open your IBM Maximo Application Suite instance, which is also known as a IBM Maximo Application Suite project.
2. From the side navigation menu, click **Administration** > **CustomResourceDefinitions**.
3. In the **CustomResourceDefinitions** page, filter for the `ManageWorkspace` CR.
4. Open the `ManageWorkspace` CR.
5. Click the **Instances** tab and open the instance that you want to work on.
6. Click the **YAML** tab.
7. In the **Find** field, enter `defaultJMS`.
8. In the results, locate the `defaultJMS` property within **spec** > **settings** > **deployment**.
9. Change the `defaultJMS` property value to `true`.
10. Click **Save**. Wait for the operator to complete the following actions and apply your changes.
 - Rebuild bundle images.
 - Apply your changes to the bundle images.
 - Stop and restart the related server bundles in **Workloads** > **Pods**.

Configuring the user interface

You can change the appearance of the user interface by adding notifications to the login page, including your own branding and custom visual style, or hiding guided tours.

Customer-managed **Updating the user interface**

Note: If you are using a Maximo Application Suite as a Service environment that is based on Maximo Application Suite 9.0, this feature is not available.

Starting in Maximo Application Suite 9.0, you can change the appearance of the user interface by relabeling the common header and apply custom styles by using advanced CSS customizations.

About this task

To update the common header and the Maximo Application Suite CSS, you can use the User interface customization option on the **Configurations** page to implement the following changes:

Customized branding

You can replace the IBM and Maximo Application Suite branding on the header and also remove IBM brand styles by using CSS overrides, applying colors, typographic styles to any visual element.

Environment identification

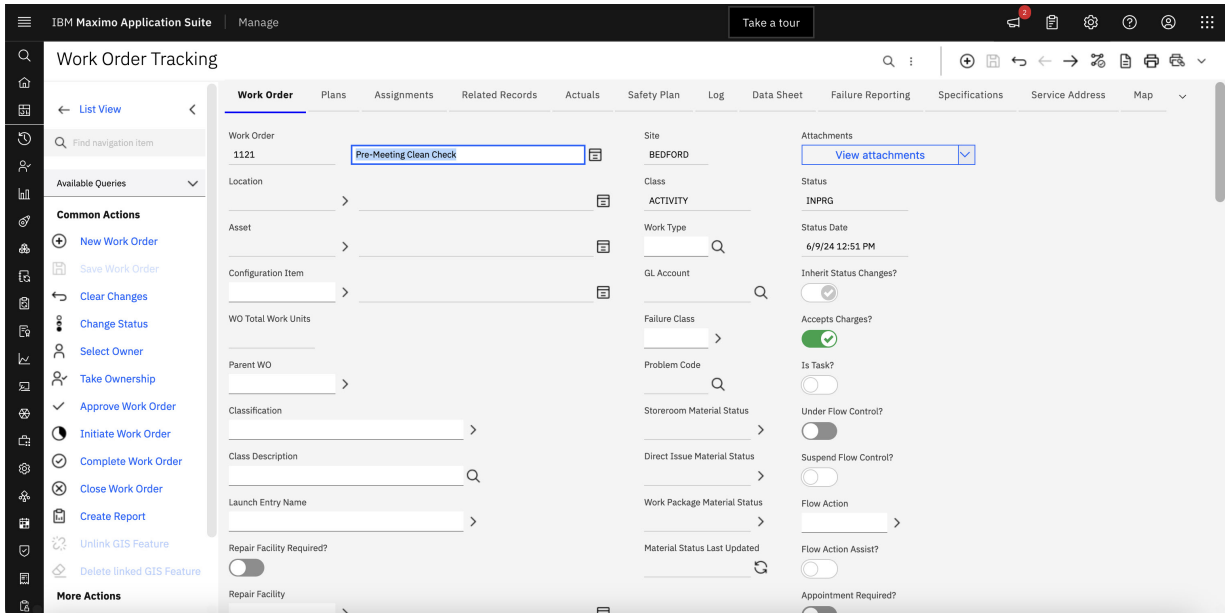
Because these customizations are instance-specific, one set of styles can be applied to development or test environments, including labeling the name of the application, to differentiate them from production environments.

Style adjustments

Similar to all IBM products, the Maximo Application Suite user interface uses the Carbon Design System, which is a visual and UX design language. For more information, see, [Carbon Design System](#).

By using the user interface customization feature, you can change the style of the current Carbon skin. For example, you can change the skin in Maximo Manage to reduce the font size and required scrolling. Previously, Maximo Manage supported several skins, but in Maximo Application Suite all applications use a single Carbon-based skin. Combined with existing UI system properties for Maximo Manage, you can change the current skin in Maximo Manage to be more similar to an older skin.

To get started, you can use the `Manage_Overrides_template.css` template file, which includes multiple key CSS classes and example customizations. The following image illustrates the updated skin for Maximo Manage and shows a smaller font size and reduced spacing.



Note: This template is an example of changes that you can implement and is not officially supported by IBM Support.

It might be necessary to customize individual applications. For example, customizing Maximo Manage does not impact Maximo Health.

You can update the user interface by applying changes on the **Header configuration** and **CSS customization** pages:

Header configuration

You can replace the IBM and Maximo Application Suite branding with your own company branding and also upload a formatted company logo.

CSS customization

Cascading Style Sheets (CSS) is code that defines the visual appearance of the user interface. With this customization, a single `override` code block is loaded on each page, which you can define to change the appearance of any element in Maximo Application Suite. These changes include colors, spacing, font sizes, and typography.

CSS classes are shared across pages and applications. Any CSS change to an application is considered a customization to code and might cause unexpected changes to any application in the Suite.

Considerations

- CSS customization is not supported by IBM. Consider whether changes that are made by CSS customization cause any issues before you open support cases. The purpose of user interface customization is only to streamline the delivery of these types of changes.
- CSS classes can change from one release to another, particularly in the newer, role-based applications. After you upgrade from Maximo Application Suite 9.0 to a later version, test your customization again when the upgrade is complete.
- CSS classes might affect multiple pages or applications.

- Customizations that are made do not impact Maximo Monitor or Maximo Mobile applications from the app stores in Maximo Application Suite 9.0.

Procedure

1. On the **Suite administration** page, from the side navigation menu, click **Configurations** and then click **User interface customization**.
2. To show your own logo and company branding in the header, update the header configuration.
 - a) On the **Header configuration** tab, enter your company name and product name.
 - b) Upload a .svg or .png file that shows before the company name in the header.
The maximum file size is 1 MB, and the optimum height is 40 pixels or less.
3. To change the appearance of the user interface by using CSS, update the current CSS.
 - a) On the **CSS customization** tab, set **Enable CSS customization** to on.
 - b) Enter your CSS changes. You can also upload the CSS changes.
To search the current CSS and find classes, you can use the developer tools in your browser.
For example, to change the color of the header class from black to dark cyan, add the following syntax:


```
.bx--header {
  background-color: darkcyan;
}
```
 - c) To disable changes and return to the default CSS, you can set **Enable CSS customization** to off.
4. Save your changes.

Example

If you want to change the appearance of the login page to show your company name and a different image, you can update the CSS and header. For more information, see [“Example: Changing the login page” on page 710](#).

What to do next

If you update the CSS and custom changes do not show, you must configure the content security policy (CSP).

Customer-managed

Example: Changing the login page

Note: If you are using a Maximo Application Suite as a Service environment that is based on Maximo Application Suite 9.0, this feature is not available.

In Maximo Application Suite 9.0, a system administrator wants to change the appearance of the login page to include their company image and replace the title heading on the login page to their company name.

Procedure

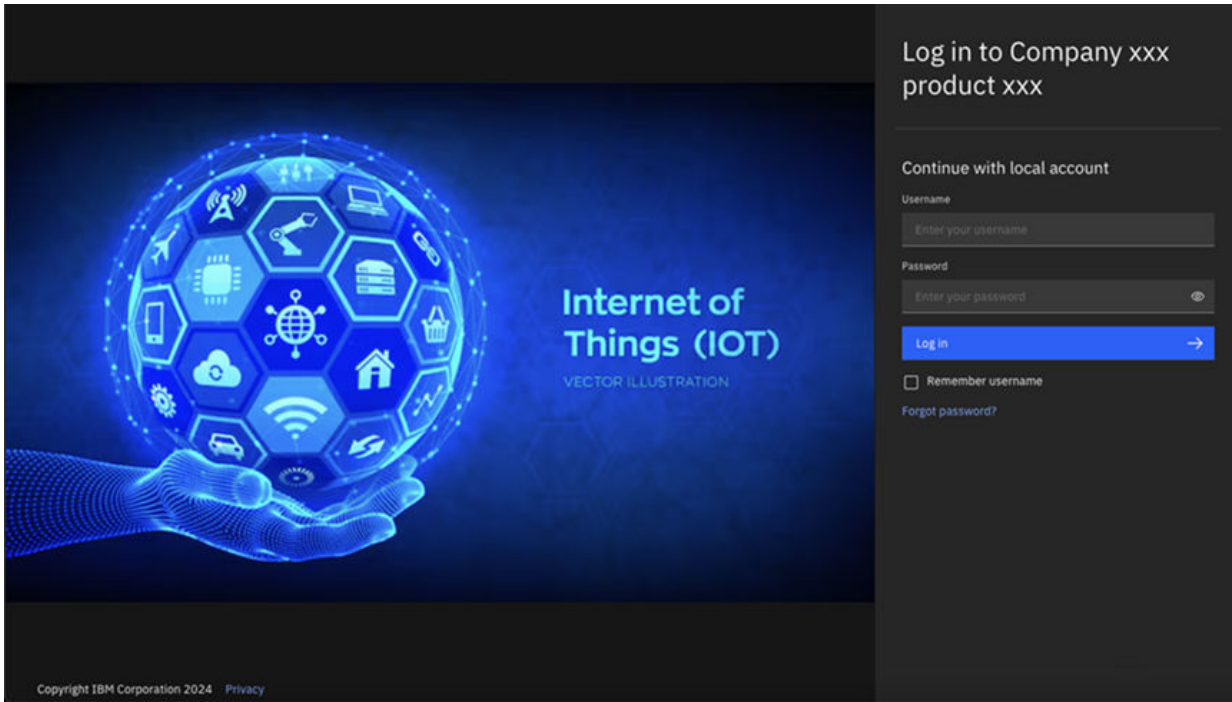
1. On navigation menu, in the suite administration page, select **Configurations > User interface customization**.
2. Select the **Header configuration** tab and enter your company name and product name in the corresponding fields.
3. Select the **CSS customization** tab and set **Enable CSS customization** to on.

4. To add an image to the login page, enter the following CSS update to add the image as the background image.

```
.login-background-image {  
    background-image: url("image_url")  
}
```

where *image_url* is the URL of the image that you want to use.

For example, after you save your changes, the login page shows the new image and the name of your company.

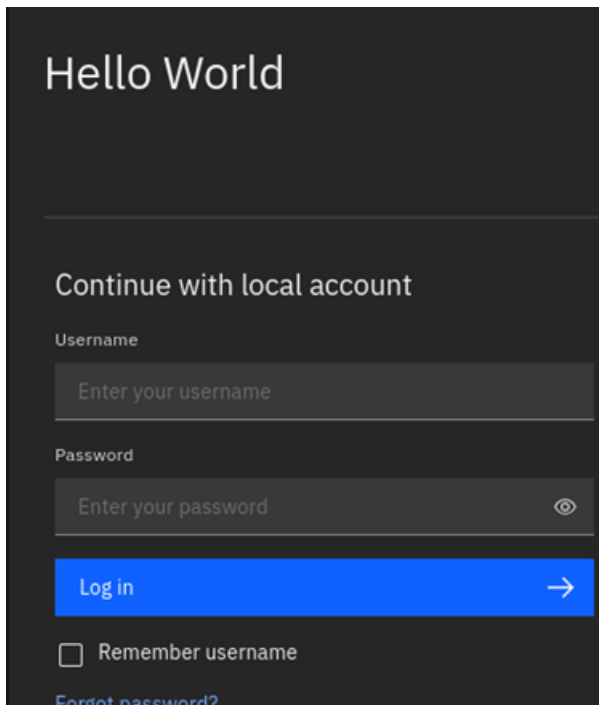


5. If you want to change the title text on the login page, enter the following CSS update

```
p#x7ynq::after { visibility: visible; content: "Hello World"; position: absolute; left: 0;  
top: 0; } p#x7ynq { visibility: hidden; position: relative; }
```

where *Hello World* is the text that you are updating to.

For example, after you save your changes, the login page shows the updated text.



Customer-managed

Updating the content security policy in Maximo Manage

Note: If you are using a Maximo Application Suite as a Service environment that is based on Maximo Application Suite 9.0, this feature is not available.

If you update the CSS for Maximo Manage 9.0 and later and custom changes do not show, you must configure the content security policy (CSP) for Maximo Manage.

Before you begin

Determine the base domain of your Maximo Manage server. To apply CSS overrides and custom images, the content security policy must include the base domain of the Maximo Manage server. For example, if your host is `main.mainage.mycompany.com` then the base domain is `mycompany.com`.

Procedure

1. In Maximo Manage, open the System Properties application.
2. Select **Filter** and search for the **mxr.sec.header.Content_Security_Policy** property.
3. In the **Global Value** field of the content security policy property, add `*.base-domain` where *base-domain* is the base domain of your host.
 - a) Add your base domain immediately after `style-src.`
 - b) Add your base domain after `img-src.`
For example, if your base domain is `mycompany.com`, add `*.mycompany.com`.
4. Save your changes.
 - a) From the **Common Actions** menu, click **Save Property**.
 - b) Select the checkbox for the property name that you updated and then click **Live Refresh**.
 - c) In the **Live Refresh** window, confirm the changes that you made and click **OK** to refresh.

Related tasks

[Updating the user interface](#)

Enabling login notification

Starting in Maximo Application Suite 8.11, you can create and display a system message on the login page to provide security and privacy information to users.

Procedure

1. Enable a login notification on the **Authentication** page.
 - In Maximo Application Suite 9.1, from the side navigation menu, select **Suite > Administration > Authentication**.
 - In Maximo Application Suite 9.0 and earlier, on the **Suite administration** page, select **Users** from the side navigation menu and then click the **Authentication** tab.
2. In the Login notification section, enable the message and enter the information.
When you save your changes, the message is shown on the login page.



Attention:

If you configure SAML as the default identity provider and enable seamless login so that users authenticate to the login page that uses the SAML identity provider, the Maximo Application Suite login page is not shown. If you need to display a security message to comply with federal regulations, ensure that seamless login is disabled. Otherwise, users do not see any system notification that might be shown on the login page. For more information, see [configuring default identity providers](#).

3. To disable the login notification, click the toggle to remove the message.
Disabling the login notification deletes the message content and removes the message on the login page.

Related tasks

[Configuring default identity providers](#)

If you configure more than one identity provider, such as LDAP or SAML, you can specify which identity provider is the primary login option for users by updating authentication options on the **Suite administration** page. Alternatively, you can update the custom resource file in Red Hat OpenShift Container Platform.

Disabling or hiding login options by using APIs

Starting in Maximo Application Suite, a system administrator or an identity provider (IdP) administrator can hide or disable identity providers so that they are not accessible to users. Administrators can continue to maintain the existing IdP configuration and account links.

An administrator may choose to hide or disable an authentication option to enhance security by removing outdated or vulnerable methods, enforce compliance with organizational or regulatory requirements, streamline user experience by reducing complexity, or centralize identity management through a preferred provider. Also, these actions can prevent phishing attacks, ensure correct logging and traceability, and avoid incompatibility with existing infrastructure. Sometimes, cost control or transitioning to newer authentication systems also necessitates disabling older or less secure options.

About this task

You can change the login options by using APIs for the supported IdP types:

- Local
- LDAP
- SAML
- OIDC

As an alternative, you can change the login options by using custom resource (CR) definitions. For more information, see [“Disabling or hiding login options in custom resource” on page 714](#).

Procedure

1. To hide or show an IdP option on the login page by using APIs, use the PUT `https://api.<MAS domain>/sso/idps/<idpId>/visible` including the `{"isVisible": }` attribute in the body.

Option	Action
To hide an IdP login option, such as default-ldap	Run the following command. <pre>curl -H "Content-Type: application/json" --location \ --request PUT https://api.<MAS domain>/sso/idps/default-ldap/visible \ --header "x-access-token: <system admin or idp admin access token>" \ --insecure \ --data '{ "isVisible": false }'</pre>
To show an IdP login option	Change the "isVisible" attribute to <code>"isVisible": true</code> and run the command.

2. To disable or enable an IdP option on the login page by using APIs, use the PUT `https://api.<MAS domain>/sso/idps/<idpId>/enable` including the `{"isEnabled": }` attribute in the body.

Option	Action
To disable an IdP login option, such as local	Run the following command. <pre>curl -H "Content-Type: application/json" --location \ --request PUT https://api.<MAS domain>/sso/idps/local/enable \ --header "x-access-token: <system admin or idp admin access token>" \ --insecure \ --data '{ "isEnabled": false }'</pre> <p>Note: When the local IdP is disabled, it is still possible to access the local dedicated page <code>https://auth.<MAS domain>/idplogin/loginpage?idp=local</code>, and log in with super user credentials.</p>
To enable an IdP option	Change the "isEnabled" attribute to <code>"isEnabled": true</code> and run the command.

Results

When you hide an IdP option, the IdP is not listed as an available login option on the login page. However, automation scripts that perform authentication through APIs can authenticate in Maximo Application Suite using that IdP.

When you disable an IdP option, the IdP is not listed as an available login option in the login page. It is also not possible to authenticate with that login option in automation scripts by using APIs. Processing can take up to 5 minutes for the login page to reflect the change in IdP visibility or disabled and enabled status.

Disabling or hiding login options in custom resource

Starting in Maximo Application Suite, a system administrator or an identity provider (IdP) administrator can hide or disable identity providers so that they are not accessible to users. Administrators can continue to maintain the existing IdP configuration and account links.

An administrator may choose to hide or disable an authentication option to enhance security by removing outdated or vulnerable methods, enforce compliance with organizational or regulatory requirements, streamline user experience by reducing complexity, or centralize identity management through a preferred provider. Also, these actions can prevent phishing attacks, ensure correct logging and traceability, and avoid incompatibility with existing infrastructure. Sometimes, cost control or transitioning to newer authentication systems also necessitates disabling older or less secure options.

About this task

You can change the login options by using APIs for the supported IdP types:

- Local
- LDAP
- SAML
- OIDC

As an alternative, you can change the login options by using APIs. For more information, see [“Disabling or hiding login options by using APIs”](#) on page 713.

Procedure

1. In the suite CR, update the `spec.settings.sso.enabled` property or the `spec.settings.sso.visible` property.

Note: When you configure for the first time, you might need to add the `enabled` and `visible` attributes to the CR before you can apply the changes.

- a) In the Red Hat OpenShift Container Platform console, from the side navigation menu, in the Administration section, select **CustomResourceDefinitions**.
- b) On the **CustomResourcesDefinitions** page, select the Suite CRD.
- c) On the **Instances** tab, select the instance that you want to update.
- d) On the **YAML** tab, change the property values.

Option	Action
To disable the login option for local users	Set <code>spec.settings.sso.enabled</code> to <code>false</code> .
To enable the login option local users	Set <code>spec.settings.sso.enabled</code> to <code>true</code> .
To hide the login option local users	Set <code>spec.settings.sso.visible</code> to <code>false</code> .
To show the login option local users	Set <code>spec.settings.sso.visible</code> to <code>true</code> .

2. In the IDPCfg CR of the specific IdP that you want to change, change login options for other IdPs, LDAP, SAMP, or OIDC by updating the `spec.enabled` property or the `spec.visible` property.

Note: When you configure for the first time, you might need to add the `enabled` and `visible` attributes to the CR before you can apply the changes.

- a) In the Red Hat OpenShift Container Platform console, from the side navigation menu, in the Administration section, select **CustomResourceDefinitions**.
- b) On the **CustomResourcesDefinitions** page, select the IDPCfg CRD.
- c) On the **Instances** tab, select the instance that you want to update.
- d) On the **YAML** tab, change the property values.

Option	Action
To disable the login option	Set <code>spec.enabled</code> to <code>false</code> .
To enable the login option	Set <code>spec.enabled</code> to <code>true</code> .
To hide the login option	Set <code>spec.visible</code> to <code>false</code> .

Option	Action
To show the login option	Set <code>spec.visible</code> to <code>true</code> .

Results

When you hide an IdP option, the IdP is not listed as an available login option on the login page. However, automation scripts that perform authentication through APIs can authenticate in Maximo Application Suite using that IdP.

When you disable an IdP option, the IdP is not listed as an available login option in the login page. It is also not possible to authenticate with that login option in automation scripts by using APIs. Processing can take up to 5 minutes for the login page to reflect the change in IdP visibility or disabled and enabled status.

Hiding guided tours

Starting in Maximo Application Suite 9.0, you can configure the user interface to hide guided tours. Maximo Application Suite provides guided tours to help users learn more about different tasks and updates in the product.

About this task

Guided tours provide content, such as interactive tours, links to videos, and what's new information. When enabled, this content is available by clicking **Take a tour**. If disabled, this content and the **Take a tour** button are not available.

Starting in Maximo Application Suite 9.1, access to this content is available from the Help menu by selecting **Show me how**. If disabled, this content and the **Show me how** option are not available.

Procedure

1. From the **Suite administration** menu, click **Configurations** and then in the Other section, click **Guided tours**.
2. Set **Guided tours** to disable.
3. Save your changes.

Results

The content is no longer available. If you want to view the guided tours later, you can set **Guided tours** to enable.

Managing user profile

Starting in IBM Maximo Application Suite 9.1, edit your user information, change password, or update locale and region settings.

Procedure

1. In the Maximo Application Suite **User profile** page, click **Edit profile**.
2. Select the following tabs to update your information or preferences.
 - In the **User information** tab, update your user details, contact information such as email and phone numbers, and address.
 - In the **Change password** tab, update your password according to your company policies.
 - In the **Language and region** tab, specify your preferred locale, language, and time zone. The preferences that you select override your browser settings.

For more information, see [“Language and locale support” on page 127](#).

Related concepts

Language and locale support

IBM Maximo Application Suite supports the use of preferred language and locale for the Maximo Application Suite user interfaces. The preferences that are applied override the language and locale settings for the browser that is used to access Maximo Application Suite.

Disabling surveys

Users can complete the customer satisfaction survey to submit feedback about their experience with Maximo Application Suite. As an administrator, you can configure the user interface to hide the survey.

About this task

When enabled, the survey is available to all users to submit feedback.

Procedure

1. From the side navigation menu, select **Suite > Administration > Configurations**.
2. In the Other section, click **Survey** and set to disable.
3. Save your changes.

Results

The survey is no longer available. If you want to allow users to submit the survey later, you can set to enable.

Customer-managed

Configuring cross-origin resource sharing (CORS)

Cross-origin resource sharing (CORS) in Maximo Application Suite enables secure access for communication between various resources from different domains. You can implement CORS in Maximo Application Suite by updating the custom resource file in Red Hat OpenShift Container Platform.

About this task

Starting in Maximo Application Suite 9.1, you can also configure the Suite custom resource to implement CORS for Maximo Application Suite identity providers and extend public APIs configuration to support the following fields:

- allowedHeaders
- allowedMethods
- exposeHeaders
- maxAge

Procedure

1. In the Red Hat OpenShift Container Platform console, in the Administration section, select **Custom Resource Definitions**.
2. In the **CustomResourcesDefinitions** window, select the Suite CR file.
3. On the **Instances** tab, select the instance that you want to update.
4. On the **YAML** tab for the instance, in the `spec.settings.cors` section, enter the domain for the `allowedOrigins` property.
For example, to include the domain `https://developer.ibm.com`, specify the following domain information.

```
apiVersion: core.mas.ibm.com/v1
kind: Suite
metadata:
```

```

...
spec:
...
  settings:
...
    cors:
      allowedOrigins:
        - https://developer.ibm.com
...

```

5. To configure CORS for an identity provider specify the values for the following properties in the `spec.settings.cors` section.

Field property	Description
<code>allowedHeaders</code>	Defines a comma-separated list of HTTP headers that can be included in requests from the allowed origins. Setting the property to <code>*</code> permits all headers.
<code>allowedMethods</code>	The HTTP methods that a client is allowed to when it makes requests to the endpoint. The value is set to GET.
<code>exposeHeader</code>	Lists HTTP headers that can be exposed to the requesting client so that certain headers are accessible when API responses are processed.
<code>maxAge</code>	Determines how long the browser can cache the results of a preflight request (OPTIONS request), improving performance by reducing unnecessary preflight requests.

For example, specify the following CORS configuration information for WebSphere Liberty server.

```

spec:
  ....
  settings:
    ....
    cors:
      allowedHeaders:
        - 'Authorization, Content-Type, Origin, Access-Control-Request-Method, Access-
Control-Request-Headers'
      allowedMethods:
        - GET
        - POST
      allowedOrigins:
        - allowedOrigin
      exposeHeaders:
        - X-Custom-Header
      maxAge: 600
    ....

```

6. Save the custom resource changes.

Using single-node Red Hat OpenShift clusters

A single-node cluster in Red Hat OpenShift consists of a single control plane node that is configured to run workloads. This configuration offers both control node and worker node functions, so you can deploy a smaller Red Hat OpenShift environment with or without dependence on a centralized management cluster. A single-node cluster can run autonomously when needed, which is useful for resource-constrained environments, demos, proofs of concept, or on-premises deployments.

By deploying a single-node cluster, you can experience the benefits of Red Hat OpenShift in a more compact environment that requires fewer resources. This cluster provides an efficient way to test new features or applications in a controlled environment. However, a single-node Red Hat OpenShift cluster lacks high availability, so it might not be suitable for mission-critical workloads that require constant uptime.

When to use single-node Red Hat OpenShift clusters

A single-node Red Hat OpenShift cluster can be implemented in a production environment when redundancy, scalability, and high availability are not critical requirements.

For production environments, edge sites, or scenarios where Red Hat OpenShift clusters are required, but high availability is not critical, a single-node Red Hat OpenShift cluster can be an appropriate solution.

The following situations might be appropriate for a single-node Red Hat OpenShift cluster:

- Small implementations that contain only IBM Maximo Application Suite and IBM Maximo Manage for up to 70 concurrent users.
- Satellite or disconnected deployments, which are possibly connected to a large deployment of Maximo Application Suite.
- Upgrading small Maximo deployments to Maximo Application Suite.
- Demos and proofs of concept.

Installation prerequisites

Before you install a single-node Red Hat OpenShift cluster, review the installation prerequisites.

Testing of a single-node Red Hat OpenShift cluster is performed with the following CPU and RAM configuration.

Red Hat OpenShift

The same version that is supported by IBM Maximo Application Suite version 8.9 or later.

vCPU

16 cores

RAM

64 GB

Primary/Ephemeral storage

Minimum 300 GB

IBM entitlement key

To get the entitlement key, log in to the [IBM Container Library](#) with a user ID that has software download rights for your company's IBM Passport Advantage entitlement.

Red Hat OpenShift pull secret file

You can download this file from <https://access.redhat.com/management>. You need a valid Red Hat account to download the file.

IBM Maximo Application Suite license file `license.dat`

Access the IBM License Key Center. From the **Get Keys** menu, select **IBM AppPoint Suites**. Select **IBM MAXIMO APPLICATION SUITE AppPOINT LIC**. For more information, see [Maximo Application Suite License File](#).

Docker or Podman

Docker or Podman can be used.

Bare metal or VMware vSphere prerequisites

For more information, see [Requirements for installing Red Hat OpenShift on a single node](#)

Preparing a Docker container

To install a single-node Red Hat OpenShift cluster, you first set up a Docker container and copy the pull secret file and license file into the container.

Procedure

1. Set up a Docker container.

```
mkdir ~/sno
cd ~/sno
```

```
docker pull quay.io/ibmmas/cli
docker run -dit --name sno quay.io/ibmmas/cli:latest bash
```

2. Log in to the Docker container, create a folder for the IBM Maximo Application Suite configuration, and then exit the container.

```
docker exec -it sno bash
mkdir masconfig
exit
```

3. Copy the pull secret and Maximo Application Suite license file into the Docker container.

```
docker cp pull-secret sno:/mascli/masconfig/pull-secret
docker cp license.dat sno:/mascli/masconfig/license.dat
```

Installing on bare metal or VMware vSphere

To install on bare metal or VMware vSphere, you can generate the discovery ISO by using the Assisted Installer, manually generate the discovery ISO, or install by using USB media. You can also monitor the installation.

For more information, see [Installing OpenShift on a single node](#).

Installing logical volume manager storage

Local storage in Red Hat OpenShift means that storage devices or file systems are available locally on a node server. You need to provide the cluster with a storage class and related provisioner.

Procedure

1. Log in to the Red Hat OpenShift web console.
2. Install the LVM storage operator.
 - a) From the side navigation menu, click **Operators > OperatorHub**.
 - b) Search for LVM.
 - c) Click the **LVM Storage provided by Red Hat** tile.
 - d) Install the operator by using the default parameters.
3. Create the cluster for the LVM storage.
 - a) When the operator is ready for use, click **Create LVMCluster**.
 - b) Keep the default values and click **Create**.
 - c) Wait for the LVMCluster status to become ready.
4. Configure the YAML for the LVM storage.
 - a) From the side navigation menu, click **Home > Search**.
 - b) Click the **Resources** menu and enter `config` in the search field.
 - c) Click **Config imageregistry.operator.openshift.io/v1**.
 - d) Click **cluster** and on the **YAML** tab, edit the YAML.
 - i) Modify `rolloutStrategy` from this text:

```
rolloutStrategy:RollingUpdate
```

to this text:

```
rolloutStrategy:Recreate
```

- ii) Save and then modify storage from this text:

```
storage: {}
```


to this text:

```
storage:
  pvc:
    claim: ''
```

iii) Save and then modify managementState from this text:

```
managementState: Removed
```

to this text:

```
managementState: Managed
```

iv) Save your changes.

5. Verify the status of the image-registry-storage PVC.

- a) From the side navigation menu, click **Storage > PersistentVolumeClaims**.
- b) Search for the image-registry-storage persistent volume claim (PVC).
- c) Check the status of the image-registry-storage PVC.

The successful status is bound. If the status is pending, you need to troubleshoot the issue. For more information, see [“Troubleshooting the image-registry-storage persistent volume claim” on page 721](#).

Installing Maximo Application Suite and Maximo Manage

You can install IBM Maximo Application Suite and IBM Maximo Manage by using the Maximo Application Suite command-line interface (CLI) utility.

This utility provides commands to manage the local Docker registry, configure policies in the Red Hat OpenShift cluster, and deploy Maximo Application Suite. You can choose a noninteractive or interactive installation. For more information, see [“Standard installation with IBM Maximo Application Suite CLI” on page 218](#).

Troubleshooting the image-registry-storage persistent volume claim

To enable building and pushing images, the image-registry-storage persistent volume claim (PVC) must be in the bound status. If the image-registry-storage PVC is in pending status, update the image-registry-storage PVC.

Procedure

1. Log in to the Red Hat OpenShift web console.
2. From the side navigation menu, click **Storage > PersistentVolumeClaims**.
3. Search for the image-registry-storage PVC.
4. If the status is pending, click the image-registry-storage PVC name and on the **YAML** tab, download the YAML.
5. On the **YAML** tab, edit the YAML.
 - a. Remove the metadata fields uid, resourceVersion, and creationTimestamp.
 - b. Remove the manageFields section.
 - c. Remove the status section.
 - d. Modify accessModes from ReadWriteMany to ReadWriteOnce.

The following example shows the YAML after the changes.

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage
```

```

namespace: openshift-image-registry

annotations:
  imageregistry.openshift.io: 'true'
  volume.beta.kubernetes.io/storage-provisioner: topolvm.cybozu.com
  volume.kubernetes.io/selected-node: 00-50-56-9d-4e-28
  volume.kubernetes.io/storage-provisioner: topolvm.cybozu.com
finalizers:
- kubernetes.io/pvc-protection
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage:
        100Gi
  storageClassName: odf-lvm-vg1
  volumeMode: Filesystem

```

6. From the side navigation menu, click **Storage > PersistentVolumeClaims**.
7. In the row for the image-registry-storage PVC, click the three-dot icon and then click **Delete PersistentVolumeClaim**.
8. Verify that the image-registry-storage PVC is deleted.
 - a) From the side navigation menu, click **Storage > PersistentVolumeClaims** and check the status of the image-registry-storage PVC.
 - b) If the status is terminating, click **image-registry-storage** and then click the YAML tab.
 - c) Delete the `finalizer` section from the YAML.
 - d) Click **Save**.
9. In the **Project** menu, confirm that the project is still openshift-image-registry.
10. Create a PVC.
 - a) Click **Create PersistentVolumeClaim**.
 - b) Click the **Edit YAML** link. Replace the content of the YAML with the YAML you edited.
 - c) Click **Create**.

The new PVC immediately goes into the bound state.

Administering Maximo Application Suite

As a suite administrator, you can manage users, their access and entitlements, and monitor AppPoints usage on the suite administration and security pages.

Starting in Maximo Application Suite 9.0 and earlier, you access the **Suite administration** page by clicking the **Administration** icon in the menu bar.

Starting in Maximo Application Suite 9.1, access the **Administration** and **Security** pages on the side navigation menu.

You can also use the URL `https://admin.<mas_domain>` to access the administration page.

Performance optimization for IBM Maximo Application Suite

IBM Maximo Application Suite depends on various components for its operations. You can find the components that are listed below can improve performance for IBM Maximo Application Suite.

Key microservices and dependencies to scale for Maximo Application Suite core

Important services that are needed for user login, authentication, application management, metrics, and licensing are contained in the Maximo Application Suite core namespace.

Scale the following key microservices in the Maximo Application Suite core namespace to meet the growth in the number of concurrent IBM Maximo Application Suite users. Increase the memory and CPU limits for the pods to scale the microservices as the number of users grow. Microservices in the MAS core namespace support podTemplates, which enable manual scaling by the administrator, but they do not support autoscaling.

- MongoDB is used extensively by the coreidp pod, api-licensing pod, adoption usage pod, or other Maximo Application Suite and Suite License Service micro services.
- Maximo Application Suite core namespace
 - Coreidp pods
 - Licensing-mediator pods
 - Coreapi pods

To decrease the load on the coreapi pods, encourage users to log in directly to IBM Maximo Application Suiteso that the **Suite navigator** page is bypassed.

- Suite License Service namespace
 - Api-licensing pods
- Kubernetes apiserver pods. The coreapi pods issue Kubernetes API calls to retrieve information from Maximo Application Suite application custom resource files and config maps.

Key MongoDB metrics

Monitor memory usage, CPU usage, and average read and write latency to identify scale problems. If MongoDB is not scaled properly, it can become a bottleneck when the number of concurrent users grows. Liveness probes are used to determine when a container needs to be restarted. For instance, they detect a deadlock situation, where an application is still running but unable to continue functioning properly. If a container continuously fails its liveness probe, the kubelet restarts it. If the liveness probes time out and results in a pod restart, your MongoDB cluster is likely undersizedFor more information, see [Liveness, Readiness, and Startup Probes](#)

Monitor the following MongoDB metrics to identify scaling problems.

Memory usage

Monitoring memory usage is crucial for MongoDB performance. By default, MongoDB tries to cache the active dataset in memory by using the Wired Tiger cache. If you notice a high number of cache evictions or if MongoDB servers are terminated by the `oomkiller` process for using too much memory, consider increasing the memory that is allocated to the MongoDB server.

CPU usage

Help ensure that the MongoDB servers are not reaching their allocated CPU limits. Consistently high CPU usage might lead to performance issues, so regularly check this metric to avoid bottlenecks.

Average read and write latency

Ideally, the average read and write latency needs to be under 50 milliseconds. If the latency exceeds this threshold, it might indicate an undersized MongoDB cluster. Check whether the MongoDB cluster has sufficient memory allocation and verify disk performance.

Lock waiters

A large number of lock waiters is a sign of contention on collections or documents in MongoDB and might lead to performance degradation. Monitor and address any excessive lock wait times.

Scaling MongoDB community edition

Scale MongoDB based on the number of concurrent users and the login rate. When you scale the MongoDB community edition, specify the CPU and memory limits in the MongoDBCommunity custom resource.

The following example shows how to use the MongoDBcommunity custom resource to scale CPU and memory limits.

- Log in to the Red Hat OpenShift Container Platform Console.
- Under **Administration**, Select **CustomResourceDefinitions**
- Select the **MongoDBCommunity** CustomResourceDefinition.
- Navigate to the Instances section and Select the **mas-mongo-ce** instance.
- On the YAML tab, edit the CPU and memory limits for name:mongod, You can use the values in the CPU and memory limit guidelines for scaling table.

```
spec:
  statefulset:
    spec:
      template:
        spec:
          containers:
            - name: mongod
              resources:
                limits:
                  cpu: <cpu limit>
                  memory: <mem limit>
```

The following table provides guidelines for scaling the MongoDB community based on the number of concurrent users and login rate

<i>Table 130. CPU and memory limit guidelines for scaling MongoDB community</i>		
Login rate (logins per minute)	MongoDB CPU limit	MongoDB memory limit (GB)
75	2	4
150	2	4
300	4	8
600	8	12
1200	12	16

Scaling coreidp service in the Maximo Application Suite core namespace

The following table provides guidelines for scaling the coreidp service based on the number of concurrent users and login rate.

<i>Table 131. CPU and memory limit guidelines for scaling coreidp service</i>			
Login rate (logins per minutes)	coreidp replicas	coreidp CPU limit	coreidp memory limit
75	1	6	1
150	1	6	1
300	1	6	1
600	2	6	2
1200	4	6	3

Scaling the licensing-mediator service in the Maximo Application Suite core namespace

The following table provides guidelines for scaling the licensing-mediator service based on the number of concurrent users and login rate.

Table 132. CPU and memory limit guidelines for scaling licensing-mediator service

Login rate (logins per minute)	licensing-mediator replicas	licensing-mediator CPU limit	licensing-mediator memory limit (GB)
75	1	1	1
150	1	1	1
300	2	2	1
600	4	3	1
1200	6	3	1

Scaling api-licensing service in the IBM Suite License Service namespace

The following table provides guidance for scaling the api-licensing service based on the number of concurrent users and login rate.

Table 133. CPU and memory limit guidelines for scaling api-licensing service

Login rate (logins per minute)	api-licensing replicas	api-licensing CPU limit	api-licensing memory limit (GB)
75	1	1	2
150	1	2	2
300	2	2	2
600	2	2	2
1200	2	2	2

Scaling coreapi service in the Maximo Application Suite core namespace

The following table provides guidelines for scaling the coreapi service based on the number of concurrent users and login rate.

Table 134. CPU and memory limit guidelines for scaling coreapi service

Login rate (login per minute)	coreapi replicas	coreapi CPU limit	coreapi memory limit (GB)
75	3	1	2
150	3	1	2
300	3	1	2
600	3	2	2
1200	3	3	2

Enabling Red Hat OpenShift Container Platform Cluster Insights Advisor

You can use the Red Hat OpenShift Container Platform Insights Advisor to identify any issues that are related to your current version, nodes, or configurations. This functionality is only present in Self-Managed Red Hat OpenShift. For more information, see [Self-managed Red Hat OpenShift subscription guide](#).

Procedure

1. Log in to the Red Hat OpenShift Container Platform console.

2. From the side navigation menu, click **Administration** > **Cluster Settings** tab.
3. On the **Cluster Settings** page, in the **Subscription** section, select Red Hat OpenShift Container Platform Cluster Manager.
You are redirected to the Red Hat OpenShift Container Platform Hybrid Cloud Console.
4. Click **Insights Advisor**

Configuring PID limits for Docker

You can enable or change the process identifier (PID) limits for the Docker container by using the command-line utility

About this task

A process identity (PID) is a distinct number that the Linux kernel assigns to each process or thread that is actively running on the system. The PID limit is a setting that can be enabled to control the number of processes that can run within one single container. If the value is too small, it can cause issues. For most workloads, a pod PID limit of 4096 is sufficient.

Platform version	Default value
IBM Cloud Kubernetes Service 4.8	231129
Red Hat OpenShift Service on Amazon Web Services	4096 in Red Hat OpenShift 4.11 and later
Microsoft Azure self-managed Red Hat OpenShift Container Platform	1024

Procedure

1. Log in to your Red Hat OpenShift Container Platform command line as an admin user.
2. Run the following command.

```
oc debug node/$NODE_NAME
chroot /host
vi /etc/crio/crio.conf
```

The `.conf` file, which you need to edit, is opened.

3. Insert the following line in the file and then save the file.

```
pids_limit =<new value>
```

4. Run the following commands to restart services and worker nodes.

```
systemctl daemon-reload
systemctl restart crio
shutdown -r now
```

Configuring the HAProxy router

The HAProxy router balances loads and routes traffic in cloud and container orchestration systems such as Red Hat OpenShift. You can scale the HAProxy router ingress controller, set the maximum number of concurrent connections, and select a different log balancing algorithm.

Scaling the Ingress controller

The peak number of concurrent incoming connections is the primary factor that influences the number of ingress pods to create. Each ingress pod in the Red Hat OpenShift HAProxy can support up to 20,000 incoming TCP connections per pod.

To scale the default ingress controller to a specific number of replicas, use the `oc patch` command in the Red Hat OpenShift command-line interface. The following example scales the default ingress controller to 3 replicas:

```
oc patch -n openshift-ingress-operator ingresscontroller/default --patch '{"spec":{"replicas":3}}' --type=merge
```

Setting the maximum number of concurrent connections

The HAProxy router can handle a default maximum of 20,000 concurrent connections, but this limit can be adjusted as needed. To configure the maximum connections, you can use the `--max-connections` option.

```
oc edit ingresscontroller default -n openshift-ingress-operator
spec:
  tuningOptions:
    maxConnections: 10000
```

Also, you can modify the **ROUTER_MAX_CONNECTIONS** environment variable in the router's deployment configuration to change the limit after which the router pods restart with the new value. If **ROUTER_MAX_CONNECTIONS** is not present, the default value of 20,000 is used.

However, it is crucial to verify that the node settings are configured to handle the specified connection limit. The `sysctl fs.nr_open` and `sysctl fs.file-max` system parameters must be large enough to support the increased max connections or HAProxy will not start.

Red Hat OpenShift Container Platform does not support modifying ingress controller deployments by using environment variables, such as **ROUTER_THREADS** and **ROUTER_DEFAULT_TUNNEL_TIMEOUT**, because these values are overwritten if the ingress operator is enabled.

Changing the load balancer algorithm

Starting in Red Hat OpenShift Container Platform 4.10, the following load-balancing algorithms are available:

- source
- roundrobin
- random, which is used by default
- leastconn

Use the random algorithm because it yields a balanced distribution of incoming connections to the IBM Maximo Application Suite pods.

To customize the load-balancing algorithm for a specific route, you can use route annotations. For example, to set the algorithm to roundrobin, run the following command:

```
oc annotate --overwrite route/example-route haproxy.router.openshift.io/balance=roundrobin
```

Node considerations

Instance types offer different combinations of CPU, memory, disk, and network configurations that you can configure to meet the needs of your users.

Consider the following guidelines when you configure nodes:

- Each worker node reserves approximately 1 core for internal services. To prevent overcommitting resources, a 16-core, 64 GB instance is a good starting point for a typical worker node. An 8-core instance can lack sufficient capacity, and a 32-core instance results in significant capacity loss if an outage occurs.
- Balance CPU-memory worker nodes, with a CPU-to-memory ratio of 1:4.
- For database nodes, use instances with a higher CPU-to-memory ratio, such as 1:8.

- A minimum of 3 worker nodes is required for high availability, which helps ensure built-in redundancy.
- For production environments, use 8-core, 32 GB instances for control plane nodes to avoid bottlenecks with internal services.
- Use a 10 GB Ethernet connection for production environments.
- A minimum of 300 GB of storage per worker node can support the build process requirements for Maximo Application Suite.

Backing up and restoring IBM Maximo Application Suite

Plan and implement backup and restore strategies for recovery from unplanned events, such as data loss or catastrophic failure of Red Hat OpenShift Container Platform.

About this task

- When you are restoring a Maximo Application Suite instance, the *instanceId* from the backups must match the *instanceId* of the new Maximo Application Suite instance.
- If you want to restore to a different Red Hat OpenShift Container Platform deployment, the *instanceId* must be the same as the *instanceId* of the backups.
- Alternatively, you can use IBM Storage Fusion with recipes for backing up and restoring Maximo Application Suite.

Overview of the components and processes for back up and restore procedures

You need to back up IBM Maximo Application Suite components on a scheduled basis to ensure that a full restore of IBM Maximo Application Suite is possible.

Maximo Application Suite is built for resilience. The main components and the way they interact in the logical architecture of Maximo Application Suite ensure the availability of your assets and information.

For more information, see [“Logical architecture” on page 213](#) and [“Resilient architecture components” on page 215](#).

Backup and restore scenarios overview

The timing and frequency of when to take backups and the adequate type of backup depends on the scenario as does restoring the persistent state of a Maximo Application Suite instance.

Persistent database failures

A persistent database failure can occur due to a hardware failure, an infrastructure failure, or a software failure. When a persistent database failure occurs, the Maximo Application Suite instance is not recoverable by using normal processes like a database roll back, starting another pod instance in Kubernetes, or rerouting the transaction by using a load balancer.

To restore the Maximo Application Suite instance, you must deploy or activate another database instance, and then restore the data from a backup.

If a single database failure occurs, you can do a partial restore for that portion of Maximo Application Suite. For more information, see [Partial recovery scenarios](#).

If multiple database failures occur, you must verify that data references are consistent. For more information, see [Persistence store relationships](#).

Data corruption induced by an application

Data corruption errors can occur due to defects or data entry errors.

For relational database management system, you can filter out these errors by using the restore process with point-in-time recovery or filtered transaction restore.

For other persistent stores, only the most recent backup point is available. To avoid losing data, you need to backup frequently. For example, if you are backing up data only once in a 24-hour period, you can lose up to 24 hours of data.

Moving data between Maximo Application Suite instances

You can restore data to a different Red Hat OpenShift cluster than the original cluster and move data between Maximo Application Suite instances.

Moving data between clusters or instances can be useful in the following scenarios:

- Restoring a production Maximo Application Suite instance to a new availability zone.
- Using production data for development and testing.
- Migrating a production Maximo Application Suite instance to a new location, such as a larger Red Hat OpenShift cluster or a different availability zone.
- Using sample data in new deployments, such as demo systems, proof-of-concept instances, model training, or new development.
- Resetting a Maximo Application Suite instance to the factory default.

Partial recoveries

A partial recovery is done when the scope of a failure is contained within one or more of the persistent data store types. To plan for a partial recovery, consider the scenario, the deployment configuration, and the wanted outcome.

This type of failure can occur when the deployed Maximo Application Suite instance is configured for higher levels of availability. For example, if you configure multiple nodes for the document database, and the nodes are distributed across zones or even regions, the document database is more resilient. Even in a complete region disaster, transactions can fail over to a replicated node and continue operations. Backups are still taken for the other scenarios but are rarely necessary for failure scenarios.

For more information, see [“High availability”](#) on page 213.

You can apply replication to the relational database management system database, the persistent volumes, and Red Hat OpenShift configuration management. IBM Cloud Object Storage uses replication to avoid backups. Turn on versioning to recover from data corruption induced by an application.

Backup and restore process overview

Learn which persistent data stores to back up and restore for each Maximo Application Suite application.

About this task

Depending on the application, you need to back up certain components. You can restore the same components for each application. The following list describes the individual application requirements.

Maximo Application Suite core

Back up MongoDB and the namespace. For more information about the back up and restore sequences, see [“Backing up Maximo Application Suite core”](#) on page 738 and [“Restoring and validating Maximo Application Suite core”](#) on page 758.

Maximo Manage

Back up MongoDB, relational data stores, cloud object stores, the namespace, and persistent volumes. For more information about the back up and restore sequences, see [“Backing up Maximo Manage”](#) on page 742 and [“Restoring and validating Maximo Manage”](#) on page 761.

The IoT tool

Back up MongoDB, relational data stores, and the namespace. For more information about the back up and restore sequences, see [“Backing up the IoT tool”](#) on page 747 and [“Restoring and validating the IoT tool”](#) on page 763.

Visual Inspection

Back up MongoDB, the namespace, and persistent volumes. For more information about the back up and restore sequences, see [“Backing up Maximo Visual Inspection” on page 748](#) and [“Restoring and validating Maximo Visual Inspection” on page 764](#).

Maximo Monitor

Back up MongoDB, and the namespace. For more information about the back up and restore sequences, see [“Backing up Maximo Monitor” on page 751](#) and [“Restoring and validating Maximo Monitor” on page 769](#).

Maximo Optimizer

Back up MongoDB. For more information about the back up and restore sequences, see [“Backing up Maximo Optimizer” on page 753](#) and [“Restoring and validating Maximo Optimizer” on page 771](#).

Maximo Collaborate

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Back up the namespace, CouchDB, Watson Discovery, and cloud object storage. For more information about the back up and restore sequences, see [“Backing up Maximo Collaborate” on page 753](#) and [“Restoring and validating Maximo Collaborate” on page 773](#).

Maximo Health or Maximo Health and Predict - Utilities

Back up Watson Studio data. For more information about the back up and restore sequences, see [“Backing up Maximo Health or Maximo Health and Predict - Utilities” on page 756](#) and [“Restoring Maximo Health or Maximo Health and Predict - Utilities” on page 776](#).

Maximo Predict

Back up Watson Studio and Watson Machine Learning data. For more information about the back up and restore sequences, see [“Backing up Maximo Predict” on page 756](#) and [“Restoring Maximo Predict” on page 777](#).

Scheduling nightly and weekly backups

MongoDB and Relational Database Management Systems (RDBMS) include back up and restore services. Decide how frequently to take the backups.

If you are using cloud services for MongoDB or RDBMS, you do not need to schedule backups.

Nightly backups

Backups of MongoDB, Red Hat OpenShift environment, and incremental backups of the RDBMS are scheduled frequently to ensure the availability of your data.

Schedule a MongoDB backup

Schedule a MongoDB backup by using one of the nodes in the cluster. For more information, see [“MongoDB overview” on page 731](#).

Enterprise MongoDB

Includes advanced tools for scheduling the backup.

Community Edition of MongoDB

Build your own automation configuration and tools.

You can set up a nonvoting replica node specifically for backup processing.

Schedule a RDBMS incremental backup

Schedule a RDBMS incremental backup and store it locally.

In the cloud object store, copy the backup to a cross-region bucket. This process increases transaction log space because the prior logs are not necessary for recovery.

For more information, see [“Backup and restore process for relational database management systems” on page 732](#).

Schedule the Maximo Application Suite Operator

Schedule the Maximo Application Suite Operator to create nightly backups of the Red Hat OpenShift environment.

You can schedule a backup by using the OADP Operator and Velero and store it to a cross-region bucket in the cloud object store.

Weekly backups

Since full RDBMS backups take longer, they are scheduled weekly.

Schedule a weekly full RDBMS backup and store it locally.

Copy the backup to a cross-region bucket in the cloud object store. This process increases transaction log space and also removes incremental backup files.

If you host multiple tenants or organizations, you must create a cross-region bucket for each tenant and organization.

MongoDB overview

The steps to back up MongoDB depend on where your MongoDB instance is and how it is configured.

Review the [mongodump MongoDB documentation](#). Determine where to store the backups and how to access the MongoDB instance.

The MongoDB instance is a replica set that includes three members. For more information, see [replica set](#).

A replica set consists of one primary and two secondary members. Assume that backups and restores are done only to the primary member.

MongoDB instance hosted by a third party

If you are using a third-party hosted MongoDB instance, review the hosting vendor's documentation for more information about backups.

You might use one of the following third-party hosted instances:

- IBM Cloud Databases for MongoDB. For more information, see [IBM Cloud Databases for MongoDB](#).
- MongoDB Atlas Database. For more information, see [MongoDB Atlas Database](#)
- Amazon Document DB with MongoDB compatibility. For more information, see [Amazon DocumentDB \(with MongoDB compatibility\)](#).

If your vendor does not provide backup services, your Maximo Application Suite administrator must decide where to store backup snapshots.

MongoDB community Kubernetes operator

If you are using a MongoDB instance that is based on the MongoDB Community Kubernetes operator, your Maximo Application Suite administrator must plan and schedule back ups.

For more information, see [MongoDB Community Kubernetes Operator](#).

Related tasks

[“Backing up MongoDB for Maximo Application Suite core” on page 738](#)

You need to back up specific MongoDB that are used by Maximo Application Suite core. Creating backups of the MongoDB databases that Maximo Application Suite core depends on, must be planned and scheduled to happen on a timely basis.

Related information

[The mongo Shell](#)

[MongoDB Shell \(mongosh\)](#)

[The MongoDB Database Tools Documentation](#)

Backup and restore process for relational database management systems

The backup and restore process for relational database management systems involves a full backup, incremental backups, and transaction logs.

You need to take a full back up which is used as the initial recovery point and has a copy of the entire database content. Because a full back up can be a slow process, make regular incremental backups that include only the most recent changes. Full backups are done less often.

Backups and the transaction logs must be kept in reliable storage locations, such as cloud object stores.

A restore is done by taking the information from the last full backup, applying each incremental backup, and replaying the transaction logs to the point of failure, which means all the committed transactions are safe.

A full backup releases previous incremental backups, and an incremental backup releases previous transaction logs. Both releases reduce data storage requirements.

Determine the best balance of operational backup processing time versus storage space and recovery processing time. Usually, each application service maintains transactional consistency. If a single relational database management system instance fails and is restored, you can complete the backup and restore operations across different relational database management system instances without concern for data inconsistency.

Cloud object storage versioning

Balance the additional feature storage costs for versioning and cross-region replication for disaster recovery and the importance of cloud object storage versioning recovery. You do not have to use versioning for all file types.

Backup file storage is normally not updated after it is saved, but you can use versioning to avoid accidental deletion. Deleting versioned files marks the files as deleted, but the files are still accessible. You can permanently delete the files after new backups are created and after a retention period.

The backup process is built into IBM Cloud Object Storage. Basic failure points do not need recovery. Loss of data can occur due to user or application errors, such as deleting an attachment. If you want to keep previous copies accessible, configure versioning. You can also use Amazon Web Services S3 cloud object storage.

Persistent volumes

A persistent volume (PV) is an API object that captures details of the storage implementation. Each application stores configuration data in the persistent volumes that are assigned to the containerized virtual machines. The application identifies which services are backed up by tagging them under the application's namespace.

Namespace in Red Hat OpenShift

Regularly back up the resources in the namespace manually or by using a script.

- For more information, see [“Backing up Maximo Application Suite core namespace” on page 738](#).
- For more information, see [“Backing up the Maximo Manage namespace” on page 743](#).

Persistence store relationships

If the recovery process is not coordinated properly, the way that data is linked across persistence stores can cause inconsistencies.

Device message flow

A typical device metric flows into the IoT tool from various sources. The IoT tool might use an MQTT message on a bidirectional connection for direct device management features and continuous connectivity. The IoT tool might use a Kafka topic or gateway connectivity. These messages are often

persisted to enable guaranteed delivery, specifically in the MessageSight or Kafka subcomponents of the IoT tool.

These subcomponents are not typically backed up or restored. The components handle reliability through replication. Each message goes through a lifecycle. Restarting the IoT tool reestablishes this lifecycle and starts processing new messages.

Device registration and metric timestamp data

The IoT common service registers new device types and device instances into both MongoDB metadata store and RDBMS before a schema is created for that device to store its timestamp metric values. The registration for the relational database management system is completed in a different transaction boundary. If a single MongoDB node fails, replication from one of the other nodes accomplishes recovery automatically. If MongoDB has a multinode loss of data before the next backup, inconsistencies can arise.

If MongoDB has a multinode loss of data before the next backup, the device is not registered in the recovered MongoDB, but the device still has the historical timestamp metric data and some authentication and metadata. The device must be re-created to authorize IoT connections and receive new metric data. To avoid a multinode loss, use a script to automate IoT device registration and before you enable connectivity to the IoT tool common service to receive time-series data, record a MongoDB backup.

Data dictionary relationships

The Asset Data Dictionary is a service that handles complex graph relationships between assets and devices, device tag hierarchies, and other advanced features. The Asset Data Dictionary uses a relational database management system for persistence and references identifiers in different application service schemas, which might be in the same relational database management system instance or different instances that have different transaction boundaries. If the application services share the relational database management system instance, back up and restore all the schemas together to keep them consistent. If the schemas are spread across multiple relational database management system instances, restoring data for Maximo Manage or Maximo Monitor independently of the Asset Data Dictionary might require performance of those asynchronous operations.

Attachments

Some Maximo Application Suite applications use attachments. Attachments can include receipts, certifications, invoices, and any files that you attach to your assets, work orders, or job plans. You need to back up the attachments regularly.

- For more information, see [“Backing up Maximo Manage attachments”](#) on page 746.

Kafka message provider overview

To avoid potential data loss, you need to replicate the Kafka storage. You will need to consult the documentation from your Kafka provider.

For more information, see the back up and restore documentation or guidelines that are provided by your Kafka provider.

For more information about AMQ Kafka, see [Understanding Red Hat AMQ Streams Components for Red Hat OpenShift and Kubernetes - Part 2](#)

If you are using a cloud-based Kafka solution, you can enable more data protection measures by enabling cloud backup.

Related concepts

[“Replicating Apache Kafka messages for Maximo Manage”](#) on page 747

If Maximo Manage is using Apache Kafka messaging platform, you must replicate the Kafka storage of messages to enable recovery from an unplanned data loss event.

Encrypting and compressing backups

To avoid unauthorized access and tampering of the backup file, you can encrypt and compress the backup.

Procedure

- To encrypt and compress the contents of the current working directory, enter the following command.

```
tar -czf - * | openssl enc -e -aes256 -out secured.tar.gz
```

- To encrypt and compress the contents of the directory for a pod container, enter the following command.

```
oc exec -it $POD -c $CONTAINER --namespace $NAMESPACE-- bash -c "tar -cz $DIRECTORY/* | openssl enc -e -aes256 -out <targzfilename> "
```

Decrypting and extracting backups

To gain access to your secured and compressed backup file, you need to decrypt and extract.

Procedure

- To decrypt and extract contents of the compressed backup, enter the following command.

```
openssl enc -d -aes256 -in <targzfilename> | tar xz
```

- To decrypt and extract contents in a pod container, enter the following command.

```
oc exec -it $POD -c $CONTAINER --namespace $NAMESPACE -- bash -c 'openssl enc -d -aes256 -in <targzfile> | tar xz'
```

Backing up Maximo Application Suite

On a scheduled and regular basis, back up both Maximo Application Suite core and Maximo Application Suite applications.

Before you begin

Ensure that both backup and restore Red Hat OpenShift clusters are running and use the same versions of Maximo Application Suite core and Maximo Application Suite applications.

Backing up MongoDB for Maximo Application Suite

Creating backups of the MongoDB databases that Maximo Application Suite depends on, must be planned and scheduled to happen on a timely basis.

Before you begin

- Ensure that you installed MongoDB by following the procedures for installing MongoDB on-premises, or the steps in [Dependency Management: MongoDB](#). If you installed MongoDB in a different way, the backup process does not work.
 - For more information, see [Installing MongoDB](#).
 - For more information, see [MongoDB](#).
- The ability to enter the **mongo-ce** pods in the mongoce namespace by using the **oc** command-line tool.
- The password for the MongoDB admin user. For more information, see [Get the credentials](#).
- The primary replica set member, its internal service hostname, and pod name. For more information, see [Determine the primary set member](#).

- If the MongoDB instance is using IBM Suite License Service, backup the IBM Suite License Service. For more information, see [“Backing up the IBM Suite License Service license file and ID”](#) on page 739.
- The MongoDB CA certificate.
 - If you are using port forwarding, a local copy is needed.
- To install MongoDB Database Tools. For more information, see [MongoDB Database Tools](#).
 - If you are using port forwarding, **mongodump** and **mongoexport** need to be installed locally.

About this task

You can back up MongoDB by using manual steps.

After you get the credentials and determine the primary replica set member, you need to choose a connection strategy: internally within the Red Hat OpenShift cluster or by using port forwarding to connect to the single replica member.

Getting the credentials and determining the primary set member

The administrator must get the credentials and determine the primary replica set member. Then, the administrator must choose an option for backing up: either using MongoDB CE pods internally, or by using port forwarding.

Before you begin

You need to log in by using **oc** and have administrative privileges for the **mongoce** namespace.

Procedure

1. Get the credentials.

The password for the admin user is contained in the secret **mas-mongo-ce-admin-admin** in the **mongoce** namespace.

- a) Get the encoded admin password and decode **data.password** from the secret.

```
oc get secret mas-mongo-ce-admin-admin -n mongoce -o yaml | yq .data.password | base64 -d
```

- b) Get the encoded **ca.crt** and decode it.

```
oc get secret mongo-ca-secret -n mongoce -o yaml | yq '.data["ca.crt"]' | base64 -d > /tmp/ca.crt
```

2. Determine the primary replica set member.

Enter the MongoDB container of any of the MongoDB replica set member pods in the pod **mas-mongo-ce** and get a copy.

By default, it can be any of these pods:

- **mas-mongo-ce-0**
- **mas-mongo-ce-1**
- **mas-mongo-ce-2**

Any of the listed pods are valid. The example uses *mas-mongo-ce-0*.

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce bash
```

3. Determine the file path to the certificate authority (CA) which depends on the MongoDB CE operator version.

MongoDB CE operator version	
For older versions of the MongoDB CE operator	the path is <code>/var/lib/tls/ca/ca.crt</code> .

MongoDB CE operator version	
For the latest version of the MongoDB CE operator	<p>find the path by using the command</p> <pre>oc exec -it mas-mongo-ce-0 -c mongod -- namespace mongoce -- bash -c "cat /data/ automation-mongod.conf"</pre>

4. Output the **mongod** config file and the path to the CA file by using this command:

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "cat /data/automation-
mongod.conf"
net:
  bindIp: 0.0.0.0
  port: 27017
  tls:
    CAFile: /var/lib/tls/ca/
10f4a08a1c4ec1c05b550811eda26dc91b8f0e8baf86c37235630373b8e13096.pem
    allowConnectionsWithoutCertificates: true
    allowInvalidCertificates: true
    allowInvalidHostnames: true
    certificateKeyFile: /var/lib/tls/server/
870b9305462bfba1006a0d2af677de0ad5df1db15307313f03902ae55cef1b09.pem
    mode: requireTLS
  replication:
    replSetName: mas-mongo-ce
  security:
    authorization: enabled
    keyFile: /var/lib/mongodb-mms-automation/authentication/keyfile
  setParameter:
    authenticationMechanisms: SCRAM-SHA-256,SCRAM-SHA-1
  storage:
    dbPath: /data
    wiredTiger:
      engineConfig:
        journalCompressor: snappy
```

5. Inside the **mongod** container, make a connection to the replica set.

Update the value of the **tlsCAFile** file path to the value that was determined in the previous step.

```
mongo "mongodb://
admin:{decodedPassword}@mas-mongo-ce-0.mas-mongo-ce-svc.mongoce.svc.cluster.local:27017,mas-
mongo-ce-1.mas-mongo-ce-svc.mongoce.svc.cluster.local:27017,mas-mongo-ce-2.mas-mongo-ce-
svc.mongoce.svc.cluster.local:27017/admin?replicaSet=mas-mongo-ce" --tls --
tlsCAFile=/var/lib/tls/ca/ca.crt
```

6. Use the mongo shell and run the **db.runCommand("ismaster");** command to make the connection to the replica set.

```
db.runCommand("ismaster");
```

7. Obtain the service hostname and the pod of the primary replica set member.

The command displays a result with an attribute that is named **primary**, which is used to identify the primary replica set member.

```
{
  ...
  ...
  ...
  "setName" : "mas-mongo-ce",
  "setVersion" : 1,
  "ismaster" : true,
  "secondary" : false,
  "primary" : "mas-mongo-ce-0.mas-mongo-ce-svc.mongoce.svc.cluster.local:27017",
  ...
  ...
  ...
}
```

8. Exit back to the container shell.


```
mas-mongo-ce:PRIMARY> exit
```

- Exit out of the container.

```
2000@mas-mongo-ce-0:/$ exit
```

Backing up by using MongoDB CE pods internally

Back up by using a MongoDB CE replica set pod internally within the Red Hat OpenShift cluster.

Before you begin

The administrator needs to get the credentials and determine the primary set member. For more information, see [“Getting the credentials and determining the primary set member”](#) on page 735.

Procedure

- Create a directory for the backups.

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "mkdir /data/backups"
```

- Back up the databases.

Example of backing up Maximo Application Suite core databases `mas_{instanceId}_core` and `mas_{instanceId}_catalog`.

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "mongodump --host={{primaryHost}} --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt -d mas_{instanceId}_core --archive=/data/backups/mas_{instanceId}_core.archive"
```

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "mongodump --host={{primaryHost}} --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt -d mas_{instanceId}_catalog --archive=/data/backups/mas_{instanceId}_catalog.archive"
```

- Copy the backup archives to the local system.

```
oc cp mongoce/mas-mongo-ce-0:/data/backups ./backups -c mongod
```

- Remove back ups from the mongod container in the `mas-mongo-ce-0` pod.

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "rm -rf /data/backups"
```

Note: The archive files are stored in `./backups` on the administrator's local computer. Copy the archives to a safe location.

Backing up by using port forwarding

Back up by using port forwarding to connect directly to the primary replica set host.

Before you begin

The administrator needs to get the credentials and determine the primary set member. For more information, see [“Getting the credentials and determining the primary set member”](#) on page 735.

Procedure

- In one terminal, start the port forwarding to the pod that is the primary replica set member.

```
oc port-forward {{podThatIsPrimary}} -n mongoce 28015:27017
```

- In another terminal, make sure that a `backups` subfolder exists and backup the databases.

- Make sure that a `backups` subfolder exists.

```
mkdir backups
```

b) Back up the databases.

```
mongodump --host=localhost --port=28015 --username=admin --password={{decodedPassword}}
--authenticationDatabase=admin --ssl --sslCAFile=/tmp/ca.crt -d mas_{{instanceId}}_core --
archive=./backups/mas_{{instanceId}}_core.archive
```

```
mongodump --host=localhost --port=28015 --username=admin --password={{decodedPassword}}
--authenticationDatabase=admin --ssl --sslCAFile=/tmp/ca.crt -d mas_{{instanceId}}_catalog
--archive=./backups/mas_{{instanceId}}_catalog.archive
```

Backing up Maximo Application Suite core

Back up the Maximo Application Suite core databases and the namespace in Red Hat OpenShift.

Before you begin

- Before you plan your backup and restore strategy, make note of the Maximo Application Suite core instance ID *{instanceId}*.
- You must have administrator privileges for the Maximo Application Suite core namespace.

About this task

You can back up Maximo Application Suite core by using manual steps.

Backing up MongoDB for Maximo Application Suite core

You need to back up specific MongoDB that are used by Maximo Application Suite core. Creating backups of the MongoDB databases that Maximo Application Suite core depends on, must be planned and scheduled to happen on a timely basis.

About this task

After you get the credentials and determine the primary replica set member, you need to choose a connection strategy: internally within the Red Hat OpenShift cluster or by using port forwarding to connect to the single replica member. To learn more, see [“Backing up MongoDB for Maximo Application Suite” on page 734](#).

Procedure

- Back up the Maximo Application Suite core databases.
 - mas_{{instanceId}}_core
 - mas_{{instanceId}}_catalog

Note: You don't need to back up the mas_{{instanceId}}_adoptionusage database.

Related concepts

[“MongoDB overview” on page 731](#)

The steps to back up MongoDB depend on where your MongoDB instance is and how it is configured.

Backing up Maximo Application Suite core namespace

Maximo Application Suite core is mostly stateless. However, to ensure that a Maximo Application Suite core instance can be restored to an expected state, you must back up specific resources regularly.

About this task

This procedure is based on the configuration and deployment that was made by using the [MAS Devops Ansible Collection](#).

Backing up the version of the IBM Maximo Operator Catalog

Back up the version of the IBM Maximo Operator Catalog that is specified in the CatalogSource that is named `ibm-operator-catalog`.

Before you begin

- You need administrator credentials to access the Red Hat OpenShift command-line interface (CLI). For more information, see [Red Hat OpenShift command-line interface \(CLI\)](#).
- Make sure that the cli command yq version 4.33.3 or earlier versions is installed in the path. For more information, see [CLI yq](#).

Procedure

Note and save the catalog version.

- a) Export the CatalogSource `ibm-operator-catalog` from the `openshift-marketplace` namespace to retrieve the currently configured version.

```
oc get CatalogSource ibm-operator-catalog -n openshift-marketplace -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-
configuration"], .status)'
```

The following example shows standard output for this command:

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: ibm-operator-catalog
  namespace: openshift-marketplace
spec:
  description: Catalog Source for IBM Maximo Application Suite (Internal Use Only)
  displayName: IBM Maximo Operators (v8-master-amd64 Dev)
  image: icr.io/cpopen/ibm-maximo-operator-catalog:v8-230414-amd64
  priority: 100
  publisher: IBM
  sourceType: grpc
  updateStrategy:
    registryPoll:
      interval: 45m
```

From the YAML in the example, notice that the IBM Maximo Operator Catalog version is 230414. The version is referenced in the `.spec.image` attribute.

```
image: icr.io/cpopen/ibm-maximo-operator-catalog:v8-230414-amd64
```

Only the version needs to be noted and saved.

Backing up the IBM Suite License Service license file and ID

If IBM Suite License Service is installed on the same Red Hat OpenShift cluster as Maximo Application Suite core, the license file and ID must be backed up.

Procedure

- Export and back up the contents of the `ibm-sls-sls-entitlement` secret in the `ibm-sls` namespace.

The secret exists in IBM Suite License Service 3.3.0 and later.

```
oc get Secret ibm-sls-sls-entitlement -n ibm-sls -o yaml
apiVersion: v1
data:
  entitlement: U0VSVkVSIglibS1zbHMgMDI0MmFjMTEwMDAyIDI3MDAwClZ ...
```

Backing up the resources in the Maximo Application Suite core namespace manually

You need to backup the resources in the `mas-{instanceId}-core` namespace on a scheduled basis manually or by using a script.

Procedure

1. Back up the subscription.

```
oc get Subscription ibm-mas-operator -n mas-{instanceId}-core -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
.resourceVersion, .metadata.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-
configuration"], .status)' > ./Subscription-ibm-mas-operator.yaml
```

2. Back up the OperatorGroup.

```
oc get OperatorGroup ibm-mas-operator-group -n mas-{instanceId}-core -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
.resourceVersion, .metadata.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-
configuration"], .status)' > ./OperatorGroup-ibm-mas-operator-group.yaml
```

3. Back up the superuser credentials secret.

```
oc get Secret $MAS_INSTANCE_ID-credentials-superuser -n mas-{instanceId}-core -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
.resourceVersion, .metadata.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-
configuration"], .status)' > ./Secret-$MAS_INSTANCE_ID-credentials-superuser.yaml
```

4. Back up the ibm-entitlement secret.

```
oc get Secret ibm-entitlement -n mas-{instanceId}-core -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
.resourceVersion, .metadata.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-
configuration"], .status)' > ./Secret-ibm-entitlement.yaml
```

5. Back up core.mas.ibm.com resource.

- a) Back up the Suite.core.mas.ibm.com resource, only one instance that is named *{instanceId}*.
- b) Back up all the Workspace.core.mas.ibm.com resources.

6. Back up all *.config.mas.ibm.com resources.

- **MongoCfg**
- **KafkaCfg**
- **JdbcCfg**
- **SlsCfg**
- **BasCfg**
- **Smtpcfg**
- **WatsonStudioCfg**
- **ObjectStorageCfg**
- **PushNotificationCfg**
- **ScimCfg**
- **IDPCfg**

7. Back up all instances of addons.mas.ibm.com.

- **AppConnect**
- **Humai**

Note: Starting in Maximo Application Suite 8.11, Parts Identifier is no longer available. If Parts Identifier is deployed and active in your environment and you are upgrading to Maximo Application Suite 8.11, you must deactivate and delete Parts Identifier before you can complete the upgrade.

- **MVIEdge**

The following manual example uses the JdbcCfg resource.

```
oc get JdbcCfg -n mas-{instanceId}-core
```

Make note of the names of each instance and back up each one:

```
oc get JdbcCfg {jdbcInstnaceName} -n mas-{instanceId}-core -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata
```

```
.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./JdbcCfg-{jdbcInstanceName}.yaml
```

- Optional: If the steps were done without the script, export any credentials that might be kept in a *Secret* that the **config** resource references.
If the **config** resource contains a section that is called `.spec.config.credentials.secretName`, back up the secret that *secretName* references.

```
oc get Secret <secretName> -n mas-{instanceId}-core -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./Secret-<secretName>.yaml
```

- Optional: Backup the config map `mas-messages` in the Maximo Application Suite core namespace.
If custom **Access denial message** is set in the license consumption admin page, you must backup the configmap `mas-messages`.

```
oc get cm mas-messages -n mas-{instanceId}-core -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./configmap-mas-messages.yaml"
```

Backing up when manual certificate management is enabled

Determine whether manual certificate management is enabled and back up.

About this task

If manual certificate management was configured, you must back up the secret *{instanceId}-cert-publicneeds*.

- For more information, see [“Enabling manual certificate management” on page 601](#).
- For more information, see [“Uploading public certificates in Red Hat OpenShift” on page 601](#).

To determine whether manual certificate management is enabled, examine the Suite resource in the Maximo Application Suite core namespace by running the following command:

```
oc get Suite $MAS_INSTANCE_ID -n mas-{instanceId}-core -o yaml
```

If `.spec.settings.manualCertMgmt` is set to `true`, then manual certificate management is enabled.

Procedure

Back up the associated secret.

- Use the following command:

```
oc get Secret $MAS_INSTANCE_ID-cert-public -n mas-{instanceId}-core -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./Secret-$MAS_INSTANCE_ID-cert-public.yaml
```

- Alternatively, use the script, which automatically detects whether manual certificate management is enabled and backs up the secret *{instanceId}-cert-public*.
 - For more information, see [“Backing up Maximo Application Suite core namespace by using a script” on page 742](#).

IBM Cloud Pak foundational services

A back up of IBM Cloud Pak foundational services is necessary when a custom ClusterIssuer or when the Let's Encrypt Integration are in use for Maximo Application Suite.

All the applicable ClusterIssuer and Secret resources need to be exported and backed up.

- For more information, see [custom ClusterIssuer](#).

If only the provided self-signed certificates are used, then a Maximo Application Suite installation generates a new self-signed certificate that is satisfactory.

Backing up a custom ClusterIssuer configuration

If Maximo Application Suite core is configured to use a custom ClusterIssuer, then the ClusterIssuer and all of its associated Secret resources must be backed up manually.

About this task

The automated script `mascore-backup-restore.sh` can detect whether a custom ClusterIssuer is configured. It does not back up the ClusterIssuer instance or any of its associated Secrets. The script can be found in [mascore-backup-restore.sh](#).

Procedure

1. To determine whether a custom ClusterIssuer is configured for Maximo Application Suite core, examine the Suite resource in the Maximo Application Suite core namespace.

```
oc get Suite $MAS_INSTANCE_ID -n mas-{instanceId}-core -o yaml
```

If the customer is using a custom ClusterIssuer, the result in the Suite CR includes a `.spec.certificateIssuer` section, as shown in the following example.

```
spec:
  certManagerNamespace: ibm-common-services
  certificateIssuer:
    duration: 8760h0m0s
    name: ivt811x-01-cis-le-prod
    renewBefore: 720h0m0s
```

2. Back up the named ClusterIssuer and all of its associated secrets.

Related information

[Cert-manager configuration](#)

Backing up Maximo Application Suite core namespace by using a script

An alternative approach to manually backing up is to use the script.

- To access the script, go to [mascore-backup-restore.sh](#)
- The script can detect whether a custom ClusterIssuer is configured, but it will not backup the ClusterIssuer instance or any of its associated Secrets.
- The script automatically detects whether manual certificate management is enabled and backs up the secret `{instanceId}-cert-public`.
 - For more information, see [manual certificate management](#).
 - For more information, see [secret {instanceId}-cert-public](#).

Related information

[Restoring the Maximo Application Suite core namespace in Red Hat OpenShift using a script](#)

Backing up Maximo Manage

Back up Maximo Manage data stores, the Maximo Manage namespace in Red Hat OpenShift, and Maximo Manage attachments.

Before you begin

- Back up Maximo Application Suite core before you back up Maximo Manage.
 - For more information, see [“Backing up Maximo Application Suite core”](#) on page 738.
- If Maximo Manage is integrated with Apache Kafka, you must replicate the Kafka message storage.

- For more information, see [“Kafka message provider overview” on page 733](#).

About this task

You can back up Maximo Manage by using manual steps.

Backing up Maximo Manage databases

Back up the database to ensure continuous availability of data. Back up the database before and after you upgrade Maximo Manage. To back up, you create a copy of the data, which can be recovered and restored if you have a data failure.

The user data that is created in Maximo Application Suite is stored in MongoDB and synced to the Maximo Manage database.

To ensure the consistency of user data, obtain a MongoDB backup first, followed by a database backup.

For more information, see [“MongoDB overview” on page 731](#).

For more information, see [“Backing up MongoDB for Maximo Application Suite” on page 734](#)

To back up the database that is used by Maximo Manage, follow the database provider's documentation. For more information, see [Backing up and restoring Db2](#).

Related concepts

[“Restoring Maximo Manage databases” on page 761](#)

You can use your database backups to restore the databases that you use for Maximo Manage.

Backing up the Maximo Manage namespace

To ensure that a Maximo Manage instance can be restored to an expected state, you must back up the resources from the `mas-{instanceId}-manage` namespace on a scheduled basis.

Before you begin

- You need the version of the operator catalog that is specified in `CatalogSource` that has the name `ibm-operator-catalog`.
- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the `oc` command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the `yq` command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

The following list is an example of the types of resources that you must back up from the `mas-{instanceId}-manage` namespace.

- `ibm-mas-manage` subscription and operator group
- `ibm-entitlement` secret
- `{workspaceId}-mange-encryptionsecret` secret
- `{workspaceId}-manage-encryptionsecret-operator` secret
- all `{workspaceId}-mange-d-sb*-asc--sn` JMS server secrets, if a JMS server is installed
- `mangeapps.apps.ibm.com` resource
- `manageworkspaces.apps.mas.ibm.com` resource

Backing up Maximo Manage namespace manually

On a scheduled basis, back up the resources in the namespace by using commands.

Procedure

1. Back up the `ibm-mas-manage` subscription and operator group and the `ibm-entitlement` secret. Enter the following commands.

```
oc get Subscription ibm-mas-manage -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./Subscription-ibm-mas-manage.yaml
```

```
oc get OperatorGroup mas-{instanceId}-manage-operator-group -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./OperatorGroup-ibm-manage-operator-group.yaml
```

```
oc get Secret ibm-entitlement -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./Secret-ibm-entitlement.yaml
```

2. Back up the `manageapps.apps.ibm.com` resource. Back up the `{instanceId}` instance.

```
oc get ManageApp {instanceId} -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./manageapp.yaml
```

3. Back up the `manageworkspaces.apps.mas.ibm.com` resource. Back up all instances.

```
oc get ManageWorkspace {instanceId}-{workspaceId} -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./{workspaceId}-manageworkspace.yaml
```

4. Back up the secrets.

The secrets that you back up depend on your environment. Run the following command for each secret that you back up.

```
oc get Secret <secret> -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./Secret-<secret>.yaml
```

- `{workspaceId}-manage-encryptionsecret`
- `{workspaceId}-manage-encryptionsecret-operator`
- `{workspaceId}-manage-d-sb*-asc-sn`, when the `jms` server is deployed. The secret names can be different. The following list of secret names is a list of default secret names. To check your secret names, open `spec.settings.deployment.serverBundles` and check the secret names in `ManageWorkspace` CR for the `additionalServerConfig` property.
 - `{workspaceId}-manage-d-sb0-asc-sn`
 - `{workspaceId}-manage-d-sb1-asc-sn`
 - `{workspaceId}-manage-d-sb2-asc-sn`
 - `{workspaceId}-manage-d-sb3-asc-sn`
 - `{workspaceId}-manage-d-sb4-asc-sn`
- `{workspaceId}-manage-d-ls3d-sk-sn`. To check your secret name, open `spec.settings.deployment.loggingS3Destination` and check the secret name in `ManageWorkspace` CR for the `secretKey` property.

5. Back up your own certificate.

This type of certificate is also known as bring your own certificate (BYO).

- a) Examine the suite resource in the Maximo Application Suite core namespace to determine whether manual certificate management is enabled.

For more information, see [manual certificate management](#).

- If a value of `false` is set for `.spec.settings.manualCertMgmt` or the value does not exist, then manual certificate management is not enabled.
- If a value of `true` is set for `.spec.settings.manualCertMgmt`, the manual certificate management is enabled.

```
$ oc get Suite $MAS_INSTANCE_ID -n mas-{instanceId}-core -o yaml
```

- b) If the manual certificate management is enabled, back up the `{instanceId}-{workspaceId}-cert-public-81` secret.

```
oc get Secret{instanceId}-{workspaceId}-cert-public-81 -n mas-{instanceId}-manage -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.generation, .metadata.resourceVersion, .metadata.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-configuration"], .status)' > ./Secret-cert-public-81.yaml
```

If you don't do a manual backup, then you might instead back up by using the `manage-backup-restore.sh` backup script. The backup script automatically detects when manual certificate management is enabled and backs up the secret `{instanceId}-{workspaceId}-cert-public-81`.

Related tasks

[“Restoring the Maximo Manage namespace in Red Hat OpenShift” on page 761](#)

You can restore your backed up Maximo Manage namespace manually or by using the `manage-backup-restore.sh` script.

Backing up Maximo Manage namespace with a script

On a scheduled basis, use the `manage-backup-restore.sh` script to back up the Maximo Manage namespace.

Before you begin

Get the script [manage-backup-restore.sh](#).

Procedure

Run the following command to back up the Maximo Manage namespace.

```
manage-backup-restore.sh -w <mas-workspace-id> -i <MAS_INSTANCE_ID> -f <BACKUP_FOLDER> -m backup
```

```
manage-backup-restore.sh -w <mas-workspace-id> -i dev -f ./ -m backup
```

-w, --mas-workspace-id

The Maximo Manage workspace ID that you are backing up.

-i, --mas-instance-id

The Maximo Application Suite instance ID that you are backing up.

-f, --backup-folder

The folder where the backup artifacts are written to.

-m, --mode

Indicates whether to backup or restore. Use the value **restore**.

Related tasks

[“Restoring the Maximo Manage namespace in Red Hat OpenShift” on page 761](#)

You can restore your backed up Maximo Manage namespace manually or by using the `manage-backup-restore.sh` script.

Backing up Maximo Manage attachments

Back up the attachments in Maximo Manage, which can include receipts, certifications, invoices, and any files that you attach to your assets, work orders, or job plans.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the `oc` command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the `yq` command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Copy the attachments from the `all` or `ui` container on the server pod to a local system. Then, archive the attachments folder from the local system to an external storage.

Procedure

1. Get the `mountPath`. In the `ManageWorkspace` custom resource (CR), go to `spec.settings.deployment.persistentVolumes`.
2. Log in to your cluster.
 - a) In Red Hat OpenShift, click the drop-down menu in the global navigation bar and click **Copy login command**.
 - b) On the new page, click **Display token**.
 - c) In the **Log in with this token** field, copy the login command.

```
oc project <your "Manage" namespace>
```

3. Get the pod name of your server `$POD` for the `all` or `ui` containers.

```
oc get pods
```

4. Back up your attachments.

- Use the following command for an `all` server deployment.

```
oc exec -it $POD -c all --namespace $NAMESPACE -- bash -c "tar -cz <tarzfile> <mountPathfolder>/"
```

- Use the following command for a `ui` server deployment.

```
oc exec -it $POD -c ui --namespace $NAMESPACE -- bash -c "tar -cz <tarzfile> <mountPathfolder>/"
```

5. Copy the compressed file of attachments to your local system.

```
oc cp $POD:/$TARFOLDER ./$LOCALFOLDER -c all
```

Related tasks

[“Encrypting and compressing backups” on page 734](#)

To avoid unauthorized access and tampering of the backup file, you can encrypt and compress the backup.

[“Restoring Maximo Manage attachments” on page 762](#)

You can restore attachments for Maximo Manage, which can include receipts, certifications, invoices, and any files that are attached to assets, work orders, or job plans.

Replicating Apache Kafka messages for Maximo Manage

If Maximo Manage is using Apache Kafka messaging platform, you must replicate the Kafka storage of messages to enable recovery from an unplanned data loss event.

- Consult the documentation of your Kafka provider.
- If you use cloud-based Apache Kafka, enable the cloud backup.
- If you have a data loss event, create a Kafka instance on a restore cluster and configure Maximo Manage to use the new Kafka instance.

For more information, see [“Kafka message provider overview” on page 733](#).

Configuring Maximo Manage for a new Kafka instance

To configure Maximo Manage for a new Kafka instance, update the Kafka configuration by using the **Add/Modify Message Providers** action in the External Systems application in Maximo Manage. If necessary, import Kafka certificates.

For more information, see [Integration by using Apache Kafka](#).

Resending data from Kafka

If you do not back up Kafka messages, you can resend message data from Apache Kafka and reset the offset. If you use this approach, you might get duplicate records because some records were already processed. To find the correct offset, detect the records that were already processed.

Backing up the IoT tool

You can back up the IoT tool if it is deployed on an IBM Cloud , Amazon Web Services, or on-premises environment.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Regularly back up the IoT tool databases, relational databases, and IoT namespace in Red Hat OpenShift.

You can back up the IoT tool by using manual steps.

Backing up MongoDB for the IoT tool

You need to back up specific MongoDB databases that are used by the IoT tool.

Back up the following MongoDB databases that are used by the IoT tool:

- `iot_{MAS_INSTANCE_ID}_cs_activity_db`
- `iot_{MAS_INSTANCE_ID}_d_actions`
- `iot_{MAS_INSTANCE_ID}_d_core`
- `iot_{MAS_INSTANCE_ID}_d_dashboard`
- `iot_{MAS_INSTANCE_ID}_d_deviceregistry`
- `iot_{MAS_INSTANCE_ID}_d_dmserver`
- `iot_{MAS_INSTANCE_ID}_d_dsc`
- `iot_{MAS_INSTANCE_ID}_d_infomgmt`
- `iot_{MAS_INSTANCE_ID}_d_provision_s2s`

- `iot_{MAS_INSTANCE_ID}_d_riskmgmtsecurity`
- `iot_{MAS_INSTANCE_ID}_organizations`

For more information, see [“Backing up MongoDB for Maximo Application Suite”](#) on page 734.

Backing up relational databases for the IoT tool

You can back up relational databases that are used by the IoT tool if the database instance is hosted within the Red Hat OpenShift cluster.

To ensure the consistency of user data, first back up the document database that is used by the IoT tool and then back up the relational database.

The user data that is created in Maximo Application Suite is stored in the IoT tool document database and synchronized into the relational database that is used by the IoT tool.

If you are not using Db2, refer to your database provider documentation.

Use one of the following methods to back up Db2.

- Use the instructions in the Db2 documentation.
 - For more information, see [Backing up and restoring Db2](#).
 - The Db2 container name is `c-mas-{MAS_INSTANCE_ID}-system-db2u-0`.
 - For more information, see [Backing up a Db2 database online](#).

Backing up the IoT tool namespace

You need to back up three encryption secrets in the namespace.

In the `mas-{MAS_INSTANCE_ID}-iot` namespace, back up the following secrets:

actions-credsenckey

The secret that is used to encrypt and decrypt data for the actions component, like the user credentials that are used by the IoT tool.

auth-encryption-secret

The secret that is used to encrypt and decrypt data for the user authentication, like api keys, users, and tokens that are used by the IoT tool.

provision-creds-enckey

The secret that is used to encrypt and decrypt data for the Maximo Application Suite workspace that is used by the IoT tool.

Use the following command to back up each secret individually:

```
oc get secret {secret_name} -n mas-{MAS_INSTANCE_ID}-iot -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadat
a.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-configuration"], .status)' > ./
{MAS_INSTANCE_ID}-{secret_name}.yaml
```

Backing up Maximo Visual Inspection

If Maximo Visual Inspection is deployed on IBM Cloud, Amazon Web Services, or an on-premises environment, back up the document data stores and persistent volumes for Maximo Visual Inspection.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the `oc` command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the `yq` command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can back up Maximo Visual Inspection by using manual steps.

Related tasks

[“Restoring and validating Maximo Visual Inspection” on page 764](#)

To restore Maximo Visual Inspection, restore the MongoDB backup and the backed up data from the persistent volume claim and then restart the service pods. You can also validate the restoration.

Backing up MongoDB for Maximo Visual Inspection

Back up the specific MongoDB databases that are used by Maximo Visual Inspection.

Back up the following MongoDB databases:

- mas-*{MAS_INSTANCE_ID}*-visualinspection
- mas-*{MAS_INSTANCE_ID}*-edgeman

For more information, see [“Backing up MongoDB for Maximo Application Suite” on page 734](#).

Backing up Maximo Visual Inspection persistent volumes

Back up Maximo Visual Inspection data that is stored in persistent volume claims.

About this task

All Maximo Visual Inspection data is stored in a persistent volume claim. This persistent volume claim can exist in any directory that is accessible by all worker nodes. You need to back up the Maximo Visual Inspection data in the data folder.

In the data folder, each user has a subdirectory that includes the data sets, dnn-sources or imported custom models, and trained models that the user created.

Procedure

1. Optional: To preserve data integrity, stop the vision-service pod.
 - a) In the Red Hat OpenShift administration console, set the value of the operator pod for the mas-visualinspection deployment to 0.
 - b) In the Red Hat OpenShift administration console, set the operator pod for the vision-service deployment to 0.
2. Back up the persistent volume claim data.

The method that you use depends on where you plan to save the backup data.

Location of backup data	Method
Cloud Object Storage	Back up the persistent volume claim data to Cloud Object Storage.
Local workstation	Download the persistent volume claim data to a local workstation.

3. Restart the vision-service pod to bring the application back online.
 - a) In the Red Hat OpenShift administration console, set the value of the operator pod for the mas-visualinspection deployment to 1.

Backing up the persistent volume claim data to Cloud Object Storage

You can copy the persistent volume claim data from the pod to Cloud Object Storage directly. Create a network policy to allow access and then create a job to provide the Cloud Object Storage credentials and back up the persistent volume claim data.

About this task

Create a Red Hat OpenShift job to copy the persistent volume claim data to Cloud Object Storage. Because the Maximo Application Suite namespace blocks the egress network by default, you need to create a network policy to allow the backup job to access Cloud Object Storage.

Procedure

1. Create a network policy.

In the Red Hat OpenShift web console, open the **Import YAML** page. Copy and paste the following NetworkPolicy content, replace the variables, and click **Create**.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: allow-vi-backup
  namespace: mas-{MAS_INSTANCE_ID}-visualinspection
spec:
  podSelector:
    matchLabels:
      job-name: vi-backup
  egress:
    - {}
  policyTypes:
    - Egress
```

2. Create a Red Hat OpenShift job to backup persistent claim volume data to Cloud Object Storage.

In the Red Hat OpenShift web console, open the **Import YAML** page. Copy and paste the following Job content, replace the variables, and click **Create**.

```
apiVersion: batch/v1
kind: Job
metadata:
  name: vi-backup
  namespace: mas-{MAS_INSTANCE_ID}-visualinspection
spec:
  parallelism: 1
  completions: 1
  backoffLimit: 6
  template:
    metadata:
      name: vi-backup
    labels:
      app: vi-backup
    spec:
      serviceAccountName: ibm-mas-visualinspection-operator
      serviceAccount: ibm-mas-visualinspection-operator
      containers:
        - name: vi-backup
          image: rclone/rclone:1.62.2
          command:
            - sh
            - '-c'
            - >-
              export RCLONE_CONFIG=/tmp/rclone.conf; rclone config create
              brcos s3 provider={S3_PROVIDER} endpoint={S3_URL} access_key_id={S3_ACCESS_KEY}
              secret_access_key={S3_SECRET_KEY} region={S3_REGION}; rclone --links --progress --no-check-
              certificate --config /tmp/rclone.conf copy /opt/powerai-vision/data brcos:{S3_BUCKET}/data
          volumeMounts:
            - name: data-mount
              mountPath: /opt/powerai-vision/data
              subPath: data
          securityContext:
            privileged: false
            readOnlyRootFilesystem: false
            allowPrivilegeEscalation: false
          restartPolicy: OnFailure
      volumes:
        - name: data-mount
          persistentVolumeClaim:
            claimName: {MAS_INSTANCE_ID}-data-pvc
```

The following values are used in the Job content:

{MAS_INSTANCE_ID}

The Maximo Application Suite instance ID.

{S3_PROVIDER}

The Amazon S3 storage provider that you are using. For more information, see [Amazon S3 Storage Providers](#).

{S3_URL}

Endpoint for S3 API.

{S3_ACCESS_KEY}

Your Amazon Web Services access key ID.

{S3_SECRET_KEY}

Your Amazon Web Services secret access key or password.

{S3_REGION}

The region to connect to.

{S3_BUCKET}

The backup data is stored in `{S3_BUCKET}/data`.

3. In the Red Hat OpenShift web console, wait for the **vi-backup** job to be completed.

You can check the log of this job for details of the progress.

What to do next

- Check that the backup data is copied to the Cloud Object Storage successfully.
- After the backup job is completed, delete the network policy and job.

Downloading the persistent volume claim data

If you are not using Cloud Object Storage, you can download the persistent volume claim data from the service pod to your local workstation. The download speed depends on the speed and status of the network that your workstation is using.

Procedure

1. Get the Maximo Visual Inspection service pod name. You need to run the backup data commands in this pod.

```
POD_NAME=$(oc get pod -n mas-${MAS_INSTANCE_ID}-visualinspection | grep ${MAS_INSTANCE_ID}-service | awk '{print $1}')
```

2. Compress the persistent volume claim data in the pod.

```
oc exec ${POD_NAME} -c ${MAS_INSTANCE_ID}-service -n mas-${MAS_INSTANCE_ID}-visualinspection -- bash -c "rm -rf /tmp/data.tar; tar -cf /tmp/data.tar -C /opt/powerai-vision/data ."
```

3. Download the backup file from the pod.

```
oc cp --retries=1 -c ${MAS_INSTANCE_ID}-service mas-${MAS_INSTANCE_ID}-visualinspection/${POD_NAME}:/tmp/data.tar /tmp
```

Backing up Maximo Monitor

If Maximo Monitor is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, you back up MongoDB and the namespace for Maximo Monitor.

Before you begin

- Before you back up Maximo Monitor, you need to back up the IoT tool. For more information, see [“Backing up the IoT tool” on page 747](#).

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Maximo Monitor depends on the IoT tool. Because they share a relational database, you don't need to back up the relational database for Maximo Monitor separately.

You can back up Maximo Monitor by using manual steps.

Related tasks

“[Restoring and validating Maximo Monitor](#)” on page 769

To restore Maximo Monitor, restore the IoT tool, the MongoDB backup, and the Maximo Monitor namespace in Red Hat OpenShift. You can also validate the restoration.

Backing up MongoDB for Maximo Monitor

You must back up the MongoDB database that is used by Maximo Monitor.

Back up the following MongoDB database:

- `mas-{MAS_INSTANCE_ID}_monitor`

For more information, see “[Backing up MongoDB for Maximo Application Suite](#)” on page 734.

Backing up the Maximo Monitor namespace

You must back up four secrets in the namespace.

In the `mas-{MAS_INSTANCE_ID}-add` namespace, back up the following secrets:

datadictionary-*<MAS_WORKSPACE_ID>*

The secret that contains the credentials to access the Asset Data Dictionary APIs.

instance-admin

The secret that contains the default admin group *userid* and *password* that is used by Asset Data Dictionary components.

Use the following commands to back up the secrets:

```
oc get secret datadictionary-{MAS_WORKSPACE_ID} -n mas-{MAS_INSTANCE_ID}-add -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadata.a.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./datadictionary-{MAS_WORKSPACE_ID}.yaml
```

```
oc get secret instance-admin -n mas-{MAS_INSTANCE_ID}-add -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadata.a.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./instance-admin.yaml
```

In the `mas-{MAS_INSTANCE_ID}-monitor` namespace, back up the following secrets:

<MAS_INSTANCE_ID>-*<MAS_WORKSPACE_ID>*-datadictionaryworkspace-workspace-binding

This secret contains the *userid* and API key to access the Asset Data Dictionary APIs.

monitor-kitt

This secret also contains the *userid* and API key to access the Asset Data Dictionary APIs.

Use the following command to back up the secrets:

```
oc get secret {MAS_INSTANCE_ID}-{MAS_WORKSPACE_ID}-datadictionaryworkspace-workspace-binding -n mas-{MAS_INSTANCE_ID}-monitor -o yaml | yq 'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadata.a.uid, .metadata.annotations["kubectl.kubernetes.io/last-applied-configuration"], .status)' > ./{MAS_INSTANCE_ID}-{MAS_WORKSPACE_ID}-datadictionaryworkspace-workspace-binding.yaml
```



```
oc get secret monitor-kitt -n mas- $\{MAS\_INSTANCE\_ID\}$ -monitor -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadat
a.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-configuration"], .status)' > ./
monitor-kitt.yaml
```

Backing up Maximo Optimizer

If Maximo Optimizer is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, back up MongoDB for Maximo Optimizer.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can back up Maximo Optimizer by using manual steps.

Backing up MongoDB for Maximo Optimizer

You must back up the MongoDB database that is used by Maximo Optimizer.

Back up the following MongoDB database:

- `mas- $\{MAS_INSTANCE_ID\}$ _optimizer`

For more information, see [“Backing up MongoDB for Maximo Application Suite” on page 734](#).

Backing up Maximo Collaborate

If Maximo Collaborate is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, you back up the namespace for Maximo Collaborate, CouchDB data, and Watson Discovery data. You can also back up the data in the cloud object storage.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

Related tasks

[“Restoring and validating Maximo Collaborate” on page 773](#)

To restore Maximo Collaborate, restore the namespace for Maximo Collaborate, CouchDB data, and Watson Discovery data. Restoring data from the cloud object storage is optional. You can also validate the restoration.

Backing up Maximo Collaborate namespace

You need to back up one secret in the namespace.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

In the `mas- $\{MAS_INSTANCE_ID\}$ -collaborate` namespace, back up the following secret:

collaborate-secret

The secret that contains the Maximo Collaborate CouchDB access credentials and other configurations, such as multitenant service API key.

Use the following command to back up the secret:

```
oc get secret collaborate-secret -n mas-{MAS_INSTANCE_ID}-collaborate -o yaml | yq
'del(.metadata.creationTimestamp, .metadata.ownerReferences, .metadata.resourceVersion, .metadat
a.uid, .metadata.annotations["kubect1.kubernetes.io/last-applied-configuration"], .status)' > ./
{MAS_INSTANCE_ID}-collaborate-secret.yaml
```

Backing up CouchDB data for Maximo Collaborate

Maximo Collaborate saves data in an internal Apache CouchDB. Use the operator that is provided by Maximo Collaborate to back up the internal CouchDB data.

About this task

Use the API resource `CustomResourceDefinitions` to create the `CollaborateBackup` custom resource. To learn more about custom resource definitions, see [CustomResourceDefinitions](#).

Procedure

1. Use the operator to back up the CouchDB data and copy it to the cloud object storage that is configured in Maximo Application Suite.
 - a) In the Red Hat OpenShift web console, in the header, click the plus icon to open the **Import YAML** page.
 - b) Copy and paste the following content:

```
apiVersion: apps.mas.ibm.com/v1
kind: CollaborateBackup
metadata:
  name: {MAS_INSTANCE_ID}
  namespace: {namespace}
  labels:
    app.kubernetes.io/instance: {MAS_INSTANCE_ID}
    app.kubernetes.io/managed-by: ibm-mas-collaborate
    app.kubernetes.io/name: ibm-mas-collaborate
    mas.ibm.com/instanceId: {MAS_INSTANCE_ID}
    mas.ibm.com/applicationId: collaborate
spec:
  backupBucketName: {mascos_bucket_name}
  backupFileName: {mascos_file_name}
```

- c) Replace the variables with the values for your environment.
 - {namespace}***
The Maximo Collaborate namespace, which is typically `mas-{MAS_INSTANCE_ID}-collaborate`.
 - {mascos_file_name}***
The name of the backup file.
 - {mascos_bucket_name}***
The path where the backup is uploaded to the cloud object storage.
 - d) Click **Create**.
2. Optional: Check the backup progress by entering the following command:

```
oc get CollaborateBackup -n {namespace}
```

Check the status in the output.

What to do next

After the backup process is completed, delete the `CollaborateBackup` instance to avoid repeatedly triggering the CouchDB backup task.

```
oc delete CollaborateBackup {MAS_INSTANCE_ID} -n {namespace}
```

Backing up Watson Discovery data for Maximo Assist

You need to back up the Watson Discovery data for Maximo Assist by following the procedure for backing up IBM Cloud Pak for Data.

Note: Starting in Maximo Application Suite 9.0, Watson Discovery, which is used to support the query and diagnose functions, is no longer available as a dependency in Maximo Assist. If Maximo Assist is already deployed and activated with Watson Discovery, and you are upgrading to Maximo Application Suite 9.0, before you can complete the upgrade, you must contact IBM® Support to help with the manual removal of Watson Discovery.

To learn more about backing up IBM Cloud Pak for Data, see [Backing up and restoring data in Cloud Pak for Data](#).

Backing up cloud object storage for Maximo Collaborate

Back up Maximo Collaborate data that is stored in cloud object storage only for the disaster recovery scenario.

Before you begin

You need to install Rclone, which is a command-line program to manage files on cloud storage. For more information, see [Rclone downloads](#).

About this task

Maximo Collaborate stores documents in the cloud object storage that is configured on the **Configuration** page in Maximo Application Suite. For the disaster recovery scenario, if your source cluster and your target cluster are in different regions and your cloud object storage does not support cross-region data replication, you must manually copy the data between regions. For more information, see [“Backup and restore scenarios overview”](#) on page 728.

The method to copy data in cloud object storage depends on the cloud object storage provider. Refer to the cloud object storage provider’s documentation for details.

Procedure

1. Get the access information for the cloud object storage.

- a) Get the URL.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.url | base64 --decode
```

- b) Get the access key.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.username | base64 --decode
```

- c) Get the access secret.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.password | base64 --decode
```

2. Configure the Rclone tool.

- For an on-premises deployment, the default Maximo Collaborate installation uses Red Hat OpenShift Ceph object storage. For more information, see [Ceph](#).
- For an IBM Cloud deployment, you can use IBM Cloud Object Storage. For more information, see [IBM Cloud Object Storage](#).
- For an Amazon Web Services cloud deployment, you can use AWS S3. For more information, see [s3 configuration for the AWS S3 provider](#).

When you set up Rclone connection for the cloud object storage provider, you specify a configuration name. Provide this configuration name in the Rclone command for connecting to the cloud object storage.

3. Download the data from cloud object storage.

a) List the buckets.

```
rclone lsf --no-check-certificate ${RCLONE_CFG_NAME}
```

b) Download the data into a bucket.

```
rclone copy --no-check-certificate --progress ${RCLONE_CFG_NAME}:${COS_BUCKET_NAME} /tmp/  
collaborate-cos/${COS_BUCKET_NAME}
```

Backing up Maximo Health or Maximo Health and Predict - Utilities

If Maximo Health or Maximo Health and Predict - Utilities is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, back up the Watson Studio data for Maximo Health or Maximo Health and Predict - Utilities.

Before you begin

- Before you back up Maximo Health or Maximo Health and Predict - Utilities, you must back up Maximo Manage. For more information, see [“Backing up Maximo Manage” on page 742](#).
- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can back up Maximo Health by using manual steps.

Maximo Health and Maximo Health and Predict - Utilities depend on Maximo Manage. Because they share a relational database, you don't need to back up the relational database for Maximo Health or Maximo Health and Predict - Utilities separately.

Procedure

Back up the Watson Studio data.

Follow the steps to export a project for Watson Studio. For more information, see [Exporting a project \(Watson Studio\)](#).

Related tasks

[“Restoring Maximo Health or Maximo Health and Predict - Utilities” on page 776](#)

You must restore Maximo Manage, and then you restore Maximo Health or Maximo Health and Predict - Utilities and also your Watson Studio data.

Backing up Maximo Predict

If Maximo Predict is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, back up Watson Studio data and Watson Machine Learning data for Maximo Predict.

Before you begin

- Before you back up Maximo Predict, you need to back up Maximo Manage. For more information, see [“Backing up Maximo Manage” on page 742](#).

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the `oc` command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the `yq` command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Maximo Predict depends on Maximo Manage. Because they share a relational database, you don't need to back up the relational database for Maximo Predict separately.

Procedure

1. Back up Watson Studio data.

Follow the steps to export a project for Watson Studio. For more information, see [Exporting a project \(Watson Studio\)](#).

2. Back up Watson Machine Learning data.

Follow the steps to export assets for Watson Machine Learning. For more information, see [Exporting space assets](#).

Related tasks

[“Restoring Maximo Predict” on page 777](#)

You must restore Maximo Manage, and then you restore Maximo Predict and also your Watson Studio and Watson Machine Learning data.

Restoring and validating Maximo Application Suite

You can restore individual IBM Maximo Application Suite applications from your backups.

Before you begin

Ensure that both, the backup and restore Red Hat OpenShift clusters, are running and using the same versions of Maximo Application Suite core and Maximo Application Suite applications.

When you restore Maximo Application Suite from a backup, use the same instance ID and workspace ID that you used for the backup.

Restoring MongoDB for Maximo Application Suite

You need to restore the MongoDB databases that are used by Maximo Application Suite.

About this task

You can restore MongoDB by using manual steps.

Restoring by using MongoDB CE pods internally

You can restore the MongoDB databases by using a MongoDB CE replica set pod that is within the Red Hat OpenShift cluster.

Procedure

1. Copy the backup archives from the local system to the `mongod` container.

```
oc cp ./backups mongoce/mas-mongo-ce-0:/data -c mongod
```

2. Restore the databases.

Use `--drop` with the restore to ensure that the collections that are restored from the backup are identical to what is in the backup archives.

Example restoring Maximo Application Suite core databases.

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "mongorestore --host={{primaryHost}} --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt --drop --archive=/data/backups/mas_{{instanceId}}_core.archive"
```

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "mongorestore --host={{primaryHost}} --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt --drop --archive=/data/backups/mas_{{instanceId}}_catalog.archive"
```

If you are resorting to a new instance of Maximo Application Suite, which is already configured:

- a) The *instanceID* must be the same as the one for the backups.
- b) Do not restore the Maximo Application Suite core database collection `OauthClient`.

To exclude the `OauthClient` collection, use the following parameter: `--nsExclude=mas_{{instanceId}}_core.OauthClient`.

3. To remove backups from the `mongod` container that is in the `mas-mongo-ce-0` pod, run the following command:

```
oc exec -it mas-mongo-ce-0 -c mongod --namespace mongoce -- bash -c "rm -rf /data/backups"
```

Restoring MongoDB by using port forwarding

You can restore the MongoDB databases by using port forwarding. When you use port forwarding, you connect directly to the primary replica set host.

Procedure

1. In a command line, run the following command to start the port forwarding for the primary replica set member pod.

```
oc port-forward {{podThatIsPrimary}} -n mongoce 28015:27017
```

2. In another command line, restore the databases.

Use `--drop` with the restore to ensure that the collections that are restored from the backup are identical to what is in the backup archives.

```
mongorestore --host=localhost --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt --drop --archive=./backups/mas_{{instanceId}}_core.archive
```

```
mongorestore --host=localhost --port=27017 --username=admin --password={{decodedPassword}} --authenticationDatabase=admin --ssl --sslCAFile=/var/lib/tls/ca/ca.crt --drop --archive=./backups/mas_{{instanceId}}_catalog.archive
```

If you are resorting to a new instance of Maximo Application Suite, which is already configured:

- a) The `{{instanceId}}` must be the same as the one for the backups.
- b) Do not restore the Maximo Application Suite core database collection `OauthClient`.

To exclude the `OauthClient` collection, use the following parameter: `--nsExclude=mas_{{instanceId}}_core.OauthClient`.

Related information

[Use port forwarding to access applications in a cluster.](#)

Restoring and validating Maximo Application Suite core

You can restore Maximo Application Suite core databases and the namespace in Red Hat OpenShift.

About this task

You can restore Maximo Application Suite core by using manual steps.

Restoring MongoDB for Maximo Application Suite core

You need to restore the MongoDB databases that are used by Maximo Application Suite core.

About this task

You need to choose a connection strategy: internally within the Red Hat OpenShift cluster or by using port forwarding to connect to the single replica member. For more information, see [“Restoring MongoDB for Maximo Application Suite” on page 757](#).

Procedure

- Restore the Maximo Application Suite core databases.
 - `mas_{{instanceId}}_core`
 - `mas_{{instanceId}}_catalog`

Related tasks

[“Backing up MongoDB for Maximo Application Suite core” on page 738](#)

You need to back up specific MongoDB that are used by Maximo Application Suite core. Creating backups of the MongoDB databases that Maximo Application Suite core depends on, must be planned and scheduled to happen on a timely basis.

Restoring the Maximo Application Suite core namespace in Red Hat OpenShift

How you restore resources in Maximo Application Suite core and which configurations you restore depend on your environment and how you created the data backup. You can apply the configurations for Maximo Application Suite core by using a command for every file or by using the script.

About this task

If a custom `ClusterIssuer` is used, you must specify the value at installation.

If IBM Suite License Service is deployed and configured, you must specify the license file and ID that is backed up.

The Maximo Application Suite installation applies the necessary IBM Cloud Pak foundational services dependencies.

If manual certificates are enabled, when you restore Maximo Application Suite core, you must apply the `.yaml` file backup of the `secret {{ instanceId }}-cert-public` file.

For more information about when to apply the `MongoCfg`, `BasCfg`, and `SlsCfg` configurations, see [“When to apply the backed up configurations” on page 759](#).

Procedure

- To restore manually, use the `oc cli -f <filename>` command to restore the resources that you backed up.
The command applies the configuration that is specified to the `*.yaml` file that is provided by using `-f <filename>`.
- To restore the Maximo Application Suite core namespace, use the script `mascore-backup-restore.sh`.
The `mascore-backup-restore.sh` script runs the `oc apply` command on all the `*.yaml` files that are in a specified folder.
To access the script, go to [mascore-backup-restore.sh](#).

When to apply the backed up configurations

When you are restoring Maximo Application Suite core, new configurations for the `config.mas.ibm.com`, `MongoCfg`, and `SlsCfg` are generated. It is important to understand whether to apply the newly generated configurations or the backed up configurations.

Only Maximo Application Suite core

If you are restoring only Maximo Application Suite core, apply all the backed up configurations.

Maximo Application Suite core, IBM Suite License Service, and MongoDB

When you restore Maximo Application Suite core to a new Red Hat OpenShift cluster, new configurations for the following resources are generated.

- `config.mas.ibm.com`
- `MongoCfg`
- `SlsCfg`

These resources are used to configure Suite License Service and MongoDB. Apply the new `MongoCfg` and `SlsCfg` configurations. Do not apply the backed up configurations.

MongoDB or Suite License Service instance	Configurations	Which configuration to apply
New MongoDB instance	The <code>MongoCfg</code> configuration and its corresponding secret	Apply the new configurations.
New Suite License Service instance	The <code>SlsCfg</code> configuration	Apply the new configuration.
Existing MongoDB instance. This instance can be hosted by a third party.	<code>MongoCfg</code> and its corresponding secret	You can apply the backed up configuration or the new one.
Existing Suite License Service instance	The <code>SlsCfg</code> configuration	Apply the backed up configuration.

Suite License Service and MongoDB instances hosted outside of the Red Hat OpenShift cluster

If your Suite License Service and MongoDB instances are hosted outside of the Red Hat OpenShift cluster that you are restoring the instances to, do not apply the backed up configurations.

Validating restoration of Maximo Application Suite core

To confirm a successful restore, log in to the Maximo Application Suite admin UI.

Using Red Hat OpenShift Container Platform or `oc get <resource type> <resource name>`, verify that the restored instances of the resources in the list are in **Ready** state and have no current errors.

- `core.mas.ibm.com`
 - Suite only one instance named `{instanceId}`
 - All instances of `Workspace`
- All the instances of `config.mas.ibm.com`
 - `MongoCfg`
 - `KafkaCfg`
 - `JdbcCfg`
 - `SlsCfg`
 - `BasCfg`
 - `SntpCfg`

- WatsonStudioCfg
- ObjectStorageCfg
- PushNotificationCfg
- ScimCfg
- IDPCfg
- All the instances of `addons.mas.ibm.com`
 - AppConnect
 - Humai
 - MVIEdge

Restoring and validating Maximo Manage

You can restore Maximo Manage and validate the restored files.

Before you begin

Before you restore Maximo Manage, you must restore Maximo Application Suite core.

About this task

You can restore Maximo Manage by using manual steps.

Restoring Maximo Manage databases

You can use your database backups to restore the databases that you use for Maximo Manage.

Restore the MongoDB, which contains the data that is created in Maximo Application Suite and synced into Maximo Manage database. For more information, see [“Restoring MongoDB for Maximo Application Suite”](#) on page 757.

To restore the database that is used by Maximo Manage, follow the database provider's documentation. For more information, see [Backing up and restoring Db2](#).

Related concepts

[“Backing up Maximo Manage databases”](#) on page 743

Back up the database to ensure continuous availability of data. Back up the database before and after you upgrade Maximo Manage. To back up, you create a copy of the data, which can be recovered and restored if you have a data failure.

[“Backup and restore process for relational database management systems”](#) on page 732

The backup and restore process for relational database management systems involves a full backup, incremental backups, and transaction logs.

Restoring the Maximo Manage namespace in Red Hat OpenShift

You can restore your backed up Maximo Manage namespace manually or by using the `manage-backup-restore.sh` script.

Before you begin

- Before you restore the Maximo Manage namespace, you must restore the Maximo Application Suite core namespace.
- Get the script [manage-backup-restore.sh](#).

Procedure

- To restore manually, use the `oc cli` command to restore the resources that you backed up.

The command applies the configuration that is specified in the file that is provided. Run the command for each *.yaml file that you created when you did the backup.

```
oc apply -f <filename>
```

When you restore the Maximo Manage namespace, you must apply the YAML backup of the `{instanceId}-{workspaceId}-cert-public-81` secret.

- To restore the Maximo Manage namespace by using the `manage-backup-restore.sh` script, run the following command.

```
manage-backup-restore.sh -w <mas-workspace-id> -i <MAS_INSTANCE_ID> -f <BACKUP_FOLDER> -m restore
```

```
manage-backup-restore.sh -w <mas-workspace-id> -i dev -f ./ -m restore
```

-w, --mas-workspace-id

The Maximo Manage workspace ID where the information is restored to.

-i, --mas-instance-id

The Maximo Application Suite instance ID where the information is restored to.

-f, --backup-folder

The folder where the backup artifacts are read from.

-m, --mode

Indicates whether to backup or restore. Use the value **restore**.

The script runs the **oc apply** command on all the *.yaml files in a specified folder, which restores the files.

Restoring Maximo Manage attachments

You can restore attachments for Maximo Manage, which can include receipts, certifications, invoices, and any files that are attached to assets, work orders, or job plans.

Procedure

1. Open the Maximo Manage namespace.

```
oc project <your "Manage" namespace>
```

2. Get the pod name of your server \$POD for the `all` or `ui` containers.

```
oc get pods
```

3. Copy the backup archive folder that contains your attachments from the local system to the server for the `all` or `ui` containers.

```
oc cp $POD:/$MOUNTPATH $LOCALFOLDER -c all
```

4. Expand the `tar` file in the `mountPath` folder.

```
oc exec -it $POD -c $CONTAINER --namespace $NAMESPACE -- bash -c "tar -xz <tarfilename>"
```

Validating restoration of Maximo Manage

To verify that the restore was successful, check the status of the Maximo Manage workspace CR in the Red Hat OpenShift console.

Procedure

1. Verify the success of the following processes.
 - The build was completed.

- The **Maxinst** command was completed.
 - The server bundles were deployed.
2. Log in to Maximo Manage.

Restoring and validating the IoT tool

By using your backups, you can restore the IoT tool databases, relational databases, and namespace in Red Hat OpenShift.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the **yq** command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can restore the IoT tool by using manual steps.

Related tasks

[“Backing up the IoT tool” on page 747](#)

You can back up the IoT tool if it is deployed on an IBM Cloud , Amazon Web Services, or on-premises environment.

Restoring MongoDB for the IoT tool

You need to restore the MongoDB databases that are used by the IoT tool.

Restore the following MongoDB databases:

- `iot_{MAS_INSTANCE_ID}_cs_activity_db`
- `iot_{MAS_INSTANCE_ID}_d_actions`
- `iot_{MAS_INSTANCE_ID}_d_core`
- `iot_{MAS_INSTANCE_ID}_d_dashboard`
- `iot_{MAS_INSTANCE_ID}_d_deviceregistry`
- `iot_{MAS_INSTANCE_ID}_d_dmserver`
- `iot_{MAS_INSTANCE_ID}_d_dsc`
- `iot_{MAS_INSTANCE_ID}_d_infomgmt`
- `iot_{MAS_INSTANCE_ID}_d_provision_s2s`
- `iot_{MAS_INSTANCE_ID}_d_riskmgmtsecurity`
- `iot_{MAS_INSTANCE_ID}_organizations`

For more information, see [“Restoring MongoDB for Maximo Application Suite” on page 757](#).

Restoring the IoT tool relational databases

If the database instance is hosted within the Red Hat OpenShift cluster, you can restore the relational databases that are used by the IoT tool.

To ensure the consistency of user data, restore MongoDB first, followed by the IoT relational database.

If you are not using Db2, refer to your database provider documentation.

- To restore Db2, use the instructions in the Db2 documentation.
 - For more information, see [Backing up and restoring Db2](#).
 - The Db2 container name is `c-mas-{MAS_INSTANCE_ID}-system-db2u-0`.
 - For more information, see [Restoring Db2 from an online backup using commands](#).

Restoring the IoT tool namespace

You need to restore three encryption secrets for the `mas-MAS_INSTANCE_ID-iot` namespace.

About this task

You need to restore each one of the encryption secrets.

- `actions-credsenckey`
- `auth-encryption-secret`
- `provision-creds-enckey`

Procedure

1. Use the command `oc apply -f` to restore each secret.

```
oc apply -f ./iMAS_INSTANCE_ID-secret_name.yaml
```

2. Restart the `datapower-datapower` pods.

```
oc get pod -n mas-MAS_INSTANCE_ID-iot | grep "datapower-datapower" | awk '{print $1}' | xargs oc delete pod -n mas-MAS_INSTANCE_ID-iot
```

3. Wait for `datapower-datapower` pods to be running and in the ready status.

To monitor the status of the `datapower-datapower` pods, use the following command:

```
oc get pod -n mas-MAS_INSTANCE_ID-iot | grep datapower-datapower
```

Validating restoration of the IoT tool

You can verify whether the IoT tool was successfully restored by checking the deployment status and then opening the IoT tool dashboard to the devices and the users.

Procedure

1. Check the IoT tool deployment status.
 - a) Log in to Maximo Application Suite as an administrator.
 - b) From the side navigation menu, click **Applications**.
 - c) Click IoT.
 - d) Review the details page for the application and workspace and check whether their statuses are **Ready**.
2. Open the IoT tool.

In Maximo Application Suite, click the **AppSwitcher** icon and then click **IoT** to open the IoT tool.
3. In the IoT tool, check the IoT tool devices.
 - a) In the IoT tool, from the side navigation menu, click **Devices**.
 - b) On the **Browse Devices** page, check the list of devices. Verify that all the devices from before the backup are listed.
4. In the IoT tool, check the IoT tool users.
 - a) From the side navigation menu, click **Members**.
 - b) On the **Browse Members** page, check the list of the users who have access to the IoT tool. Verify that all the IoT tool users from before the backup are listed.

Restoring and validating Maximo Visual Inspection

To restore Maximo Visual Inspection, restore the MongoDB backup and the backed up data from the persistent volume claim and then restart the service pods. You can also validate the restoration.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the **yq** command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can restore Maximo Visual Inspection by using manual steps.

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Visual Inspection are based on different restore scenarios.

For more information about restore scenarios, see [“Backup and restore scenarios overview”](#) on page 728.

To learn more about how to deploy and undeploy Maximo Visual Inspection, see [“Deploying IBM Maximo Visual Inspection”](#) on page 391 and [“Deactivating and deleting applications”](#) on page 472.

Procedure

- For the crash recovery scenario, restore the Maximo Visual Inspection application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Visual Inspection deployment.
 - b) Redeploy the Maximo Visual Inspection application.
 - c) Restore the MongoDB data and the persistent volume claim data from the backups.

For more information about crash recovery, see [Crash recovery](#).
- For the disaster recovery scenario, restore the Maximo Visual Inspection application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Visual Inspection.
 - d) Restore the MongoDB data and the persistent volume claim data from the backups.

For more information about disaster recovery, see [Disaster recovery](#).
- For the restore data scenario, you can restore Maximo Visual Inspection backup data only to the source cluster.
 - a) Select and restore the backed up data from the MongoDB and persistent volume claim backups.

You don't need to undeploy the existing Maximo Visual Inspection deployment.

For more information about restoring data, see [Restore data](#).

Related tasks

[“Backing up Maximo Visual Inspection”](#) on page 748

If Maximo Visual Inspection is deployed on IBM Cloud, Amazon Web Services, or an on-premises environment, back up the document data stores and persistent volumes for Maximo Visual Inspection.

Restoring MongoDB for Maximo Visual Inspection

You need to restore the MongoDB databases that are used by Maximo Visual Inspection.

Restore the following MongoDB databases:

- `mas- $\{MAS_INSTANCE_ID\}$ -visualinspection`
- `mas- $\{MAS_INSTANCE_ID\}$ -edgeman`

For more information, see [“Restoring MongoDB for Maximo Application Suite” on page 757](#).

Restoring Maximo Visual Inspection persistent volumes

Restore the Maximo Visual Inspection data that is stored in the persistent volume claims.

Procedure

1. Get the *VI* user and *group ID*.

```
oc get deployment ${MAS_INSTANCE_ID}-service -n mas-${MAS_INSTANCE_ID}-visualinspection -o
yaml | yq .spec.template.spec.securityContext.runAsUser
```

2. Restore the persistent volume claim data.

The method that you use depends on where you stored the backup data.

- If you use IBM Cloud Object Storage, restore the persistent volume claim data directly from IBM Cloud Object Storage. For more information, see [Backing up the persistent volume claim data to IBM Cloud Object Storage](#).
- If you use your local workstation, upload the persistent volume claim data from a local workstation. For more information, see [“Downloading the persistent volume claim data” on page 751](#).

What to do next

Set the correct owner of the data folders.

Restoring persistent volume claim data from Cloud Object Storage

You can copy the persistent volume claim data from Cloud Object Storage directly to the pod. Create a network policy to allow access and then create a job to provide the Cloud Object Storage credentials and restore the persistent volume claim data.

Before you begin

Get the *VI* user and *group ID*.

```
oc get deployment ${MAS_INSTANCE_ID}-service -n mas-${MAS_INSTANCE_ID}-visualinspection -o yaml
| yq .spec.template.spec.securityContext.runAsUser
```

About this task

You need to create a Red Hat OpenShift job to copy the persistent volume claim data to the pod. Because the Maximo Application Suite namespace blocks the egress network by default, you need to create a network policy to allow the restore job to access Cloud Object Storage.

Procedure

1. Create a network policy.

In the Red Hat OpenShift web console, open the **Import YAML** page. Copy and paste the following NetworkPolicy content, replace the variables, and click **Create**.

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: allow-vi-restore
  namespace: mas-${MAS_INSTANCE_ID}-visualinspection
spec:
  podSelector:
    matchLabels:
      job-name: vi-restore
  egress:
    - {}
  policyTypes:
    - Egress
```

2. Create a Red Hat OpenShift job to restore the persistent claim volume data from Cloud Object Storage to a pod.

In the Red Hat OpenShift web console, open the **Import YAML** page. Copy and paste the following Job content, replace the variables, and click **Create**.

```
apiVersion: batch/v1
kind: Job
metadata:
  name: vi-restore
  namespace: mas-{MAS_INSTANCE_ID}-visualinspection
spec:
  parallelism: 1
  completions: 1
  backoffLimit: 6
  template:
    metadata:
      name: vi-restore
    labels:
      app: vi-restore
    spec:
      serviceAccountName: ibm-mas-visualinspection-operator
      serviceAccount: ibm-mas-visualinspection-operator
      securityContext:
        runAsUser: {VI_USER}
        runAsGroup: {VI_USER}
        runAsNonRoot: true
      containers:
        - name: vi-restore
          image: rclone/rclone:1.62.2
          command:
            - sh
            - '-c'
            - >-
              rm -rf /opt/powerai-vision/data/*; export RCLONE_CONFIG=/tmp/rclone.conf;
              rclone config create brcos s3 provider={S3_PROVIDER} endpoint={S3_URL}
              access_key_id={S3_ACCESS_KEY} secret_access_key={S3_SECRET_KEY} region={S3_REGION}; rclone
              --links --progress --no-check-certificate --config /tmp/rclone.conf copy brcos:{S3_BUCKET}/
              data /opt/powerai-vision/data;
          volumeMounts:
            - name: data-mount
              mountPath: /opt/powerai-vision/data
              subPath: data
          securityContext:
            privileged: false
            readOnlyRootFilesystem: false
            allowPrivilegeEscalation: false
          restartPolicy: OnFailure
          volumes:
            - name: data-mount
              persistentVolumeClaim:
                claimName: {MAS_INSTANCE_ID}-data-pvc
```

The following values are used in the Job content:

{MAS_INSTANCE_ID}

The Maximo Application Suite instance ID.

{S3_PROVIDER}

The Amazon S3 storage provider that you are using. For more information, see [Amazon S3 Storage Providers](#).

{S3_URL}

Endpoint for S3 API.

{S3_ACCESS_KEY}

Your Amazon Web Services access key ID.

{S3_SECRET_KEY}

Your Amazon Web Services secret access key or password.

{S3_REGION}

The region to connect to.

{S3_BUCKET}

The backup data is stored in **{S3_BUCKET}**/data.

{VI_USER}

The owner of the data folders.

3. In the Red Hat OpenShift web console, wait for the **vi-restore** job to be completed.
You can check the log of this job for details of the progress.

What to do next

After the job is completed, delete the network policy and job.

Uploading persistent volume claim data

If you are not using Cloud Object Storage, you can upload the persistent volume claim data from your local workstation to the service pod. The upload speed depends on the speed and status of the network that your workstation is using.

Before you begin

Get the *VI user* and *group ID*.

```
oc get deployment ${MAS_INSTANCE_ID}-service -n mas-${MAS_INSTANCE_ID}-visualinspection -o yaml | yq .spec.template.spec.securityContext.runAsUser
```

Procedure

1. Get the Maximo Visual Inspection service pod name. Run the restore data commands in this pod.

```
POD_NAME=$(oc get pod -n mas-${MAS_INSTANCE_ID}-visualinspection | grep ${MAS_INSTANCE_ID}-service | awk '{print $1}')
```

2. Upload the backup file to the pod.

```
oc cp --retries=-1 -c ${MAS_INSTANCE_ID}-service /tmp/data.tar mas-${MAS_INSTANCE_ID}-visualinspection/${POD_NAME}:/tmp/data.tar
```

3. Extract the backup file in the pod.

```
oc exec ${POD_NAME} -c ${MAS_INSTANCE_ID}-service -n mas-${MAS_INSTANCE_ID}-visualinspection -- bash -c "rm -rf /opt/powerai-vision/data/*; tar -xf /tmp/data.tar -C /opt/powerai-vision/data; chown -R ${VI_USER}:${VI_USER} /opt/powerai-vision/data"
```

Restarting pods for Maximo Visual Inspection after restoration

After you restore the data for Maximo Visual Inspection, you need to restart several pods for the new configuration to take effect.

Procedure

1. Restart the pods.

```
oc get pods -n mas-${MAS_INSTANCE_ID}-visualinspection | grep ${MAS_INSTANCE_ID}-service | awk '{print $1}' | xargs oc delete pod -n mas-${MAS_INSTANCE_ID}-visualinspection
```

2. Wait for pods to be ready.

```
watch "oc get pods -n mas-${MAS_INSTANCE_ID}-visualinspection | grep ${MAS_INSTANCE_ID}-service"
```

Validating restoration of Maximo Visual Inspection

You verify certain resources in Maximo Visual Inspection to validate the restoration of Maximo Visual Inspection.

Procedure

1. Log in to Maximo Application Suite as a Maximo Visual Inspection user.

2. Open the Maximo Visual Inspection application.
3. Verify that you have access to the following resources.
 - The restored data sets.
 - The restored trained models.
 - The deployed models and that you are able to test them.

Restoring and validating Maximo Monitor

To restore Maximo Monitor, restore the IoT tool, the MongoDB backup, and the Maximo Monitor namespace in Red Hat OpenShift. You can also validate the restoration.

Before you begin

- Restore the IoT tool. For more information, see [“Restoring and validating the IoT tool” on page 763](#).
- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the `oc` command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the `yq` command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can restore Maximo Monitor by using manual steps.

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Monitor are based on different scenarios. For more information, see [“Backup and restore scenarios overview” on page 728](#).

To learn more about how to deploy and undeploy Maximo Monitor, see [“Deploying IBM Maximo Monitor” on page 370](#) and [“Deactivating and deleting applications” on page 472](#).

Procedure

- For the crash recovery scenario, restore the Maximo Monitor application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Monitor deployment.
 - b) Redeploy the Maximo Monitor application.
 - c) Restore the MongoDB data and the data from the namespace backups.
- For the disaster recovery scenario, restore the Maximo Monitor application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Monitor.
 - d) Restore the MongoDB data and the data from the namespace backups.
- For the restore data scenario, you can restore Maximo Monitor backup data only to the source cluster.
 - a) Select and restore the backed up data from the MongoDB and namespace backups.

You don't need to undeploy the existing Maximo Monitor deployment.

Related tasks

[“Backing up Maximo Monitor” on page 751](#)

If Maximo Monitor is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, you back up MongoDB and the namespace for Maximo Monitor.

Restoring MongoDB for Maximo Monitor

You must restore the MongoDB database that is used by Maximo Monitor.

Restore the following MongoDB database:

- mas-`{MAS_INSTANCE_ID}_monitor`

For more information, see [“Restoring MongoDB for Maximo Application Suite” on page 757.](#)

Restoring Maximo Monitor namespace

You must restore four secrets in the namespace and restart four pods.

In the mas-`{MAS_INSTANCE_ID}`-add namespace, restore the following secrets:

datadictionary-`<MAS_WORKSPACE_ID>`

The secret that contains the credentials to access the Asset Data Dictionary APIs.

instance-admin

The secret that contains the default admin group user ID and password that are used by the Asset Data Dictionary components.

Use the following commands to restore the secrets:

```
oc apply -f ./datadictionary-{MAS_WORKSPACE_ID}.yaml
```

```
oc apply -f ./instance-admin.yaml
```

In the mas-`{MAS_INSTANCE_ID}`-monitor namespace, restore the following secrets:

`<MAS_INSTANCE_ID>`-`<MAS_WORKSPACE_ID>`-datadictionaryworkspace-workspace-binding

This secret contains the user ID and API key that are used to access the Asset Data Dictionary APIs.

monitor-kitt

This secret also contains the user ID and API key that are used to access the Asset Data Dictionary APIs.

Use the following commands to restore the secrets:

```
oc apply -f ./{MAS_INSTANCE_ID}-{MAS_WORKSPACE_ID}-datadictionaryworkspace-workspace-binding.yaml
```

```
oc apply -f ./monitor-kitt.yaml
```

After you restore the secrets, you must restart several pods for the new configuration to take effect.

```
oc get pod -n mas-{MAS_INSTANCE_ID}-add | grep "user-store" | awk '{print $1}' | xargs oc delete pod -n mas-{MAS_INSTANCE_ID}-add
```

```
oc get pod -n mas-{MAS_INSTANCE_ID}-add | grep "series-store" | awk '{print $1}' | xargs oc delete pod -n mas-{MAS_INSTANCE_ID}-add
```

```
oc get pod -n mas-{MAS_INSTANCE_ID}-add | grep "graph-store" | awk '{print $1}' | xargs oc delete pod -n mas-{MAS_INSTANCE_ID}-add
```

```
oc get pod -n mas-{MAS_INSTANCE_ID}-monitor | grep "{MAS_INSTANCE_ID}" | awk '{print $1}' | xargs oc delete pod -n mas-{MAS_INSTANCE_ID}-monitor
```

Validating restoration of Maximo Monitor

You can verify whether Maximo Monitor was successfully restored by opening Maximo Monitor and validating that the data for the devices is displayed.

Procedure

1. Log in to Maximo Application Suite and open Maximo Monitor by using the **AppSwitcher** icon.
2. Validate that the data is displayed for the devices.
 - a) From the side navigation menu, click **Setup**.
 - b) Click **Devices** and validate that all the data is displayed.
 - c) On the **Devices tab**, in the **Device types** section, click a device and verify that all the data is displayed for both raw and calculated metrics.

Restoring and validating Maximo Optimizer

To restore Maximo Optimizer, restore the MongoDB backup. You can also validate the restoration.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can restore Maximo Optimizer by using manual steps.

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Optimizer are based on different scenarios. For more information, see [“Backup and restore scenarios overview”](#) on page 728.

To learn more about how to deploy and undeploy Maximo Optimizer, see [“Deploying IBM Maximo Optimizer”](#) on page 410 and [“Deactivating and deleting applications”](#) on page 472.

Procedure

- For the crash recovery scenario, restore the Maximo Optimizer application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Optimizer deployment.
 - b) Redeploy the Maximo Optimizer application.
 - c) Restore the MongoDB data from the backup.
- For the disaster recovery scenario, restore the Maximo Optimizer application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Optimizer.
 - d) Restore the MongoDB data from the backup.
- For the restore data scenario, you can restore Maximo Optimizer backup data only to the source cluster.
 - a) Select and restore the backed up data from the MongoDB backup.

You don't need to undeploy the existing Maximo Optimizer deployment.

Restoring MongoDB for Maximo Optimizer

You must restore the MongoDB database that is used by Maximo Optimizer.

Restore the following Maximo Optimizer database:

- mas_{MAS_INSTANCE_ID}_optimizer

For more information, see [“Restoring MongoDB for Maximo Application Suite” on page 757.](#)

Validating restoration of Maximo Optimizer

You can verify whether Maximo Optimizer was successfully restored by checking the deployment status and then opening Maximo Optimizer to check the jobs, models, and projects.

Procedure

1. Check the Maximo Optimizer deployment status.
 - a) Log in to Maximo Application Suite as an administrator.
 - b) From the side navigation menu, click **Applications**.
 - c) Click **Optimizer**.
 - d) Review the details page for the application and workspace and check whether they are in the **Ready** status.
2. Open Maximo Optimizer by clicking the **AppSwitcher** icon and then clicking **Optimizer**.
3. From the side navigation menu, click **Jobs** and validate that all the data in the dashboard is displayed.
 - a) Select a job.
 - b) Click the **More actions** icon and download each item.
Verify that you can download each item
 - Scenario
 - Attachment
 - Solution
 - Log
4. From the side navigation menu, click **Models**. Validate that all the models are restored.
 - a) Find a model that has attached files and click **Download** to verify that it works.
5. On the side navigation menu, click **Projects**. Validate that the projects are restored.
 - a) Click a project and verify that the fields contain the expected values, especially the **API Key** field.
6. Download and submit a job and validate that it works.
 - a) From the side navigation menu, click **Jobs**.
 - b) On the row for a job, click the **More actions** icon and select **Download scenario**.
 - c) Click **Add job** and provide the following information:
 - Project**
Select the same type of project as the job that you downloaded.
 - Model**
Select the same type of model as the job that you downloaded.
 - Add Scenario File**
Upload the downloaded job.
 - d) Click **Create Job**.
 - e) When the job is completed, verify that the **Job Status** and the **Solve Status** fields are the same as the job that you downloaded.

Restoring and validating Maximo Collaborate

To restore Maximo Collaborate, restore the namespace for Maximo Collaborate, CouchDB data, and Watson Discovery data. Restoring data from the cloud object storage is optional. You can also validate the restoration.

Before you begin

- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Collaborate are based on different scenarios. For more information, see [“Backup and restore scenarios overview”](#) on page 728.

To learn more about how to deploy and undeploy Maximo Collaborate, see [“Deploying IBM Maximo Collaborate”](#) on page 292 and [“Deactivating and deleting applications”](#) on page 472.

Procedure

- For the crash recovery scenario, restore the Maximo Collaborate application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Collaborate deployment.
 - b) Redeploy the Maximo Collaborate application.
 - c) Using the backups, restore the namespace for Maximo Collaborate and then restore the CouchDB data and Watson Discovery data.
 - d) Optional: Restore the data from the cloud object storage backup.
- For the disaster recovery scenario, restore the Maximo Collaborate application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Collaborate.
 - d) Using the backups, restore the namespace for Maximo Collaborate and then restore the CouchDB data and Watson Discovery data.
 - e) Optional: Restore the data from the cloud object storage backup.
- For the restore data scenario, select the data that you want to restore from the Maximo Collaborate backups. You can restore data only to the source cluster.

You don't need to undeploy the existing Maximo Collaborate deployment.

Related tasks

[“Backing up Maximo Collaborate”](#) on page 753

If Maximo Collaborate is deployed on IBM Cloud, Amazon Web Services, or an on-premises environment, you back up the namespace for Maximo Collaborate, CouchDB data, and Watson Discovery data. You can also back up the data in the cloud object storage.

Restoring Maximo Collaborate namespace

You need to restore the secret for Maximo Collaborate in the namespace.

In the `mas-{MAS_INSTANCE_ID}-collaborate` namespace, restore the following secret:

collaborate-secret

The secret that contains the access credentials for the CouchDB that is used by Maximo Collaborate and other access credentials, such as the multitenant service API key.

Use the following command to restore the secret:

```
oc apply -f ./{MAS_INSTANCE_ID}-collaborate-secret.yaml
```

Restoring CouchDB data for Maximo Collaborate

Use the operator that is provided by Maximo Collaborate to restore the CouchDB data.

About this task

Use the API resource `CustomResourceDefinitions` to create the `CollaborateRestore` custom resource. To learn more about custom resource definitions, see [CustomResourceDefinitions](#).

Procedure

1. Use the operator to copy the specified backup file from the cloud object storage that is configured in Maximo Application Suite and restore the data to CouchDB.
 - a) In the Red Hat OpenShift web console, in the header, click the plus icon to open the **Import YAML** page.
 - b) Copy and paste the following content:

```
apiVersion: apps.mas.ibm.com/v1
kind: CollaborateRestore
metadata:
  name: {MAS_INSTANCE_ID}
  namespace: {namespace}
  labels:
    app.kubernetes.io/instance: {MAS_INSTANCE_ID}
    app.kubernetes.io/managed-by: ibm-mas-collaborate
    app.kubernetes.io/name: ibm-mas-collaborate
    mas.ibm.com/instanceId: {MAS_INSTANCE_ID}
    mas.ibm.com/applicationId: collaborate
spec:
  backupBucketName: {mascos_bucket_name}
  backupFileName: {mascos_file_name}
```

- c) Replace the variables with the values for your environment.

{namespace}

The Maximo Collaborate namespace, which is typically `mas-{MAS_INSTANCE_ID}-collaborate`.

{mascos_file_name}

The name of the backup file.

{mascos_bucket_name}

The path to download the backup from cloud object storage.

- d) Click **Create**.
2. Optional: Check the restore progress by entering the following command:

```
oc get CollaborateRestore -n {namespace}
```

Check the status in the output.

What to do next

After the restore process is completed, delete the CollaborateRestore instance to avoid repeatedly triggering the CouchDB restore task.

```
oc delete CollaborateRestore ${MAS_INSTANCE_ID} -n ${namespace}
```

Restoring Watson Discovery data for Maximo Collaborate

You need to restore the Watson Discovery data for Maximo Collaborate by following the procedure for restoring IBM Cloud Pak for Data.

To learn more about restoring IBM Cloud Pak for Data, see [Backing up and restoring data in Cloud Pak for Data](#).

Restoring cloud object storage for Maximo Collaborate

Restore the Maximo Collaborate data that is stored in cloud object storage.

Before you begin

You need to install Rclone, which is a command-line program to manage files on cloud storage. For more information, see [Rclone downloads](#).

About this task

Maximo Collaborate stores documents in the cloud object storage that is specified on the **Configuration** page in Maximo Application Suite. For more information, see [“Backup and restore scenarios overview” on page 728](#).

The method to copy data in cloud object storage depends on the cloud object storage provider. Refer to the cloud object storage provider’s documentation for details.

Procedure

1. Get the access information for the cloud object storage.

- a) Get the URL.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.url | base64 --decode
```

- b) Get the access key.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.username | base64 --decode
```

- c) Get the access secret.

```
oc get secret ${MAS_INSTANCE_ID}-objectstorage-secret -n mas-${MAS_INSTANCE_ID}-collaborate -o yaml | yq .data.password | base64 --decode
```

2. Configure the Rclone tool.

- For an on-premises deployment, the default Maximo Collaborate installation uses Red Hat OpenShift Ceph object storage. For more information, see [Ceph](#).
- For an IBM Cloud deployment, you can use IBM Cloud Object Storage. For more information, see [IBM Cloud Object Storage](#).
- For an Amazon Web Services cloud deployment, you can use AWS S3. For more information, see [s3 configuration for the AWS S3 provider](#).

When you set up the Rclone connection for the cloud object storage provider, you specify a configuration name. Provide this configuration name in the Rclone command for connecting to the cloud object storage.

3. Upload the data to cloud object storage.

a) Upload bucket data.

```
rclone copy --no-check-certificate --progress /tmp/collaborate-cos/${COS_BUCKET_NAME} $  
{RCLONE_CFG_NAME}:${COS_BUCKET_NAME}
```

b) List the buckets.

```
rclone lsf --no-check-certificate ${RCLONE_CFG_NAME}
```

Validating restoration of Maximo Collaborate

You verify specific resources in Maximo Collaborate to validate the restoration of Maximo Collaborate.

Procedure

1. Log in to Maximo Application Suite as an administrator.
2. Open the Collaborate application.
3. From the side navigation menu, click **Manage documents** and verify that all the documents are listed.
 - a) Click any document and verify that the preview works.
4. From the side navigation menu, click **Manage diagnosis libraries** and verify that the libraries are listed.
5. From the side navigation menu, click **Collaborate for technicians** and verify that you can search for and find any document.

Restoring Maximo Health or Maximo Health and Predict - Utilities

You must restore Maximo Manage, and then you restore Maximo Health or Maximo Health and Predict - Utilities and also your Watson Studio data.

Before you begin

- Restore Maximo Manage. For more information, see [“Restoring and validating Maximo Manage” on page 761](#).
- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

You can restore Maximo Health by using manual steps.

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Health or Maximo Health and Predict - Utilities are based on different scenarios. For more information, see [“Backup and restore scenarios overview” on page 728](#).

To learn more about how to deploy and undeploy Maximo Health, see [“Deploying IBM Maximo Health” on page 296](#) and [“Deactivating and deleting applications” on page 472](#).

To learn more about how to deploy and undeploy Maximo Health and Predict - Utilities, see [“Deploying IBM Maximo Health and Predict - Utilities” on page 405](#) and [“Deactivating and deleting applications” on page 472](#).

Procedure

- For the crash recovery scenario, restore the Maximo Health or Maximo Health and Predict - Utilities application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Health or Maximo Health and Predict - Utilities deployment.
 - b) Redeploy the Maximo Health or Maximo Health and Predict - Utilities application.
 - c) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).
- For the disaster recovery scenario, restore the Maximo Health or Maximo Health and Predict - Utilities application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Health or Maximo Health and Predict - Utilities.
 - d) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).
- For the restore data scenario, you can restore Maximo Health or Maximo Health and Predict - Utilities backup data only to the source cluster.
 - a) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).

You don't need to undeploy the existing Maximo Health or Maximo Health and Predict - Utilities deployment.

Related tasks

[“Backing up Maximo Health or Maximo Health and Predict - Utilities” on page 756](#)

If Maximo Health or Maximo Health and Predict - Utilities is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, back up the Watson Studio data for Maximo Health or Maximo Health and Predict - Utilities.

Restoring Maximo Predict

You must restore Maximo Manage, and then you restore Maximo Predict and also your Watson Studio and Watson Machine Learning data.

Before you begin

- Restore Maximo Manage. For more information, see [“Restoring and validating Maximo Manage” on page 761](#).
- You need access to the Red Hat OpenShift command-line interface (CLI). Make sure that the **oc** command is in your path and that you have administrator credentials. For more information, see [Getting started with Open Shift CLI](#).
- You need to install the yq command-line interface, which is a lightweight and portable command-line processor for YAML, JSON, and XML. For more information, see [yq Quick usage guide](#).

About this task

Source cluster

The original Red Hat OpenShift cluster from which you took the backup data.

Target cluster

Any Red Hat OpenShift cluster that is different than the source cluster.

The three different ways to restore Maximo Predict are based on different scenarios. For more information, see [“Backup and restore scenarios overview” on page 728](#).

To learn more about how to deploy and undeploy Maximo Predict, see [“Deploying IBM Maximo Predict”](#) on page 371 and [“Deactivating and deleting applications”](#) on page 472.

Procedure

- For the crash recovery scenario, restore the Maximo Predict application to the source cluster from which you took the backup data.
 - a) Undeploy the existing Maximo Predict deployment.
 - b) Redeploy the Maximo Predict application.
 - c) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).
 - d) Restore the Watson Machine Learning data.

Import spaces and projects in Watson Machine Learning. For more information, see [Importing spaces and projects into existing deployment spaces in Watson Machine Learning](#).
- For the disaster recovery scenario, restore the Maximo Predict application to the target cluster, which is different than the source cluster.
 - a) Install and configure the cluster.
 - b) Install Maximo Application Suite core.
 - c) Deploy Maximo Predict.
 - d) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).
 - e) Restore the Watson Machine Learning data.

Import spaces and projects in Watson Machine Learning. For more information, see [Importing spaces and projects into existing deployment spaces in Watson Machine Learning](#).
- For the restore data scenario, you can restore Maximo Predict backup data only to the source cluster.
 - a) Restore the Watson Studio data.

Import a project in Watson Studio. For more information, see [Importing a project \(Watson Studio\)](#).
 - b) Restore the Watson Machine Learning data.

Import spaces and projects in Watson Machine Learning. For more information, see [Importing spaces and projects into existing deployment spaces in Watson Machine Learning](#).

You don't need to undeploy the existing Maximo Predict deployment.

Related tasks

[“Backing up Maximo Predict”](#) on page 756

If Maximo Predict is deployed on IBM Cloud , Amazon Web Services, or an on-premises environment, back up Watson Studio data and Watson Machine Learning data for Maximo Predict.

Backing up and restoring with IBM Storage Fusion

IBM Storage Fusion and its recipes provide a fully automated solution for backing up and restoring Maximo Application Suite. IBM Storage Fusion provides an enterprise-grade backup solution with scheduling capabilities, which are tailored for managing multiple clusters. It integrates seamlessly with IBM products such as Cloud Pak for Data and Db2 and supports efficient resource deployment within or across clusters.

Before you begin

The benefits of using IBM® Storage Fusion with recipes for backing up and restoring Maximo Application Suite are:

- The recipes identify the components and perform the backup.
- IBM Storage Fusion provides a cohesive end-to-end automated backup capability.

- IBM Storage Fusion offers an intuitive user interface and scheduling capabilities.
- IBM Storage Fusion is more suited for enterprise usage where multiple clusters need to be managed.
- If you have other IBM products that use Db2 such as IBM Cloud Pak for Data, then IBM Storage Fusion already has recipes for such an environment.

A few considerations need to be looked at when you are evaluating to use IBM Storage Fusion:

- IBM Storage Fusion does not backup external databases, which are outside of the Red Hat OpenShift cluster.
- IBM Storage Fusion requires additional resources, which can either be in the cluster or in another cluster.
- IBM Storage Fusion license that is provided with Maximo Application Suite does not include backup and restore, so additional licenses need to be acquired.

You must have the following software:

- Maximo Application Suite version 8.11 and later.
- Maximo Manage
 - Make sure that the production storage that is hosting Maximo Application Suite and Maximo Manage is compatible with Container Storage Interface (CSI). If you intend to use IBM Storage Fusion for backup and restore, Maximo data must be on storage that supports CSI. For more information, see [Fusion System Requirements](#).
- IBM Storage Fusion 2.8.x or later.
 - The license for Maximo Application Suite provides access to IBM Storage Fusion for internal deployment mode only. This excludes disaster recovery, backup components, data cataloging, and advanced encryption with KMS. You can also upgrade your license. For more information, see [IBM Terms > License information](#).

Note: If you are planning to configure Kafka with IBM Maximo Application Suite along with IBM Storage Fusion Backup and Restore operator, you require AMQ-Streams operator for configuring the Kafka cluster. Do not install an open source Strimzi operator because it causes potential issues between both operators. Ensure that you match the version of Fusion with the AMQ-Streams operator that you are planning to use for IBM Maximo Application Suite. To refer to Red Hat's own recommendations, see [Redhat's recommendation](#)

About this task

IBM Storage Fusion is the leading storage solution for Red Hat OpenShift.

Using IBM Storage Fusion allows to automate the backup process that uses the recipes to perform all the required steps. The recipe has the components that are identified and performs the backup based on this definition. It does not backup an external database, that task is left to the database administrator. The IBM Storage Fusion tools and user interface are provided to simplify the process. For more information, see [IBM Storage Fusion](#).

Procedure

1. Install IBM Storage Fusion 2.8.x or later.
 - a) Get the entitlement key.
For more information, see [Obtaining entitlement key](#).
 - b) Create an image pull secret.
For more information, see [Creating image pull secret](#).
 - c) Install the IBM Storage Fusion Operator and deploy IBM Storage Fusion.
Follow the specific instructions based on your type of Red Hat OpenShift Container Platform deployment. For more information, see [Deploying IBM Storage Fusion](#).
 - d) Deploy the IBM Storage Fusion backup and restore service.

- For more information, see [IBM Storage Fusion Backup & Restore](#).
- From the IBM Storage Fusion user interface, configure general backup and restore.
 - Create a backup storage location.

Backup storage location

The object storage is where backups are stored.

It can be any S3 compatible storage or Microsoft Azure Blob storage.

For more information, see [Backup storage locations](#).

- Create a backup policy.

Backup policy

Defines the frequency, the retention period, and the location of the backups.

For more information, see [Creating backup policy](#).

- Configure the specific IBM Storage Fusion recipes for Maximo Application Suite backup and restore.
For more information, see [Using Maximo Fusion Recipes](#).

Restoring Maximo Application Suite with IBM Storage Fusion

You can restore to the same Red Hat OpenShift cluster or a different Red Hat OpenShift cluster.

Before you begin

- The target namespaces must be the same.
- The target namespaces must be empty.

Procedure

- Restore the Maximo Application Suite platform instance.
 - In the IBM Storage Fusion user interface, specify **Use the same project the application is already using**.
For more information, see [Restoring an application with IBM Storage Fusion Software](#).
IBM Storage Fusion reinstalls the Maximo Application Suite platform instance, restores the data, and recovers the instance to a working state.
- Optional: From the IBM Storage Fusion user interface, restore to a spoke cluster.
 - Establish a connection between the hub and spoke clusters.
Generate a connection snippet from the hub cluster and copy it. For more information, see [Establishing connection between hub and spoke](#).
 - On the hub cluster, start the installation.
 - On the spoke cluster, select where the recipe needs to be deployed.
For more information, see [IBM Storage Fusion>Backup & Restore](#).

Debugging backed up and restored resources

In IBM Storage Fusion, you can debug by obtaining the backup or restore recipe log information and a list of the resources that were backed up or restored.

Procedure

- Download the `getRecipeWorkflow.sh` shell script from the IBM Storage Fusion repository.
For more information, see [IBM Storage Fusion backup-restore recipes](#).
- Obtain the job uid for the backup or restore job by using either the IBM Storage Fusion user interface or from the command line.

Method	Steps
IBM Storage Fusion user interface	<p>a. From the left navigation panel, select the Applications tab.</p> <p>b. From the table, select the wanted application.</p> <p>c. Select the Backups tab.</p> <p>d. On the wanted backup, select the ... menu and click Details.</p> <p>e. Click the hyperlink of the job name.</p> <p>For example, <code>filebrowser-demo-daily-apps.spparch.spp-ocp.tuc.stglabs.ibm.com-202311220800</code></p> <p>f. Note the Job ID in the details panel.</p> <p>For example, <code>63eb6ea5-dc24-41ef-9fb7-199c48df1508</code></p>
Command line	<p>a. Open the YAML file for the wanted IBM Storage Fusionbackup CR.</p> <p>For example,</p> <pre>oc get fbackup filebrowser-demo-daily-apps.spparch.spp-ocp.tuc.stglabs.ibm.com-202311220800 -n ibm-spectrum-fusion-ns -o yaml</pre> <p>b. Locate the <code>uid</code> field in the <code>spec: metadata</code>.</p> <p>For example, <code>uid: 63eb6ea5-dc24-41ef-9fb7-199c48df1508</code></p> <p>c. Run the <code>getRecipeWorkflow.sh</code> shell script with the job <code>uid</code>.</p> <p>For example,</p> <pre>getRecipeWorkflow.sh backup 63eb6ea5-dc24-41ef-9fb7-199c48df1508</pre>

3. Get a list of the resources that were backed up or restored.
 - a) Download the `getResources.sh` shell script from the IBM Storage Fusion Git repository. For more information, see [IBM Storage Fusion backup-restore recipes](#).
 - b) Obtain the job `uid` for the backup or restore job as indicated in the previous steps.
 - c) Run the `getResources.sh` shell script with the job `uid`.

```
getResources.sh
backup 63eb6ea5-dc24-41ef-9fb7-199c48df1508
```

What to do next

Information about protecting the backup and restore service can be found in the service protection section of the IBM Storage Fusion documentation. For more information, see [Data protection](#).

Administering users and user access in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#).

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#). Starting in Maximo Application Suite 9.1, users and security groups are created and managed at the suite level.

User access and entitlements in Maximo Application Suite 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#).

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

User access to suite applications is granted through associated access rights and entitlements. When you create a user, you use security groups to assign access and specify entitlements.

You create a user by selecting **Suite > Security > Users** from the side navigation menu. For more information, see [Creating users](#).

Access type

You assign an access type to the user for access to Maximo Application Suite. Depending on the access type that you assign, AppPoints are deducted either permanently or per-user session.

Concurrent access type

The AppPoint cost is applied when the user is logged in to Maximo Application Suite. When a user starts a session, the number of AppPoints for the assigned entitlement are checked out. When the user session ends, the AppPoints are returned.

Authorized access type

The AppPoints are reserved permanently from the organization pool when the user is created. With reserved AppPoints, the user can log in without depending on your organization’s current AppPoint balance, and no additional AppPoints are checked out when the user logs in. If you change the user's access type to concurrent, the reserved AppPoints are returned to the pool. Users who need administration access must be granted authorized access.

For more information, see [“Licensing in Maximo Application Suite 9.1” on page 78](#).

Access entitlement

Access entitlement is managed by security groups. Security groups are used to control user access to applications, data, and actions by grouping users and assigning permissions. When you create a security group, you can choose the applications with specific permissions to give users access. You can add users to one or more security groups. A user's entitlement is then set to the highest entitlement tier that is required to give them access to the capabilities and applications that are defined by the combined security group authorizations. For example, users who are added to a security group to access Maximo Health are assigned the base entitlement, unless they are also assigned to a security group where the premium entitlement is required.

Note: Security groups are not applicable to Maximo Real Estate and Facilities. For more information, see [“Maximo Real Estate and Facilities access” on page 783](#).

The following entitlements are applicable for the suite applications:

Self-service

The self-service entitlement is used to grant minimal access to Maximo Manage.

Limited

With the limited entitlement, a user can work with the core Maximo Application Suite applications, which includes Maximo Manage, and Maximo Monitor.

Base

With the base entitlement, the user has the same application access as the Limited entitlement and also access to the following applications:

- Maximo Collaborate
- Maximo Health

Premium

With the premium entitlement, the user has the same application access as the base entitlement and also has access to the following applications:

- Maximo Predict
- Maximo Visual Inspection

Premium entitlement also includes access to Maximo Manage industry solutions and add-ons. For more information, see [“Licensing in Maximo Application Suite 9.1” on page 78](#).

Maximo Real Estate and Facilities access

When you create users in Maximo Application Suite for Maximo Real Estate and Facilities, you assign them one of the following access types for initial access to Maximo Real Estate and Facilities.

- Self-service
- Base
- Limited
- Premium

Important: For the first login to Maximo Real Estate and Facilities, you must create a mandatory initial administrator user with a user ID of FACILITIESADMIN. Regardless of your selection, the FACILITIESADMIN user is assigned Base access in Maximo Real Estate and Facilities.

Users with Maximo Real Estate and Facilities access are automatically synchronized to the Maximo Real Estate and Facilities application.

After users have initial access, Maximo Real Estate and Facilities administrators can give users access to security groups and other permissions by logging in to Maximo Real Estate and Facilities administration. For more information, see [Administering user access and permissions](#).

As a Maximo Real Estate and Facilities administrator, you can view the entitlements that each access type provides on the **License viewer** page in Maximo Real Estate and Facilities. For more information, see [Viewing user licenses](#).

Administrative access

Users who have an administration entitlement can be granted combinations of Maximo Application Suite administration access.

The following administration access options are available.

System configuration

The user has edit access to the catalog, configurations, and license consumption functionality in Suite administration.

IDP management

A user with IDP management access can configure OIDC, LDAP, and SAML authentication, SMTP, and user registry synchronization. This user can also customize the user interface and manage certificates. IDP management access is a subset of system configuration with fewer access privileges. For example, this access type does not include the ability to deploy or activate applications or configure database connections.

API key management

Users with API key management access can create and manage API keys in Maximo Application Suite. To assign specific access to API keys, the user must also have the same access type as the API key. For example, to create an API key with system configuration access, the user must also have system configuration access.

To set up a user as an administrator who can create and manage general users, you can assign them user management permissions. For more information, see [“Creating administrators for user management” on page 787](#)

Related tasks

[Creating users in Maximo Application Suite 9.1](#)

[Creating security groups in Maximo Application Suite in 9.1](#)

Administering users in Maximo Application Suite 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier”](#) on page 796.

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

Starting in Maximo Application Suite 9.1, you can create users and centrally manage their authorizations and security privileges. You can also synchronize user data.

When users are created in Maximo Application Suite, the data is synchronized that the user information is up-to-date and consistent for each suite application, such as Maximo Manage. You can view the user synchronization status in the **User details** page by selecting **Suite > Security > Users**. The synchronization statuses for users are Pending, Complete, or Error. For more information, see [“User synchronization in Maximo Application Suite 9.1”](#) on page 789.

Related tasks

[Administering security groups in Maximo Application Suite in 9.1](#)

Creating users in Maximo Application Suite 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier”](#) on page 796.

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

Starting in Maximo Application Suite 9.1 to create a user record, you add the users details, specify the authentication method, and assign them to security groups for access to applications and capabilities.

Before you begin

Before you can create other users, you need user management permissions. For more information, see [“Creating administrators for user management”](#) on page 787

About this task

To create multiple user records simultaneously, you can import the user data by using a .csv file. For more information, see [Importing users](#).

Procedure

1. From the side navigation menu, select **Suite > Security > Users** and then click **Create user**.
2. Specify the identity information for the user.
 - a) On the **Identity** tab, specify the user ID and user's display name, and if required, the primary email.

Note: For the first login to Maximo Real Estate and Facilities, you must create a mandatory initial administrator user with a user ID of FACILITIESADMIN and Base entitlement.
 - b) Specify the account status to automatically deactivate the account at a specified date or a specified number of days of inactivity.
 - c) Specify other identity information, such as account status, user details, and contact information.

d) In the Preferences section, specify the locale and time zone for user. These preferences override language settings for the browser that they use to access the suite.

The locale and time zone is applied at the suite level for users, who can view the settings from their profile. Users can also change their locale setting by selecting **Profile > Manage profile > Language and region**, and making changes that are automatically applied to their user record.

3. Specify the authentication type and login details for user authentication.

Depending on your configuration, the following authentication types are supported:

- Local
- LDAP
- SAML
- OIDC

4. Create the user record.

Option	Description
Create & continue to access	Continue to grant access by assigning access types and security groups to the user.
Create & exit	You create the user record but without access to any applications. You can come back later to grant access.

5. On the **Access and license** tab, in the Access type section, specify whether to grant concurrent access or authorized access. If you are creating an administrator user, select **Authorized** as the access type.

Concurrent access type

The AppPoint cost is applied when the user is logged in to Maximo Application Suite. When a user starts a session, the number of AppPoints for the assigned entitlement are checked out. When the user session ends, the AppPoints are returned.

Authorized access type

The AppPoints are reserved permanently from the organization pool when the user is created. With reserved AppPoints, the user can log in without depending on your organization’s current AppPoint balance, and no additional AppPoints are checked out when the user logs in. If you change the user's access type to concurrent, the reserved AppPoints are returned to the pool. Users who need administration access must be granted authorized access.

6. **Note:** This step is not applicable to Maximo Real Estate and Facilities.

In the Security group section, assign the security group to grant the user access to the applications that they need.

Default groups are available for each suite application to assign specific access for both administrators and general users.

Maximo Collaborate uses the following default groups:

ASSISTADMIN

Maximo Collaborate administrators can create and manage users.

ASSISTSTUDIOUSERS

Maximo Collaborate studio users can create and manage information, such as building diagnosis libraries and managing collaboration sessions.

ASSISTTECHUSERS

Maximo Collaborate technical users, such as a field technician or maintenance worker, can diagnose and resolve equipment problems.

Maximo Monitor uses the following default groups:

MONITOR_ADMIN

Maximo Monitor administrators can create and manage users.

MONITOR_USERS

Maximo Monitor users can create device types and metrics, add devices, send data, import hierarchies, and create device dashboards.

Maximo Optimizer uses the following default groups:

OPTIMIZERADMIN

Maximo Optimizer administrators can manage user access and permissions.

OPTIMIZERUSERS

Maximo Optimizer users can run optimization models to schedule or assign work in a scenario.

Maximo Visual Inspection uses the following default groups:

MVI_ADMIN

Maximo Visual Inspection administrators can manage user access and permissions.

MVI_USERS

Maximo Visual Inspection users can create and deploy models.

7. To set up users as administrators, assign specific administrative privileges.

System configuration

The user has edit access to the catalog, configurations, and license consumption functionality in Suite administration.

IDP management

A user with IDP management access can configure OIDC, LDAP, and SAML authentication, SMTP, and user registry synchronization. This user can also customize the user interface and manage certificates. IDP management access is a subset of system configuration with fewer access privileges. For example, this access type does not include the ability to deploy or activate applications or configure database connections.

API key management

Users with API key management access can create and manage API keys in Maximo Application Suite. To assign specific access to API keys, the user must also have the same access type as the API key. For example, to create an API key with system configuration access, the user must also have system configuration access.

If a user needs administrative permissions to create other users, you must assign them to the USERMANAGEMENT security group. For more information, see [“Creating administrators for user management”](#) on page 787.

8. To give users access to Maximo Real Estate and Facilities, select one of the following access types from the Maximo Real Estate and Facilities list.

- Self service
- Base
- Limited
- Premium

9. Save your changes.

The license summary provides the details of the user's access and permissions.

Tip: If you change the user details, you can click **Re-calculate license summary** to view the latest information.

Results

When users are created in Maximo Application Suite, the data is synchronized so that the user information is up-to-date and consistent for each suite application, such as Maximo Manage. You can view the user synchronization status on the **User** details page.

A user's email address is optional. If the user does not have a specified email address, system-generated emails, such as welcome or password emails, are sent to the email address for the default suite system

that is specified in the Simple Mail Transfer Protocol (SMTP) configuration. For more information, see [Suite system email address](#).

What to do next

If you need to provide further specific information for Maximo Manage users, you can update the user record in Maximo Manage by selecting the user in **Manage > Security > Users (Manage)**.

Maximo Real Estate and Facilities administrators can give users access to security groups and other permissions by logging in to Maximo Real Estate and Facilities administration. For more information, see [Administering user access and permissions](#).

Related concepts

[User access and entitlements in Maximo Application Suite 9.1](#)

[Default groups in Maximo Application Suite in 9.1](#)

Related tasks

[Importing users in Maximo Application Suite 9.1](#)

[Creating security groups in Maximo Application Suite in 9.1](#)

Creating administrators for user management

Starting in Maximo Application Suite 9.1, to set up a user as an administrator who can grant other users access to the suite applications, you assign them user management permissions.

About this task

The USERMANAGEMENT security group provides specific administrative permissions to create users.

Procedure

1. From the side navigation menu, select **Suite > Security > Users** and then click **Create user**.
2. On the **Identity** tab, enter the user's details.
 - a) Specify the user identity information, language preferences, and authentication access.
 - b) Click **Create & continue to access**.
3. On the **Access and license** tab, specify the access type and add the users to the USERMANAGEMENT security group to grant the administration entitlement to create users.
 - a) In the Access type section, select **Authorized** so that the user is authorized for continuous access.
 - b) In the Security group section, click **Select groups** and assign the USERMANAGEMENT security group.
4. Save the user record.
5. Authorize this administrator user to assign other users to security groups.
 - a) On the **Users** page, select the user that you created.
 - b) From the **More actions** menu, select **Authorize group reassignment**.
 - c) To add the groups to the assignment list for the user, click **Select groups**.
6. Save your changes.

Importing users in Maximo Application Suite 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#).

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

To create multiple users in Maximo Application Suite 9.1, you can import user data by using a .csv file. The .csv file also ensures that the format for the user information adheres to the import processing rules.

Before you begin

To create and import new user data, you can export existing user data as a .csv file to use as a template. To export user data, on the side navigation menu, select **Suite > Security > Users** and then click the download icon.

Procedure

1. In the .csv file, provide the user information to create the user records, such as identity details, contact information, access entitlements, and account status.
2. On the side navigation menu, select **Suite > Security > Users** and then click the **Import data** icon.
3. In the **Import data** window, upload the completed .csv file.
4. Specify the **Delimiter** and **Text Qualifier**.
5. Add **Object Structure**.
Use the search option to select an object structure.
6. To validate that the user data in the file adheres to the import processing rules, click **Validate**.
If there are any issues, you can address the issues in the file and upload the file again.
7. Import the file.

Results

After you import the completed .csv file, the data is processed, and a record is created for each user in the file.

After the initial import, you can modify user information or delete users by updating the file and importing the changes.

Related tasks

[Creating users in Maximo Application Suite 9.1](#)

[Importing data](#)

You can import data from a .csv file to update the database and simultaneously create multiple records, such as user records. You can use this file to add the data and ensure that the format adheres to the import processing rules.

[Exporting data](#)

Authorizing users to assign other users to security groups in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier”](#) on page 796.

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

When you create a security group, you are authorized to assign users to that group. As an administrator, you can also authorize a user to assign other users to security groups.

Procedure

1. On the side navigation menu, select **Suite > Security > Users**.
2. Select the user for whom you want to authorize.
3. In more actions menu of the **View user** page, select **Authorize group reassignment**.

4. To authorize the user to assign other users to security groups, click **Select groups** to add the groups to the users assignment list.
5. Click OK and save your changes.

Results

The user has permission to add other users to the selected groups.

User synchronization in Maximo Application Suite 9.1

Users that are created in Maximo Application Suite 9.1 are added by authorization services for data synchronization so that user information is up to date for Maximo Application Suite and its applications such as Maximo Manage.

The user synchronization is bidirectional. The synchronization occurs from Maximo Application Suite to Maximo Manage and from Maximo Manage to Maximo Application Suite

Synchronization statuses

You can view the user synchronization status in the **User details** page by selecting **Suite > Security > Users**.

The synchronization statuses for users are Pending, Complete, or Error.

- Pending indicates that the data is being processed.
- Complete indicates that the data is synchronized.
- Error indicates that the data is synchronization is not completed and has errors.

Synchronization errors

If you are creating users with LDAP or SAML, from the user registry, or from user management APIs, an error might occur and the synchronization might not complete. An error also might occur during upgrade when the users are synchronizing from MongoDB to the authorization service.

To resolve these synchronization errors, you can resynchronize by selecting the action **View user synchronization status** from **Suite > Security > Users**.

An error might occur when you manage users by using Create, Read, Update, and Delete (CRUD) or import actions in the new user application or when you upgrade from IBM Maximo Asset Management to Maximo Manage.

You can also resynchronize user updates to initiate immediate changes rather than waiting on a scheduled crontask. For example, if you change an entitlement for a user and you need to apply the update immediately. You can resynchronize these changes for a user by selecting **More actions > Resync**. Otherwise, the MASUSERSYNC crontask regularly synchronizes user entitlement updates.

Administering security groups in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#).

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

Starting in Maximo Application Suite 9.1, you can provide authorization to suite applications, actions, and data to a group of users by using security groups. Security groups are used to control user access to applications, data, and actions by grouping users and assigning permissions.

Security groups can provide broad authorizations to many applications, or you can take a modular approach by adding users to multiple groups that grant fewer access privileges. You can specify different levels of authorization, which can be a combination of read, insert, save and delete, or all levels.

Related tasks

[Administering users in Maximo Application Suite 9.1](#)

Default groups in Maximo Application Suite in 9.1

Starting in Maximo Application Suite 9.1, authorization to all suite applications, actions, and data is provided by assigning users to one or more security groups. To give users the correct level of access to the suite applications, predefined groups are available that you can add users to.

The following default groups are available for each suite application to assign specific access for both administrators and general users.

Maximo Collaborate uses the following default groups:

ASSISTADMIN

Maximo Collaborate administrators can create and manage users.

ASSISTSTUDIOUSERS

Maximo Collaborate studio users can create and manage information, such as building diagnosis libraries and managing collaboration sessions.

ASSISTTECHUSERS

Maximo Collaborate technical users, such as a field technician or maintenance worker, can diagnose and resolve equipment problems.

Maximo Monitor uses the following default groups:

MONITOR_ADMIN

Maximo Monitor administrators can create and manage users.

MONITOR_USERS

Maximo Monitor users can create device types and metrics, add devices, send data, import hierarchies, and create device dashboards.

Maximo Optimizer uses the following default groups:

OPTIMIZERADMIN

Maximo Optimizer administrators can manage user access and permissions.

OPTIMIZERUSERS

Maximo Optimizer users can run optimization models to schedule or assign work in a scenario.

Maximo Visual Inspection uses the following default groups:

MVI_ADMIN

Maximo Visual Inspection administrators can manage user access and permissions.

MVI_USERS

Maximo Visual Inspection users can create and deploy models.

All new users are also added to the following groups, including Maximo Manage:

MAXDEFLTREG

A group that grants basic access to users, which allows them to change their passwords.

MAXEVERYONE

A global group that specifies global settings for all users.

You can assign users to additional security groups to meet your business needs.

Related tasks

[Creating security groups in Maximo Application Suite in 9.1](#)

[Creating users in Maximo Application Suite 9.1](#)

Creating security groups in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0

and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796.](#)

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups.](#)

Starting in Maximo Application Suite 9.1, you create security groups to manage authorization and entitlement to suite applications. Until you create security groups to specify permissions and add users to the group, users cannot access the suite applications or complete any actions.

About this task

You use security groups to grant read, insert, save, or delete access to the applications, actions, and data that users can access. When you create a security group, you can choose the applications with specific permissions to give user access. For more information, see [“Access entitlement” on page 782.](#)

Note: Initial access to Maximo Real Estate and Facilities is applied in each user's record and not by using security groups. Maximo Real Estate and Facilities administrators can give users access to security groups and other permissions by logging in to Maximo Real Estate and Facilities administration. For more information, see [Administering user access and permissions.](#)

Procedure

1. From the side navigation menu, select **Suite > Security > Security groups** and then click **Create group**.
2. On the **Identity** tab, specify a name and a default application and click **Create group**.
After the group is created with the identity information, you specify the applications, restrictions, users, and operational dashboard to associate with the group.
3. On the **Applications** tab, select the applications that you want to authorize and grant the appropriate types of access to each and click **Next**.

Tip: Each application must be granted READ permission so that it is visible in the side navigation menu.

If you save and exit, you can continue adding security groups information. Search for the group name that you created and then click edit.

4. On the **Restrictions** tab, you can add restrictions to restrict access to objects, attributes, and collections and click **Next**.
5. On the **Users** tab, add users to the group and click **Next**.
6. On the **Operational dashboard** tab, add a dashboard for the group.
7. Save your changes.

Results

The security group is created and associated with the users record. You can edit the group to add more users, restrictions, and dashboards if required.

Related concepts

[User access and entitlements in Maximo Application Suite 9.1](#)

[Default groups in Maximo Application Suite in 9.1](#)

Related tasks

[Creating users in Maximo Application Suite 9.1](#)

Applying data restrictions in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796.](#)

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups.](#)

Starting in Maximo Application Suite 9.1, you can specify restrictions on which records can be accessed by security groups on the application level for suite applications. You can create a data restriction for all users at an object level for Maximo Manage and at an application level for the other suite applications.

About this task

You can apply the following types of data restrictions:

- Global data restriction that is applied to the entire system and not associated with a specific security group.
- Group data restriction that is applied to the users in a specified group.

Global data restrictions restrict access to specific resource instances based on one filter criterion of Qualified. Qualified is the only type of restriction that is used. Data that meets the condition is retrieved from the database. It can be applied only to primary objects and not child objects or attributes.

Restriction criteria depend on the suite application that you are applying restrictions to. For example, the criteria for applying restrictions to Maximo Manage differs from the criteria for other suite applications.

Procedure

1. From the side navigation menu, select **Suite > Security > Security groups**.
2. Select the type of restriction that you want to apply.
 - To apply a restriction to the system, select **System actions > Global data restrictions**.
 - To apply a restriction to users in a specific group, select the group and on the **Restrictions** tab, and then click **Add restriction**.
3. Select the application that you want to apply the restriction to.
4. Set **Reevaluate** to specify when the restriction condition is evaluated..

Description	Option
To reevaluate the restriction condition when the user tabs to another field	Set to Yes
To evaluate the restriction conditions after changes to the field are saved	Set to No

5. If you select an object for the Maximo Manage application, specify the object and condition.
The condition class and the expression are applied based on the condition that you selected.
6. For other suite applications, specify the expression condition.

Authorizing security group assignments for users in Maximo Application Suite in 9.1

Note: Starting in Maximo Application Suite 9.1, you can manage users and assign security groups to users in suite security. For information about managing users and user access in Maximo Application Suite 9.0 and earlier, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796](#).

To manage users and security groups in Maximo Manage 9.0 and earlier, see [Managing users and groups](#).

When you create a security group, you assign users to the group. For security purposes, you can give users permission to add or remove other users from a security group by granting them the authorization to assign. Authorizing designated users to manage user permissions in specific groups, maintains a controlled level of user access.

Procedure

1. From the side navigation menu, select **Suite > Security > Security groups**.

2. Select the group for which you want to authorize group reassignment.
3. Click the **More actions** menu and select **Authorize group reassignment**.
4. To authorize users, click **Select users** and add the user to the assignment list.
5. Save your changes.

Results

The user has permission to add other users to the group.

User management APIs

Starting in Maximo Application Suite 9.1, new user management APIs replace the existing user management APIs. The new user management APIs provide greater granularity for user authorization to allow for a more flexible and secure way to manage access within the application.

In Maximo Application Suite 9.0 and earlier, user access is managed through rigid roles, such as `user` or `administrator`, without the ability to fine-tune access levels within those roles. With the new APIs, administrators can now assign specific permissions within each area of the application, which enables the creation of more precise and context-specific roles and accurately reflects real-world responsibilities.

Deprecated APIs

The following APIs for user management are deprecated in 9.1.

User creation APIs

The endpoints that are responsible for registering new users.

Workspace assignment APIs

The endpoints that are used to link users to specific workspaces.

Role assignment APIs

The endpoints that are used to assign fixed roles, such as `user` and `administrator`, within applications.

These APIs continue to function for a limited time but no longer receive functional updates and operate with reduced capabilities. Plan to migrate to the new APIs.

Comparing API models

In the original API model, assigning a role such as `administrator` grants full access to the application, without the ability to define more specific permissions. With this model, the ability to have multiple types of administrators with distinct scopes of action was not possible.

The new APIs enable granular role-based access control (RBAC) for more precise access management, better security, and improved scalability and enable administrators to grant the following options.

- Assign specific permissions rather than relying on broad, predefined roles.
- Compose custom roles by combining multiple permissions.
- Build access models that are aligned with the actual responsibilities of each user.

Technical considerations

In addition to the changes to the authorization model and APIs, the following data persistence updates must also be considered.

- The deprecated APIs used to store information directly in MongoDB.
- The new APIs follow a transient model where data is temporarily stored in MongoDB, then migrated to relational databases and becomes the single source of truth.
- After the data persists in the relational database, the original entries in MongoDB are ignored.

- After the initial synchronization with Maximo Manage, a flag that is called `foundationSyncComplete` is added to each user for every application that uses the compatibility mode.

Some data continues to persist in MongoDB due to the following application-specific requirements.

- When you assign access to the IoT tool.
- When you assign access to Maximo Real Estate and Facilities.

What's changed

In Maximo Application Suite 9.0 and earlier, user roles are assigned in suite applications by using the following methods.

- API calls
- LDAP synchronization
- Self-registration
- SCIM integration

Starting in Maximo Application Suite 9.1, user authorization is managed by the foundation service for most of the suite applications. Public API role updates are disabled for the following suite applications that are administered by Maximo Manage.

- Maximo Collaborate
- Maximo Monitor
- Maximo Optimizer
- Maximo Predict
- Maximo Visual Inspection

These applications support an initial user configuration. However, after a user is synchronized with Maximo Manage, role changes are no longer available, which can ensure a successful migration without disrupting existing automation.

Related information

[Maximo Application Suite 9.1 APIs](#)

Administering user sessions in Maximo Application Suite 9.1

Starting in Maximo Application Suite 9.1, as an administrator, you can log out individual users so that inactive sessions are available for other users, or all users, to perform maintenance tasks, and prevent them from accessing the system.

Configuring force user log out

Starting in Maximo Application Suite 9.1, as an administrator, you can automatically log out a user to perform system maintenance or to ensure that inactive sessions are available for other users.

About this task

Note: Force user log out is available to all suite applications except Maximo Real Estate and Facilities.

Procedure

1. On the side navigation menu, select **Suite > Administration > User sessions**.
2. Select the user for whom you want to log out.
3. Click **More actions > Force logout**.
4. Specify the following options.

Field	Description
Time in minutes before user is logged out	The amount of time before the forced logout is applied. The time is displayed to the user in a system message. If you specify the time as zero minutes, the user is logged out immediately.
Time in minutes before user can log in again	The amount of time before the user can log in to system after a forced logout.

Tip: If a user ends a session without logging out, the concurrent AppPoints are not returned to the pool of AppPoints for 30 minutes. To ensure that all sessions are available, you can log out that user immediately if you set the **Time in minutes before user is logged out** to zero.

Results

The user is logged out. A message is shown in the user interface every minute to notify the user before they are logged out and specifies the time that you defined. Users can select to turn off the message.

Enabling maintenance mode by using APIs

Starting in Maximo Application Suite 9.1, to prevent all users from accessing the system during critical maintenance or when changes are being made, you can enable maintenance mode by using APIs. After you enable maintenance mode, active users are logged out and login access is restricted for all users. Only system administrators, identity provider (IdP) management administrators, and super users can log in during maintenance mode.

About this task

For more information about REST APIs, see [Maximo Application Suite 9.1 APIs](#).

Note: Force user log out is available to all suite applications except Maximo Real Estate and Facilities.

Procedure

1. To enable maintenance mode, use the PUT `https://api.<MAS domain>/forcelogout/global` API and specify values for the following attributes.

Payload attribute	Description
authenticationLockDurationMinutes	The amount of time that all users are logged out and before they can log in to system again.
logoutMessage	For users that are logged in, a message is shown in the user interface to notify them that they will be logged out.
timeToLogoutMinutes	The amount of time before users are logged out. The message is shown to notify the user before they are logged out.

For example, you can enter the following command in a command line tool that logs out all users for 60 minutes and gives them a 5-minute warning.

```
curl -H "Content-Type: application/json" --location \
--request PUT https://api.<mas domain>/forcelogout/global \
--header "x-access-token: <system admin or idp admin access token>" \
--insecure \
--data '{
```

```
"authenticationLockDurationMinutes": 60,
"logoutMessage": "You will be logged out in 5 minutes. Save your work.",
"timeToLogoutMinutes": 5
}'
```

2. To change the time for maintenance mode, you can update the time for the **authenticationLockDurationMinutes** attribute in the payload.

Setting a new duration overwrites the previous duration and does not add to the current duration time. For example, if the duration was initially 60 minutes and you change to 70 minutes, then the new duration is 70 minutes. If you change the time from 60 minutes to 10 minutes, the new duration is just 10 minutes.

3. To end maintenance mode early so that users can log in again, you can delete the API with DELETE `https://api.<MAS domain>/forcelogout/global`.

Results

All users are logged out. A message is shown in the user interface every minute to notify the user before they are logged out and specifies the time that you defined. Users can select to turn off the message.

Administrators can log in during maintenance mode.

Customer-managed **Administering users and user access in Maximo Application Suite in 9.0 and earlier**

In Maximo Application Suite in 9.0 and earlier, user records are created and managed at the suite level and are made available for application-level access with application-specific role assignments, such as administrator or user. If required, an application administrator can set detailed application privileges for each individual application.

Before you begin

You can watch this video that shows how to manage users in Maximo Application Suite 8.11.

Related concepts

[User entitlement and access in Maximo Application Suite in 9.0 and earlier](#)
User access is granted through entitlements and associated access rights.

Related reference

[Importing users in Maximo Application Suite in 9.0 and earlier](#)

To create multiple users in Maximo Application Suite, use the template file to import new users and ensure that the format for the user information adheres to the import processing rules. After you import users, you can also use the template file to modify user information and delete users.

Creating users in Maximo Application Suite in 9.0 and earlier

You can add users to Maximo Application Suite by creating a record for each user. You can also create multiple user records simultaneously by using a template file.

About this task

Maximo Application Suite user records are stored in the user collection in the MongoDB core database. The internal user registry can be configured to use external authentication providers for user authentication. For more information, see [Authentication methods](#).

In Maximo Application Suite 9.0, a user's email address is required by default. Starting in Maximo Application Suite 9.0.14, by updating the Suite custom resource in Red Hat OpenShift Container Platform, you can change this setting so that the email address is optional. For more information, see ["Configuring emails as optional in Maximo Application Suite 9.0.14" on page 636](#).

Procedure

1. To add a single user, create a user record and enter the users' details, authentication type, and entitlementment.
 - a) In **Suite administration**, click **Users** and then click **Create user**.
 - b) Specify the user's display name, user ID, and primary email.
 - c) Specify other identity information, such as account status, user details, contact information, locale, and time zone.

For more information, see [“Setting user account status in Maximo Application Suite in 9.0 and earlier”](#) on page 806 and [“Setting language and time zone preferences for users”](#) on page 807.
 - d) Specify the authentication type and login details for user authentication.
 - e) Select entitlementments to set the application and administration access levels.

For more information, see [“User entitlement and access in Maximo Application Suite in 9.0 and earlier”](#) on page 797.
 - f) In the Application access section, give the user access to applications that they need by selecting the role or entitlementment.
2. To create multiple user records, import user records by using the .csv template file.
 - a) In **Suite administration**, select **Users** and then click the **Import users** icon.
 - b) Download the .csv template.
 - c) In the .csv template, provide the user information to create the user records, such as identity details, contact information, access entitlementments, and account status. For more information, see [“Importing users in Maximo Application Suite in 9.0 and earlier”](#) on page 801.
 - d) Upload the .csv file and click **Import**.

After you import the completed .csv file, the data is processed, and a record is created for each user in the file. To modify user information or delete users, you can update the file and import the changes again.

Results

If SMTP is enabled for your environment, an email is sent to the email address that is associated with the user ID. If the SMTP email security is set to [email passwords to users](#), a second email is sent with the user's password. If SMTP is not set up, an administrator must contact the user to provide the password.

Users can self-manage their accounts to update the display name, change their password, and set their preferred language and region from the profile menu in the user interface.

Starting in Maximo Application Suite 9.0.15 and 8.11.26, you can search for multiple users by entering a comma to separate each value in the search field. You can also search by email information on the user page by default, which eliminates the need to configure the search view.

Customer-managed **User entitlement and access in Maximo Application Suite in 9.0 and earlier**

User access is granted through entitlementments and associated access rights.

You can specify the Maximo Application Suite access rights when you create users or later by using a combination of the following user access dimensions:

Entitlementment

The user can be assigned a specific administration or application access right or both.

Access

Entitled users can be granted specific access rights at the administrator and application level.

For sample entitlementment and access assignment combinations, see [“Example of user roles configurations”](#) on page 800.

Application Point cost

Application Points (AppPoints) are used to track and enforce user entitlement and access across Maximo Application Suite. At creation time, you can assign each user an entitlement level of self-service, limited, base, or premium. Each level is associated with an AppPoint cost.

Self-service

Application access requires 0 AppPoints. Administration access is not available.

Limited

Application access requires 5 AppPoints. Administration access is not available.

Base

Application access requires 10 AppPoints. Administration access requires 10 AppPoints.

Premium

Application access requires 15 AppPoints. Administration access requires 15 AppPoints.

Access type

Depending on the access type that you assign the user, AppPoints are deducted either permanently or per user session.

Concurrent access type

The AppPoint cost is applied when the user is logged in to Maximo Application Suite. When a user starts a session, AppPoints corresponding to the assigned entitlement are checked out. When the user session ends, the AppPoints are returned.

Authorized access type

The AppPoints are reserved permanently from the organization pool when the user is created. With reserved AppPoints, the user can log in without depending on your organization's current AppPoint balance and no additional AppPoints are checked out when the user logs in. If you change the user's access type to concurrent, the reserved AppPoints are returned to the pool. By default, Administrator users are granted authorized access.

For more information, see [“Licensing in Maximo Application Suite 9.0 and earlier” on page 105](#).

Administration entitlement and access

Administrative access to Maximo Application Suite and its applications is granted by using administration entitlement. The following entitlements are available:

None

If administration entitlement is not granted, the user is a regular Maximo Application Suite user without administration rights. If the user has an application entitlement other than None, the user can log in and access applications in Maximo Application Suite.

Base

Base administration entitlement gives the user application administrator rights for each application to which the user has access and for which the user was given certain administrative rights at the Maximo Application Suite level: workspace management and user management. The base administrator user has access rights to the **Suite administration** page as set by the administration access selection.

Premium

Premium administration entitlement gives the user administrator rights at the Maximo Application Suite level. In addition, the user has all rights of the Base entitlement.

Users who have an administration entitlement other than None can be granted combinations of Maximo Application Suite administration access. The following administration access options are available:

User management

Requires base administration entitlement. The user has edit access to the Users section of the **Suite administration** page and can grant users access to applications.

System configuration

Requires premium administration entitlement. The user has edit access to the catalog, Configurations, and license consumption sections of the **Suite administration** page.

API key management

Requires premium administration entitlement. Users with API key management access can create and manage API keys in Maximo Application Suite. To assign specific access to API keys, the user must also have the same access type as the API key. For example, to create an API key with system configuration access, the user must also have system configuration access.

Application entitlement and roles

The application entitlement entitles the user access to the applications and tools that make up Maximo Application Suite. The following entitlements are available:

None

If no application entitlement is granted, the user does not have access to any applications. However, the user has access to the Maximo Application Suite user interface, and if the user has administrator entitlement, they also have access to the **Suite administration** page.

Self-service

The self-service entitlement is used to grant minimal access to Maximo Manage. Further access is then granted in Maximo Manage by assigning users to one or more security groups. Depending on the authorization set by those groups, the user's application entitlement might be automatically upgraded.

Limited

With the limited entitlement, a user can work with the core Maximo Application Suite applications, which include Maximo Monitor and Maximo Manage.

Base

With the base entitlement, the user has the same application access as the Limited entitlement and also access to the following applications:

- Maximo Collaborate
- Maximo Health

Premium

With the premium entitlement, the user has the same application access as the base entitlement and also has access to the following applications:

- Maximo Predict
- Maximo Visual Inspection
- Maximo Health and Predict - Utilities

Note: The application roles are controlled by the applications and might differ from application to application. By default, users with administration entitlement are given the administrator role for all available applications. A user with no administration access is given the user role for the application that the user is entitled to.

The following access levels are available for users to access Maximo Monitor:

None

If no role is granted, the user does not have access to the applications.

User

The user has regular user access rights to the application.

Administrator

The user has administrator user access rights to the application. A user needs Administration Entitlement to be granted the administrator role.

For more information about the available user roles for specific applications, see the corresponding documentation:

- [Maximo Collaborate](#)
- [Maximo Health](#)
- [Maximo Manage](#)
- [Maximo Monitor](#)
- [Maximo Predict](#)

Example of user roles configurations

The following table lists generic Maximo Application Suite user roles and their accompanying entitlements and access settings:

<i>Table 137. Example of user roles configurations</i>				
Role	Administration entitlement	Administration access	Application entitlement	Application role
Maximo Application Suite administrator The Maximo Application Suite administrator manages overarching system configuration settings from the Suite administration page.	Premium	System configuration	None	None
<i>Application administrator</i> The application administrator administers one or more applications, adds and assigns users to these applications, and uses the application-specific user interfaces to manage further user privileges.	Base	User management	Limited	Example: Maximo Monitor with Administrator role

Table 137. Example of user roles configurations (continued)

Role	Administration entitlement	Administration access	Application entitlement	Application role
<p><i>Application user (Maximo Manage)</i></p> <p>The Maximo Manage application user has access to the Suite navigator and to Maximo Manage at the access level that is set by security groups in the application.</p>	None	None	Self-service or higher	Example: Maximo Manage with User role
<p><i>Application user</i></p> <p>The application user has access to the Suite navigator and to one or more applications at an access level that is set by the application.</p>	None	None	Limited	Example: Maximo Monitor with User role

Related tasks

Customer-managed

[Administering users and user access in Maximo Application Suite in 9.0 and earlier](#)
 In Maximo Application Suite in 9.0 and earlier, user records are created and managed at the suite level and are made available for application-level access with application-specific role assignments, such as administrator or user. If required, an application administrator can set detailed application privileges for each individual application.

Importing users in Maximo Application Suite in 9.0 and earlier

To create multiple users in Maximo Application Suite, use the template file to import new users and ensure that the format for the user information adheres to the import processing rules. After you import users, you can also use the template file to modify user information and delete users.

You can download the .csv template from the Suite administration user interface.

1. In the Suite administration, select **Users** and then click the **Import users** icon.
2. Download the .csv template.
3. Enter users details.
4. In the **Import users**, import the file.
5. To modify user information or delete users, you can update the file and import the changes.

The following information describes the column names and values that you provide in the .csv template for each user. Each column corresponds to a field entry in the user record.

Identity

The following information describes the values that you specify to import the user's identity details.

id

The user ID is the internal identification for the user and is a required field. The default maximum length is 100 alphanumeric characters.

After the user record is created, this ID cannot be changed.

Note: If you enabled Maximo Application Suite to include the use of all special characters for user ID and username and use double bytes characters, save the file as CSV UTF-8 format. For more information, see [“Enabling special characters for user ID and username” on page 647](#).

username

The username is the name that the user types to log in and is a required field. The default maximum length is 100 uppercase alphabetic characters.

The username can be the same as the user ID. The username can be changed.

Note: If you enabled Maximo Application Suite to include the use of all special characters for user ID and username and use double bytes characters, save the file as CSV UTF-8 format. For more information, see [“Enabling special characters for user ID and username” on page 647](#).

displayName

The display name that is shown when the user is logged in. This field is a required field.

The default maximum length is 82 alphanumeric characters.

givenName

The given name of the user. The default maximum length is 30 alphanumeric characters.

familyName

The surname of the user. The default maximum length is 50 alphanumeric characters.

title

The title of the user, such as Mr. or Ms.

Password and authentication

The following information describes the values that you specify to import the user's password and authentication details.

generatePassword

Enter either TRUE or FALSE.

Enter TRUE to automatically generate a password. You can generate passwords only for new users. Passwords cannot be generated for current users.

sendPasswordToEmail

Enter either TRUE or FALSE.

Enter TRUE to email the password credentials to the user after the user record is created.

password

Enter NONE or specify a custom password.

If **generatePassword** is TRUE, then enter NONE for **password**. Otherwise, enter a custom password.

forcePasswordChange

Enter either TRUE or FALSE.

Enter TRUE if you want users to change their password during their first login.

issuer

The issuer is the authentication type and is a required field.

Enter local, ldap, or saml.

Starting in Maximo Application Suite 8.11, the **issuer** field is deprecated. For 8.11 and later versions, use the **Identities** field.

User entitlement and application access

The following information describes the values that you can specify to import the user's entitlement and application access details.

permissions__systemAdmin

Enter either TRUE or FALSE. This field is a required field.

Enter TRUE if the user needs system configuration permissions. A user with system configuration access has administrative privileges for the core Maximo Application Suite settings and can deploy applications, update configurations, and manage license files.

permissions__userAdmin

Enter either TRUE or FALSE. This field is a required field.

Enter TRUE if the user needs user management permissions. A user with user management access can create and manage users and assign entitlements and access levels.

entitlement__application

Enter one of the following values:

- NONE
- SELF_SERVICE
- LIMITED
- BASE
- PREMIUM

This field is a required field.

If the user needs access to applications, specify the application entitlement for the user. For more information, see [“Application entitlement and roles” on page 799](#)

entitlement__admin

Enter one of the following values:

- NONE
- ADMIN_LIMITED
- ADMIN_BASE
- ADMIN_PREMIUM

This field is a required field.

Specify the administrator entitlement for the user. For more information, see [“Administration entitlement and access” on page 798](#).

authorizedUser

Enter either TRUE or FALSE.

Enter TRUE if the user needs authorized access. AppPoints are permanently reserved for authorized users, so they can log in to Maximo Application Suite at any time.

Users who have administrator entitlement, such as BASE or PREMIUM in the **entitlement__admin** column, can be granted authorized access.

For more information, see [“Access type” on page 798](#).

workspaceId

Specify the name of the workspace that you are adding the user to.

appId

Specify the name of the application that the user needs access to:

- manage
- monitor
- health
- visual inspections
- assist
- predict

If the user needs access to multiple applications, create another row and specify the value for the application.

appId__access

Enter either ADMIN, USER, or MANAGEUSER.

To give users administrator access to an application, specify ADMIN. A user needs an administration entitlement of BASE or PREMIUM in the **entitlement__admin** column to be granted the administrator role.

To give users regular user access rights to an application, specify USER.

To give users access to Maximo Manage, specify MANAGEUSER.

Contact information

The following information describes the values to import the user's contact information.

emails__value

The user's email address. The default maximum length is 100 alphanumeric characters.

emails__type

Indicates whether the email address is a work or home email address. Enter either HOME or WORK.

emails__primary

Enter either TRUE or FALSE. This field is a required field.

If the email address is the main address to contact the user, specify TRUE. If the user's information includes multiple email addresses, one email address must be set to TRUE.

phoneNumbers__value

The user's phone number. The default maximum length is 20 alphanumeric characters.

phoneNumbers__type

Enter either WORK or MOBILE.

phoneNumbers__primary

Enter either TRUE or FALSE.

If the phone number is the main number to contact the user, specify TRUE. If the user's information includes multiple phone numbers, one phone number must be set to TRUE.

addresses__streetAddress

The street address. The default maximum length is 169 alphanumeric characters.

addresses__locality

The name of the city. The default maximum length is 36 alphanumeric characters.

addresses__region

The name of the region. The default maximum length is 36 alphanumeric characters.

addresses__postalCode

The postcode. The default maximum length is 12 alphanumeric characters.

addresses__country

The country. The default maximum length is 36 alphanumeric characters.

addresses__primary

Enter either TRUE or FALSE.

If the address is the main address of the user, specify TRUE.

Account status information

The following information describes the values that you can specify to import the user's account status information.

status__active

Enter either TRUE or FALSE.

Enter TRUE to activate or FALSE to deactivate the user account.

status__activationUpdateReason

Enter in plain text that explains why the user account status is updated. The maximum is 512 characters, which include spaces.

status__expiresAt

Specify the date in YYYY-MM-DD format that the account is automatically deactivated.

status__inactivityTimeout

Specify the duration after which if the user does not log in, the account is automatically deactivated.

Enter alphanumeric text in the following format:

P[n]DT[n]H[n]M[n]S

or

P[n]W

For example, P4DT12H30M5S represents a duration of four days, twelve hours, thirty minutes, and five seconds.

Identities for seamless user authentication

Starting in 8.11, the following information describes the values that you specify to import identities for seamless user authentication.

local_user

Indicates that the user account is local within the system. This field is a required field.

Enter either TRUE or FALSE. Ensure that you use accurate values in the field.

identities_saml

The particular Security Assertion Markup Language (SAML) authentication configuration to use for the respective user.

Enter `default-saml`.

If this authentication is not required, leave this field empty or enter NONE.

identities_saml_id

The critical linkage between the user's system ID and their SAML-based login credentials.

Enter `<user SAML ID>.saml`

Use the specific SAML ID that is assigned to a user, for example `Jsmith.saml`.

identities_ldap

The type of Lightweight Directory Access Protocol (LDAP) authentication that is used for the user.

Enter `default-ldap`.

If this authentication is not required, leave this field empty or enter NONE.

Note: Maximo Application Suite supports LDAP as the default configuration.

identities_ldap_id

The bridge between the user's system identity and their LDAP-based login credentials.

Enter `<user LDAP ID>.ldap`.

Use the specific LDAP ID that is assigned to a user, for example `Jsmith.ldap`.

User deletion

If you need to delete users from Maximo Application Suite, specify TRUE in the **delete** column. The user record is removed from the Maximo Application Suite user registry. When you are adding users, specify FALSE.

Troubleshooting

If the import process has errors, you can download the processed file. The processed file provides the data that was not imported and includes an error message that explains the problem. For example, if an error is related to the issuer field, the following message is shown:

```
Failed validating field: enum, reason: 'SAML' is not one of ['local', 'ldap', 'saml']
```

Fix the error and upload the file again.

If you cannot download the error file, check the importuser pod in the Maximo Application Suite core namespace.

If the importuser pod is not created or if the process does not start or shows a generic error, check the coreapi pod in the Maximo Application Suite core namespace.

Related tasks

Customer-managed

[Administering users and user access in Maximo Application Suite in 9.0 and earlier](#)
In Maximo Application Suite in 9.0 and earlier, user records are created and managed at the suite level and are made available for application-level access with application-specific role assignments, such as administrator or user. If required, an application administrator can set detailed application privileges for each individual application.

SaaS

[Adding users in Maximo Application Suite as a Service](#)
To add users to Maximo Application Suite as a Service, you create a record to specify user details, such as username and contact information. You can also specify access and entitlement for suite applications. You can also create multiple user records simultaneously by importing user information in a .csv file template. .

Related information

[Troubleshooting Import user template](#)

Starting in IBM Maximo Application Suite 8.11, when you upgrade the IBM Maximo Application Suite, you might find differences in the Import user template .csv headers.

Setting user account status in Maximo Application Suite in 9.0 and earlier

The application administrator can change the user account status by activating or deactivating it on the **User management** page. You can also deactivate the user account by a specific date or after a duration of inactivity from that account. You can also change the user account status in bulk from the **User management** page or by importing .csv file.

About this task

To change account status of multiple user accounts from the **User management** page, select multiple user records from the table, and then click **Activate** to activate or **Deactivate** to deactivate. Also, you can change user account status in bulk by importing .csv file. For more information, see [“Importing users in Maximo Application Suite in 9.0 and earlier” on page 801](#).

To change account status of a user account from the **User management** page, click the **Overflow menu** and then refer to the following choices to take action.

Procedure

- To activate an account, click **Activate**.
- To deactivate an account, click **Deactivate**.
- To deactivate an account automatically by date, click **Edit**, in the **Account status** section, set **Deactivate on date** to **Enabled** and then select a date on which the user account must be deactivated.
- To deactivate an account automatically due to inactivity, click **Edit**, in the **Account status** section, set **Deactivate after user inactivity** to **Enabled** and then specify the number of days after which if the user does not log in, the account is deactivated.

Note: Change in user account status due to LDAP synchronization does not affect the user account's entitlement.



Trouble: If reactivation of an existing authorized or administrator user fails due to the lack of AppPoints, the LDAP synchronization shows an error. An administrator must manually update the entitlement for such user records to remove the administrator permissions or change to a concurrent user. The LDAP synchronization must be triggered again to update the manual changes in synchronized data.

Setting language and time zone preferences for users

You can specify the preferred locale and time zone for users in the user record for the suite user interfaces. The preferences that are applied override the language settings for the browser that is used to access the suite.

For example, if the user's preferred language is set to German and Maximo Application Suite is accessed from a browser set to English, the user interface is displayed in German. Similarly, if the time zone for the user is CET, the user interface date and time is based on CET regardless of the time zone that is set for the user's system or browser.

About this task

Important: Not all applications support these preferences and instead use application-specific settings. For more information, see [Language support table](#).

Procedure

1. From the **Suite administration** page, click **Users > Add user** or select the user record.
2. In the Additional information section, click **Locale and time zone**.
3. Select the preferred locale and time zone for the user.
 - The locale preference specifies the language and date and time representation.
 - The time zone preference sets your approximate geographical location. Date and time information is displayed by using this setting.
4. Apply your changes.

Results

The locale and time zone is applied for users, who can view the settings from their profile. Users can also change their locale setting by selecting **Profile > Manage profile > Language and region**, which is automatically applied to their user record.

Related concepts

[Language and locale support](#)

IBM Maximo Application Suite supports the use of preferred language and locale for the Maximo Application Suite user interfaces. The preferences that are applied override the language and locale settings for the browser that is used to access Maximo Application Suite.

Deleting and anonymizing user data

When you delete a user in IBM Maximo Application Suite, the user data is retained in the database by default. Starting in Maximo Application Suite 8.11.11, you can anonymize personal information, such as username, emails, and display name, before you delete users.

About this task

You can anonymize user data by using one of the following options:

- In the user interface (UI), you can anonymize user data locally on the user management page when you select to delete a user. The **Anonymize user data** checkbox in the conformation dialog of user deletion can be managed for each user in the application.
- Before you delete users, you can globally anonymize user data by updating the custom resource in the Red Hat OpenShift web console.

If you anonymize the user data, you can reuse a deleted user's unique username and email when you create users.

Note: If the username is the same as the user ID, then you cannot reuse the username.

Procedure

- Anonymize user data globally in the Red Hat OpenShift web console.
 - a) In the Red Hat OpenShift Container Platform console,, click **Administration > CustomResourceDefinitions**.
 - b) Search for Suite.
 - c) Select the custom resource definition for the suite.
 - d) Click the **Instances** tab.
 - e) Select your Maximo Application Suite instance ID.
 - f) On the **YAML** tab, add the following to the `spec.settings` of the Maximo Application Suite instance custom resource.

```
---
userDataObfuscation:
  obfuscateDataOnDeletion: true
```

- g) Save the information.
 - h) In **Suite administration**, click **Users** and then click the **Delete user**.

When you delete users in the UI, the **Anonymize personal information** checkbox in the confirmation dialog of user deletion is read-only because you have globally set this option.
- Anonymize user data locally in Maximo Application Suite user interface.
 - a) In **Suite administration**, click **Users** and then click the **Delete user**.
 - b) In the deletion confirmation dialog, select the **Anonymize personal information** checkbox.

The personal information is deleted from the user record that is retained in the database.

Note: The user ID is not deleted. If that user ID contains any personal information, it is stored in the database.

If you do not anonymize the user data, the data is stored in the Maximo Application Suite database.

Results

When you delete users, the user ID is removed from the Maximo Application Suite user registry. If you are using LDAP or SAML authentication, the user account remains on the identity provider server but is no longer associated with a Maximo Application Suite account. If you use user registry synchronization, you must delete the user on the LDAP server and then synchronize to remove the Maximo Application Suite user.

Upgrade of users from Maximo Asset Management to Maximo Application Suite

Users are added in Maximo Application Suite and their profiles, login information, and application entitlements are managed in Maximo Application Suite. If you are upgrading from Maximo Asset Management to Maximo Application Suite with Maximo Manage, user data is upgraded.

If you are upgrading from Maximo Asset Management and use an integration to create or update users, you must send user data to Maximo Application Suite and not to Maximo Manage.

Note:

When you upgrade Maximo Asset Management to Maximo Application Suite, user and administrator records in the Maximo database are imported by Maximo Application Suite if they do not exist there or if they are not managed by an identity provider. Information includes language, time zone, and locale. Imported users are considered as already synchronized to Maximo Manage. User records also remain in the database that was upgraded.

If SMTP is enabled in Maximo Application Suite, a welcome email is sent to the email address that is associated with the user ID when the user or administrator account is created. If the SMTP email security is set to [email passwords to all users](#), a second email is sent with the user or administrator password. If SMTP is not set up, or if passwords are not emailed to users or administrators, a Maximo Application Suite administrator must supply the password.

If you have users with a different user ID and login ID and you are upgrading to Maximo Application Suite 8.7 or earlier, you must re-create your existing users. For more information, see [Troubleshooting user migration](#).

For the user synchronized between Maximo Application Suite and Maximo Manage, you can set different owner and issuer for different cron task instances. For more information, see [Owner and Issuer for cron task instance](#).

For more information about the process of synchronizing and troubleshooting the migrated users, see [“Managing users post upgrade” on page 560](#).

Managing users in Maximo Manage

Although, user details created in IBM Maximo Application Suite are synchronized with the applications under Maximo Application Suite. There are some details and access that you must manage at application level for IBM Maximo Manage user.

By using the **Security** module of the Maximo Manage, you can secure and manage users which includes the following list of activities.

Note: Following features are accessible based on the user entitlement.

- Configure and manage user settings
- Configure and manage user groups
- Reset E-Signature Key
- Manage user sessions
- Configure and manage security groups for users
- Configure user authorizations

To learn more about managing users and security groups in Maximo Manage, see [Managing users and groups](#).

Local user account settings

For local user authentication, you can configure password settings and you can enable account lockout .

Configuring password settings

As a suite administrator, you can configure password settings for local authentication to meet the requirements of your organization's security policy.

About this task

In Maximo Application Suite 8.10 and earlier versions, configuring passwords for local authentication is available by selecting **Users > Password settings** in Suite administration. Starting in Maximo Application Suite 8.11, you configure passwords for local authentication by selecting **Users > Authentication**.

Procedure

1. Open the **Password setting** page.
 - a) In **Suite administration**, click **Users** and then the **Authentication** tab.
 - b) In the Identity providers section, for the local identity provider, click the **More actions > Configure password**.
2. In the Password requirements section, specify the password length and the minimum number of uppercase, lowercase, numeric, and special characters.

Password length must be 6–35 characters.
3. Specify whether users can include their username in their password.
4. Specify whether users must change their password on first login.
5. In the Allowed placement of password characters section, specify the placement of numbers or special characters.
6. Save your settings.

Results

These password rules are enforced when you create a password for new users or when users change their password by selecting **Profile > Manage profile > Change password**.

Setting password expiration for local users

Starting in Maximo Application Suite 9.1, you can enable password expiration to define when a user is required to change their password.

About this task

You can specify how long the password is valid before a user must change their password and include a grace period that users can still log in after that password expires. You can also send an email to users to notify when a password is due to expire.

Procedure

1. From the navigation menu, select **Suite > Administration > Authentication**.
2. In the Identity providers section, for the local identity provider, click the **More actions > Configure password**.
3. In the password expiration section, specify the password expiration details.

- a) Enable password expiration.
 - b) Specify the number of days that the password is valid before a user must change their password.
 - c) Specify the grace period if you want to provide additional time that the user can still log in before expiration.
 - d) To notify the user that their password is due to expire, select the **Email user about password expiration** check box and specify the number of days to notify before the actual expiry date.
4. Save your updates.

Results

The users' password is set to expired based on your configuration. When the password expires or due to expire, the user is prompted to change the password on the login page.

Enabling account lockout

Starting in Maximo Application Suite 9.0, account lockout is available. For local authentication, you can enable account lockout to define the conditions that prevent users from logging in after consecutive unsuccessful login attempts.

Procedure

1. In Suite administration, from the side navigation menu, click **Users** and then click the **Authentication** tab.
2. In the Identity providers section, for the local identity provider, click **More actions > Configure password**.
3. In the Account lockout section, specify the number of consecutive password attempts before the users account is locked.
4. Choose whether to lock the account by duration or until an administrator unlocks the account.
5. Save your changes.

Results

When the lockout conditions are met, the account of the user is locked and a message is shown on the login page. If SMTP is configured, an email is also sent to the user to inform them that their account is locked. The user must either wait a predetermined amount of time before they can log in again or contact a system administrator to unlock the account.

What to do next

Customer-managed You can specify the length of time that elapses after a failed password attempt before the number of consecutive password attempts resets by updating the password policy setting in Red Hat OpenShift Container Platform. The default is that the number of consecutive password attempts resets only when the user logs in or after the users account is unlocked.

Related tasks

[Resetting login attempts](#)

Customer-managed Resetting login attempts

Starting in Maximo Application Suite 9.0, account lockout is available. For local authentication, you can specify the length of time that elapses after a password attempt before the number of consecutive password attempts is reset. You can define the reset setting by updating the password policy setting in Red Hat OpenShift Container Platform.

About this task

To reset the number of consecutive password attempts after a predefined interval, you change the duration of the **loginAttemptsResetSeconds** key from the default value of 0 to a specific length of time in seconds. By default, the number of consecutive password attempts is reset only when the user logs in or after the user's account is unlocked.

For example, as an administrator, in Maximo Application Suite, you set the maximum number of password attempts for users to 3 before they are locked out of their account. In Red Hat OpenShift Container Platform, you set the value of the **loginAttemptsResetSeconds** key to 600, which is 10 minutes. If a user enters their credentials incorrectly three consecutive times, then they are locked out until their account is unlocked. If a user enters incorrect credentials two times and 10 minutes elapse before they try to log in again, the number of remaining login attempts resets from 1 remaining attempt to 3.

Procedure

1. In the Red Hat OpenShift Container Platform console, from the side navigation, select **Home > Project** and search for the mas-*<instance id>*-core namespace where Maximo Application Suite is deployed.
2. From the side navigation menu, click **Workloads > Secrets**.
3. Search for a secret that is named *<instance id>*-password-policy and update the duration of the **loginAttemptsResetSeconds** key.

Option	Action
If the secret exists, update the duration of the loginAttemptsResetSeconds key.	<ol style="list-style-type: none">a. From the Actions menu, click Edit Secret.b. Find the loginAttemptsResetSeconds key.c. For the loginAttemptsResetSeconds key, in the Value field, specify the duration in seconds.
If the secret does not exist, create the secret and then update the duration of the loginAttemptsResetSeconds key.	<ol style="list-style-type: none">a. Click Create and select Key/value secret.b. In the Secret name field, enter <i><instance id></i>-password-policy.c. In the Key field, enter loginAttemptsResetSeconds.d. In the Drag and drop file with your value here or browse to upload it field, enter the duration in seconds.e. Click Create.

Results

The change is applied immediately. To validate, you can enter invalid credentials up until the final attempt that will lock you out. After the duration that you set in the **loginAttemptsResetSeconds** key elapses, enter the incorrect credentials again. The account remains unlocked.

Related tasks

[Enabling account lockout](#)

Starting in Maximo Application Suite 9.0, account lockout is available. For local authentication, you can enable account lockout to define the conditions that prevent users from logging in after consecutive unsuccessful login attempts.

The user access limit for your Maximo Application Suite organization is set by your contracted AppPoint entitlement. By monitoring your AppPoint usage, you can periodically review whether you need to adjust your entitlement and update your license file.

Note: The following information pertains to license management for customer-managed Maximo Application Suite. For Maximo Application Suite Dedicated, uploading of license keys is handled by your Maximo Application Suite representative. For more information, see your Welcome letter.

From the license and reporting page, you can do the following tasks:

- [“View usage reports” on page 813](#)
- [Configure session idle timeout](#)
- [Configure licenses and reporting](#)

Related concepts

[AppPoints for customer-managed Maximo Application Suite](#)

Understand AppPoint allocation

The Maximo Application Suite license usage is managed by the Suite License Service (SLS). Each Maximo Application Suite instance can be connected to a unique SLS instance, or multiple Maximo Application Suite instances can share an SLS and the corresponding license file.

For shared Suite License Service, the following conditions apply to the usage report AppPoint data.

- The pool of AppPoints granted by your license is shared among all Maximo Application Suite instances in the order of requests.
- Reserved and concurrent user logins from multiple Maximo Application Suite instances only result in a single checkout of AppPoints on condition that the entitlement of the logged in user is the same.
- Each Maximo Application Suite licensing report shows the total AppPoint usage across all Maximo Application Suite instances that share the Suite License Service.

View usage reports

The **Report** page provides an overview of the AppPoint entitlement for your Maximo Application Suite contract and license information, such as start and expiration dates. To view usage reports in Suite administration, click **License consumption** from the side navigation menu and select the **Report** tab.

To get an overview of your organization's AppPoint usage over time, use the built-in reports to see how well your contracted AppPoints meet your environment requirements.

Use the report information to identify AppPoint overages to see whether they occur regularly and can be managed by adjusting your contract and [uploading a new license](#).

The following report information is available and is based on the selected time period that you choose, such as the current month or day or choose a custom range. You can specify the period of time by selecting **Show report for**.

Overview

The Overview section contains AppPoint usage information for the selected time period, including entitled capacity and peak usage.

Details

The Details section provides a summary of the configurations that you set to manage AppPoints consumption, such as compliance enforcement, idle timeout, and entitlements in your license. You can click **View** for more details of each configuration or also click the **Configuration** tab.

Report chart

The report chart outlines how your AppPoints consumption varied over the selected period of time. The chart is divided into reserved and concurrent user AppPoints and you can hover over a usage data

point in the chart for details. For a more detailed breakdown of the AppPoint consumption by user entitlement, click **View breakdown**.

Tip: You can also download the reports in .csv format for further analysis.

Understand AppPoint consumption

AppPoints usage is divided into reserved and concurrent usage.

Reserved usage

The AppPoints are reserved permanently from the organization pool when an authorized access user is created. Reserved AppPoints are also consumed when you deploy applications or capabilities that have a reserved cost. Nonproduction installations do not incur reserved costs for applications or capabilities.

Concurrent usage

The AppPoint cost is applied when a concurrent access user is logged in to Maximo Application Suite. When a user starts a session, AppPoints corresponding to the assigned entitlement are checked out. When the user session ends, the AppPoints are returned.

Understand denial and overage usage

Denial and overage usage shows the cumulative number of AppPoints that were blocked or that were checked out beyond your entitlement during the report window.

If your environment is configured for compliance enforcement, concurrent users are blocked from logging in if the AppPoint entitlement is exceeded. If enforcement is not configured, the overage might incur extra costs.

Configure session idle timeout

You can define how long users can stay logged in by enabling session idle timeout and specifying how long a web browser session can be idle before that session is automatically logged out.

When users are logged out, their AppPoint cost is returned to the license pool. If idle timeout is disabled, users stay logged in until they actively log out. By default, sessions are timed out after 30 minutes of inactivity.

1. From the **Suite administration** page, select **license consumption**.
2. On the **Configuration** tab, enable **Idle timeout**.
3. Specify the length of time a session can be idle before that session is automatically logged out.
4. Save your changes.

Configure licenses

On the **Configuration** tab you can enable or disable compliance enforcement, set a default login message, set automatic reporting, and upload a new license file.

Set compliance enforcement

The number of AppPoints that are allocated to each Maximo Application Suite workspace is set by the product license. These AppPoints are consumed as you deploy applications, add-ons, and industry solutions that come with an AppPoint cost, and also when users log in and use the product. You can configure your environment to enforce the AppPoints limit and block users from logging in if the AppPoint entitlement is exceeded. For more information about AppPoint costs, see [AppPoints](#).

When enforcement is enabled, you can provide a login message that is displayed if a user's login is denied. Modify the sample message to fit your environment and situation.

The maximum number of simultaneous users are already logged in. Try logging in again in a little while.

Update your license key file

If changes are made to your IBM Maximo Application Suite license, you must upload a new license file that reflects that change.

The license update process includes two main tasks.

1. Acquire your updated Maximo Application Suite license key.

The license key is provided with your purchase of Maximo Application Suite. You can download the file in the [License Key Center](#) by logging in to your IBM Rational license Key Center account. The login information is provided with the license Key Center welcome letter. For more help on licensing, see the [IBM Support - Licensing page on IBM.com](#).

To create the license file, you must provide the following license server parameters that were provided in the license step of the [setup process](#):

Parameter	Value
Configuration	Single
Host ID	The server <i>MAC address</i>
Hostname	The server <i>Hostname</i>
Port	The server <i>port</i> Default: 27000
Host ID Type	The server <i>Ethernet address</i>

2. Update the license.

You can replace your existing license key file by using the **Suite administration** page in Maximo Application Suite.

- a. On the **Suite administration** page, select **license consumption**.
- b. On the **Configuration** tab, click **Replace license file**.
- c. Upload the new license file.

The new license is applied on commit.

Note: If the license file update fails, Maximo Application Suite automatically rolls back to the previous license.

Generating and managing API keys

Starting in 8.11, generate and manage application programming interface (API) keys to use in automation processes for user management in IBM Maximo Application Suite.

About this task

As a suite administrator, you can create, edit, or delete API keys in Maximo Application Suite. You can activate, deactivate, or refresh authentication tokens for the API keys. The API keys are used in user management for automation processes.

Log in to Maximo Application Suite as an administrator user.

Procedure

- In IBM Maximo Application Suite, from the side navigation menu, select **API keys**.
 - Create an API key.
 1. On the **API keys** page, click **Create API key**.
 2. On the **Create API key** page, enter the following details:

Description

Text that provides context about the API key, such as the purpose of the API key.

Authentication token expiry

The number of days or provide custom days for the authentication token to expire. This token can be used for the duration that the authentication token for the API key is valid before you need to refresh. The token is automatically inactive and no longer available to use after the expiration date.

3. Select the suite administrative privileges, such as **System configuration** or **User management**, that are applicable to the API key.
4. Click **Submit**.

Tip: Copy the authentication token details. If authentication token details are lost, you cannot recover the details. To create a token, you must create an API key.

5. Verify that the key is saved and available in the list of API keys.
- Edit an API key.
 1. On the **API keys** page, select the API key that you need to edit.
 2. From the **More actions** icon, select **Edit**.

Alternatively, you can edit the API key after you open an API key to view the key details.
 3. On the **Edit API key** page, change details, such as the description, authentication token expiry duration, and administrative privileges.
 4. Click **Submit**.
 5. Verify that the key is saved and available in the list of API keys.
 - Delete an API key.
 1. On the **API keys** page, select a single key or multiple keys to delete.
 2. Click **Delete**.
 3. In the **Confirm delete API key** dialog, click **Delete**.
 4. Verify that the key or keys are removed from the list of API keys.
 - Activate an API key.
 1. On the **API keys** page, select a single key or multiple keys.
 2. Click **Activate**.
 3. In the list of API keys, verify that the status of the API key changes to active.
 - Deactivate an API key.
 1. On the **API keys** page, a select single key or multiple keys.
 2. Click **Dectivate**.
 3. In the list of API keys, verify that the status of the API key changes to inactive.
 - Refresh an authentication token.
 1. Select a single key or multiple keys and click **Refresh authentication token**.
 2. In the **Confirm refresh authentication tokens** dialog, click **Confirm**.

When you refresh the authentication tokens for the API keys, new tokens are generated, and the current tokens expire.

Related information

[Maximo Application Suite APIs 9.1](#)

[Maximo Application Suite Admin APIs 9.0](#)

[Maximo Application Suite Admin APIs 8.11](#)

Audit logging in Maximo Application Suite

As a Suite administrator, starting in 8.11, you can forward all logging from Red Hat OpenShift into an external system so that logs can be aggregated and securely stored. The IBM Maximo Application Suite logs can be aggregated from an Red Hat OpenShift Cluster to third-party systems.

Multiple resources within Maximo Application Suite must be considered for audit events and creating a cohesive audit trail.

For more information about logging subsystem provided by Red Hat OpenShift, see the following Red Hat OpenShift documentation.

- [Understanding the logging subsystem for Red Hat OpenShift](#)
- [Installing the logging subsystem for Red Hat OpenShift](#)
- [Forwarding logs](#)

Note: Audit logs are unavailable for forwarding for IBM Cloud when you use the Red Hat OpenShift logging subsystem. For Red Hat OpenShift clusters hosted by IBM, see [Reviewing service, API server, and worker node logs](#).

For container and application logs, and not for audit logs, IBM Log Analysis works out of the box with Red Hat OpenShift clusters hosted in IBM Cloud. You need not install the Red Hat OpenShift logging subsystem unless logs are exported somewhere other than IBM Log Analysis.

Extracting audit events from the Red Hat OpenShift Container Platform audit logs

Many audit events that are captured by the Red Hat OpenShift Container Platform audit logging feature are relevant to Maximo Application Suite. Therefore, it is recommended to include audit events captured by Red Hat OpenShift into any audit trail. For more information, see [Red Hat OpenShift Container Platform audit logging](#).

Note: Red Hat OpenShift audit logs may not be directly accessible - even from the Red Hat OpenShift logging subsystem. Depending on the hosting environment for the Red Hat OpenShift cluster special steps might be required to access the audit logs. For example, the following documentation must be followed to obtain Red Hat OpenShift audit logs when the hosting environment is IBM Cloud [Reviewing service, API server, and worker node logs](#).

After the Red Hat OpenShift Container Platform audit logs are forwarded to a third-party system, they can easily be queried by using the applicable Maximo Application Suite and Manage namespaces. For example, consider the following sample audit log event:

Example of audit event

```
{
  "level": "Metadata",
  "auditID": "1234e7f-32cb-44ab-85e6-1ef57d901ff5",
  "stage": "RequestReceived",
  "requestURI": "/apis/config.mas.ibm.com/v1/namespaces/mas-pfvttjq-core/mongocfgs/pfvttjq-mongo-system/status",
  "verb": "update",
  "user": {
    "username": "system:serviceaccount:mas-pfvttjq-core:ibm-mas-entitymgr-mongocfg",
    "uid": "01a234f5-f678-9101-a3c6-1a87c20a3d1a",
    "groups": ["system:serviceaccounts", "system:serviceaccounts:mas-pqrstuv-core"],
    "system:authenticated": true,
    "extra": {
      "authentication.kubernetes.io/pod-name": ["pqrstuv-entitymgr-mongocfg-6b9d6bc5c8-4sv2s"],
      "authentication.kubernetes.io/pod-uid": ["1ac345aa-f6a7-8c90-a1a1-f3eda571f72b"]
    }
  },
  "sourceIPs": ["123.18.456.789"],
  "userAgent": "ansible-operator/v0.0.0 (linux/amd64) kubernetes/$Format",
  "objectRef": {
    "resource": "mongocfgs",
    "namespace": "mas-pfvttjq-core",
    "name": "pvfvtjq-mongo-system",
    "apiGroup": "config.mas.ibm.com",
    "apiVersion": "v1",
```

```
    "subresource": "status"
  },
  "requestReceivedTimestamp": "2023-07-07T03:07:42.949453Z",
  "stageTimestamp": "2023-07-07T03:07:42.949453Z"
}
```

To refer to events related to Maximo Application Suite, see [About the API audit log](#). The `requestURI` and the `objectRef` can be used for querying namespace, Custom Resource type or both.

Detailed documentation describing the entries that appear in an Red Hat OpenShift audit log event can be found at [Red Hat OpenShift Container Platform audit logging](#).

Extracting audit events from Maximo Application Suite container logs

Maximo Application Suite container logs (that is, pods in the core namespace) contains audit log entries for specific actions. For example, audit events that might appear in container logs include - but not limited to:

- API Request and Responses
- Authentication and Authorization events
- User management

Maximo Application Suite audit events can be queried from container logs that have been forwarded to third-party systems by looking for log messages that are logged at the custom level AUDIT or containing the following string AIUEV. All Maximo Application Suite audit events are logged at the custom log level AUDIT. Any log messages that are associated with a log message logged at the level AUDIT has a message that uses a message ID prefixed with AIUEV. The particular source or container can also be used when harvesting Maximo Application Suite audit log events. The source container should be considered and recorded when creating the audit log trail.

An example of an audit log event from a Maximo Application Suite container can be:

```
2023-07-05 17:54:15,341 api.passwordPolicy AUDIT AIUEV1001I: Response GET /utils/
passwordpolicy? from 172.30.219.135 by user mDVVqsrszXWCPuUAdXIG1omcS6701t9cv resulted in 200 OK
2023-07-05 17:54:15,859 api.passwordPolicy AUDIT AIUEV1000I: Request GET /utils/passwordpolicy?
from 172.30.219.135 by user mDVVqsrszXWCPuUAdXIG1omcS6701t9cv
2023-07-05 17:54:15,874 api.passwordPolicy AUDIT AIUEV1001I: Response GET /utils/
passwordpolicy? from 172.30.219.135 by user mDVVqsrszXWCPuUAdXIG1omcS6701t9cv resulted in 200 OK
2023-07-05 17:54:15,977 api.users.v2userAPIs AUDIT AIUEV1000I: Request POST /v2/users? from
172.30.219.135 by user mDVVqsrszXWCPuUAdXIG1omcS6701t9cv
```

Keeping track of the container that produced the event that a trail can be established around the following:

- The time of the event
- The container that produced the event
- The event message

Audit events for Maximo Manage or Maximo Asset Management

For audit events specific to Maximo Manage or Maximo Asset Management, see [Viewing Maximo Manage Logs](#).

Importing data

You can import data from a .csv file to update the database and simultaneously create multiple records, such as user records. You can use this file to add the data and ensure that the format adheres to the import processing rules.

Procedure

1. On the side navigation menu, select **Suite > Security > Users** and then click the **Import data** icon.

2. In the .csv file, provide the user information to create the user records, such as identity details, contact information, access entitlements, and account status.
3. Select the **Import data** action.
4. In the **Import data** window, upload the completed file.
5. Specify the **Delimiter** and **Text Qualifier**.
6. Add **Object Structure**.
Use the search option to select an object structure.
7. To validate that the data in the file adheres to the import processing rules, click **Validate**.
If there are any issues, you can address the issues in the file and upload the file again.
8. Import the file.

Related tasks

[Importing users in Maximo Application Suite 9.1](#)

Exporting data

Starting in IBM Maximo Application Suite 9.1, you can export or download data to an external file system for backup or importing to another Maximo Application Suite instance. For example, to view or edit user profiles, export user data to a file such as a .csv file.

Procedure

1. On the side navigation menu, select **Suite > Security > Users** and then click the download icon.
2. In **Export data**, export either visible columns or custom columns data.
When you select custom columns, you can select a template to export data, and specify **Delimiter** and **Text Qualifier** values.
3. Export the file.

Related tasks

[Importing users in Maximo Application Suite 9.1](#)

SaaS **Managing SaaS users**

You can create and manage SaaS users at the suite-level and also give users access and entitlement to the applications that they need. After the user is created, you can give users specific permissions to access individual applications. To connect an external server for authentication or user sync registry, you can submit a request with IBM Support.

Related concepts

[Getting started as a SaaS suite administrator](#)

Get started in Maximo Application Suite as a Service applications by adding users and specifying their entitlements, requesting server authentication and user synchronization, and monitoring application usage.

SaaS **Requesting external authentication and user registry synchronization in Maximo Application Suite as a Service**

Maximo Application Suite as a Service supports external Lightweight Directory Access Protocol (LDAP) authentication and Security Assertion Markup Language (SAML) authentication. By using user registry synchronization, you can automate your user management by synchronizing users and groups between an LDAP server and your local user registry. To connect an external server for authentication or user sync registry, you open a case with IBM Support.

Procedure

1. On the **Suite administration** page, click **Users > Request LDAP SAML or user sync configuration**.
2. Select the options that you need to configure:
 - LDAP authentication
 - SAML authentication
 - User registry synchronization
3. In the template, enter the parameter information for the configuration options that you selected.
4. Copy the information and then click **Open IBM Support portal**.
5. In the support portal, click **Open a case** and specify the following information:

Field	Case information
Case title	Request Maximo Application Suite as a Service configuration
Product	Maximo Application Suite as a Service
Severity information	Select the severity for the case.
Case description	Paste the configuration and parameter information from the template.

You might need to log in with your IBM credentials.

6. Submit the case.

What to do next

After you submit the case, IBM Support will be in contact to help you complete the configuration.

If access and entitlement are not specified in this support case, you can user give users access and entitlement to suite applications by updating their user records.

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Related tasks

[Adding users in Maximo Application Suite as a Service](#)

To add users to Maximo Application Suite as a Service, you create a record to specify user details, such as username and contact information. You can also specify access and entitlement for suite applications. You can also create multiple user records simultaneously by importing user information in a .csv file template. .

SaaS

Adding users in Maximo Application Suite as a Service

To add users to Maximo Application Suite as a Service, you create a record to specify user details, such as username and contact information. You can also specify access and entitlement for suite applications. You can also create multiple user records simultaneously by importing user information in a .csv file template. .

About this task

You can create individual users or synchronize user creation. Maximo Application Suite as a Service supports external Lightweight Directory Access Protocol (LDAP) authentication and Security Assertion Markup Language (SAML) authentication. By using user registry synchronization, you can automate your user management by synchronizing users and groups between an LDAP server and your local user

registry. To connect an external server for authentication or user sync registry, you open a case with IBM Support.

Procedure

1. To create a user record, on navigation menu, in the suite security page, and click **Users > Create users**.
 - a) In the Identity section, specify the user's display name, user ID, and primary email.
You can also enter user information, such as contact information and define the user's locale and time zone.
 - b) If the user needs administrator access, in the Suite administration access section, select **User management** to set their administration privileges.
A user with user management access can create and manage users, assign entitlement and access levels, and view usage data.
 - c) In the Application access section, give the user access to applications that they need by selecting the role or entitlement.
2. To create multiple user records, import user information by using the .csv file template.
 - a) On the navigation menu in the suite security page, select **Users**, and click the **Import users** icon.
 - b) Download the .csv template.
 - c) In the .csv template, provide the information for users who are added to the user records.
User information includes identity details, contact information, and access entitlements. For more information about completing the template, see [“Importing users in Maximo Application Suite in 9.0 and earlier”](#) on page 801.
 - d) Upload the .csv file and click **Import**.
When you import the completed .csv file, the data that you provided is processed and creates a record for each user in the file.

What to do next

Users are sent an email with their login details and the URL to access to their environment.

For users that are assigned a Maximo Manage entitlement, you can provide security group permissions in Maximo Manage to define the Maximo Manage applications, options, and data that the user can access.

Related concepts

[AppPoints for Maximo Application Suite as a Service](#)

Related tasks

[Requesting external authentication and user registry synchronization in Maximo Application Suite as a Service](#)

Maximo Application Suite as a Service supports external Lightweight Directory Access Protocol (LDAP) authentication and Security Assertion Markup Language (SAML) authentication. By using user registry synchronization, you can automate your user management by synchronizing users and groups between an LDAP server and your local user registry. To connect an external server for authentication or user sync registry, you open a case with IBM Support.

Related reference

[Importing users in Maximo Application Suite in 9.0 and earlier](#)

To create multiple users in Maximo Application Suite, use the template file to import new users and ensure that the format for the user information adheres to the import processing rules. After you import users, you can also use the template file to modify user information and delete users.

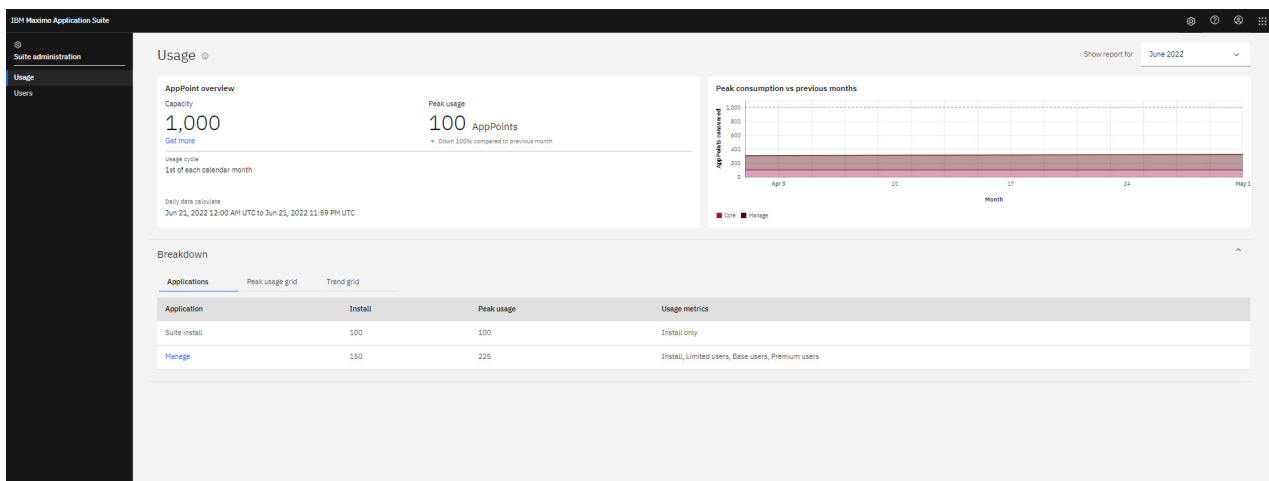
SaaS **Monitoring AppPoint usage for Maximo Application Suite as a Service**

As a SaaS suite administrator, you can use the Usage dashboard to monitor the peak usage trends of how AppPoints are consumed in each application, such as when users use the Maximo Manage application. You can manage the daily usage of AppPoints for each application and ensure that your organization stays within their AppPoint capacity.

For more information about AppPoints, see [AppPoints for Maximo Application Suite as a Service](#).

AppPoints are spent on installing and using capabilities and applications within the suite. Each application has its own charging model, which is defined in the catalog entry for that capability. For example, the Maximo Manage application is charged AppPoints based on users, while Maximo Visual Inspection is charged usage AppPoints based on models trained and inferences made.

The total AppPoint usage is the sum of usage across the suite and individual applications, and can be viewed in the Usage dashboard.



To review your organization's AppPoint usage, on the **Suite administration** page, select **Usage** from the side navigation menu and then on the Usage dashboard, select **Show report for** the month that you want to review.

The main page of the Usage dashboard provides an overview of AppPoints that shows your subscription capacity and peak usage for the selected month. The **Capacity** shows the number of subscribed AppPoints for your organization. For example, if your organization purchase a subscription of 1,000 AppPoints, then your subscription capacity is 1,000 AppPoints. **Peak usage** shows the highest amount of AppPoints that were used at a given moment within that calendar month and a comparison to the previous month, which is a helpful indicator if the peak usage is close to exceeding the subscription capacity.

The graphical view that compares your peak consumption with previous months also helps you analyze usage trends and determine whether you might exceed your current capacity. If you determine that you need to purchase more AppPoints, on the **Suite administration** page, you can select **Usage** from the side navigation menu and on the AppPoint overview card, select **Get more**.

For detailed information of usage per application, select the application in the Breakdown section.

Related concepts

[Getting started as a SaaS suite administrator](#)

Get started in Maximo Application Suite as a Service applications by adding users and specifying their entitlements, requesting server authentication and user synchronization, and monitoring application usage.

[AppPoints for Maximo Application Suite as a Service](#)

Purchasing AppPoints for Maximo Application Suite as a Service

If your AppPoint usage or number of users increases in Maximo Application Suite as a Service, you can adjust your AppPoint capacity by purchasing more AppPoints.

Procedure

1. On the **Suite administration** page, from the side navigation menu, select **Usage**.
2. On the AppPoint overview card, click **Get more**.
3. In the **Request AppPoints** dialog, enter the information in the template, including your customer ID, contact email, and the number of AppPoints that you want to purchase.
4. Copy the information and then click **Open IBM Support portal**.
5. In the support portal, click **Open a case** and specify the following information:

Field	Case information
Case title	Purchase Maximo Application Suite as a Service AppPoints
Product	Maximo Application Suite as a Service
Severity information	Select the severity for the case.
Case description	Paste the details of the AppPoints request from the template.

You might need to log in with your IBM credentials.

6. Submit the case.

What to do next

After you submit the case, IBM Support will contact you to complete your request.

Related concepts

[AppPoints for Maximo Application Suite as a Service](#)

Scenario: Monitoring AppPoint usage for Maximo Manage in Maximo Application Suite as a Service

Kelly is a SaaS suite administrator and is monitoring usage of the Maximo Manage application in Maximo Application Suite as a Service for her organization.

Reviewing usage and analyzing the data

To review the AppPoint usage, on the **Suite administration** page, Kelly selects **Usage** from the side navigation menu. On the Usage dashboard, Kelly sees that the peak is higher than she expected and that the Maximo Manage application is responsible for most of that usage. In the Breakdown section, on the **Applications** tab, she selects Manage to view the detailed usage report. On the **Daily usage** card, she can see that the peak usage is consistently higher than she anticipated every day in the month, but the average usage was much closer to her original estimation. In the Breakdown of the last 24 hours section, Kelly views the hourly usage to clarify the cause of these high peaks.

Identifying the cause of high peak usage

By analyzing the hourly usage, Kelly noticed that at both 3 PM and 9 PM a substantially higher number of users are accessing Manage than during the rest of the day. To investigate further, Kelly selects 3 PM from the **Select hour** list to show the names and license types of all users who access the application at that time. Because 3 PM is the shift changeover, she realizes that users from both shifts are logged in, which increases the peak usage.

Resolving the issue

Kelly can now choose to either contact her IBM representative or submit a request with IBM Support to discuss an AppPoint subscription that has the capacity to contain these usage increases or make business process changes within her own organization to ensure a prompt access handover at shift changes.

Customer-managed

Monitoring Maximo Application Suite

Configure Red Hat OpenShift cluster monitoring and install Grafana to help you monitor Maximo Application Suite.

Customer-managed

Configuring Red Hat OpenShift cluster monitoring

Maximo Application Suite applications provide application level metrics and dashboards for monitoring various aspects for application health and performance. Maximo Application Suite uses the Prometheus monitoring stack within OCP for storing application level metrics. Maximo Application Suite also uses Grafana for rendering application level metrics in integrated dashboards.

Red Hat OpenShift Container Platform (OCP) is preconfigured with a Prometheus based monitoring stack that collects resource level metrics from compute nodes in the cluster. Some examples of the metrics that are collected by OCP are compute node CPU, memory, disk, and I/O metrics. Maximo Application Suite applications cannot use the preconfigured Prometheus cluster for collecting Maximo Application Suite application metrics, as it is reserved for OCP cluster metrics. Instead, a second Prometheus cluster can be enabled and configure to collect metrics from user-defined projects.

For more information about the Red Hat OpenShift Monitoring stack, see [Red Hat OpenShift Container Platform : Monitoring overview](#).

Tip: This task maps to the following Ansible role: `cluster_monitoring`. For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

Before you begin

Consider how many days to store Prometheus metrics. The number of retention days determines how much storage to configure for both the base and user workload Prometheus clusters. Allocate 5 GB - 10 GB of storage for each retention day. The amount of storage that is required by Prometheus depends on the number of compute nodes in the cluster, the number of Maximo Application Suite applications installed, and the number of retention days.

About this task


Use the following storage classes to configure Prometheus storage, according to the Cloud Service Provider hosting your Red Hat OpenShift cluster:

Cloud Service Provider	Prometheus Storage Classes - <code>\$_{PROMETHEUS_STORAGE_CLASS}</code>
On-premises	<code>ocs-storagecluster-ceph-rbd</code>
AWS	<code>ocs-storagecluster-ceph-rbd</code>
Azure	
IBM Cloud	<code>ibmc-block-bronze</code>

Procedure

Install by using the Red Hat OpenShift Container Platform web console.

1. Update the cluster-monitoring-config and user-workload-monitoring-config ConfigMaps.

- a) Click **Import YAML** ()
- b) Enter the following YAML to configure both the base and user workload Prometheus clusters.
 - Replace `${PROMETHEUS_STORAGE_CLASS}` with the corresponding Prometheus Storage Class from the preceding table according to your Cloud Service Provider hosting your installation.
 - Any storage class that supports RWO access mode and file system volume mode is sufficient. The I/O requirements for the Prometheus persistent volumes are not significant.
 - In the example YAML, both the base and user workload Prometheus clusters are configured to retain metrics for 15 days.

```
---
apiVersion: v1
kind: ConfigMap
data:
  config.yaml: |
    prometheusOperator:
      baseImage: quay.io/coreos/prometheus-operator
      prometheusConfigReloaderBaseImage: quay.io/coreos/prometheus-config-reloader
      configReloaderBaseImage: quay.io/coreos/configmap-reload
    prometheusK8s:
      retention: "15d"
      baseImage: openshift/prometheus
      volumeClaimTemplate:
        spec:
          storageClassName: "${PROMETHEUS_STORAGE_CLASS}"
          resources:
            requests:
              storage: "150Gi"
    alertmanagerMain:
      baseImage: openshift/prometheus-alertmanager
      volumeClaimTemplate:
        spec:
          storageClassName: "${PROMETHEUS_STORAGE_CLASS}"
          resources:
            requests:
              storage: "20Gi"
    enableUserWorkload: true
    nodeExporter:
      baseImage: openshift/prometheus-node-exporter
    kubeRbacProxy:
      baseImage: quay.io/coreos/kube-rbac-proxy
    kubeStateMetrics:
      baseImage: quay.io/coreos/kube-state-metrics
    grafana:
      baseImage: grafana/grafana
    auth:
      baseImage: openshift/oauth-proxy
metadata:
  name: cluster-monitoring-config
  namespace: openshift-monitoring
```

```
---
apiVersion: v1
kind: ConfigMap
data:
  config.yaml: |
    prometheus:
      retention: "15d"
      volumeClaimTemplate:
        spec:
          storageClassName: "${PROMETHEUS_STORAGE_CLASS}"
          resources:
            requests:
              storage: "150Gi"
metadata:
```

```
name: user-workload-monitoring-config
namespace: openshift-user-workload-monitoring
```

c) Click **Create**.

2. On the **Workloads > StatefulSets** page, switch to the `openshift-user-workload-monitoring` project and wait for the `prometheus-user-workload` statefulset to indicate that there are two pods in `Running` state.

No configuration required in MAS. PodMonitor and ServiceMonitor resources that are created by Maximo Application Suite and Maximo Application Suite applications will automatically be registered with the user workload Prometheus cluster. MAS metrics are scraped by Prometheus.

Customer-managed **Installing Grafana**

Red Hat OpenShift Container Platform is preconfigured with a Grafana instance for visualizing Prometheus metrics from compute nodes in the cluster. This Grafana instance is reserved for OCP cluster metrics, such as compute node CPU, memory, disk, and I/O metrics. Maximo Application Suite applications cannot use the base Grafana instance. You can install another Grafana instance to host dashboards for Maximo Application Suite applications.

For more information about Red Hat OpenShift monitoring, see [Red Hat OpenShift Container Platform : Accessing third-party UIs](#).

Not required to install Maximo Application Suite, but required for monitoring Maximo Application Suite.

Note:

- Starting in 9.0.5 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM System/390x architecture, the Grafana operators are not supported.
- Starting in 9.0.12 and 9.1, if Maximo Application Suite core and Maximo Manage base are configured on IBM Power (ppc64le) architecture, the Grafana operators are not supported.

Before you begin

Ensure that the user workload monitoring Prometheus cluster is enabled and configured.

About this task

Tip: This task maps to the following Ansible role: `cluster_monitoring`. For more information, see [“IBM Maximo Application Suite installation with Ansible collection”](#) on page 276.

Use the following storage classes to configure Grafana storage, according to the Cloud Service Provider hosting your Red Hat OpenShift cluster:

Cloud Service Provider	Grafana Storage Classes - <code>\$_{GRAFANA_STORAGE_CLASS}</code>
On premises	<code>ocs-storagecluster-cephfs</code>
Amazon Web Services	<code>ocs-storagecluster-cephfs</code>
Microsoft Azure	
IBM Cloud	<code>ibmc-block-bronze</code>

Procedure

Install by using the Red Hat OpenShift Container Platform web console.

1. Configure role-based access control for Grafana.

The Grafana operator requires permission to scan Maximo Application Suite application namespaces for GrafanaDashboard custom resources.

a) In the banner, click **Import YAML** (.

b) Enter the following YAML.

```
---
apiVersion: operators.coreos.com/v1alpha2
kind: OperatorGroup
metadata:
  name: grafana-operator
  namespace: openshift-user-workload-monitoring
spec:
  targetNamespaces:
  - openshift-user-workload-monitoring
```

c) Enter the following YAML.

```
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: grafana-operator
  namespace: openshift-user-workload-monitoring
```

d) Enter the following YAML.

```
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: grafana-operator
rules:
  - apiGroups:
    - ""
    resources:
      - pods
      - services
      - endpoints
      - persistentvolumeclaims
      - configmaps
      - secrets
      - serviceaccounts
      - configmaps
    verbs:
      - get
      - list
      - create
      - update
      - delete
      - deletecollection
      - watch
  - apiGroups:
    - ""
    resources:
      - events
    verbs:
      - get
      - list
      - watch
      - create
      - delete
      - update
      - patch
  - apiGroups:
    - apps
    resources:
      - deployments
      - deployments/finalizers
      - daemonsets
      - replicaset
      - statefulsets
    verbs:
      - get
      - list
      - create
      - update
      - delete
      - deletecollection
      - watch
```

```

- apiGroups:
  - route.openshift.io
resources:
  - routes
  - routes/custom-host
verbs:
  - get
  - list
  - create
  - update
  - delete
  - deletecollection
  - watch
  - create
- apiGroups:
  - extensions
resources:
  - ingresses
verbs:
  - get
  - list
  - create
  - update
  - delete
  - deletecollection
  - watch
- apiGroups:
  - integreatly.org
resources:
  - grafanas
  - grafanas/status
  - grafanas/finalizers
  - grafanadashboards
  - grafanadatasources
  - grafanadatasources/status
verbs:
  - get
  - list
  - create
  - update
  - delete
  - deletecollection
  - watch
- apiGroups:
  - networking.k8s.io
resources:
  - ingresses
verbs:
  - get
  - list
  - create
  - update
  - delete
  - deletecollection
  - watch
  - create

```

e) Enter the following YAML.

```

---
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: aggregate-grafana-admin-edit
  labels:
    rbac.authorization.k8s.io/aggregate-to-admin: "true"
    rbac.authorization.k8s.io/aggregate-to-edit: "true"
rules:
- apiGroups:
  - "integreatly.org"
resources:
  - grafanas
  - grafanas/status
  - grafanas/finalizers
  - grafanadashboards
  - grafanadatasources
  - grafanadatasources/status
verbs:
  - "get"
  - "list"
  - "watch"

```

```
- "create"
- "update"
- "patch"
- "delete"
- "deletecollection"
```

f) Enter the following YAML.

```
---
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: aggregate-grafana-view
  labels:
    rbac.authorization.k8s.io/aggregate-to-view: "true"
    rbac.authorization.k8s.io/aggregate-to-cluster-reader: "true"
rules:
- apiGroups:
  - "integreatly.org"
  resources:
  - grafanas
  - grafanas/status
  - grafanas/finalizers
  - grafanadashboards
  - grafanadatasources
  - grafanadatasources/status
  verbs:
  - "get"
  - "list"
  - "watch"
```

g) Enter the following YAML.

```
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: grafana-operator
roleRef:
  name: grafana-operator
  kind: ClusterRole
  apiGroup: ""
subjects:
- kind: ServiceAccount
  name: grafana-operator
  namespace: openshift-user-workload-monitoring
```

h) Click **Create**.

2. Install the Grafana Operator.

Follows these steps to install Grafana Operator v5.

a) In the banner, click **Import YAML** (). Enter the following YAML.

```
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: grafana-operator
  namespace: openshift-user-workload-monitoring
  labels:
    operators.coreos.com/grafana-operator.openshift-user-workload-monitoring
spec:
  channel: v5
  installPlanApproval: Automatic
  name: grafana-operator
  source: community-operators
  sourceNamespace: openshift-marketplace
config:
  env:
  - name: "WATCH_NAMESPACE"
    value: ""
  - name: "DASHBOARD_NAMESPACES_ALL"
    value: "true"
```

b) Click **Create**.


c) Verify that the Grafana operator installed successfully.

```
oc get csv -n openshift-user-workload-monitoring -l operators.coreos.com/grafana-operator.openshift-user-workload-monitoring=""
```

Sample output for v5

NAME	DISPLAY	VERSION	REPLACES	PHASE
grafana-operator.v5.6.3	Grafana Operator	5.6.3	grafana-operator.v5.6.2	Succeeded

3. Create the Grafana instance resource.

a) In the banner, click **Import YAML** (). Enter the following YAML to create the Grafana instance resource.

Note:

- Replace `${GRAFANA_STORAGE_CLASS}` by the corresponding Grafana Storage Class from the preceding table according to your Cloud Service Provider hosting your installation.
- Any storage class that supports RWX access mode and file system volume mode is sufficient. The I/O requirements for the Grafana persistent volumes are not significant.

```
---
apiVersion: grafana.integreatly.org/v1beta1
kind: Grafana
metadata:
  name: mas-grafana
  namespace: openshift-user-workload-monitoring
labels:
  dashboards: "grafanav5"
spec:
  config:
    auth:
      disable_login_form: "false"
      disable_signout_menu: "true"
  log:
    level: warn
    mode: console
  dataStorage:
    accessModes:
      - ReadWriteOnce
    class: ${GRAFANA_STORAGE_CLASS}
    size: 20Gi
  deployment:
    strategy:
      type: Recreate
  spec:
    replicas: 3
    template:
      spec:
        containers:
          - name: grafana
            readinessProbe:
              httpGet:
                path: /api/health
                port: 3000
                scheme: HTTP
              failureThreshold: 5
              initialDelaySeconds: 30
              periodSeconds: 10
              successThreshold: 1
              timeoutSeconds: 20
            livenessProbe:
              httpGet:
                path: /api/health
                port: 3000
                scheme: HTTP
              failureThreshold: 5
              initialDelaySeconds: 60
              periodSeconds: 10
              successThreshold: 1
              timeoutSeconds: 20
```

b) Click **Create**.

On the **Workloads Deployments** page, switch to the `openshift-user-workload-monitoring` project and wait for the `grafana-deployment` deployment to indicate that three pods are in Ready state.

4. Add the `cluster-monitoring-view` cluster role to the Grafana service account.

```
oc adm policy add-cluster-role-to-user cluster-monitoring-view -z grafana-serviceaccount -n openshift-user-workload-monitoring
```

Sample output


```
clusterrole.rbac.authorization.k8s.io/cluster-monitoring-view added: "grafana-serviceaccount"
```

5. Get the Bearer token from the `grafana-serviceaccount` service account.

```
oc sa get-token grafana-serviceaccount -n openshift-user-workload-monitoring
```

Save this Bearer token. You need it to create the Prometheus data source in the next step.

6. Create the Prometheus data source and configure it as the default data source in Grafana.

- a) In the banner, click **Import YAML** (). Enter the following YAML to create the GrafanaDataSource resource:

Replace `${TOKEN}` with the Bearer token from the previous step.

```
---
apiVersion: grafana.integreatly.org/v1beta1
kind: GrafanaDataSource
metadata:
  name: mas-prom-grafanadatasource
  namespace: openshift-user-workload-monitoring
spec:
  instanceSelector:
    matchLabels:
      dashboards: "grafanav5"
  datasource:
    name: prometheus
    type: prometheus
    access: proxy
    url: https://thanos-querier.openshift-monitoring.svc.cluster.local:9091
    isDefault: true
    editable: true
  jsonData:
    httpHeaderName1: Authorization
    timeInterval: 5s
    tlsSkipVerify: true
  secureJsonData:
    httpHeaderValue1: Bearer ${TOKEN}
```

- b) Click **Create**.


7. Get the Grafana admin credentials from the `grafana-admin-credentials` secret.

```
oc get secret grafana-admin-credentials -n openshift-user-workload-monitoring -o jsonpath='{.data.GF_SECURITY_ADMIN_USER}' | base64 -d ; echo
```

```
oc get secret grafana-admin-credentials -n openshift-user-workload-monitoring -o jsonpath='{.data.GF_SECURITY_ADMIN_PASSWORD}' | base64 -d ; echo
```

8. Log in to the Grafana console

`http://grafana-route-openshift-user-workload-monitoring.apps.cluster1.example-cluster.com`

- a) Click the login icon . Enter the Grafana admin credentials from the previous step and log in.

- b) On the side navigation, click the dashboards icon  then click **Manage**.

Grafana resources that are created during application installation are imported by the Grafana operator. The GrafanaDashboard scans resources across all namespaces and the resources are now visible. Dashboards are organized into folders that correspond to namespaces. Expand a folder to see dashboards.

9. Loading the Maximo Manage dashboard into Grafana.

- The IBM Maximo Manage application does not include a dashboard however you can create a dashboard for Manage by using the following steps:
 - a. To load the Grafana Dashboard for Manage, in the Grafana dashboard, download the `maximo-dashboard.json` file from: [maximo-dashboard.zip](#).
 - b. Go to Import it into Dashboards/Manage.
 - c. Click the **Import** button. Note that you see the Import option only when you log in as the admin user that you entered when you created the Grafana instance.
 - d. In the **Import via panel json** field, enter the JSON data from the file downloaded in step “10.a” on page 832.
 - e. Click the **Load** button. On the next screen, set the folder name to the namespace of your Manage instance and click the **Import** button.
 - f. You will now see your Manage dashboards.
- The IBM Maximo Manage and IBM Maximo Visual Inspection applications do not include a dashboard. However, you can create dashboards for Maximo Manage and Maximo Visual Inspection by using the following steps:
 - a. To load the Grafana Dashboard for Maximo Manage and Maximo Visual Inspection, in the Grafana dashboard, download the `maximo-manage-dashboard.json` and `maximo-mvi-dashboard.json` files from: [maximo-dashboard.zip](#).
 - b. In the Grafana web interface, from **Dashboards > Manage**, click **Import**.
 - i) In the **Import via panel json** field, enter the JSON data from the file `maximo-manage-dashboard.json` downloaded in step “10.a” on page 832.
 - ii) Click **Load**. On the next screen, set the folder name to the namespace of your Maximo Manage instance and click **Import**.
 - c. In the Grafana web interface, from **Dashboards > Manage**, click **Import**.
 - i) In the **Import via panel json** field, enter the JSON data from the file `maximo-mvi-dashboard.json` downloaded in step 10.a
 - ii) Click **Load**. On the next screen, set the folder name to the namespace of your IBM Maximo Visual Inspection Edge instance and click **Import**.
 - d. You will now see your Maximo Manage and Maximo Visual Inspection dashboards.

Note: Starting in Maximo Visual Inspection 8.8, the Grafana dashboard is installed automatically.

What to do next

No configuration required in Maximo Application Suite. PodMonitor and ServiceMonitor resources that are created by Maximo Application Suite and Maximo Application Suite applications are automatically registered with the user workload Prometheus cluster. Maximo Application Suite metrics are scraped by Prometheus.

Using the serviceability dashboard

You can enable or disable the serviceability dashboard. By using the serviceability dashboard, you can monitor the health and performance of some applications that are deployed in IBM Maximo Application Suite: IBM Maximo Collaborate, IBM Maximo Health, IBM Maximo Optimizer 8.5.0 and later, and IBM Maximo Predict.

The serviceability dashboard is implemented with Grafana, OpenTelemetry, and Prometheus.

Installing Grafana, OpenTelemetry, and Prometheus

Before you can enable the serviceability dashboard, you must install Grafana, OpenTelemetry, and Prometheus.

Procedure

1. Install Python. Any in-support 3.9 or later version of Python can be used.
2. Install Ansible and check the version by using the following commands.

```
python3 --version
python3 -m pip install ansible junit_xml pymongo xmljson jmespath
kubernetes==12.0.1 openshift==0.12.1
ansible --version
ansible-playbook --version
```
3. Install the Ansible collection. Run the following command to install the [ibm.mas_devops](#) collection directly from Ansible Galaxy.

```
ansible-galaxy collection install ibm.mas_devops
```
4. Login to the Red Hat OpenShift cluster. If your cluster is in the IBM Cloud, log in to the Red Hat OpenShift cluster in IBM Cloud Red Hat OpenShift Kubernetes Service with the IBM Cloud command-line interface. If your cluster is not in the IBM Cloud, log in to the Red Hat OpenShift cluster by using the `oc` command line tool.
5. Run the Ansible [cluster_monitoring](#) role by running a playbook. Use the following command:

```
ansible-playbook playbook.yml
```

You can create a `playbook.yml` file by using the following example:

```
- hosts: localhost
  any_errors_fatal: true
  vars:
    prometheus_storage_class: "ibmc-block-gold"
    prometheus_alertmgr_storage_class: "ibmc-file-gold-gid"
    grafana_instance_storage_class: "ibmc-file-gold-gid"
    cluster_monitoring_include_opentelemetry: true

  roles:
    - ibm.mas_devops.cluster_monitoring
```

Installing the OpenTelemetry operator manually

When you installed Grafana and Prometheus, if you ran the `cluster_monitoring` role in the Red Hat OpenShift cluster and `cluster_monitoring_include_opentelemetry` was set to the default value of `false`, then you need to install the OpenTelemetry operator manually.

Procedure

1. Log in to the web console of the Red Hat OpenShift cluster.
2. From the side navigation menu, click **Operators > OperatorHub** to open the **OperatorHub** page.
3. Search for `opentelemetry`.
4. Click the **Community OpenTelemetry Operator** tile.
5. On the **Community OpenTelemetry Operator** page, click **Install**. Keep the default values for all the input fields.
6. Click **Install** and then wait until the OpenTelemetry operator installation is completed and the OpenTelemetry operator is ready.

Enabling the serviceability dashboard

You can enable the serviceability dashboard for IBM Maximo Collaborate, IBM Maximo Health, IBM Maximo Optimizer 8.5.0 and later, and IBM Maximo Predict.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

About this task

You can enable and disable the serviceability dashboard by editing the YAML of the custom resource instance for the application. For example, you can set `spec.settings.enableOpenTelemetry` to `true` to enable the serviceability dashboard and set `spec.settings.enableOpenTelemetry` to `false` to disable it.

Procedure

1. In Red Hat OpenShift Container Platform, from the side navigation menu, click **Administration > CustomResourceDefinitions**.
2. Search for the custom resource definition for your application.

Table 140. Applications, custom resource definitions, and namespaces

Application	Custom resource definition	Namespace
Maximo Collaborate	CollaborateApp	mas- <i><instanceID></i> -collaborate
Stand-alone Maximo Health	HealthWorkspace	mas- <i><instanceID></i> -health
Maximo Health as part of Maximo Manage	ManageWorkspace	mas- <i><instanceID></i> -manage
Maximo Optimizer	OptimizerApp	mas- <i><instanceID></i> -optimizer
Maximo Predict	PredictWorkspace	mas- <i><instanceID></i> -predict

3. Click the custom resource definition.
4. On the **Instances** tab, click the instance name.
5. On the **YAML** tab, enable OpenTelemetry.

Application	To enable
Stand-alone Maximo Health and Maximo Health as part of IBM Maximo Manage	Set <code>spec.settings.opentelemetry.enabled</code> to <code>true</code> .
Maximo Collaborate, Maximo Optimizer, and Maximo Predict	Set <code>spec.settings.enableOpenTelemetry</code> to <code>true</code> .

6. Click **Save**.
7. From the side navigation menu, click **Workloads > Pods**.
8. Wait until the operator for the application reconciles the changes, and the `otel-collector-0` or `<instanceID>-<workspaceID>-otel-collector-0` pod is displayed in the `mas-<instanceID>-<application>` project.

Viewing metrics on the serviceability dashboard

After the serviceability dashboard is enabled, you can view the service-level objective metrics for IBM Maximo Collaborate, IBM Maximo Health, IBM Maximo Optimizer 8.5.0 and later, and IBM Maximo Predict.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Procedure

1. In a browser, open the Grafana user interface by opening the location of the grafana-route within the project in which the Grafana instance is installed.
For example, if the Grafana instance is installed in the grafana namespace, the location might be `https://grafana-route-grafana.<your_cluster_domain>`.
2. In the Grafana user interface, click **Dashboards** > **Manage** to open the Grafana dashboards user interface.
3. Click the **mas-<instanceID>-<application>** folder to expand it and click **MAS <application> Metrics** to open the serviceability dashboard and view metrics on the dashboard.

You can switch among different IBM Maximo Application Suite instance IDs. You can also select the time range within which you want to view the metrics.

Server-level monitoring metrics

The server-level monitoring metrics provide the response time and response code for IBM Maximo Collaborate, IBM Maximo Health, IBM Maximo Optimizer, and IBM Maximo Predict APIs.

Note: Starting in Maximo Application Suite 9.1, IBM Maximo Assist is now named IBM Maximo Collaborate.

Maximo Collaborate

The server-level monitoring metrics for Maximo Collaborate are described in the following table.

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Query API - query documents	mas_collaborate_query_querydocsapi_duration_seconds	https://<workspaceId>.collaborate.<instanceID>.<domain>/api/v2/document-query/query	GET
Query API - build filter criteria	mas_collaborate_query_buildfiltercriteriaapi_duration_seconds	https://<workspaceId>.collaborate.<instanceID>.<domain>/api/v2/document-query/build-filter	POST
Query API - list query configurations	mas_collaborate_query_listallqueryconfigapi_duration_seconds	https://<workspaceId>.collaborate.<instanceID>.<domain>/api/v2/document-query/queryConfigs	GET
Query API - get query configuration	mas_collaborate_query_getqueryconfigapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v2/document-query/queryConfigs/<queryConfigID>	GET

Table 141. Server-level monitoring metrics for Maximo Collaborate (continued)

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Query API - list field values	mas_collaborate_query_listfieldsvaluesapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v2/document-query/field-values	GET
Diagnosis API - list diagnosis libraries	mas_collaborate_diagnosis_listalldiagnosislibrariesapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v2/diagnosis/libraries	GET
Diagnosis API - get diagnosis library count	mas_collaborate_diagnosis_countdiagnosislibrariesapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v2/diagnosis/librariesCount	GET
Diagnosis API - trigger diagnosis flow	mas_collaborate_diagnosis_triggerdiagnosisflowapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v2/diagnosis/diagnosis/<libraryID>	POST
Document management API - create document collection	mas_collaborate_docmgmt_createdoccollectionapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v1/document-management/collections	POST
Document management API - bind collection with Watson Discovery (WD) collection	mas_collaborate_docmgmt_bindwdcollectionapi_duration_seconds	https://<workspaceID>.collaborate.<instanceID>.<domain>/api/v1/document-management/discovery/bindings	POST

Maximo Health

The server-level monitoring metrics for Maximo Health are described in the following table.

Table 142. Server-level monitoring metrics for Maximo Health

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Get matrix count for each cell	mas_health_matrix_matrixcountapi_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiasset?action=getmatrixcount...	POST
List assets for matrix drill in of current scores	mas_health_matrix_get_matrixdrillincurrent_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiasset?oslc.select=...ahscore_xxx.value...	GET
List assets for matrix drill in of future scores	mas_health_matrix_get_matrixdrillinfuture_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiasset?oslc.select=...ahimedfuturescore...	GET
View charts statistic information	mas_health_charts_get_charts_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiasset?action=GETAHINSIGHTSDATA...	POST
List AIO projects	mas_health_aio_list_projects_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiaioproject	GET
View AIO strategies for one project	mas_health_aio_get_strategiesforoneproject_duration_seconds	https://<workspaceID>.health.<instanceID>.<domain>/maximo/oslc/os/mxapiaiostrategy?...oslc.where=aioprojectnum...	GET

Table 142. Server-level monitoring metrics for Maximo Health (continued)

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
List AIO asset replace plans for one strategy	mas_health_aio_get_as sets_plans_per_strategy _duration_seconds	https:// <workspaceID>.health.<instanceID>.<d omain>/maximo/ oslc/os/ mxapiaioschedasset ?...domaininternal where=action=REPLA CE...&aioprojectnu m=...AND aiostrategynum=...	GET
List AIO schedules for one strategy	mas_health_aio_get_sc hedules_per_strategy_d uration_seconds	https:// <workspaceID>.health.<instanceID>.<d omain>/maximo/ oslc/os/ mxapiaioschedule?. ...domaininternalwh ere=action=REPLACE ...&aioprojectnum= ...AND aiostrategynum=...	GET
Run AIO analysis start	mas_health_aio_run_an alysis_start_duration_se conds	https:// <workspaceID>.health.<instanceID>.<d omain>/maximo/ oslc/os/ mxapiaiostrategy/. ...?...action=wsmet hod=startAnalysis. ..	POST

Maximo Optimizer

The server-level monitoring metrics for Maximo Optimizer are described in the following table.

Table 143. Server-level monitoring metrics for Maximo Optimizer

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Create job API	mas_optimizer_api_crea te_job_duration_second s	https:// <optimizer_REST_UR L>/jobs	POST
Submit job API	mas_optimizer_api_exec ute_job_duration_secon ds	https:// <optimizer_REST_UR L>/jobs/<ID>/ execute	POST

Table 143. Server-level monitoring metrics for Maximo Optimizer (continued)

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Upload scenario data API	mas_optimizer_api_upload_scenario_data_duration_seconds	https://<optimizer_REST_URL>/jobs/<ID>/attachments/<attID>/blob	PUT
Get job execution status API	mas_optimizer_api_get_job_status_duration_seconds	https://<optimizer_REST_URL>/jobs/<ID>/execute	GET
Download job attachment API	mas_optimizer_api_download_job_attachment_duration_seconds	https://<optimizer_REST_URL>/jobs/<ID>/attachments/<attID>/blob	GET
List job API	mas_optimizer_api_list_job_duration_seconds	https://<optimizer_REST_URL>/jobs	GET
List project API	mas_optimizer_api_list_projects_duration_seconds	https://<optimizer_REST_URL>/projects	GET

Maximo Predict

The server-level monitoring metrics for Maximo Predict are described in the following table.

Table 144. Server-level monitoring metrics for Maximo Predict

Row title in the serviceability dashboard	Metric name	Monitored URL template	Monitored URL method
Load prediction scoring details	mas_predict_asset_predictions_load_duration_seconds	https://<workspaceID>.predict.<instanceID>.<domain>/ibm/pmi/service/prediction/result	POST

Detailed metrics on the serviceability dashboard

The serviceability dashboard provides detailed metrics data for service-level objective (SLO) metrics.

Expand a row in the serviceability dashboard to view the detailed metrics data for an SLO metric. For each SLO metric, eight cards show the metrics data for latency and error rate. Five cards provide information about the response time, or latency, for each SLO metric. Three cards provide information about the error rate for each SLO metric.

Note: For each service-level objective metric on the serviceability dashboard, the first data point, which is the metric data of the first incoming call, is not counted or shown on the dashboard because the underlying Prometheus Query Language (PromQL) query, `increase(... [$__range])`, calculates the increase in the time series in the range vector.

Bucket heat maps

Bucket heat maps indicate the distribution of duration of the response time in seconds of incoming HTTP requests.

The underlying query expression, which is written in Prometheus Query Language (PromQL), calculates the increase of HTTP requests that are served during the selected time range by histogram bucket. For more information, see [Histograms and summaries](#).

Note: The value of the label `le` denotes the inclusive upper bound of the bucket.

Query expression for Maximo Collaborate

```
sum(round(increase(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}[$__range]))) by (le)
```

Query expression for Maximo Health

```
sum(round(increase(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}[$__range]))) by (le)
```

Query expression for Maximo Optimizer

```
sum(round(increase(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}[$__range]))) by (le)
```

Query expression for Maximo Predict

```
sum(round(increase(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}[$__range]))) by (le)
```

Response time in last 30 mins

Response time in the last 30 minutes indicates the 90th, 95th, and 99th percentile of request durations over the last 30 minutes.

The underlying query expression calculates the 90th, 95th, and 99th percentile of request durations aggregated by the `le` label over the last 30 minutes by using the histogram quantile function. For more information, see [histogram_quantile](#).

Query expression for Maximo Collaborate

```
histogram_quantile(.9,
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}
[30m:])) by (le))

histogram_quantile(.95,
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}
[30m:])) by (le))

histogram_quantile(.99,
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}
[30m:])) by (le))
```

Query expression for Maximo Health

```
histogram_quantile(.9,
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}
[30m:])) by (le))

histogram_quantile(.95,
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\ "$instanceId\"}
[30m:])) by (le))
```



```

    histogram_quantile(.99,
    sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\"$instance
    ce
    Id\"}[30m:])) by (1e))

```

Query expression for Maximo Optimizer

```

    histogram_quantile(.9,
    sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId
    \"}
    [30m:])) by (1e))

    histogram_quantile(.95,
    sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId
    \"}
    [30m:])) by (1e))

    histogram_quantile(.99,
    sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId
    \"}
    [30m:])) by (1e))

```

Query expression for Maximo Predict

```

    histogram_quantile(.9,
    sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\"$insta
    nceId\"}[30m:])) by (1e))

    histogram_quantile(.95,
    sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\"$insta
    nceId\"}[30m:])) by (1e))

    histogram_quantile(.99,
    sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\"$insta
    nceId\"}[30m:])) by (1e))

```

Average response time

Average response time indicates the average response time in milliseconds of incoming requests.

The underlying query expression calculates the average response time of requests that are served during the selected time range.

Query expression for Maximo Collaborate

```

    sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_sum{masinstanceid=\"$instanceId
    \"}[
    $__range]))/
    sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=\"$i
    nstanceId\"}[
    $__range]))

```

Query expression for Maximo Health

```

    sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_sum{masinstanceid=\"$instanceId\
    \"}[
    $__range]))/
    sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid
    =\"$instanceId\"}[
    $__range]))

```

Query expression for Maximo Optimizer

```

    sum(rate(mas_optimizer_api_create_job_duration_seconds_sum{masinstanceid=\"$instanceId\"}[
    $__
    range]))/sum(rate(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=\"$inst
    anceId\"}[
    $__range]))

```

Query expression for Maximo Predict

```

    sum(rate(mas_predict_asset_predictions_load_duration_seconds_sum{masinstanceid=\"$instanceId\"
    }
    [
    $__range]))/
    sum(rate(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid=\"$instanceI
    d\"}
    [
    $__range]))

```

% of requests which broke SLO

% of requests which broke SLO indicates the percentage of request durations that exceed the threshold in the SLO metric definition, for example, 4 seconds.

The underlying query expression calculates the percent of request durations that exceed the threshold in the SLO metric definition during the selected time range.

Query expression for Maximo Collaborate

```
(1-  
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"10\"}[$__range]))/  
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range]))) * 100
```

Query expression for Maximo Health

```
(1-  
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"4\"}[$__range]))/  
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range]))) * 100
```

Query expression for Maximo Optimizer

```
(1-  
sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"4\"}[$__range]))/  
sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range]))) * 100
```

Query expression for Maximo Predict

```
(1-  
sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"5\"}[$__range]))/  
sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range]))) * 100
```

% of requests which met SLO

% of requests which met SLO indicates the percentage of request durations that met the threshold in the SLO metric definition, for example, 4 seconds.

The underlying query expression calculates the percent of request durations that met the threshold in the SLO metric definition during the selected time range.

Query expression for Maximo Collaborate

```
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"10\"}[$__range]))/  
sum(rate(mas_collaborate_query_querydocsapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range])) * 100
```

Query expression for Maximo Health

```
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"4\"}[$__range]))/  
sum(rate(mas_health_matrix_matrixcountapi_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"+Inf\"}[$__range])) * 100
```

Query expression for Maximo Optimizer

```
sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\"$instanceId\",le=\"4\"}[$__range]))
```

```
sum(rate(mas_optimizer_api_create_job_duration_seconds_bucket{masinstanceid=\ "$instanceId\",le="+Inf"}[$__range]))*100
```

Query expression for Maximo Predict

```
sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid="\ $instanceId",le="5"}[$__range]))/sum(rate(mas_predict_asset_predictions_load_duration_seconds_bucket{masinstanceid="\ $instanceId",le="+Inf"}[$__range]))*100
```

Error rate

Error rate indicates the percentage of erroneous calls, for example, calls with HTTP status code in the 500 range.

The underlying query expression calculates the percent of requests that are served during the selected time range that return a status code that starts with 5, for example, 500, 502, 503.

Query expression for Maximo Collaborate

```
sum(increase(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=\ "$instanceId",code=~"5.."}[$__range]))/sum(increase(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range]))*100
```

Query expression for Maximo Health

```
sum(increase(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid=\ "$instanceId",code=~"5.."}[$__range]))/sum(increase(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range]))*100
```

Query expression for Maximo Optimizer

```
sum(increase(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=\ "$instanceId",code=~"5.."}[$__range]))/sum(increase(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range]))*100
```

Query expression for Maximo Predict

```
sum(increase(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid=\ "$instanceId",code=~"5.."}[$__range]))/sum(increase(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid="\ $instanceId"}[$__range]))*100
```

Calls

Calls indicates the ratio of calls with different HTTP status codes, for example, 200, 400, 401, 403, 500, and so on.

The underlying query expression calculates the count of requests aggregated by the status code.

Query expression for Maximo Collaborate

```
sum(round(increase(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range])) by (code)
```

Query expression for Maximo Health

```
sum(round(increase(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range])) by (code)
```

Query expression for Maximo Optimizer

```
sum(round(increase(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=\ "$instanceId"}[$__range])) by (code)
```

Query expression for Maximo Predict

```
sum(round(increase(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid=$instanceId}{$__range}))) by (code)
```

Error rate in last 30 mins

Error rate in last 30 minutes indicates the percentage of erroneous calls over the last 30 minutes.

The underlying query expression calculates the percent of requests that are served over the last 30 minutes that return the status code that starts with 5, for example, 500, 502, 503.

Query expression for Maximo Collaborate

```
(sum(increase(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=$instanceId},code=~"5.."[30m])))/  
(sum(increase(mas_collaborate_query_querydocsapi_duration_seconds_count{masinstanceid=$instanceId}{30m}))*100
```

Query expression for Maximo Health

```
(sum(increase(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid=$instanceId},code=~"5.."[30m])))/  
(sum(increase(mas_health_matrix_matrixcountapi_duration_seconds_count{masinstanceid=$instanceId}{30m}))*100
```

Query expression for Maximo Optimizer

```
(sum(increase(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=$instanceId},code=~"5.."[30m])))/  
(sum(increase(mas_optimizer_api_create_job_duration_seconds_count{masinstanceid=$instanceId}{30m}))*100
```

Query expression for Maximo Predict

```
(sum(increase(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid=$instanceId},code=~"5.."[30m])))/  
(sum(increase(mas_predict_asset_predictions_load_duration_seconds_count{masinstanceid=$instanceId}{30m}))*100
```

Monitoring Maximo Manage data

IBM Maximo Manage includes a monitoring agent that collects statistics about application usage and performance. You can use Prometheus and Grafana to view usage data that the agent collects.

For more information, see [Monitoring agent for Maximo Manage](#).

Developing

You can develop new Maximo Application Suite applications and modify existing applications.

Extending applications

Extensibility is the capability of extending the base functions of an application to meet business-specific needs.

Extensibility is built into the design of Maximo Application Suite. Maximo Application Suite includes powerful and flexible tools to create and customize applications to suit your business needs. You can:

- Add or modify applications, business objects, and schemas.
- Add or modify business rules and processes.
- Inject new analytic rules for anomaly detection.

- Add formulas to calculate the health and criticality of resources.
- Integrate with external services.
- Orchestrate new business flows by combining internal and external services

Extensibility overview

Extensibility enables you to adapt an application for a specific business need.

Maximo Application Suite extensibility is achieved by extending Maximo Application Suite components, documenting adaptations, and educating users. Extensions can be packaged and offered in the Red Hat Marketplace or within Maximo Application Suite itself using the tools and applications that are provided.

For more information, see [“Maximo Application Suite accelerators” on page 76](#).

Extensions

Logic and trigger points define a Maximo Application Suite application extension.

Logic forms the core of the extension and defines the value that it brings to an application.

Trigger points define how logic is initiated. Logic can be initiated directly through a client or indirectly through another application component. Logic can be initiated synchronously or asynchronously, on-demand, or on a schedule.

Extensions are often expressed as metadata that is stored in a persistent repository that is controlled by the application that hosts it.

Roles

Administrators have access to the applications and APIs that enable them to deploy and activate an extension. They can develop, configure, and manage extensions. Extensions can be managed through a configuration application or a configuration API. Administrators are also responsible for packaging and distributing extensions.

Users directly or indirectly use an extension to achieve a business result. Users can directly start a service of an extension. Users can also indirectly start a service of an extension when they use a default service that itself relies on an extension. Users are not typically involved with the development or configuration of extensions.

Extension types

Each Maximo Application Suite component includes extension points. These extension points can be broadly classified into certain extension types.

Business object schema extensions

Business object schema extensions extend the base application object schema to store more business data. This extension type can modify business objects by adding a property or modifying the metadata of an existing property. All modifications must be compatible with existing base application function. Business object schema extensions can add new objects to the schema. New objects might be related to a base object. It can be used to store relevant business data for business practices that are related to the business domain of the application.

Business rule extensions

Business rule extensions can dictate business object or application behaviors.

A business rule extension might affect a business object property by applying logic to a business object properties lifecycle event. These rules can impact the data and metadata of a property.

A business rule extension might affect business object rules that apply to an object's lifecycle events.

A business rule extension might affect a business process by defining new business processes that can span multiple business objects.

Business rule extensions can be developed to configure scheduled jobs to perform actions on behalf of a business.

Business rule extensions can be developed to configure audits and electronic signatures for compliance requirements.

Integration extensions

Integration extensions use application APIs to extend object schemas and object rules. Any extensions that are developed need to be accessible through APIs.

User interface extensions

User interface extensions extend application UI to show new data that is introduced through the extension of business schema and business rules extensibility. User interface extensions are also used to display base object information in a different view.

User interface extensions can update an application UI to display newly added business objects. They can also be used to update an application UI to accommodate a new business process. User interface extensions are commonly used to modify an application UI by displaying data in different formats or by adding quick actions.

Reporting extensions

Reporting extensions are used to create custom reports for the business data that is generated by an application. Reports include data that is relevant to a business process and define how that data is presented and formatted.

Characteristics of extensions

Maximo Application Suite extensions share a common set of characteristics. However, Maximo Application Suite extensions are not required to include every characteristic listed.

Hot deployment support

Deploying an extension must not require any full-service downtime.

Configurable

Extension definitions must be configurable through APIs and applications. Extensions must be flexible enough to mature with the evolving needs of a business.

Ideally, an extension is designed to maintain its definition in a metadata format. Metadata is the perfect solution for data storage, portability, and updatability without requiring any downtime. All applications must allow for the configuration of extensions through user interfaces for administrative users and APIs for external tools.

Pluggable

Applications must allow the activation or deactivation of an extension component without requiring a restart of the Docker container.

Security

Applications need to provide as much security for extensions as is provided for default application features and functions. Security must be applied to the configuration user interface (UI), API, and runtime

of the extension. For security reasons, the configuration UI and API of the extension must be accessible only to the administrator role.

Upgradeable

Application upgrades must not negatively impact the configuration and behavior of an extension.

Discoverable

Extensions must be distinguishable from base application features and functions. Applications must provide tools to discover extensions. Having the capacity to recognize extensions helps you determine the extent of customization for an application. Extensions can be activated and deactivated for troubleshooting an application.

Version control

Extensions must maintain release information and communicate the version of the current installation.

Distribution

Extensions must be packaged so that they can be offered in a marketplace. An extension must provide a command-line interface (CLI) or graphical method to package, distribute, and import itself among environments.

Documentation

Extensions must be documented extensively, including the following documentation:

- A feature or reference document
- Tutorials
- Best practice information
- Educational courses

Extension documentation must focus on both experienced and inexperienced audiences.

Serviceability

Extensions must be serviceable by supporting the following items:

- Artifact logging capability
- Monitoring
- MustGather tools and documentation
- Tools for testing the extension

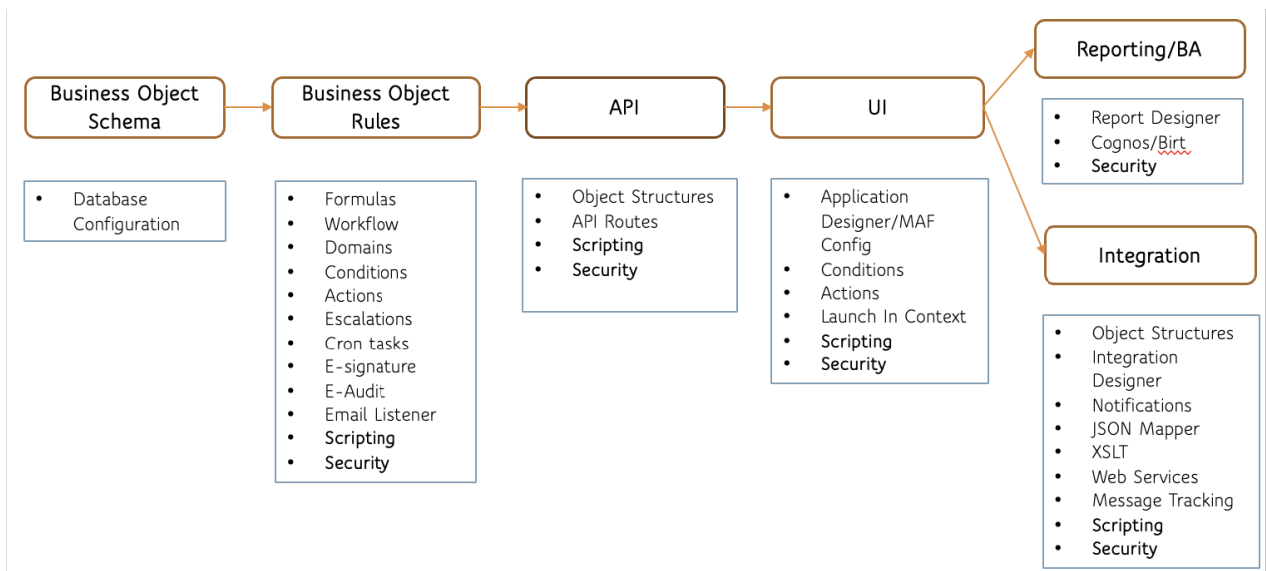
Extending Maximo Application Suite applications

Every Maximo Application Suite application can be extended to accommodate unique business needs.

Each Maximo Application Suite application can be customized and configured through various extension points.

Maximo Manage extensibility

Maximo Manage includes extension points for most of its modules and services.



Extending Maximo Manage can require you to work across several extension types. As an example, supplementing asset data with a new data type can require the following steps:

1. Create a business object to store the additional data, for example, ASSETVENDORINFO. Use the Database Configuration application to extend the database schema to include the new business object and its object attributes, attribute types, and attribute constraints. Include a list of supported values, and designate them as required or optional, and define relationships to other objects, for example, ASSET or LOCATIONS.
2. Use scripting to define the behavior of the business object through business rules on events, such as on add, update, and delete. You can also apply business rules at the attribute level by configuring conditional constraints, conditional defaults, conditional actions, and conditional validations to perform based on user input. As an example, you can define a business rule that the REFERREDBY attribute becomes required when the RATING attribute value is less than 50.
3. Use the Object Structures application to create one or more APIs for the business object. Include it as part of an existing object structure or create one. Object structures are sets of one or more related business objects. An API allows a client to perform create, retrieve, update, and delete operations on an object structure in the scope of an atomic transaction. Set the Role Based Access Control and Row Level Access Control rules for the API by using the Security configuration application.
4. Use the Application Designer application to create a UI for users to interact with the business object. As an example, you can modify an existing Maximo Manage UI to add a menu option, or you can create a completely new UI.
5. Use the Report Administration application to create reports for the new data type. Maximo Manage includes an embedded BIRT-based reporting engine that integrates with Cognos Analytics. You can import custom BIRT report designs to support ad-hoc reporting.
6. Use the integration applications to integrate with other systems. For example, in the Enterprise Services application, you can configure an inbound queue with the Maximo Manage enterprise service so that the business object receives data from other systems. You can configure an outbound queue with the Maximo Manage publish channel to send data from the business object to an external system. You might need to transform data to a format that is compatible with integration exits, XSL, or a JSON Mapper. You can configure protocol mediation and endpoint handlers, and you can store and forward messages by using JMS and Kafka.

Business object schema extensibility resources

Resources are available to learn more about extending the extension points for business object schemas in Maximo Manage.

Based on your business needs, you might want to create a schema extension. For example, a schema extension could be used to extend Maximo Manage objects to store additional data. You can also modify an attribute's data type or length to suit your business data requirements. Use the Maximo Manage Database Configuration application to draft and apply your schema changes.

Database configuration

The Maximo Manage database offers multiple extension points that can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Configuring databases](#)
- [Asset & Facilities Management community database](#)

Business object rule extensibility resources

Resources are available to learn more about extending the extension points for business object rules in Maximo Manage.

You can use Maximo Manage business rule tools to customize the behavior of business objects, attributes, and processes. You might want to modify the business rules for an existing object or attribute to suit your business needs. For example, you could mark an item as required or read only by using the Database Configuration application. You could create complex conditional logic using the Scripting application. You could introduce a new business process for approvals by leveraging workflows.

Database configuration

The Maximo Manage database offers multiple extensibility points that can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Configuring databases](#)
- [Asset & Facilities Management community database](#)

.

Formulas

Maximo Manage formula calculations can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Configuring databases](#)
- [Maximo Formula Documentation](#)
- [Asset & Facilities Management community forumlas](#)

.

Workflows

Maximo Manage workflows can be customized and configured to meet your business needs.

For more information, see [Implementing workflow processes](#).

Domains

Maximo Manage domains can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Managing domains](#)
- [Associating domain values with conditions](#)

Conditions

Maximo Manage conditions can be customized and configured to meet your business needs.

For more information, see [Conditional Expression Manager](#).

Actions

Maximo Manage actions can be customized and configured to meet your business needs.

For more information, see [Action types](#).

Escalations

Maximo Manage escalations can be customized and configured to meet your business needs.

For more information, see [Managing escalations](#).

Communication templates

Maximo Manage communication templates can be customized and configured to meet your business needs.

For more information, see [Managing communication templates](#).

Cron tasks

Maximo Manage cron tasks can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Managing cron tasks](#)
- [Cron task loggers](#)

.

E-signatures

Maximo Manage e-signatures can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Specifying the electronic signature key](#)
- [Electronic signature properties](#)

.

E-audit

Maximo Manage e-audits can be customized and configured to meet your business needs.

For more information, see [Electronic audit records](#).

Email listener

Maximo Manage email listeners can be customized and configured to meet your business needs.

More information is available in the following resources:

- [E-mail Listeners](#)
- [Email listener properties](#)

.

Scripting

Maximo Manage scripting can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Automating routine tasks with automation scripts](#)
- [Automation scripts loggers](#)
- [Maximo Autoscripting](#)
- [Asset & Facilities Management community scripting](#)

.

Security

Maximo Manage security can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Securing](#)
- [Security settings](#)

API extensibility resources

Resources are available to learn more about extending the extension points for APIs in Maximo Manage.

APIs are used to integrate and extract data, develop role-based applications, and develop custom 3rd part applications. APIs form the basis of stateless interaction with Maximo Manage business objects. Anytime a new business object type is created, a business object attribute is added or modified, or new business logic is introduced, APIs must be updated. APIs can be extended using the Maximo Manage Object Structures and Scripting apps. They can also be used to create entirely new APIs.

Object structures

Maximo Manage object structures offer multiple extensibility points that can be customized and configured to meet your business needs.

For more information, see [Object structures](#).

API routes

The Maximo Manage API routes offer multiple extensibility points that can be customized and configured to meet your business needs.

More information is available in the following resources:

- [IBM Maximo REST API Guide](#)
- [Sample REST API requests for browsing Kafka queues](#)[Asset & Facilities Management community integration](#)

Scripting

Maximo Manage scripting can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Automating routine tasks with automation scripts](#)
- [Automation scripts loggers](#)
- [Maximo Autoscripting](#)
- [Asset & Facilities Management community scripting](#)

Security

Maximo Manage security can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Securing](#)
- [Security settings](#)

User interface extensibility resources

Resources are available to learn more about extending the extension points for user interfaces in Maximo Manage.

User interface extensions can be used to create a user interface for a new application, or modify an existing one. User interface extensions often associated with business object extensions, business rules extensions, API extensions, and in certain cases, Integration extensions. Modifying a user interface can take many forms:

- Change layouts by using the Application Designer.
- Add new sections to an application or highlight an attribute of a custom business object by using the Application Designer.
- Implement conditional user interface elements by using the Application Designer, configuring sigoptions in the Security application. Add further logic with the Condition Expression Builder and the Scripting application.
- Add new actions to the user interface by using the Application Designer, configuring sigoptions in the Security application, and the Scripting application.
- Set access rights to an application or specific areas of an application by using the Security application.

In addition to the Maximo Manage user interface, role-based user interfaces can also be extended by using the appropriate application designer.

Application Designer

The Application Designer in Maximo Manage offers an environment to customize and configure applications to meet your business needs.

More information is available in the following resources:

- [Application development](#)
- [Asset & Facilities Management community application designer](#)

Maximo Application Framework Configuration application

The Maximo Application Framework Configuration application offers a lightweight tool to customize and configure applications to meet your business needs.

More information is available in the following resources:

- [Developing applications with the Maximo Application Framework Configuration application](#)
- [Asset & Facilities Management community application MAF Configuration application](#)

Conditions

Maximo Manage conditions can be customized and configured to meet your business needs.

For more information, see [Conditional Expression Manager](#).

Actions

Maximo Manage actions can be customized and configured to meet your business needs.

For more information, see [Action types](#).

Launch in context

Maximo Manage launch in context functions can be customized and configured to meet your business needs.

For more information, see [Launch in Context](#).

Scripting

Maximo Manage scripting can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Automating routine tasks with automation scripts](#)
- [Automation scripts loggers](#)
- [Maximo Autoscripting](#)
- [Asset & Facilities Management community scripting](#)

.

Security

Maximo Manage security can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Securing](#)
- [Security settings](#)

.

Reporting extensibility resources

Resources are available to learn more about extending reports in Maximo Manage.

Maximo Manage includes an embedded BIRT reporting engine and integrates with business intelligence tools like Cognos. You can use existing Maximo Manage BIRT reports to generate and review ad-hoc business data. You can also use the stand alone Report Designer tool to create your own reports and then import them into Maximo Manage by using the Report Administration app.

Use the KPI Administration app to configure KPIs to display performance information in the role-based dashboards or start centers.

Report Designer

Maximo Manage reports can be customized and configured in the Report Designer to meet your business needs.

More information is available in the following resources:

- [Report Administration](#)
- [Configuring the BIRT report development environment](#)

Cognos Analytics

Maximo Manage reports can be integrated with Cognos Analytics to generate analytics reports.

For more information, see [Integrating with Cognos Analytics server](#).

Security

Maximo Manage security can be customized and configured to meet your business needs. More information is available in the following resources:

- [Securing](#)
- [Security settings](#)

Integration extensibility resources

Resources are available to learn more about extending integration points in Maximo Manage.

As an enterprise asset management system, Maximo Manage integrates with complementary solutions for purchasing, inventory management, and supply chain operations. You can customize Maximo Manage integrations or use the Maximo integration framework to further extend the integration capabilities of Maximo Manage. For example, you can use the Maximo integration framework to load initial data, extract that data, and then integrate with an external weather service.

The Maximo integration framework provides a comprehensive set of extension tools that can be used in many ways:

- Define the data set to import and export by using the Object Structures application and scripting capabilities.
- Define ways to publish data by using data transformation and protocol mediation to external services by using the Publish Channel application, endpoints, scripting, JSON mapping, and XSL.
- Define ways to interface synchronously with data transformation and protocol mediation to external services and return a response by using invoke channels, endpoints, scripting, JSON mapping, and XSL.
- Define Maximo Manage business objects as services with data transformation capabilities by using enterprise services, scripting, JSON mapping, and XSL.
- Define security policies for services and interactions by using the Security Groups application.

Object structures

Maximo Manage object structures offer multiple extensibility points that can be customized and configured to meet your business needs.

For more information, see [Object structures](#).

Integration tools

The Maximo Manage integration tools offer an environment to customize and configure integration points to meet your business needs.

More information is available in the following resources:

- [Integrating data with external applications](#)
- [Integration module](#)

Notifications

Maximo Manage notifications can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Managing notifications](#)
- [Push notifications](#)

Scripting

Maximo Manage scripting can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Automating routine tasks with automation scripts](#)
- [Automation scripts loggers](#)
- [Maximo Autoscripting](#)
- [Asset & Facilities Management community scripting](#)

Security

Maximo Manage security can be customized and configured to meet your business needs.

More information is available in the following resources:

- [Securing](#)
- [Security settings](#)

Application Configuration

Application Configuration is used to customize and configure Maximo Application Suite applications.

Application Configuration overview

Maximo Application Suite customers who deployed and activated Maximo Manage can download Application Configuration from an unentitled registry.

Application Configuration supports IBM Maximo Mobile, IBM Maximo Health, and IBM Maximo Manage. Additionally, you can also configure non-mobile, role-based applications.

Check out the video for a tour of Application Configuration:

Application Configuration v9.0 and earlier works through Docker container technology and Maximo Application Framework tools. Maximo Application Framework tools are used to create a configuration environment within a Docker image on your local system. Application Configuration v9.1 is deployed directly in the Maximo Application Suite user interface. You can download, modify, preview, and publish your configured IBM Maximo Manage applications with minimal setup.

Application list

The **Application list** page displays the applications that you are authorized to run. Each application listed includes description information, the last time it was modified, and the current revision. Each application listed has an edit menu that you can use to duplicate the application, re-download it from the server, delete it, or view the history of each past revision of the application. Click the name of an application to open it in the **Editor** page.

XML editor

The application XML code is available for editing in the XML editor. When you select a component in the XML editor, it renders in the **Canvas**.

If your changes require custom JavaScript to implement, add code to the `AppCustomizations.js` file, and then reference it in the application XML code.

Canvas

You can use the Canvas area to check the appearance of the component in the application. You can also modify XML code by altering the rendered components in the Canvas.

Navigator

Use the Navigator tree to locate specific components in the XML code. Select an entry to locate the component in the application XML file. For example, select the **pages** component to show all the pages that are defined in the application.

Component search

You can use the component search field to filter entries in the Navigator. For example, if you search for `dialog`, the Navigator displays only dialog components that are found in application pages. Alternatively, you can also search for a component ID. By default, searches return results that are like the term you entered. To search for an exact term, add an equals sign before the search entry. For example, `=id771`.

Properties editor

As an alternative to the XML editor, you can use the Properties editor to update component values in the XML file. Select a component from the Navigator or in the XML editor to view the properties for that component. Required properties are marked with an asterisk.

Property names must be declared using lowercase letters, for example, `item.somename`. Properties must always be declared and used in lowercase because the Maximo application server provides the data field in JSON as lowercase.

Application Configuration provides a lightweight tool that you can use to edit applications built using the Maximo Application Framework. This includes many of the Maximo Application Suite applications. Any configuration changes made to applications based on the previous Maximo Asset Management framework cannot be migrated to Maximo Application Framework applications. Any configuration changes made using a product such as Maximo Anywhere must be re-implemented in a Maximo Application Framework application.

For additional configuration, customization, and upgrade discussion and information, including examples, refer to the [IBM Asset & Facilities Management Community](#).

Getting started for Application Configuration users

You can download Application Configuration versions 9.0 and earlier from an entitled or non-entitled IBM registry. Starting with Maximo Application Suite 9.1, Application Configuration is integrated into the Maximo Application Suite user interface. When you open the application for the first time, you must enter the server address for Maximo Application Suite, not the server address for Maximo Manage. In Application Configuration, you can configure or customize applications, such as adding a query to a service request, and adding custom fields to a work order page.

Authorizing with the Maximo Manage server

Grant Application Configuration v9.0 and earlier access to the app data object structure.

1. Log in to Maximo Manage
2. Open the Object Structures application and select the OSLCMAFAPPDATA object structure.
3. Select **Configure Object Structure Security** and then select **Use Object Structure for Authorization Name**.
4. Open the Security Group application and choose a security group or create a new group.
5. Grant Application Configuration user group access to the OSLCMAFAPPDATA object structure.
6. Assign the user group to the user you want to authorize to use Application Configuration..

Repeat this process for the MXAPIWPEDITSETTING object structure.

If you add an object structure in a Maximo data source to an existing or duplicated application, you need to grant access to that object structure.

Connecting to the Maximo Manage server

You can connect to the Maximo Manage server from Application Configuration by using an API key, Maximo Application Suite credentials, or a username and password for Maximo Asset Management.

Procedure

1. Open a browser and open `http://localhost:3001`.

2. On the Application Configuration login page, select an authentication method.

Option	Description
Log in using an API Key	<p>Using an API key is the default authentication method. Enter the Maximo Asset Management or Maximo Manage server URL, for example, <code>https://host:port/maximo</code> and the API key value that is assigned to the user.</p> <p>An API key can be obtained from your system administrator, the API Keys application, or by using the REST API to generate one.</p> <p>To create an API key by using the REST API, log in and issue the following REST API POST command:</p> <pre data-bbox="521 489 1472 531">POST /oslc/apitoken/create { "expiration":-1 }</pre> <p>The JSON response includes an API key value for the logged-in user. You can use the same REST API POST command to regenerate the API key.</p> <p>For more information, see REST API guide.</p>
Log in using Maximo Application Suite credentials	<p>To log in using the Maximo Application Suite credentials, the Maximo Application Suite server must be configured to accept an OIDC log in. For more information, see https://www.ibm.com/docs/en/mas-cd/continuous-delivery?topic=authentication-methods.</p> <p>Enter the URL for the Maximo Manage server, for example, <code>https://Maximo Application Suite host/maximo</code>, and then select Login. You are redirected to the Maximo Application Suite log in page. After you log in to Maximo Application Suite, you are redirected back to the Maximo Application Framework Configuration application.</p> <p>After you log in, the application list includes the Category column, which indicates whether the application is stored in Maximo Application Suite core.</p>
Log in using Maximo Asset Management username and password	<p>Enter the URL for the Maximo Asset Management server, for example, <code>https://host:port/maximo</code>, and then enter your username and password.</p>

Related reference

“Troubleshooting Application Configuration login issues” on page 858

Several scenarios might prevent you from logging in to a server from the Maximo Application Framework Configuration application.

Troubleshooting Application Configuration login issues

Several scenarios might prevent you from logging in to a server from the Maximo Application Framework Configuration application.

REST API

Application Configuration uses the REST API to get and post data. Ensure that the REST API is available by making a simple REST API request. Maximo Manage Work Centers also use the REST API. If the Work Centers are accessible, that is a good indication that the REST API is available.

File system permissions

When you log in to a server from Application Configuration, it creates a `credentials.json` file in the workspace directory. When you open an application file for editing, that JSON file is copied to the application directory. If the Application Configuration container does not have write permission to the workspace directory, an error occurs.

```
[Error: EACCES: permission denied, open '../..../workspace/credentials.json']
```

Grant all users write permission to workspace folders.

Misconfigured web.xml

Log in to a server from Application Configuration, open another browser tab, and then verify that a `/maximo/api/whoami` REST API request returns a valid result. If a `Servlet not found` error is returned, then the `web.xml` file might not be properly configured. A `<servlet-mapping>` entry is needed for the API URL pattern.

```
<servlet-mapping>
  <servlet-name>OSLCServlet</servlet-name>
  <url-pattern>/api/*</url-pattern>
</servlet-mapping>
```

Modifying applications

After you log in to Application Configuration, you can create or modify applications.

About this task

You can modify the code of an application by using the XML editor, properties panel, and the canvas area.

You can make changes to an application by adding, deleting, or modifying XML elements directly in the XML editor. Alternatively, you can select an XML component in the XML editor, and add, remove, and modify component property values in the properties panel.

Besides editing app XML in the XML editor, you can select the `AppCustomizations.js` file and add custom JavaScript to the application. In addition, you can select the `i18n/labels.json` file to update labels shown in the application.

You can also modify the appearance of the application UI through the canvas area. Preview an application XML component in the canvas, for example, an application page, and then drag and drop elements to rearrange the display. You can also delete or copy and paste elements in the UI. The XML code displayed in the XML editor updates as you change the appearance of an element from the canvas area.

Procedure

1. If Maximo Application Suite v9.0 or earlier is installed, run the Application Configuration Docker image. If Maximo Application Suite v9.1 is installed, access Application Configuration through the Maximo Application Suite UI.
2. Log in to Maximo Manage.
The **Application list** page displays the applications that you are authorized to configure.
3. Click the name of an application to open it in the **Editor** page.
4. Select `app.xml` from the file section drop-down menu.
5. Make a change to the XML of the application.
For example, you can add an owner group reference on a work order.
Do not change component ID values in the application XML files of default Maximo Mobile applications. Before an upgrade, if you changed default Maximo Mobile applications, you can generate a delta file that contains your configuration changes. The delta file relies on the component IDs of the default applications to identify changes.
6. Click **Save**.
7. Select **Preview** and verify that the change you made works as intended.
8. Select **Publish**.

Do not publish an application immediately after you start editing the application. When you start editing an application, the local React preview server starts and can take a few minutes to complete the startup process. Errors can occur if you publish at the same time that the React server is starting. View the console for the message that the preview server is available.

After an application is published, mobile devices can then download the new version of the application. If the application is being published for the first time, assign it to the appropriate security groups. If you edit application files outside of Application Configuration, then you must ensure the integrity of those files. Use a virus or malware scanner before you publish the files to the Maximo application server.

Related information

[“Running the Docker image” on page 878](#)

After the image download is completed, you can run the Docker image to start the Maximo Application Framework Configuration application.

Drag and drop UI elements

A subset of application UI elements can be dragged and dropped on the canvas to configure the UI.

Elements that can be dragged and dropped

These elements can be dragged to different spaces on the canvas and then dropped to reposition them in the UI.

- adaptive-row
- box
- button-group
- button
- data-list
- field
- label
- list-item
- navigator
- panel
- rich-text-viewer
- smart-input
- text-input
- wrapped-text

Elements that can contain other elements

Other elements can be dropped into these elements on the canvas.

- box
- button-group
- data-list
- field
- navigator
- panel

Application revisions

You can view the revision history of an application from Application Configuration.

You can view revision details for an application in the Application list. Click the Application's three dot menu and select **History**.

The revision history includes read-only information about any updates made to that application. Details include the publish date, the user that published it, the version of the application, and the status of each

revision. A new revision is created whenever an application is published or when a new version of Maximo Asset Management or Maximo Manage is deployed or patched.

Application preview

You can preview changes that you make in an application before you publish it.

When you change the XML of an application, you can preview your changes in two ways.

Canvas preview

After you change a component of the application in the XML editor, you can preview how it renders in the **Canvas** panel. The component is not functional in this preview, but it can help you evaluate its design. Not all application components are available to preview. Components such as pages and menus can be previewed, while datasources cannot. By default, when you select a component, the entire page that contains the component is displayed in the **Canvas** panel. You can choose to display only the component by selecting the option from the **Canvas** panel. You can also enable the ability to highlight the XML code for a component in the XML editor when you hover over a component in the canvas.

When you click components displayed in the **Canvas** panel, the XML code for that component is highlighted in the XML editor. The properties for the component are displayed in the **Properties** panel. The **Navigator** also opens to the entry for that component.

Application preview

You can preview the entire application using the **Preview** button. When you preview the application, a React server starts and the application is loaded in a new browser tab. The application is functional, so you can test changes that you make before you publish the application.

Duplicating applications

If you don't want to modify a default application, you can duplicate it, give it a new name, and then change the new application.

About this task

When you duplicate an application, it inherits the security settings from the original application. You can change security settings for an application with the Security Groups application. For more information, see [Creating security groups](#).

Custom or cloned applications cannot be used to circumvent licensing terms. You cannot use custom applications to grant users functions or access beyond your license agreement.

Procedure

1. If Maximo Application Suite v9.0 or earlier is installed, run the Application Configuration Docker image. If Maximo Application Suite v9.1 is installed, access Application Configuration through the Maximo Application Suite UI.
2. Log in to Maximo Manage.
3. Highlight an application from the application list and then select **Duplicate** from the More actions menu.
4. Enter a unique name for the application and a description.
5. Select the option to enable mobile download if the application is mobile.
6. Click **Duplicate**.

Re-downloading applications

You can download the most recently published version of an application from the Maximo Manage server anytime by selecting **Redownload** from the More actions menu.

About this task

When you re-download an application, any unpublished updates that you made to an application are discarded.

Procedure

1. If Maximo Application Suite v9.0 or earlier is installed, run the Application Configuration Docker image. If Maximo Application Suite v9.1 is installed, access Application Configuration through the Maximo Application Suite UI.
2. Log in to Maximo Manage.
The **Application list** page displays the applications that you are authorized to configure.
3. Highlight an application from the application list and then select **Redownload** from the More actions menu.

Downloading and modifying application source

You can download a zip file of the entire set of source files for an application.

About this task

There might be times when your application customization includes the addition of a new file or a rewrite of an existing file. You can download all of the source files for an application, make modifications, and then re-upload the entire zip file, or individual files back in the Application Configuration editor.

Procedure

1. Log in to Maximo Application Suite and open Application Configuration.
2. Select an application from the application list.
3. From the drop down menu, select **Download app src**.
A zip file of the application's source files are downloaded to your local system.
4. Unzip the source files.
5. Make modifications.
For example, add a new JavaScript file.
6. From the drop down menu, select **Upload**.
7. Click **Add file**, select the file you want to upload, and then click **OK**.

Deleting applications

You can delete applications that you have created through the duplication process from the More actions menu.

About this task

Default applications cannot be deleted.

Procedure

1. If Maximo Application Suite v9.0 or earlier is installed, run the Application Configuration Docker image. If Maximo Application Suite v9.1 is installed, access Application Configuration through the Maximo Application Suite UI.
2. Log in to Maximo Manage.
The **Application list** page displays the applications that you are authorized to configure.
3. Highlight an application from the application list and then select **Delete** from the More actions menu.

Application upgrade using Application Configuration

After you configure an application, create an upgrade plan for fix packs and upgrades.

Process

Application Configuration allows you to configure a default Maximo Application Framework application, such as the Maximo Mobile Technician application, which is named TECHMOBILE. You can also duplicate a default application, name it something like MYTECHMOBILE, and then configure that duplicated application. When you upgrade Maximo Mobile or Maximo Manage, or apply a fix pack, the default applications are automatically updated to new versions.

If you configured a default application, it is archived during an upgrade or after you apply a fix pack, and the new application becomes the active version of the application. Alternatively, if you configure a duplicate application, the application is not impacted by an upgrade or fix pack. In either scenario, you can bring forward your configurations to the latest version or fix pack.

At a high level, application upgrade is a two step process. The first step is creating a delta file that contains the configuration changes that you implemented in a previous release. The second step is applying the changes contained in the delta file to latest version or fix pack of the application.

Environments

It is beneficial to complete an upgrade in a test environment before you upgrade other environments. Features and functions that are included with a new release might prevent you from applying your existing configuration values without some modification. When you plan for an upgrade, ideally you would have three environments:

- An environment that includes your existing application. For example, version 8.10.
- An environment with the new default version of the application. For example, version 8.11.
- An environment with the new version of the application that includes your configuration changes.

Having these three environments can help you diagnose any configuration issues that result from the upgrade. Any changes included in a new version of Maximo Manage, such as automation scripts or domain configurations, can impact your configured application. This can limit your ability to test the new version of the application.

Applications

Applications are stored as compressed files in the Maximo Manage database. The compressed file contains application XML files and related JavaScript files. When a new version of an application is deployed or if a fix pack is applied, a new row is created in the database table. That row contains the active version of the application. The previous application is disabled and is no longer available to use or configure.

Workspace files

When you open an application in Application Configuration, the application is downloaded to a local workspace directory on your system. The workspace directory name includes information about the server and credentials that are used to download the application. In the workspace directory there is a subdirectory that is named after the application. For example, if you downloaded the Technician application, the directory, TECHMOBILE, is created in your workspace.

For Maximo Application Suite 8.11.0, the standard XML files for base applications are stored in the / `ibm-config` subdirectory. These files are used to create the delta file for future upgrades from Maximo Application Suite 8.11.0.

Use different local workspace folders for different versions of Maximo Application Suite applications, including patch versions. Application files can be overwritten if you use the same workspace folder for multiple versions.

File management

You manage multiple versions of configured application files to facilitate the upgrade. You can track versions of application files by using a hierarchy of folders.

```
/810to811 (Main folder)
  /810ORIG (Contains the original version 8.10 files before configuration changes.)
  /TECHMOBILE (Folder for each application to upgrade. Only modified files listed.)
    app.xml
    wo-card-group.xml
  /810CONF (Contains the configured version 8.10 files.)
  /810DIFF (Contains the version 8.10 delta files.)
  /811ORIG (Contains the original versions of the version 8.11 files before configuration
changes.)
  /811ONF (Contains the upgraded files for version 8.11)
```

To upgrade an application from a previous release of Maximo Application Suite to the application in Maximo Application Suite 8.11, you need to create a delta file. The delta file identifies any configuration changes that were made to the application XML files. Move the original application files from a previous version of Maximo Application Suite into the ORIG folder. Place your configured application XML files in the CONF folder. You then can use the **diff** command to generate a delta file by comparing the two files. You need to repeat this process for every application XML file that you have configured.

To move changes made to default applications to the next version, you can use the **Upgrade** action in the Maximo Application Suite Configuration application. To create the diff file for a duplicated application, you must manually run the **diff** command to generate a delta file and then manually run the **apply** command to migrate your changes to the next version.

Upgrade example

This example demonstrates how to configure an application and then upgrade it to the next version of Maximo Mobile.

You can upgrade from one version of Maximo Mobile to another. For example, Maximo Mobile 8.10 to 8.11 or Maximo Mobile 8.11.0 and 8.11.1. You need separate local workspace folders for each version. If you don't use separate folders, you can lose the local version of an application when you connect to a server with a newer version of Maximo Mobile. Using the same local workspace for multiple versions can complicate the application upgrade process.

Configuring an application

In this example, you add a **Risk** field to the **Create work order** page of the Technician application.

Procedure

1. Open Application Configuration and log in to a Maximo Manage server.
2. Select **TECHMOBILE** from the list of applications.
3. From the XML editor, in the `app.xml` file, locate the **Create work** page.

```
<page id="createwo" title="Create work order" path="/createwo" comp-group-valid-
chg="{page.state.cgvc=event}"
controller="WorkOrderCreateController">
```

4. Add a **risk** attribute after the **wopriority** attribute for the **dsCreateWo** datasource.

```
<maximo-datasource id="dsCreateWo" object-structure="mxapiwodetail"
where="wonum='&quot;0&quot;";"
item-url="{app.device.isMaximoMobile ? 'nohref' : undefined}" auto-save="false">
  <schema id="w224k">
    <attribute name="workorderid" searchable="true" unique-id="true" id="bj6qp"/>
    <attribute name="wonum" searchable="true" id="n52n8"/>
    <attribute name="description" max-length="100" id="y7g99"/>
    <attribute name="description_longdescription" id="r8gpn"/>
    <attribute name="wopriority" id="ye34g"/>
    <attribute name="risk" id="re34g"/>
```


5. Add the **Risk** field and display it under the **Priority** field.

```
<box direction="row" children-sizes="100" fill-parent="true" fill-child="true" padding-top=".5" padding-bottom=".5" id="a5ezy">
<smart-input label="Priority" hide-step-buttons="true" placeholder="Enter {page.state.minPriority} to {page.state.maxPriority}" value="{dsCreateWo.item.wopriority}" on-blur="validateFields" min="{page.state.minPriority}" max="{page.state.maxPriority}" id="j8265" />
</box>
<box children-sizes="100" direction="row" fill-child="true" fill-parent="true" padding-bottom=".5" padding-top=".5">
<smart-input hide-step-buttons="true" label="Risk" value="{dsCreateWo.item.risk}" />
</box>
```

6. Save the application and then click **Preview**.
7. From the preview of the Technician application, open the **Create work order** page and confirm that the **Risk** field is displayed it under the **Priority** field.
8. Publish the configured application to the Maximo Manage server.

Results

After you publish the configured application, a copy of the updated `app.xml` file is written to the `TECHMOBILE/src` workspace folder. For Maximo Mobile 8.10 and earlier, a copy of the unmodified `app.xml` file is located in the `TECHMOBILE_backup` folder. For Maximo Mobile 8.11, a copy of the unmodified `app.xml` file is located in the `TECHMOBILE/ibm-config` folder.

What to do next

The next step is to create a delta file by using these two files.

Creating a delta file

Create a delta file that contains application configuration changes made to a default application.

Before you begin

For this example, the following folder structure exists under the `Users/myuser/Documents` file path.

```
/810to811
  /810ORIG
  /TECHMOBILE
  /810CONF
  /TECHMOBILE
  /810DIFF
  /TECHMOBILE
  /811ORIG
  /TECHMOBILE
  /811CONF
  /TECHMOBILE
```

The workspace folder that is used is `/aconfig810/MaximoUser-Maximohost:port/TECHMOBILE`.

About this task

In Maximo Mobile 8.11, Application Configuration automatically creates delta files, such as `app.delta.xml`. This automatically generated file is used for upgrading to future versions such as 8.11.1. or the next major release. To upgrade Maximo Mobile 8.10 and earlier releases, you must manually create the delta file.

Procedure

1. Copy the `/aconfig810/MaximoUser-Maximohost:port/TECHMOBILE/src/app.xml` file to the `/810to811/810CONF/TECHMOBILE` folder.
2. Copy the `TECHMOBILE` application `app.xml` file to the `810to811/810ORIG/TECHMOBILE` folder.
3. Generate the diff file.

```
docker run -it --privileged --entrypoint "/graphite/scripts/graphite-tools.js" --workdir /
graphite/app
-v /Users/myuser/Documents/810to811:/graphite/app cp.icr.io/cp/manage/maf-tools:8.11 diff --
original 810ORIG/TECHMOBILE/app.xml
--modified 810CONF/TECHMOBILE/app.xml --deltaOutput 810DIFF/TECHMOBILE/app.xml -d
```

```
docker run -it --privileged --entrypoint "\graphite\scripts\graphite-tools.js" --workdir
\graphite\app
-v c:\Users\myuser\Documents\810to811:\graphite\app cp.icr.io\cp\manage\maf-tools:8.11 diff
--original 810ORIG\TECHMOBILE\app.xml
--modified 810CONF\TECHMOBILE\app.xml --deltaOutput 810DIFF\TECHMOBILE\app.xml -d
```

The **-d** parameter is for debug mode. Debug mode produces the detailed output. Copy the diff command output from the console and save to a file. Refer to this output information to troubleshoot during testing.

The diff command generates a delta file with content similar to the following example.

```
<?xml version="1.0" encoding="UTF-8"?>
<updatescript version="">
  <app apptype="maf" id="TECHMOBILE">
    <add before="xd5k8" container="w224k" control="attribute">
      <set property="id" value="jyggb"/>
      <set property="name" value="risk"/>
    </add>
    <add before="dr9pg" container="gvvyv" control="box">
      <set property="children-sizes" value="100"/>
      <set property="direction" value="row"/>
      <set property="fill-child" value="true"/>
      <set property="fill-parent" value="true"/>
      <set property="id" value="ww34n"/>
      <set property="padding-bottom" value=".5"/>
      <set property="padding-top" value=".5"/>
    </add>
    <add container="ww34n" control="smart-input">
      <set property="hide-step-buttons" value="true"/>
      <set property="id" value="xb269"/>
      <set property="label" value="Risk"/>
      <set property="value" value="{dsCreateWo.item.risk}"/>
    </add>
  </app>
</updatescript>
```

Applying the delta file for Maximo Mobile 8.10 and earlier

Apply configuration changes that are recorded in the delta file to the new version of the application. In this example, you merge configuration changes that are made to version 8.10 of the Technician application to the newer 8.11 version of the application.

About this task

For Maximo Mobile 8.10 and earlier versions, the delta file is applied from the command line.

Procedure

1. Open Application Configuration and log in to the Maximo Manage server.
2. Select **TECHMOBILE** from the list of applications.
3. Copy the `/aconfig89/MaximoUser-Maximohost:port/TECHMOBILE/src/app.xml` file to the `/89to810/810ORIG/TECHMOBILE` folder.
4. Apply the delta file to create a version 8.10 Technician application that includes your version 8.9 configuration data.

```
docker run -it --privileged --entrypoint "/graphite/scripts/graphite-tools.js" --workdir /
graphite/app
-v /Users/myuser/Documents/89to810:/graphite/app cp.icr.io/cp/manage/maf-tools:8.10 apply
--delta 89DIFF/TECHMOBILE/app.xml --input 810ORIG/TECHMOBILE/app.xml --output 810CONF/
TECHMOBILE/app.xml -d
```

```
docker run -it --privileged --entrypoint "/graphite/scripts/graphite-tools.js" --workdir /
graphite/app
-v C:/Users/myuser/Documents/89to810:/graphite/app cp.icr.io/cp/manage/maf-tools:8.10 apply
--delta 89DIFF/TECHMOBILE/app.xml --input 8100RIG/TECHMOBILE/app.xml --output 810CONF/
TECHMOBILE/app.xml -d
```

Copy the apply command output from the console and save to a file. Refer to this output information to troubleshoot during testing.

The apply command merges configuration changes to the version 8.10 Technician application. The apply command also generates output statements based on the configuration changes.

```
--      ADD: label (emnww) to e_bm6 before b7w3z
        set id to emnww
        set label to Inspection Summary:
        set wrap to true

--      REMOVE: y4_eq from rzvz4
--      RemoveNode: Removing: y4_eq

--      MODIFY: aw99_
--      clear slot

-- Warning: (techmobile) Cannot MOVE control (brq_e) to CONTAINER (jjwgb). NOT_FOUND_ERR: An
attempt is made to
reference a node in a context where it does not exist.
.
.
.

--      MOVE: brq_e to CONTAINER (jjwgb) before bynje
```

The apply process attempts to resolve any warnings encountered earlier in the process. You might encounter an error that is related to a configuration change you made to an application component that was deprecated in the latest version of the application. In this case, you must reevaluate the change to determine whether it is still needed. If the configuration change is required, you must reimplement it using a different component in the `app.xml` file.

5. Copy the `app.xml` file to the `TECHMOBILE/src` folder of your local workspace. If your configuration introduced new XML files, they can be copied from the version 8.9 workspace folder. If you created JavaScript code, the `AppCustomizations.js` file can be copied to the version 8.10 `TECHMOBILE/src` folder.

Applying the delta file for Maximo Mobile 8.11

Apply configuration changes that are recorded in the delta file to the new version of the application. In this example, you merge configuration changes that are made to version 8.10 of the Technician application to the newer 8.11 version of the application.

About this task

For Maximo Mobile 8.11, you can apply the delta file using the Application Configuration **Upgrade** action. The **Upgrade** action upgrades all delta XML files under the `/src`, including subfolders. Use the **Upgrade** action to upgrade default applications only. Duplicated applications must be upgraded from the command line.

Procedure

1. Run the Docker image to start the Maximo Application Framework configuration application.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v maximo_8_11_workspace_directory:/
graphite/.workspace -it cp.icr.io/cp/manage/maf-tools:8.11
```

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v "C:/home/
```

```
core/maximo_workspace_directory:/graphite/.workspace" -it -e CHOKIDAR_USEPOLLING=true
cp.icr.io/cp/manage/maf-tools:8.11
```

2. Log in to the Maximo Manage server.
3. Select **TECHMOBILE** from the list of applications.
The **TECHMOBILE** application is downloaded to your Maximo Mobile 8.11 workspace folder.
4. Close the application by returning to the **Application List** window.
5. From your local system, copy the delta XML files from your previous version, including all subfolders content, to the `/src` folder of the Maximo Mobile 8.11 workspace. If the modified XML files are located in a subfolder, copy the delta file to that subfolder.
6. Rename the delta files.
Delta files must include `delta` in the file name.
For example, `app.xml` files must be rename `app.delta.xml`. The `wo-card-group.xml` file must be renamed `wo-card-group.delta.xml`.
7. Copy the `AppCustomizations.js` file and any new XML files to the `/src` folder of the Maximo Mobile 8.10 workspace.
8. Select **TECHMOBILE** from the list of applications.
9. From the **Action** menu, select **Upgrade**.

Results

After the upgrade process is complete, the file menu lists a log file for each upgraded XML file. Review the log files to confirm that your changes were applied to the Maximo Mobile 8.11 application.

Upgrading duplicated applications for Maximo Mobile 8.11

There are deviations and additional steps to the upgrade process if you are upgrading a duplicated application.

Before you begin

Ensure that you are familiar with the creation and application of delta files for default Maximo Mobile 8.10 and earlier applications.

About this task

The **Upgrade** action of Application Configuration is not compatible with duplicated applications. Duplicated applications must be upgraded from the command line. For example, when upgrading from Maximo Mobile 8.9 applications to version 8.10, or upgrading from Maximo Mobile 8.10 applications to version 8.11.

Procedure

1. Run the Docker image to start Application Configuration.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v /home/core/
maximo_8_10_workspace_directory:/graphite/.workspace -it cp.icr.io/cp/manage/maf-tools:8.11
```

If you run the Maximo Configuration v8.11 Docker image on a Windows system, you must include the **-e** parameter.

The following code is an example of the command with the parameter included.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v "C:/home/core/
maximo_8_10_workspace_directory:/graphite/.workspace" -it -e CHOKIDAR_USEPOLLING=true
cp.icr.io/cp/manage/maf-tools:8.11.0
```

2. Log in to the Maximo Manage server.
3. After you create the delta file, open the version 8.10 **TECHMOBILE** and the **TECHDUP** applications to download their files to your local workspace.

4. Copy the contents of the TECHMOBILE/src folder and overwrite the contents of the TECHDUP/src folder.
5. Open the app.xml file of the TECHDUP application in a text editor. Change the ID attribute of the maximo-application element from **techmobile** to **techdup**.

```
<maximo-application navigator-title-order="100" controller="AppController" theme="touch"
product-name="Maximo"
product-name-adaptive="Maximo" title="Maximo" id="techdup" version="8.10.0.0" default-log-
level="error"
user-menu-enabled="true" nav-initial-open-state="false" mas-enabled="false">
```

6. Copy the TECHDUP app.xml file to the 8100RIG folder.
7. Run the **apply** command.

[“Configuring an application” on page 864](#)

In this example, you add a **Risk** field to the **Create work order** page of the Technician application.

[“Creating a delta file” on page 865](#)

Create a delta file that contains application configuration changes made to a default application.

[“Applying the delta file for Maximo Mobile 8.10 and earlier” on page 866](#)

Apply configuration changes that are recorded in the delta file to the new version of the application.

In this example, you merge configuration changes that are made to version 8.10 of the Technician application to the newer 8.11 version of the application.

Upgrading duplicated applications for Maximo Mobile 9.0

There are unique steps to upgrade a duplicated application from Maximo Mobile 8.11 to Maximo Mobile 9.0.

About this task

The Application Configuration 9.0 **Upgrade** action is compatible with duplicated applications from Maximo Mobile 8.11.

The upgrade process requires the use of one or more delta files, for example, app.delta.xml, that contain configuration changes made to a duplicated Maximo Mobile 8.11 application. Starting with Maximo Mobile 8.11, the Maximo Application Framework configuration application automatically creates these delta files.

Procedure

1. Run the Docker image to start the Application Configuration.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v /home/core/
maximo_9_0_workspace_directory:/graphite/.workspace -it cp.icr.io/cp/manage/maf-tools:9.0
```

If you run the Maximo Configuration v9 or later Docker image on a Windows system, you must include the **-e** parameter.

The following code is an example of the command with the parameter included.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v "C:/home/core/
maximo_9_0_workspace_directory:/graphite/.workspace" -it -e WATCHPACK_POLLING=true
cp.icr.io/cp/manage/maf-tools:9.0
```

This command starts Application Configuration 9.0 and creates the workspace for Maximo Mobile 9.0 applications.

2. Log into Application Configuration and then download the version of the application that you duplicated in Maximo Mobile 8.11 from the application list.
For this example, the application duplicated from the Technician application is called TMT01.
The TMT01 application is downloaded into the local /home/core/maximo_9_0_workspace_directory directory.

3. From your local system, copy the `app.delta.xml` delta file from the Application Configuration 8.11 workspace to the 9.0 workspace.
For example, copy `/home/core/maximo_8_11_workspace_directory/TMT01/src/app.delta.xml` to `/home/core/maximo_9_0_workspace_directory/TMT01/src/app.delta.xml`.
4. Select the TMT01 application from the application list.
The TMT01 application should display a **Version** value similar to 8.11.0.
5. From the action menu, select **Upgrade**.
6. To apply changes from the delta file, click **Yes**.
7. Review and then close the **Upgrade results** window.
8. Click the file selection list and select the `app.xml.log` file to review the changes that were made to upgrade TMT01 to version 9.0.
9. Publish the TMT01 application to the server.

Upgrading Maximo Application Suite v9.0 applications to v9.1

Upgrading previously customized Maximo Application Suite applications to v9.1 using Application Configuration v9.1 is an automated process.

About this task

Customized application files are kept both locally on the computer where the customizations were implemented and in the database when it was published. Previous releases of Application Configuration relied on the local files to upgrade. Application Configuration v9.1 uses the zip file contained in the database.

Procedure

1. Log into Maximo Application Suite and open Application Configuration v9.1.
2. Open an application that was customized in a previous release. For example, open an application that you last customized in Maximo Application Suite v8.11.
3. From the drop-down menu, select **Migrate prior configurations**.
A dialog box appears and informs you that a newer version of Maximo Application Suite has been deployed. You are prompted to import the customizations you made in the Maximo Application Suite v8.11 application into the newer Maximo Application Suite deployment.
4. Click **Migrate prior configurations**.
Customizations are applied to the latest version of the application.

Application Configuration

Use Application Configuration to update applications that are built with the Maximo Application Framework.

Prerequisites

Get your environment ready to deploy Application Configuration v9.0 or earlier.

- A minimum of 8 GB RAM on the development system (16 GB recommended).
- Docker Desktop (<https://docs.docker.com/desktop/>) or Docker Engine (<https://docs.docker.com/engine/>) installed on your development system. Docker engine can be installed natively on Linux systems. For Windows, Docker must be installed on a Windows Subsystem for Linux (WSL) version 2 backend. For Mac OS, a Docker engine-supported operating system must be installed on VM. Ensure that your system meets all Docker requirements.

Set the Memory Resources preference in Docker Desktop to 8 GB of memory.

- A working directory on your local system. The working directory is used to run Docker commands to start Application Configuration. Make a note of the location of your working directory for future configuration efforts.
- Ensure that API Key is enabled on the Maximo Manage application server. You must configure the `mxe.oslc.webappurl` system property with the name of your Maximo Asset Management or Maximo Manage server.
- If you are using Windows, you must have the Windows Subsystem for Linux (WSL) version 2 installed on your system.
- To configure Application Configuration as an ODIC client for Maximo Application Suite, see [Manage OpenID client for Configuration Tool](#).

New applications

Duplicate a Maximo Application Framework application and modify it to create a new application you can configure to suit your business needs.

When you duplicate an application, a new application is created on the server and several tables are populated with application metadata to support the new application. This data includes the definition of the application, and it defines where the module is accessible in the Maximo Application Suite menus. Metadata also allows for the application to be configured for authorization by using the Security Groups application.

You cannot change the module of a duplicated application. You also cannot change the security template for the application. You might have to manually update the authorization set for object structures, and then grant authorization to each object structure individually.

Default applications support a mobile deployment. The application can run on a mobile device in connected or disconnected mode. The application can also be run in a browser in connected mode. Default role-based applications can run in a browser in connected mode, but they cannot be run on a mobile device. Duplicate an application that aligns with the deployment capabilities you are looking for in your new application.

Duplicate an application that has the features that you want in your new application, for example, a List page with a Details page. You can repurpose elements of the duplicated application.

If you intend to create a new work management application, you can duplicate the Maximo Mobile Technician application. Starting from the Technician application gives you insights into what data sources are available and it also provides examples of how to use those data sources in the new application.

If you plan to develop a new application that is not similar to any existing applications, you might want to duplicate a simple application, such as APIKEY, and then delete most of its contents before you start.

After you duplicate an application, you can edit the application XML code. You can add JavaScript code to the `AppCustomizations.js` file with Application Configuration. If your new application has significant functionality, you might have multiple JavaScript Controller files that are located under the application's `/src` directory. You can create and edit those files outside of the configuration application with your preferred editor. You still use Application Configuration to build, test, and publish your application.

Related tasks

[Duplicating applications](#)

If you don't want to modify a default application, you can duplicate it, give it a new name, and then change the new application.

Development documentation

A Maximo Manage component reference is available from help menu. This information describes the properties and XML code of Maximo Manage application components.

In addition, visit the [MAF Configuration Practices playbook](#) for application configuration examples and the [IBM Maximo REST API Guide](#) for API information.

Configuration training can be found at [Course overview](#).

Deprecated Maximo Application Framework components

Maximo Application Framework application components are sometimes deprecated and replaced.

Container

The Maximo Application Framework container component that was used to configure application layout is deprecated in Maximo Application Suite v8.10. The box and border-layout components are now used for application layout.

<i>Table 145. Layout component properties.</i> Property information for migrating from the container component to the box component				
container		box		Remarks
Property in the container	Value in the container	Property in the box	Value in the box	
layout	vertical horizontal	direction	column row	
align-self	start end			The box component does not support the align-self property. Use the horizontal-align and vertical-align properties as an alternative.
hidden	true false	hidden	true false	
halign-items	start center end between around evenly	horizontal-align	start center end	
valign-items	start center end between around evenly	vertical-align	start center end	
min-width	<i>number#/%</i>	min-width	<i>number#/%</i>	min-width is a string and can be used with PX or %.

Table 145. Layout component properties. Property information for migrating from the container component to the box component (continued)

container		box		Remarks
grow	<i>number</i> true false	children-sizes + fill-parent		
shrink	<i>number</i> true false			
padding	internal internal- external external	padding padding-top padding-bottom padding-left padding-right		The box component padding property applies the margin value in fraction of REM.
border	true false	border	true false	
border-color	Color code	border-color	Color code	
border-top	true false	border-top	true false	
border-bottom	true false	border-bottom	true false	
border-start	true false	border-start	true false	
border-end	true false	border-end	true false	
border-width	<i>number</i>	border-width	<i>number</i>	
background-color	Color code	background-color	Color code	
max-width	Pixel size	children-sizes	Percentage	
disable-auto-layout				
children-fill-parent-width	true false	fill-parent	true false	
display	flex inline-flex grid			
overflow		children-hide-overflow	true false	
size		size	<i>number</i>	

Table 145. Layout component properties. Property information for migrating from the container component to the box component (continued)

container		box		Remarks
		manage-children	true false	When set to true, children divide and equal the width of the parent. When set to false, children shrink to fit their width.

page-header

The Maximo Application Framework page-header component manages the header for a page is deprecated in Maximo Application Suite v8.11. The header-template component is now used to manage the header for a page. Refer to the graphite component reference guide for information about Maximo Application Framework components and properties.

Table 146. Page header component properties. Property information for migrating from the page-header component to the header-template

page-header		header-template		Remarks
Property	Value	Property	Value	
title	<string>	title	<string>	
back	true false			The header-template component doesn't support back.
busy	true false			The header-template component doesn't support busy.
sticky	true false	mode	sticky static dynamic condensed	
back-label	<string>			The header-template component doesn't support back-label.
background	grey normal			The header-template component doesn't support background.
on-back-clicked	<event>			

Related information

[IBM Maximo Application Framework Component Reference](#)

Serviceability

The Application Configuration Docker image includes network tools that are commonly used to troubleshoot deployment issues. IBM Support can direct you to use tools such as **ping**, **tracpath**, **nano**, **ifconfig**, **netstat**, **traceroute**, and **dig** to collect information about your deployment to assist them in troubleshooting issues.

Removing Docker containers through the command line

Downloading Docker images for Application Configuration v9.0 and earlier and instantiating multiple containers can use significant disk space. Remove unneeded Docker containers periodically.

To remove a Docker container from your system, complete the following steps.

1. List all Docker containers.

```
docker container ls -a
```

The output lists all running containers and their numeric IDs.

2. Stop the container.

```
docker container stop container_id
```

3. Remove the stopped container.

```
docker container rm container_id
```

Commands that use a *container_id* value can typically accept a *container_name* value instead. Refer to Docker documentation for more information..

Command variable legend

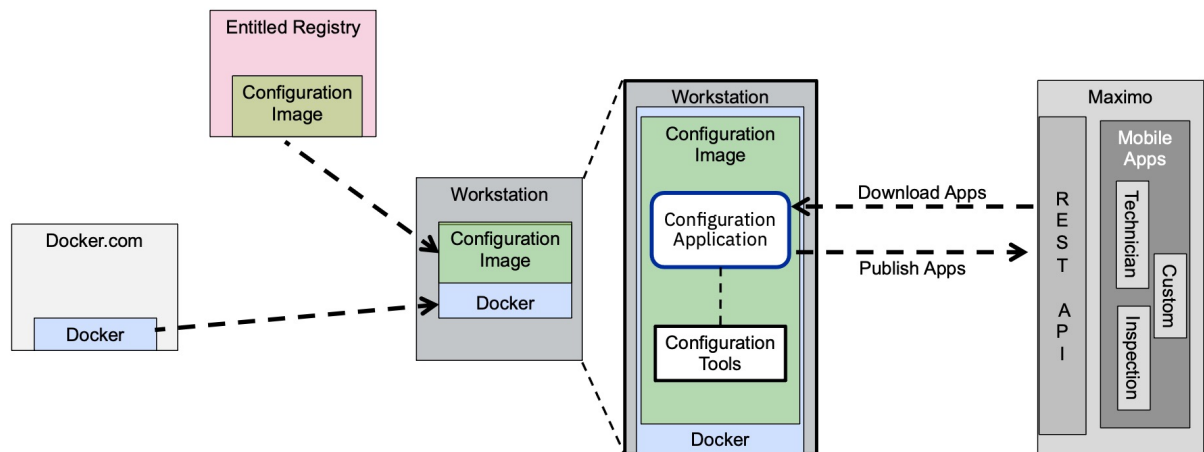
These variables are used in examples.

Variable	Description
<i>`\${maximo_workspace_directory}`</i>	The working directory that you create on your local system.
<i>`\${app_directory}`</i>	The directory on your local system that contains application code when you download an application for configuring.
<i>`\${maximo_app_name}`</i>	The name of the Maximo Mobile application that you are configuring. For example, Inspections or Technician.

Configuration and customization summary

Configuration refers to updates to the display data contained in application XML code. *Customization* refers to the creation of JavaScript code to support application functions. For example, you might configure an application by adding a field to an application page. If the field requires JavaScript to retrieve data, you would customize the application by creating JavaScript to support that field.

Deployment - (Maximo 7.6.1.2 & MAS 8.x)



To configure and customize an IBM Maximo Application Suite v9.0 and earlier application, complete the following steps:

1. Install Docker on your local development system.
2. Download the Application Configuration Docker image from the IBM Entitlement Registry.
3. Configure the [IBM Maximo application server application authorization settings](#).
4. Start the Application Configuration Docker container.
5. Log in to the Configuration application which requires an IBM Maximo Manage server and related credentials.
6. Select a Maximo Application Suite application to modify.
7. Modify the application.
8. Save your changes.
9. You can view changes to components with the component preview pane, or view the entire updated application locally by using the Preview function.

Previewing an application rebuilds the application and starts a React server.

10. Publish the application to the IBM Maximo Manage server.

Configuring Application Configuration v9.1

Configure Application Configuration in your OpenShift environment to make it available to users in Maximo Application Suite v9.1.

Procedure

1. Log into Maximo Application Suite as an administrator.
2. Configure Application Configuration in Maximo Application Suite.
 - a) Open the Admin dashboard. Under Administration, select **Configurations**.
 - b) Scroll down to the Other section, and select **Application configuration**.
 - c) Click **Configure**.
 - d) Click the option button to enable Application Configuration.
 - e) Fill in the Persistent Volume details.

Persistent volume name

Add a name for the persistent volume. For example, `config-workspace`

Size

Provide the amount of space you need. For example, 8Gi

Storage class name

Provide the name of an existing storage class in the cluster. For example, `nfs-client`

f) Click **Save**.

Application Configuration is now available from the application navigation panel, under Suite applications, select **Application Administrator > Application configuration**.

3. Add users to the security group.

a) Open the Admin dashboard. Under Security, select **Security Groups**.

b) Search for the MAFCONFIGADMIN security group and then select it.

You have the option to configure an alternative security group if needed.

c) Choose the Users tab and click **Add users**.

d) Search for and select users you want to add to the MAFCONFIGADMIN security group and then click **OK**.

Users belonging to the MAFCONFIGADMIN security group can now access Application Configuration from their navigation menu when the next log on.

Application Configuration v9.0 and earlier deployment

Download the Application Configuration Docker v9.0 and earlier image and start it as a new Docker container.

Environment configuration

With IBM Maximo Manage v9.0 or earlier, you can either host your entire Application Configuration environment in a single virtual machine (VM) or you can host the IBM Maximo Manage application Server in a VM and your Maximo Application Framework Configuration application and browser on a local system.

Before you begin, see [“Prerequisites” on page 870](#) for information on environment requirements.

1. Host your entire environment in a single VM. Your environment includes the IBM Maximo Manage application server, the Maximo Application Framework Configuration application, and your browser.
2. Host a IBM Maximo Manage server in a VM and host your Maximo Application Framework Configuration application and browser on your local system. Ensure that you can access the IBM Maximo Manage server VM from your local system.

Do not host Application Configuration and your browser in one VM and deploy the IBM Maximo Manage server in a separate VM. This configuration causes performance issues.

Container network accessibility can cause access issues for the IBM Maximo Manage server if the following conditions occur:

- The Application Configuration Docker image is hosted in a VM.
- You cannot access the IBM Maximo Manage server during the configuration login.
- You can access the IBM Maximo Manage server from a browser that is hosted in the same VM.

Maintain separate workspaces for each version of IBM Maximo Manage deployed in your environment. For example, establish separate workspaces for version 8.7, 8.8 and 8.9x` .

Downloading the Docker container image

Application Configuration is provided in a Docker container image.

Before you begin

Refer to the [Configuring and customizing IBM® Maximo® Mobile support page](#) to determine the version of the container image that is appropriate for your version of Maximo Application Suite.

About this task

You can download the Application Configuration Docker container image from the IBM® Entitled Registry. You need an entitlement key and access to IBM Passport Advantage to download the container image.

Alternatively, you can download the same image from the IBM Open Registry. If you download the image from the Open Registry, you do not need an entitlement key and access to IBM Passport Advantage is not required.

Procedure

1. Download the Application Configuration Docker container image from the IBM® Entitled Registry.
 - a) Log in to the [My IBM container software library](#) by using the IBMid and password.
 - b) In the Entitlement keys section, click **Get entitlement key**.
 - c) Click **Copy key**.
 - d) Paste the key into an empty file. You need the key to deploy the Maximo Application Framework Configuration application with Docker.
 - e) Open a command prompt on your system.
 - f) Run the following command to log in to the IBM Entitled Registry.
Replace the *entitlement_key* variable with the key that you copied.

```
docker login cp.icr.io --username cp --password entitlement_key
```

- g) Run the following command to download the Docker container image:

```
docker pull cp.icr.io/cp/manage/maf-tools:9.0.0
```

2. Download the Application Configuration Docker container image from the IBM® Open Registry.
 - a) Open a command prompt on your system.
 - b) Run the following command to download the Docker container image:

```
docker pull icr.io/cpopen/maf-tools:9.0.0
```

Running the Docker image

After the image download is completed, you can run the Docker image to start the Maximo Application Framework Configuration application.

Note: Before you run any commands, ensure that the Maximo Application Framework Configuration application container has write permission to the workspace directory.

Run the Docker image downloaded from the IBM Entitled Registry.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v /home/core/maximo_workspace_directory:/graphite/.workspace -it cp.icr.io/cp/manage/maf-tools:9.0.0
```

Run the Docker image downloaded from the IBM Open Registry.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v /home/core/maximo_workspace_directory:/graphite/.workspace -it icr.io/cpopen/maf-tools:9.0.0
```

Replace the *maximo_workspace_directory* variable with a working directory on your local system. Specify a port number with a value in the range 3000 - 3050. If you are using a Microsoft Windows system, you have to enable file sharing for your workspace directory.

If you run the Maximo Configuration v8.11 Docker image on a Windows system, you must include the **-e** parameter.

The following code is an example of the command with the parameter included.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v "C:/home/core/  
maximo_workspace_directory:/graphite/.workspace" -it -e CHOKIDAR_USEPOLLING=true cp.icr.io/cp/  
manage/maf-tools:8.11.0
```

If you run the Maximo Configuration v9 or later Docker image on a Windows system, you must include the **-e** parameter.

The following code is an example of the command with the parameter included.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v "C:/home/core/  
maximo_workspace_directory:/graphite/.workspace" -it -e WATCHPACK_POLLING=true cp.icr.io/cp/  
manage/maf-tools:9.0.0
```

Maximo Manage requires a valid security certificate to retrieve and publish applications from the Maximo Application Framework Configuration application in production environments. In development or demonstration environments, you can use the **NODE_TLS_REJECT_UNAUTHORIZED** parameter to bypass system checks for a valid security certificate.

The following code is an example of the command with the parameter included.

```
docker run -it --privileged --env NODE_TLS_REJECT_UNAUTHORIZED=0 -p 3001:3001 -p  
3006:3006 -v /home/core/maximo_workspace_directory:/graphite/.workspace -it cp.icr.io/cp/manage/  
maf-tools:9.0.0
```

Alternatively, for a secure option, you can use a self-signed certificate if you do not want to bypass validation.

The following code is an example of the command using a self-signed certificate.

```
docker run -it --privileged -p 3001:3001 -p 3006:3006 -v /Users/maximo/IBM/  
MAF:/graphite/.workspace -v /home/core//Users/maximo/IBM/MAF/certs:/etc/ssl/certs --env  
NODE_EXTRA_CA_CERTS=/etc/ssl/certs/maximo.crt -it cp.icr.io/cp/manage/maf-tools:9.0.0
```

The `-v /Users/maximo/IBM/MAF/certs:/etc/ssl/certs` parameter defines the directory where the root certificate can be found.

The `--env NODE_EXTRA_CA_CERTS=/etc/ssl/certs/maximo.crt` parameter sets an environment variable in the Docker container that instructs the node to use the root certificate that was mounted in the directory.

Do not forget to stop the Application Configuration Docker container when you are done with your work to free up system resources.

Related information

[“Command variable legend” on page 875](#)

Migrating applications

Migrate a Maximo Application Suite v9.0 or earlier application from one environment to another.

About this task

Maximo Migration Manager does not support migrating applications from one environment to another/. However, you can manually migrate applications. For example, if you configured the Maximo Mobile Technician application in a development environment you can move that application to a test or production environment.

If you are migrating a duplicated application between environments, you must duplicate the application in both environments before you can migrate files.

Procedure

1. Open Application Configuration and download the TECHMOBILE application from a development environment.

2. Use Application Configuration to make some changes and save them.

When you edit the TECHMOBILE application, a development workspace is created on your system that contains the application files.

```
C:/myDEVworkspace/user-maximoserver/TECHMOBILE/src/  
  app.xml  
  AppCustomizations.js
```

3. Download the TECHMOBILE application from a test environment.

When you edit the TECHMOBILE application, a test workspace is created on your system that contains the application files.

```
C:/myTESTworkspace/user-maximoserver/TECHMOBILE/src/  
  app.xml  
  AppCustomizations.js
```

4. From Application Configuration, return to the application list page.

5. Copy the `app.xml` and `AppCustomizations.js` files from the development workspace and overwrite those files on the test workspace.

6. Verify that the changes were made to the test workspace files by editing the TECHMOBILE application and searching for the changes in the XML editor.

7. Use the REST API to publish from a development environment to a production environment.

This process works well when performing a full migration because the ZIP file used is not environment specific. For example:

```
POST /maximo/api/os/OSLCMAFAPPDATA?  
appid=INSPECTION&action=wsmethod:uploadFile&version=9.0.0.0  
  
Body: binary of the zip  
{workspace}  
  
{user}{appname}\build\app\build directory. This  
is the local workspace on your computer that builds the app.
```

Language support for Maximo Application Framework applications v9.0 and earlier

Text that you add or change in the UI for an application can be globalized so it appears translated when you select a non-English base language.

When you open an application to edit, it is downloaded to your local workspace. In the `/src` folder of the application, there is a subfolder named `i18n`. For example, `workspace/TECHMOBILE/src/i18n`. In this folder is a `labels.json` file that is built from the application. This file contains the default labels used by the application. The `labels.json` contains English labels by default.

Application XML files define elements that includes text.

```
<field label="Current status" id="y7nv5" swap-position="false" label-class-name="header-small">
```

The text is also contained in the `labels.json` file.

```
"y7nv5_label": "Current status"
```

In a multi-language deployment, non-English label files are located in the `public/i18n/` folder. For example, `workspace/TECHMOBILE/src/public/i18n`. The naming convention for these files are determined by using ISO language codes. For example, for French, the labels file is `labels-fr.json` and German is `labels-de.json`.

An application determines the appropriate labels file to display based on the locale of the current user. If there is no translated label file for the user's locale, the application defaults to the English `labels.json` file.

If a new component with a label is added to the application XML file, the `labels.json` file is updated when the application is rebuilt. If the label requires translation to another language, for example, French, you must manually add the label to `labels-fr.json` file and then translate the file.

Customer-managed **Troubleshooting**

To identify and resolve problems with IBM Maximo Application Suite, you can use the troubleshooting and support information. For information regarding the service and support, see the [IBM Support Guide](#) on IBM.com.

Continuous delivery support

Starting with Maximo Application Suite 8.7, Maximo Application Suite introduced a continuous delivery support model that incrementally delivers regular product fix packs. For more information, see [IBM® Maximo® Application Suite product lifecycle](#).

Release notes

A maintenance fix might be available to resolve your problem. Review the readme files before you apply fixes to ensure that the version, release, modification, and fix level is appropriate for your requirements. For more information, see [Maximo Application Suite release notes](#).

Product-specific help guides

The following product-specific help guides are available:

- [Maximo Application Suite support](#)
- [Product Support Details for Maximo Application Suite](#)

Application-specific troubleshooting

Application-specific troubleshooting documentation is available to help:

- [Maximo Health](#)
- [IBM Maximo Monitor](#)
- [Maximo Predict](#)
- [IBM Maximo Visual Inspection](#)

Self-help

Before you report your problem to IBM Support, see whether the self-help support options resolve your problem:

- Ensure that your service is available and is not undergoing any maintenance work.
 - Maximo Application Suite
 - Contact your local Maximo Application Suite administrator.
- Search the IBM Support Community [knowledge bases](#) and [forums](#) for answers to your question or issue.
- Look through open technotes that cover [troubleshooting Maximo Application Suite issues](#).

Prometheus and Grafana support

A number of components of the IoT tool and the suite licensing service generate Prometheus-based metrics. The metrics can be accessed directly or by using a set of sample Grafana dashboards.

For more information about how to use Prometheus and Grafana with Maximo Application Suite, see:

- The [“Monitoring Maximo Application Suite” on page 824](#) documentation.
- The [Monitoring samples](#) on GitHub.

Product, application, and tools version

In communications with your local Maximo Application Suite administrator or with IBM Support, you might need to provide information about the product, application, or tool version that you are inquiring about.

Complete information about your Maximo Application Suite setup is available from the Maximo Application Suite user interface **Help > About** menu. From there, you can see the version number of each component and also export the full version listing as a .JSON file.

Collecting MustGather data

Maximo Application Suite provides a MustGather script that you must use to collect information before you raise an issue that requires IBM support.

For more information, see the [Collect MustGather data](#) technote.

Troubleshooting Amazon Web Services

To get help and support for your IBM Maximo Application Suite product, see the IBM Support Guide on IBM.com.

Product-specific help guides are available here:

- [Maximo Application Suite support](#)
- [Product Support Details for Maximo Application Suite](#)

You might find useful troubleshooting information in the following topics:

- [Troubleshooting installation problems](#)
- [Retrieving the installation source code version](#)

Troubleshooting installation problems for Amazon Web Services

An unsuccessful Maximo Application Suite installation has many possible causes, such as missing or invalid installation parameters, Bootnode creation failures, or cluster creation problems.

Related tasks

[Deleting the Maximo Application Suite stack on Amazon Web Services](#)

Failure points

When you start a Maximo Application Suite installation on Amazon Web Services, the CloudFormation stack template that you configured is used to create the stack. During this process, a Bootnode is created. The Bootnode completes the rest of the Maximo Application Suite installation.

To create the Red Hat OpenShift cluster, the Bootnode starts a bootstrap process. This process creates a bootstrap node that uses the Red Hat OpenShift installer to create master and worker nodes.

A Maximo Application Suite installation can fail at the following points:

- The stack creation process. If the installation failed during this process, the following indicators apply:
 - The stack is not created.
 - The Bootnode is not created.
 - In the **CloudFormation->Stacks** page, the stack status is **ROLLBACK_COMPLETE**.
- The bootstrap process. If the installation failed during this process, the following indicators apply:
 - The stack is created.

- The Bootnode is created.
- In the **CloudFormation->Stacks** page, the stack status is **CREATE_COMPLETE**.
- In the **Output** tab, the DeploymentStatus parameter displays an installation failure message that indicates the cause of the failure, for example: ID-aws-small-NA:FAILURE#The provided ER key is not valid. It does not have access to download the MAS images

Common causes of failure

The installation might fail for one of the following reasons:

- Mandatory installation parameters are missing or invalid optional parameters are specified.
- An unsupported AWS region is selected.
- The Red Hat OpenShift cluster installer times out after it waits for virtual infrastructure resources to be created.

Missing or invalid installation parameters

If you do not enter all of the mandatory parameters, the installation fails. In addition, for groups of optional parameters, such as the Maximo Manage database configuration parameters, you must either enter all of the group's parameters or leave all of them empty.

The following table indicates the failure points for missing mandatory parameters and invalid optional parameters.

Parameter/Group	Mandatory/optional	Failure point	Further information
SSHKey	Mandatory	Stack creation process	
BootnodeSGIngressCidr Ip	Mandatory	Stack creation process	
EntitledRegistryKey	Mandatory	Bootstrap process	
OpenShiftPullSecret	Mandatory	Bootstrap process	
MASLicenseUrl	Mandatory	Bootstrap process	
PublicHostedZone	Optional	Bootstrap process	If you want to create a new Red Hat OpenShift cluster, you must provide this parameter.
Existing OCP cluster details	Optional	Bootstrap process	If you want to reuse an existing Red Hat OpenShift cluster, you must provide the following parameters: <ul style="list-style-type: none"> • OpenshiftClusterApiUrl • OpenShiftUser • OpenShiftPassword

Parameter/Group	Mandatory/optional	Failure point	Further information
Maximo Manage database configuration details	Optional	Bootstrap process	If you want the Maximo Manage application to use a preconfigured database, you must provide the following parameters: - <ul style="list-style-type: none"> • MASLicenseUrl • MASManageDBUser • MASManageDBPassword • MASManageDBJdbcUrl • MASManageDBCertificateUrl
Existing SLS details	Optional	Bootstrap process	If you want to reuse an existing Suite License Service instance, you must provide the following parameters: <ul style="list-style-type: none"> • SLSEndpointUrl • SLSRegistrationKey • SLSPublicCertificateUrl
Existing IBM User Data Services details	Optional	Bootstrap process	If you want to reuse an existing User Data Services instance, you must provide all three parameters. If only few are provided, the deployment will fail.

Important: Maximo Application Suite is configured with existing Suite License Service and consumes licenses configured with Suite License Service because the licensing is managed by Suite License Service.

Note: MASLicenseUrl is applicable for BYOL AWS product only.

Unsupported AWS region

If you chose an unsupported region when you subscribed to the Maximo Application Suite in Amazon Web Services Marketplace, the installation fails in the stack creation process. An error message is displayed in the CloudFormation template, such as the following message: `Template error: Unable to get mapping for RegionMap::us-west-1::HVM64`

Cluster installer timeout

If the Maximo Application Suite installation process takes too long to create the network resources that the Red Hat OpenShift cluster installer waits for, the cluster installer might time out. In this case, the Maximo Application Suite installation fails in the bootstrap process and the following error message is displayed in the CloudFormation stack console: `Failure in creating OCP cluster.`

Failure messages

If an installation fails in the bootstrap process, the failure message that is displayed in the **CloudFormation->Stacks->Output** page indicates the cause of the failure. Further details on failures

and causes are available in the installation log file, which is `mas-provisioning.log`. You can retrieve this file in the following ways:

- [Connect to the Bootnode by using Secure Shell \(SSH\)](#) and retrieve the file from the `/root/mas-on-aws` directory.
- In the Amazon S3 service, connect to the `<cluster-name>-bucket-<region>` bucket and retrieve the file from the `ocp-cluster-provisioning-deployment-context` directory. For more information, see [Buckets overview](#) in the AWS documentation.

For information on installation failures, their possible causes, and the recommended next steps, see the following table:

Failure message	Cause of failure	Next steps
Failure in creating OCP cluster.	The AWS resources that the cluster requires, such as subnets or NAT gateways, are not created.	This issue is intermittent. Uninstall and then reinstall the Maximo Application Suite.
	A resource quota is reached. For example, no more EIPs or NAT gateways can be created because their resource quotas are reached.	- Delete the Maximo Application Suite stack. - Increase the service quotas that apply to the resources or clean up the existing resources. For more information, see AWS service quotas in the AWS documentation. - In your AWS account, verify that similar resources are not used by an existing failed installation. - reinstall the Maximo Application Suite.
	The user has insufficient permissions to install the Maximo Application Suite.	- Delete the Maximo Application Suite stack. - Ensure that the IAM user who installs the Maximo Application Suite has the correct permissions . - reinstall the Maximo Application Suite.
Failure in configuring OCP cluster.	The Red Hat OpenShift cluster is not configured because of an error. For example, the cluster configuration fails because the IBM Operator Catalog cannot be accessed.	- Connect to the Bootnode by using Secure Shell (SSH) and retrieve the <code>`mas-provisioning.log`</code> file from the <code>`/root/mas-on-aws`</code> directory. Or from the S3 bucket. - Contact IBM Support
Failure in creating Bastion host.	The bastion host is not created in the Red Hat OpenShift cluster because of an error. For example, the bastion host creation fails because your AWS account reaches its service limits.	- Connect to the Bootnode by using Secure Shell (SSH) and retrieve the <code>`mas-provisioning.log`</code> file from the <code>`/root/mas-on-aws`</code> directory. Or from the S3 bucket. - Contact IBM Maximo Application Suite support .

Failure message	Cause of failure	Next steps
Failed in the Ansible playbook execution.	An error occurs when the Bootnode uses Ansible automation to deploy Maximo Application Suite prerequisites or the Maximo Application Suite itself in the Red Hat OpenShift cluster.	- Connect to the Bootnode by using Secure Shell (SSH) and retrieve the <code>`mas-provisioning.log`</code> file from the <code>`/root/mas-on-aws`</code> directory. Or from the S3 bucket . - Contact IBM Maximo Application Suite support
This region is not supported for MAS deployment.	In the Maximo Application Suite configuration page on AWS, in the Region field, you selected an unsupported region.	- Delete the Maximo Application Suite stack. - When you reinstall the Maximo Application Suite , in the configuration page, select a supported region. For the list of supported regions, see the “Preparing to install Maximo Application Suite on Amazon Web Services” on page 146 topic.
The provided ER key is not valid. It does not have access to download the MAS images.	The Maximo Application Suite images cannot be downloaded from the IBM Entitled Registry by using the key that was provided in the <code>`EntitledRegistryKey`</code> installation parameter.	<p>- Delete the Maximo Application Suite stack. - Ensure that your ER key has the entitlement to download the Maximo Application Suite images from the registry.</p> <p>You can run a docker command to pull the image of the operator of the Maximo Application Suite version that you want to install to perform the testing.</p> <p>For example, if you are willing to install Maximo Application Suite version 8.8.0, run the following docker commands:</p> <pre data-bbox="1068 1287 1464 1388">docker login cp.icr.io -u cp - p <er-key></pre> <pre data-bbox="1068 1398 1464 1499">docker pull icr.io/cpopen/ibm-mas:8.8.0</pre> <p>- Reinstall the Maximo Application Suite.</p>
Please provide OCP pull secret.	When you use the Red Hat OpenShift IPI or UPI option, if <code>`OpenShiftPullSecret`</code> installation parameter is not specified.	- Delete the Maximo Application Suite stack. - When you reinstall the Maximo Application Suite , provide the pull secret in the <code>`OpenShiftPullSecret`</code> installation parameter.

Failure message	Cause of failure	Next steps
Please provide a valid MAS license URL.	In the MASLicenseUrl installation parameter, the HTTP or S3 location of the Maximo Application Suite license file is not provided.	- Delete the Maximo Application Suite stack. - When you <u>reinstall</u> the Maximo Application Suite, provide the HTTP or S3 location of the Maximo Application Suite license file in the MASLicenseUrl installation parameter.
Please provide all the inputs to use existing SLS.	Invalid parameters are provided for an existing SLS instance.	- Delete the Maximo Application Suite stack. - When you <u>reinstall</u> the Maximo Application Suite, either provide all of the parameters for the `Existing SLS details` group or leave them all empty. If you leave all of these parameters empty, a new SLS instance is created in the cluster.
Please provide all the inputs to use existing OCP.	Invalid parameters are provided for an existing Red Hat OpenShift cluster.	- Delete the Maximo Application Suite stack. - When you <u>reinstall</u> the Maximo Application Suite, either provide all of the parameters for the Existing OCP cluster details group or leave them all empty. If you leave all of these parameters empty, a new Red Hat OpenShift cluster is created.
The provided Hosted zone is not a public hosted zone. Please provide a public hosted zone.	When you use the Red Hat OpenShift IPI option, in the `PublicHostedZone` installation parameter, a private hosted zone is selected.	- Delete the Maximo Application Suite stack. - When you <u>reinstall</u> the Maximo Application Suite, select a public hosted zone in the `PublicHostedZone` installation parameter.
The JDBC details for Maximo Manage are missing or invalid.	The Maximo Manage database configuration parameters are not provided, or the provided parameters are invalid.	- Delete the Maximo Application Suite stack. - Provide valid JDBC parameters to connect to the Maximo Manage database. By using a database connectivity tool, such as <u>dbeaver</u> , and your Maximo Manage database configuration credentials, verify that you can connect to the database. - <u>reinstall</u> the Maximo Application Suite.

Failure message	Cause of failure	Next steps
<p>Error occurs when running the uninstall script (cleanup-mas-deployment.sh), when ran using the following command:</p> <pre>./cleanup-mas-deployment -s <stack name> -r <region></pre> <p>The error message displayed is:</p> <pre>...line 79: \$ {SUPPORTED_REGIONS,,}: bad substitution.</pre>	<p>The installed bash version is shown after the required version.</p>	<p>Update your bash version and try again.</p>
<p>Error occurs after shutting down the cluster within 24 hours of creation and then restarting after that time:Unable to connect to the server: EOF while oc login.</p>	<p>Node certificate is expired on master nodes.</p>	<p>Renew master node certificate. For more information, see https://access.redhat.com/solutions/5953441.</p>

Note: The MAS License URL is applicable for BYOL AWS product only.

Retrieving the installation source code version

You can retrieve the version of the Maximo Application Suite source code that was used in your installation from the AWS CloudFormation console.

The source code for the automated Maximo Application Suite installation process for Amazon Web Services is available in the following public GitHub repository:

- <https://github.com/ibm-mas/multicloud-bootstrap>

For the Maximo Application Suite, the versioning system that is used is the release-semantic versioning system, that is: <major>.<minor>.<patch>, for example 8.8.0.

Procedure

1. In your Amazon Web Services account, open the CloudFormation console.
2. Open the stack that you created when you installed the Maximo Application Suite.
3. In the **Output** page of the CloudFormation console, retrieve the code version from the MASCloudAutomationVersion parameter.

Troubleshooting installation problems for Microsoft Azure

An unsuccessful IBM Maximo Application Suite installation has many possible causes, such as missing or invalid installation parameters, boot node creation failures, or cluster creation problems.

An unsuccessful Maximo Application Suite installation has many possible causes, such as missing or invalid installation parameters, boot node creation failures, or cluster creation problems.

Failure points

When you start a Maximo Application Suite installation on Microsoft Azure, the Microsoft Azure Resource Manager template that you configured is used to create the deployment in the bootnode's resource group. During this process, a boot node is created. The boot node then completes the rest of the Maximo Application Suite installation.

To create the Red Hat OpenShift cluster, the boot node starts a bootstrap process. This process creates a bootstrap node that uses the OpenShift installer to create master and worker nodes.

A Maximo Application Suite installation can fail during the Microsoft Azure deployment initiation process or at bootstrap process.

If the installation failed during the Microsoft Azure deployment initiation process, the following indicators apply:

- The deployment submission is not successful.
- The boot node is not created.
- In the resource group's deployments section, the deployment associated with the Maximo Application Suite deployment is shown as failed.

If the installation failed during the bootstrap process, the following indicators apply:

- The deployment submission is successful.
- The boot node is created.
- In the resource group's deployments section, the deployment associated with the Maximo Application Suite deployment is shown as successful.
- The **MAS provisioning status** field in the email notification displays an installation failure message that indicates the cause of the failure, for example:

```
FAILURE#The provided ER key is not valid. It does not have access to download the MAS images
```

Common causes of failure

The installation might fail for one of the following reasons:

- Mandatory installation parameters are missing, or invalid optional parameters are specified.
- An unsupported Microsoft Azure region is selected.
- The Red Hat OpenShift cluster installer times out after it waits for virtual infrastructure resources to be created.

Missing or invalid installation parameters

If you do not enter all of the mandatory parameters, the installation fails. In addition, for groups of optional parameters, such as the Maximo Manage database configuration parameters, you must either enter all of the group's parameters or leave all of them empty.

The following table indicates the failure points for missing mandatory parameters and invalid optional parameters.

Parameter/Group	Mandatory/optional	Failure point	Further information
SSH public key	Mandatory	Deployment initiation process	
Bootnode NSG Ingress CIDR IP range	Mandatory	Deployment initiation process	
Microsoft Azure service principal client Id	Mandatory	Deployment initiation process	
Microsoft Azure service principal client secret	Mandatory	Deployment initiation process	
Entitled registry key	Mandatory	Deployment initiation process	

Parameter/Group	Mandatory/optional	Failure point	Further information
MAS license URL	Mandatory	Deployment initiation process	
OpenShift pull secret	Optional	Bootstrap process	If you want to create a new Red Hat OpenShift cluster, you must provide this parameter.
Public domain	Optional	Bootstrap process	
Public domain resource group	Optional	Bootstrap process	
Red Hat OpenShift cluster API URL	Optional	Bootstrap process	If you want to reuse an existing Red Hat OpenShift cluster, you must provide all three parameters. If only few are provided, the deployment will fail.
OpenShift user			
OpenShift password			
SLS endpoint URL	Optional	Bootstrap process	If you want to reuse an existing SLS instance, you must provide all three parameters. If only few are provided, the deployment will fail.
SLS registration key			
SLS public certificate URL			
DRO endpoint URL	Optional	Bootstrap process	If you want to reuse an existing DRO instance, you must provide all three parameters. If only few are provided, the deployment will fail.
DRO API key			
DRO public certificate URL			
Maximo Manage DB user	Optional	Bootstrap process	These parameters are mandatory only for Maximo Application Suite core and Maximo Manage offering.
Maximo Manage DB password			
Maximo Manage DB JDBC URL			
Maximo Manage DB certificate URL			
SMTP host	Optional	Bootstrap process	These parameters are mandatory only if Email notification is set to true.
SMTP port			
SMTP username			
SMTP password			

Unsupported Microsoft Azure region

If you chose an unsupported region when you subscribed to the Maximo Application Suite in Microsoft Azure Marketplace, the installation fails in the bootstrap process. An error message is shown in the log file and in the email notification as:

You are using unsupported region for deploying

Maximo Application Suite.

Cluster installer timeout

If the Maximo Application Suite installation process takes too long to create the network resources that the Red Hat OpenShift cluster installer waits for, the cluster installer might time out. In this case, the Maximo Application Suite installation fails in the bootstrap process and the following error message is displayed in the log file:

```
Failure in creating OCP cluster.
```

Failure messages

For information on installation failures, their possible causes, and the recommended next steps, see the following table:

Failure message	Cause of failure	Next steps
Failure in creating OCP cluster.	The Microsoft Azure resources that the cluster requires, such as VNet, subnets or gateways, are not created.	This issue is intermittent. Uninstall and then Reinstall the Maximo Application Suite.
Failure in configuring OCP cluster.	The Red Hat OpenShift cluster is not configured because of an error. For example, the cluster configuration fails because the IBM Operator Catalog cannot be accessed.	Connect to the boot node by using Secure Shell (SSH) and retrieve the <code>mas-provisioning.log</code> file from the <code>/root/mas-on-aws</code> directory. Contact IBM Maximo Application Suite support.
Failure in creating Bastion host.	The bastion host is not created in the Red Hat OpenShift cluster because of an error. For example, the bastion host creation fails because your AWS account reaches its service limits.	Connect to the boot node by using Secure Shell (SSH) and retrieve the <code>/root/ansible-devops/multicloud-bootstrap/mas-provisioning.log</code> file. Contact IBM Maximo Application Suite support.
Failed in the Ansible playbook execution.	An error occurs when the boot node uses Ansible automation to deploy Maximo Application Suite prerequisites or the Maximo Application Suite itself in the Red Hat OpenShift cluster.	Connect to the boot node by using Secure Shell (SSH) and retrieve the <code>/root/ansible-devops/multicloud-bootstrap/mas-provisioning.log</code> file. Contact IBM Maximo Application Suite support.

Failure message	Cause of failure	Next steps
This region is not supported for MAS deployment.	During the Maximo Application Suite deployment from Microsoft Azure Marketplace, you selected an unsupported region.	<p>Uninstall the Maximo Application Suite.</p> <p>When you Reinstall the Maximo Application Suite, in the configuration page, select a supported region. For the list of supported regions, see the “Preparing to installing Maximo Application Suite on Microsoft Azure” on page 167 topic.</p>
The provided ER key is not valid. It does not have access to download the MAS images.	The Maximo Application Suite images cannot be downloaded from the IBM Entitled Registry by using the key that was provided in the Entitled registry key installation parameter.	<p>Uninstall the Maximo Application Suite.</p> <p>Ensure that your ER key has the entitlement to download the Maximo Application Suite images from the registry. For example, in a Docker command shell, use the following commands to verify that you can manually download a Maximo Application Suite image:</p> <pre data-bbox="1068 947 1466 1087">docker login cp.icr.io -u cp -p <ER-key> docker pull cp.icr.io/cp/mas/admin- dashboard:5.1.27</pre> <p>Reinstall the Maximo Application Suite.</p>
Please provide OCP pull secret.	In the OpenShift pull secret installation parameter, the OpenShift pull secret is not provided.	<p>Uninstall the Maximo Application Suite.</p> <p>When you Reinstall the Maximo Application Suite, provide the pull secret in the OpenShiftPullSecret installation parameter.</p>
Provide a valid MAS license URL.	In the MAS license URL installation parameter, the HTTP or S3 location of the Maximo Application Suite license file is not provided.	<p>Uninstall the Maximo Application Suite.</p> <p>When you Reinstall the Maximo Application Suite, provide the HTTP or S3 location of the Maximo Application Suite license file in the MASLicenseUrl installation parameter.</p>

Failure message	Cause of failure	Next steps
Please provide all the inputs to use existing SLS.	Invalid parameters are provided for an existing SLS instance.	<p><u>Uninstall the Maximo Application Suite.</u></p> <p>When you <u>Reinstall the Maximo Application Suite</u>, either provide all of the parameters for SLS details in the <u>Existing Infrastructure</u> group or leave them all empty. If you leave all of these parameters empty, a new SLS instance is created in the cluster.</p>
Please provide all the inputs to use existing .DRO	Invalid parameters are provided for an existing DRO instance.	<p><u>Uninstall the Maximo Application Suite.</u></p> <p>When you <u>Reinstall the Maximo Application Suite</u>, either provide all of the parameters for DRO details in the <u>Existing Infrastructure</u> group or leave them all empty. If you leave all of these parameters empty, a new DRO instance is created in the cluster.</p>
Please provide all the inputs to use existing OCP.	Invalid parameters are provided for an existing Red Hat OpenShift cluster.	<p><u>Uninstall the Maximo Application Suite.</u></p> <p>When you <u>Reinstall the Maximo Application Suite</u>, either provide all of the parameters for Red Hat OpenShift cluster details in the <u>Existing Infrastructure</u> group or leave them all empty. If you leave all of these parameters empty, a new Red Hat OpenShift cluster is created.</p>
The JDBC details for Maximo Manage are missing or invalid.	The Maximo Manage database configuration parameters are not provided, or the provided parameters are invalid.	<p><u>Uninstall the Maximo Application Suite.</u></p> <p>Provide valid JDBC parameters to connect to the Maximo Manage database. By using a database connectivity tool, such as <u>dbeaver</u>, and your Maximo Manage database configuration credentials, verify that you can connect to the database.</p> <p><u>Reinstall the Maximo Application Suite.</u></p>

Retrieving the installation source code version

You can retrieve the IBM Maximo Application Suite source code version that was used in your installation from a VM image.

About this task

The source code for the installation process is available in the following public GitHub repository:

<https://github.com/ibm-mas/multicloud-bootstrap>

For IBM Maximo Application Suite, the versioning system that is used is the release-semantic version system that is: <major>.<minor>.<patch>. For example, 8.8.0.

Procedure

1. In your Microsoft Azure portal, open the bootnode resource group.
2. Look for the succeeded deployments and open the deployment that is related to the bootstrap process.
3. In the **Outputs** section, retrieve the code version from the `masCloudAutomationVersion` parameter.

Customer-managed Troubleshooting installation and configuration issues

Installation and configuration troubleshooting information for Maximo Application Suite is available in the Support knowledge base and throughout the documentation.

In the Support knowledge base

Use the following dynamic search link to find Maximo Application Suite technotes: [IBM Support](#)

Throughout the documentation

The “[Installing Maximo Application Suite](#)” on page 218 and “[Deploying Maximo Application Suite applications](#)” on page 292 documentation contains hints and tips throughout the end-to-end tasks for each installation path. These hints and tips include a mapping of Ansible roles to each documentation task, sample inputs and outputs for common and complex commands, and path-specific troubleshooting for Amazon Web Services (AWS) and Microsoft Azure installations.

Important: First log in to Maximo Real Estate and Facilities with the mandatory initial FACILITIESADMIN user and set up other users. If you try to log in to Maximo Real Estate and Facilities with any other user before you log in with the FACILITIESADMIN user, you get a blank screen and an error message that says the user is invalid, see “[Administering Maximo Real Estate and Facilities users](#)” on page 385.

Troubleshooting Microsoft Active Directory user synchronization

When you synchronize users from Microsoft Active Directory into IBM Maximo Application Suite, the number of users that are created in Maximo Application Suite might be fewer than the expected number of users.

Symptoms

Microsoft Active Directory might limit the Maximo Application Suite user search.

Causes

The Maximo Application Suite `customMaxSearchResults` configuration is set to 10000 by default. This configuration can be used to increase the number of users that can be synchronized from the LDAP server into Maximo Application Suite. However, if the limit of the user search value of the Microsoft

Active Directory **MaxPageSize** property is less than the **customMaxSearchResults** property value, the **MaxPageSize** property takes precedence. If the value of **customMaxSearchResults** is greater than **MaxPageSize**, the page size is used to limit the number of users in the search result.

Resolving the problem

Starting in IBM Maximo Application Suite 8.9, you can use two methods to resolve the limitation in user registry synchronization when you use Microsoft Active Directory.

- Increase **MaxPageSize** in Microsoft Active Directory.
- Alternatively, change the **ldapType** setting to **Microsoft Active Directory** in the ScimConfig custom resource (CR) in Red Hat OpenShift. The **customMaxSearchResults** property takes precedence over **MaxPageSize** of Microsoft Active Directory.

The Maximo Application Suite user registry synchronization settings **customMaxSearchResults** and **ldapType** are not shown in the user interface. The settings must be configured in Red Hat OpenShift.

1. In Red Hat OpenShift, from the side navigation menu, click **Administration > CustomResourceDefinitions**.
2. On the **CustomResourceDefinitions** page, search for and open the ScimCFG custom resource definition.
3. On the **Instances** tab, search for and open the CR that starts with the Maximo Application Suite instance ID.

For example, `scimcfgs.config.mas.<your_company_name>.com`

4. On the **YAML** tab, in the `spec:` under `config:`, add `ldapType: 'Microsoft Active Directory'`.

By adding the **ldapType** property, the maximum number of users that are synchronized matches the value of the **customMaxSearchResults** property even if the Microsoft Active Directory limit is less.

For example, in the following ScimCFG custom resource, the **customMaxSearchResults** value is set to 30000 to increase the maximum number of synchronized users even if **MaxPageSize** in the Microsoft Active Directory is set to 10000 users.

```
...
spec:
  config:
    ldapType: 'Microsoft Active Directory'
    customMaxSearchResults: '30000'
```

LDAP user mapping from Microsoft Active Directory

When synchronizing LDAP users from Microsoft Active Directory into the Maximo Application Suite database, some LDAP user properties might not match to the corresponding Maximo Application Suite user properties.

Symptoms

LDAP user properties do not match to the Maximo Application Suite user properties. For example, an LDAP user called *John Doe* might have the **givenName** property set to John Doe and the **displayName** property set to John Doe Doe.

Resolving the problem

A Maximo Application Suite administrator with access to the Red Hat OpenShift cluster can configure the properties from LDAP to map to Maximo Application Suite by updating the **mappings** property in the ScimCFG custom resource.

The following Maximo Application Suite user properties can be mapped to LDAP user properties:

- **title**

- **givenName**
- **familyName**
- **displayName**

In the ScimCFG custom resource, you configure the **mappings** property in the **spec.config.userSync** section by using the following format:

```
spec:
  ...
  config:
    ...
    userSync:
      ...
      mappings:
        standard:
          <Maximo Application Suite user property name>: <LDAP user property name>
          <Another Maximo Application Suite user property name>: <Another LDAP user property
name>
      ...
```

For example, when synchronizing Microsoft Active Directory LDAP users into Maximo Application Suite so that the LDAP user John Doe is stored with the **givenName** set to John and the **displayName** set to John Doe, configure the **mappings** property in the ScimCFG custom resource.

```
spec:
  ...
  config:
    ...
    userSync:
      ...
      mappings:
        standard:
          givenName: givenName
          displayName: displayName
```

When Maximo Application Suite connects to LDAP systems, the SCIM specification is processed internally. This uses a set of standard properties from the LDAP registry to form the given name and display name for users who are created in Maximo Application Suite. By using the mapping configuration in the ScimCFG custom resource, an administrator can synchronize the **givenName** and **displayName** to the same attributes within the LDAP directory. This configuration prevents the need for complex assembly naming formats.

When you save the mapping changes, the configuration is processed and applied to the **status.config.userSync** section. In the next scheduled synchronization cron job, the user synchronization changes are applied.

Related concepts

[LDAP user registry synchronization](#)

User registry synchronization simplifies Maximo Application Suite user management by synchronizing users and groups between an LDAP server and your local Maximo Application Suite user registry.

Customer-managed **Troubleshooting upgrade issues**

Troubleshooting information is available for IBM Maximo Application Suite upgrade issues, such as expired API keys, dropped database connections, or data migration.

After you upgrade Maximo Application Suite, problems might occur in Maximo Application Suite components or applications. Use the following links to view troubleshooting information for specific Suite applications:

Suite

Use this information and these solutions to resolve general upgrade issues that you might encounter with the Maximo Application Suite.

Suite applications cannot connect to Db2 Warehouse databases

After Cloud Pak for Data and the Db2 Warehouse service are upgraded, Suite applications cannot connect to databases because the security certificates are no longer valid.

Symptom

After you upgrade Cloud Pak for Data and Db2 Warehouse to the versions that are required by version 8.7.0 of the Suite, some Suite applications cannot connect to Db2 Warehouse databases.

Cause

When you upgrade Cloud Pak for Data and the corresponding Db2 Warehouse service, the Db2 Warehouse SSL certificates are refreshed. As a result, the Suite's Java database connection (JDBC) configuration is invalidated.

Resolution

1. In the Cloud Pak for Data dashboard menu, download the new SSL certificate by clicking **Databases > Db2 instance details > Download SSL Certificate**.
2. In the Maximo Application Suite administration dashboard, click **Configurations > > Database connection** and select the connection that you want to update.
3. Click **Edit** and paste the content of the new certificate.
4. Click **Confirm**, reenter the connection's username and password, and save your changes. The Suite operator identifies this update and ensures that the new certificate is added to the Suite's truststore.
5. After the certificate is added to the truststore, connect to your Red Hat OpenShift cluster by using the web console.
6. In the `openshift-operators` namespace, delete the `service-binding-operator` pod. This step ensures that the new SSL certificate is updated across all Suite applications.

Maximo Application Suite cannot be upgraded during data migration

Starting in 8.11, if IBM Maximo Application Suite is configured for Security Assertion Markup Language (SAML) authentication, a data migration conflict might occur when you upgrade. To resolve, you must create a `ConfigMap` in the Maximo Application Suite core namespace in Red Hat OpenShift.

A data migration conflict occurs when the environment contains any user whose `userid` is not the same as their `username` and authentication is configured to use SAML. Create a `ConfigMap` to migrate data by using either the SAML `username` or SAML `userid`.

About this task

If you see the following error message in Red Hat OpenShift when you upgrade Maximo Application Suite, create a `ConfigMap`:

```
Unable to upgrade from MAS Core {{ currentVersion }} to {{ targetVersion }}:  
Conflict detected when running Data Migration job.  
This error occurs when you have users of type SAML with userid != username.
```

Procedure

Use either a YAML file in Red Hat OpenShift or Red Hat OpenShift commands to create the ConfigMap.

Check with your application administrator to determine which field is being used for the link between Maximo Application Suite and the SAML identity provider. If user ID is used, use `_id` in the ConfigMap, otherwise use `username`.



Attention: If the wrong field is used, access to Maximo Application Suite might be lost and the fields must be updated manually.

- Create a ConfigMap by using a YAML file.
 - a) In Red Hat OpenShift from the side navigation menu, click **Workloads > ConfigMaps** and then click **Create ConfigMap**.
 - b) Click **YAML view**.
 - c) Add configmap `name` in **metadata** and `saml_id_field_link` in **data**.
You can add either `_id` or `username` in the `<your_saml_id_link>` field.
 - d) Click **Create**.
For example, the following YAML uses the SAML userid to migrate data:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: 'datarepair-config-811'
  namespace: <your_mas_core_namespace>
data:
  saml_id_field_link: <your_saml_id_link>
```

- Create a ConfigMap by using Red Hat OpenShift commands.
In your command window, run the following commands:

```
oc project mas-<instance_id>-core

echo 'apiVersion: v1
kind: ConfigMap
metadata:
  name: 'datarepair-config-811'
data:
  saml_id_field_link: <your_saml_id_link>' | oc apply -f -
```

What to do next

In Red Hat OpenShift, from the side navigation menu, click **Workloads > Pods > Logs**, verify that all users are migrated.

Troubleshooting Import user template

Starting in IBM Maximo Application Suite 8.11, when you upgrade the IBM Maximo Application Suite, you might find differences in the Import user template `.csv` headers.

Symptoms

You might see the following error when you import users that indicate a mismatch between the columns in your data and the required template.

```
Invalid input violation:
['The column titles in the .csv file do not match the column titles in the .csv file template.
View and fix the errors, and then upload the file']
```

Causes

Starting in Maximo Application Suite 8.11, in Import user template, the **issuer** field is replaced with **local_user**. Four new fields such as **identities_saml**, **identities_saml_id**, **identities_ldap**, and **identities_ldap_id** are added.

Resolving the problem

Resolve the problem of header discrepancies by using any of the following methods.

- – For Maximo Application Suite 9.1, open the csv file from the **Suite > Security > Users** page.
 - For Maximo Application Suite 9.0 and earlier, download the updated template from the **Suite administration** page, by selecting **Users > Import users**.
- Modify the existing template to include the new fields by replacing the **issuer** field with **local_user** field. Add the new fields **identities_saml**, **identities_saml_id**, **identities_ldap**, and **identities_ldap_id**.

Related reference

[Importing users in Maximo Application Suite in 9.0 and earlier](#)

To create multiple users in Maximo Application Suite, use the template file to import new users and ensure that the format for the user information adheres to the import processing rules. After you import users, you can also use the template file to modify user information and delete users.

Blank or unresponsive user interface after upgrade

When you upgrade to IBM Maximo Application Suite 8.11 without deactivating and uninstalling IBM Maximo Health and Predict - Utilities, you might have issues with the user interface.

Symptoms

The Maximo Application Suite user interface is not displayed or is unresponsive.

Causes

Maximo Application Suite was upgraded to 8.11 without deactivating and uninstalling Maximo Health and Predict - Utilities. In Red Hat OpenShift, the custom resource in the HPUapp and HPUworkspace custom resource definitions was deleted, and the Maximo Health and Predict - Utilities namespace was deleted.

Resolving the problem

1. In the Red Hat OpenShift web console, from the side navigation menu, click **Workloads > Pods**, and then find and open the internalapi pod in the mas-*<instanceid>*-core namespace or project.
2. On the **Terminal** tab of the internalapi pod, run the following command.

```
curl -v -X DELETE https://internalapi.mas-<your_instanceid>-core.svc:443/v1/workspaces/<your_workspaceid>/applications/hputilities  
--cert /etc/pki/tls/certs/mascore-cert/tls.crt --key /etc/pki/tls/certs/mascore-cert/tls.key  
--cacert /etc/pki/tls/certs/mascore-cert/ca.crt
```

Tip: In the command example, replace *<your_instanceid>* and *<your_workspaceid>* values to match your environment.

3. Log in to the MongoDB instance that is used by your Maximo Application Suite.
4. Find the mas_*<instance_id>*_catalog database and select the reservations collection.
5. Delete the HPUutilities document from the reservations collection.
6. In the Red Hat OpenShift web console, from the side navigation menu, click **Workloads > Pods** and find and restart the coreapi pods.

Customer-managed Troubleshooting Suite upgrade issues in Maximo Health , Maximo Predict, and Maximo Health and Predict - Utilities

Use this information and these solutions to resolve Suite upgrade issues that you might encounter in Maximo Health , Maximo Predict, and Maximo Health and Predict - Utilities.

Customer-managed

Updates of the Maximo Manage and Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities applications cause integration issues

After Maximo Manage and Maximo Health, Maximo Predict, or Maximo Health and Predict - Utilities are updated, Maximo Manage cannot fulfill requests from the other applications.

Symptom

After Maximo Manage is updated to version 8.3.0 and Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities are updated to their latest versions, Maximo Manage cannot fulfill requests from the other applications to complete the following tasks:

- Create service requests for assets and locations.
- Create work orders for assets and locations.
- Edit asset and location records.

Cause

HTTP requests from Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities to API endpoints in Maximo Manage fail because the API key token is invalid. Either the token is expired or it was revoked by the administrator.

Resolution

In Maximo Health, Maximo Predict, or Maximo Health and Predict - Utilities, complete the following steps:

1. Add an API key for the integration user by completing step 1 of the procedure in the following topic: [Configuring the integration users](#)
2. Reconfigure the external EAM system properties by completing step 1 of the procedure in the following topic: [Configuring the application or industry solution](#)

Customer-managed

Update of Maximo Monitor breaks functionality in Maximo Predict

After Maximo Monitor is updated, predictive and anomaly detection functionality no longer works in Maximo Predict because of a software dependency mismatch.

Symptom

After Maximo Monitor is updated to version 8.6.2, model predictions and anomaly detection no longer work correctly in version 8.4.1 of Maximo Predict. In the Maximo Monitor API log, the following error message appears:

```
ModuleNotFoundError: No module named 'pandas.io.formats.string'
```

Cause

Maximo Monitor 8.6.2 and Maximo Predict 8.4.1 use different versions of the pandas Python module.

Resolution

Update Maximo Monitor to version 8.6.3.

Customer-managed

Upgrade of Cloud Pak for Data makes some notebooks unusable in Maximo Health and Predict - Utilities

After Cloud Pak for Data is upgraded to version 4.0, the Watson Studio notebooks and jobs that are used in Maximo Health and Predict - Utilities no longer work correctly.

Symptom

After Cloud Pak for Data is upgraded to version 4.0, the asset-class specific notebooks and jobs in Watson Studio that are associated with Cloud Pak for Data 3.0 cannot be used in scoring groups in Maximo Health and Predict - Utilities.

Cause

The notebooks and jobs in Watson Studio are not associated with version 3.7 of the Python runtime environment.

Resolution

1. Log in to the Cloud Pak for Data dashboard as an administrator.
2. Click **Quick navigation->All projects**.
3. In your Maximo Health and Predict - Utilities project, click **Assets->Notebooks**.
4. Associate each notebook that includes 3.0.0 in its name with the Python 3.7 runtime environment.
 - a. Click on the notebooks, for example IBM-SCFF-Cables-3.0.0.ipynb, and click **Change Environment**.
 - b. In the drop down, click **Default Python 3.7 (1 vCPU and 2GB RAM)** and click **Associate**.
 - c. Edit the notebook and comment out the following lines by inserting a hash character, that is '#', before each line:

```
from ctypes import *  
lib1=cdll.LoadLibrary('/opt/conda/envs/Python-3.7-main/lib/libpython3.7m.so.1.0')
```

- a. Click **Save Notebook** and **Save Version**.
5. In your Maximo Health and Predict - Utilities project, click **Jobs**.
 6. Associate each job that includes 3-0-0 in its name with the Python 3.7 runtime environment.
 - a. Click on the jobs, for example Run-IBM-SCGF-Cables-3-0-0, and click **Edit Configuration->Next**.
 - b. In the dropdown, click **Default Python 3.7**.
 - c. Click **Next** until you reach the Review and save step.
 - d. Review the details and save your changes.

Customer-managed

Updates to JDBC configuration are not applied in the Maximo Predict application

After Suite-level JDBC configuration is updated, the Service Binding Operator does not apply these updates in the Maximo Predict application.

Symptom

Asset predictions that relate to downtime, degradation, and failures are not displayed in the application's user interface.

Cause

The Service Binding Operator , which is responsible for applying Suite-level configuration updates to Suite applications, does not apply these updates in the Maximo Predict application.

Resolution

1. Log in to the Red Hat OpenShift web console as an administrator.
2. In the mas-<instanceName>-predict project, click **Workloads->Pods**.
3. Search for the pods that are named predict-api-*
4. Delete the pods and allow them to fully restart. After the pods restart, the JDBC configuration updates are applied in the Maximo Predict application, and asset descriptions are correctly displayed in the user interface.

Customer-managed **Troubleshooting Suite upgrade issues in IoT tool**

Use this information and these solutions to resolve Suite upgrade issues that you might encounter in IoT tool.

Customer-managed **Updates to SSL certificates in JDBC and Kafka configuration are not applied in the IoT tool**

After SSL certificates are updated in JDBC or Kafka configuration, the Service Binding Operator does not apply these updates in the IoT tool .

Symptom

Suite-level JDBC and Kafka configuration changes, such as updates to SSL certificates, are not applied immediately in the IoT tool .

Cause

The Service Binding Operator , which is responsible for applying Suite-level configuration changes to Suite applications, does not apply the changes to the IoT tool.

Resolution

After an SSL certificate is updated in JDBC or Kafka configuration and the new certificate is added to the Suite's truststore, complete the following steps:

1. Connect to your Red Hat OpenShift cluster by using the web console.
2. In the openshift-operators namespace, delete the service-binding-operator pod. This step ensures that the Service Binding Operator is restarted and that it applies the new SSL certificate to the binding secrets of all Suite applications.
3. In the IoT tool namespace, that is mas-<instanceid>-iot, click **Workloads->Secrets** and sort the secrets by descending created date.
4. Verify that the new JDBC or Kafka binding secrets are created and that they include the contents of the new SSL certificate.
5. In the pod information, verify that truststore-worker jobs are created. These jobs ensure that the new SSL certificate is added to the IoT tool truststore.
6. Restart the relevant pods for the Suite-level configuration that was updated.

- a. To apply JDBC configuration updates, restart the auth and state component pods by running the following commands:

```
oc get pod -n mas-$MAS_INSTANCE_ID-iot | grep state- | awk '{print $1}' | xargs oc delete pod -n mas-$MAS_INSTANCE_ID-iot
oc get pod -n mas-$MAS_INSTANCE_ID-iot | grep auth- | awk '{print $1}' | xargs oc delete pod -n mas-$MAS_INSTANCE_ID-iot
```

- a. To apply Kafka configuration updates or other configuration updates, restart all IoT tool pods by running the following command:

```
oc get pod -n mas-$MAS_INSTANCE_ID-iot | awk '{print $1}' | xargs oc delete pod -n mas-$MAS_INSTANCE_ID-iot
```

Customer-managed **Troubleshooting Suite upgrade issues in Maximo Manage**

Use this information and these solutions to resolve Suite upgrade issues that you might encounter in Maximo Manage.

Customer-managed **Updates to JDBC configuration and other data resources are not applied in the Maximo Manage application**

After JDBC configuration or other data configuration that affects the Maximo Manage application is updated, the Service Binding Operator does not apply these updates in the application.

Symptom

Suite-level updates to JDBC configuration and other configuration that affects the Maximo Manage application, such as the Smtplib or Data Dictionary resources, are not applied immediately in the application.

Cause

The Service Binding Operator, which is responsible for applying Suite-level configuration changes to Suite applications, does not apply the changes in the Maximo Manage application.

Resolution

The Service Binding Operator and the Maximo Manage workspace operator need to be restarted. For more information, see [this IBM Support technote](#).

Customer-managed **Updates of the Maximo Manage and Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities applications cause integration issues**

After Maximo Manage and Maximo Health, Maximo Predict, or Maximo Health and Predict - Utilities are updated, Maximo Manage cannot fulfill requests from the other applications.

Symptom

After Maximo Manage is updated to version 8.3.0 and Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities are updated to their latest versions, Maximo Manage cannot fulfill requests from the other applications to complete the following tasks:

- Create service requests for assets and locations.
- Create work orders for assets and locations.
- Edit asset and location records.

Cause

HTTP requests from Maximo Health , Maximo Predict, or Maximo Health and Predict - Utilities to API endpoints in Maximo Manage fail because the API key token is invalid. Either the token is expired or it was revoked by the administrator.

Resolution

In Maximo Health, Maximo Predict, or Maximo Health and Predict - Utilities, complete the following steps:

1. Add an API key for the integration user by completing step 1 of the procedure in the following topic: [Configuring the integration users](#)
2. Reconfigure the external EAM system properties by completing step 1 of the procedure in the following topic: [Configuring the application or industry solution](#)

Customer-managed **Troubleshooting Suite upgrade issues in Maximo Monitor**

Use this information and these solutions to resolve Suite upgrade issues that you might encounter in Maximo Monitor .

Customer-managed **Update of Maximo Monitor breaks functionality in Maximo Predict**

After Maximo Monitor is updated, predictive and anomaly detection functionality no longer works in Maximo Predict because of a software dependency mismatch.

Symptom

After Maximo Monitor is updated to version 8.6.2, model predictions and anomaly detection no longer work correctly in version 8.4.1 of Maximo Predict. In the Maximo Monitor API log, the following error message appears:

```
ModuleNotFoundError: No module named 'pandas.io.formats.string'
```

Cause

Maximo Monitor 8.6.2 and Maximo Predict 8.4.1 use different versions of the pandas Python module.

Resolution

Update Maximo Monitor to version 8.6.3.

Customer-managed **Updates to JDBC configuration are not applied in the Maximo Monitor application**

After Suite-level JDBC configuration is updated, the Service Binding Operator does not apply these updates in the Maximo Monitor application.

Symptom

Suite-level updates to JDBC configuration are not applied immediately in the Maximo Monitor application.

Cause

The Service Binding Operator , which is responsible for applying Suite-level configuration changes to Suite applications, does not apply the changes in the Maximo Monitor application.

Resolution

1. Log in to the Red Hat OpenShift web console as an administrator.
 2. In the `mas-<instanceName>-monitor` project, click **Workloads** > **Pods**.
 3. Locate the pod that is named `ibm-mas-monitor-operator-*` and delete it.
 4. In the web console or a command shell, delete all of the cron jobs in the project.
- In the web console, click **Workloads** > **Cron Jobs** and delete all of the cron jobs in the project.
 - Alternatively, in a command shell, perform a bulk delete of all of the project's cron jobs by running the following command:

```
for j in $(kubectl get cj --namespace mas-<instanceName>-monitor -o custom-columns=:.metadata.name); do kubectl delete cj $j; done
```

1. Open a connection to the Maximo Monitor database and delete the row in the `IOTANALYTICS.KPI_IMAGE_MANAGEMENT` table where the value in the `TENANT_ID` column matches your tenant identifier.
2. Click **Workloads** > **Pods** and locate the Maximo Monitor workspace pod, which is `entitymgr-ws`.
3. Delete the pod and allow it to fully restart.

Reference

Installing Red Hat OpenShift Container Platform and Maximo Application Suite on a Windows system

You can install Red Hat OpenShift Container Platform and IBM Maximo Application Suite on a single computer, such as a laptop for evaluation or development purpose. You can use this installation to evaluate a compact Maximo Application Suite and Maximo Manage that are installed on a single computer. The following example uses a Windows system.

Before you begin



Attention: Installing on Red Hat OpenShift Container Platform and IBM Maximo Application Suite on a single computer or locally is not supported. Any issues encountered with the installation scripts or setup process is outside of the scope of support.

Before you install IBM Maximo Application Suite on a Windows computer, review the virtual environment and software requirements.

This deployment requires 14 virtual CPUs and 30 GiB of memory in the virtual environment.

This deployment requires Red Hat OpenShift Local 2.19, formerly known as Red Hat CodeReady Containers, that includes Red Hat OpenShift Container Platform 4.12.

Windows 11 is the required operating system for this installation.

Important: Maximo Application Suite does not support file systems that are not POSIX-compliant. For example, some NFS file system implementations are not POSIX-compliant. If you want to use an NFS file system for storage, verify with the vendor that the NFS implementation is fully POSIX-compliant.

About this task

In this installation, IBM Suite License Service and IBM Db2 are installed locally and the slim version of IBM User Data Services is used. A slim User Data Services is required because a full User Data Services deployment does not fit in the configured size of the locally installed Red Hat OpenShift Container Platform.

Installing Red Hat OpenShift Container Platform Local

Before you install IBM Maximo Application Suite and IBM Maximo Manage on your computer, you must install and run Red Hat OpenShift Container Platform Local 4.12. You can install this product by installing Red Hat CodeReady Containers 2.19.

Before you begin



Attention: Installing on Red Hat OpenShift Container Platform and IBM Maximo Application Suite on a single computer or locally is not supported. Any issues encountered with the installation scripts or setup process is outside of the scope of support.

If you don't already have a Red Hat account, register with Red Hat at <https://sso.redhat.com/auth/>.

Procedure

1. Download the `crc-windows-installer.zip` file from <https://developers.redhat.com/content-gateway/rest/mirror/pub/openshift-v4/clients/crc/2.19.0>.
2. Decompress the compressed file.
3. Install Red Hat OpenShift Container Platform Local by running the `crc-windows-amd64.msi` file.
4. Open the Red Hat Hybrid Cloud Console at <https://console.redhat.com/openshift>.
5. Click **Create cluster**.
6. Click the **Local** tab.
7. Download or copy your pull secret.
8. Open a command prompt and run `crc setup`.
9. Run the following commands to set the CPUs, memory, and disk size.

```
crc config set consent-telemetry no
crc config set cpus 14
crc config set memory 30720
crc start
crc stop
crc config set disk-size 200
crc start
```

10. Make a note of the login credentials so you can log in to the cluster.

Results

Red Hat OpenShift Container Platform Local is running and ready to host Maximo Application Suite and Maximo Manage.

Installing Maximo Application Suite and Maximo Manage

You can now install IBM Maximo Application Suite and IBM Maximo Manage.

Before you begin



Attention: Installing on Red Hat OpenShift Container Platform and IBM Maximo Application Suite on a single computer or locally is not supported. Any issues encountered with the installation scripts or setup process is outside of the scope of support.

Before you install Maximo Application Suite and Maximo Manage, complete these prerequisite tasks.

1. Create a directory for the contents of the GitHub repository.
2. Download the contents of the GitHub repository by running the following command from the directory that you created.

```
git clone https://github.com/evilADevil/mas-local
cd mas-local
```

3. If you have never done so, create a registry key. To create a registry key, go to <https://myibm.ibm.com/dashboard/>. Click Container Software & Entitlement Keys, then select Add new key.
4. Download the entitled registry key. This key must be enabled to get the Maximo Application Suite and IBM Cloud Pak for Data images. Log in to <https://myibm.ibm.com/dashboard/> and click **Container Software & Entitlement Keys**. Click **Copy** and save the entitlement key to a text file.
5. Download a `license.dat`, which is Maximo Application Suite license file, from the license key server. For more information, see <https://www.ibm.com/support/pages/ibm-support-licensing-start-page>. Put this file in the `mas-local` directory.
6. Get a license ID that matches the Maximo Application Suite license file. You can get the license ID by opening the license file in a text editor and checking the first line. The license ID is the second-to-last number. For example, if your first line is `SERVER sls-rlks-0.rlks 0272bc344002 27000` then your license ID is `0272bc344002`.
7. Replace the `masocpl.yml` file with the `masocpl.suds.yml` file by deleting `masocpl.yml` and then renaming `masocpl.suds.yml` to `masocpl.yml`.
8. Customize `masocpl.yml` by using the information that you collected.
 - a. Replace `<<your ER key>>` with the entitled registry key that you saved.
 - b. Replace `<<your license id>>` with the license ID that you got from the license file `license.dat`.
 - c. Replace the `uds_contact` information.

About this task

Your working directory must include the following files:

- `masdevops.yaml`
- `masocpl.yml`
- `license.dat`
- `masinst.bat`

Procedure

1. Log in to the Red Hat OpenShift Container Platform Local cluster as an administrator. If you haven't collected the credentials previously, get the credentials by running the following command:

```
crc start
```

In the output for this command, make a note of the log-in credentials.

2. At a command prompt, run the following commands by using the password that your environment provided:

```
@FOR /f "tokens=*" %i IN ('crc oc-env') DO @call %i
oc login -u kubeadmin -p <password> https://api.crc.testing:6443
```

3. Install Maximo Application Suite and Maximo Manage by running the `masinst` batch file and waiting for it to finish.

Pay attention to the log and record the user ID and password of the Maximo Application Suite superuser. You can find your username and password by checking for the following message in the log:

```
Maximo Application Suite is Ready, use the superuser credentials to authenticate
```

If you don't have the log anymore, you can retrieve the credentials from the `masdemo-credentials-superuser` secret in the `mas-masdemo-core` namespace.

Maximo Application Suite core services

Refer to the details about IBM Maximo Application Suite architecture, topology, components and how the components are distributed as micro services solutions.

Kubernetes control plane

Controller managers

ibm-mas-operator

`ibm-mas-operator` watches `Suite.core.mas.ibm.com`, acts as the primary controller manager for an installation of the Maximo Application Suite core services, installing all required entity managers and provisioning the primary resources detailed on this page.

ibm-truststore-mgr-controller-manager

`ibm-truststore-mgr-controller-manager` watches `Truststore.ibm-truststore-mgr.ibm.com` and manages all of the truststores in use in the core services namespace.

Entity managers

entitymgr-addons

`entitymgr-addons` add ons configuration.

entitymgr-bascfg

`entitymgr-bascfg` watches `BASCfg.config.mas.ibm.com`, manages the DRO integration with Maximo Application Suite.

entitymgr-coreidp

`entitymgr-coreidp` watches `CoreIDP.internal.mas.ibm.com`, manages the Core IDP component.

entitymgr-idpcfg

`entitymgr-idpcfg` watches `IDPCfg.config.mas.ibm.com`, manages IDP integration with Maximo Application Suite.

entitymgr-jdbccfg

`entitymgr-jdbccfg` watches `JDBCCfg.config.mas.ibm.com`, manages JDBC integration with Maximo Application Suite, performing configuration validation.

entitymgr-kafkacfg

`entitymgr-kafkacfg` watches `KafkaCfg.config.mas.ibm.com`, manages Kafka integration with Maximo Application Suite, performing configuration validation.

entitymgr-jdbccfg

`entitymgr-jdbccfg` watches `MongoCfg.config.mas.ibm.com`, manages Mongo integration with Maximo Application Suite, performing configuration validation.

entitymgr-objectstorage

`entitymgr-objectstorage` watches `ObjectStorageCfg.config.mas.ibm.com`, manages ObjectStorage integration with Maximo Application Suite, performing configuration validation.

entitymgr-pushnotificationcfg

`entitymgr-pushnotificationcfg`, watches `PushNotificationCfg.config.mas.ibm.com`, manages PushNotification integration with Maximo Application Suite, performing configuration validation.

entitymgr-scimcfg

`entitymgr-scimcfg` watches `SCIMCfg.config.mas.ibm.com`, manages SCIM (LDAP User sync) integration with Maximo Application Suite, performing configuration validation and resources creation such as `scimsync-agent` job and `scimsync` `liberty` pod.

entitymgr-slscfg

`entitymgr-slscfg` watches `SLSCfg.config.mas.ibm.com`, manages SLS integration with Maximo Application Suite, performing configuration validation and resources creation such as `licensing-mediator` pod. This pod is also responsible to register the SLS client in the SLS server.

entitymgr-smtpcfg

entitymgr-smtpcfg watches SMTPCfg.config.mas.ibm.com, manages SMTP integration with Maximo Application Suite, performing configuration validation.

entitymgr-watsonstudiocfg

entitymgr-watsonstudiocfg watches WatsonStudioCfg.config.mas.ibm.com, manages Watson Studio integration with Maximo Application Suite, performing configuration validation.

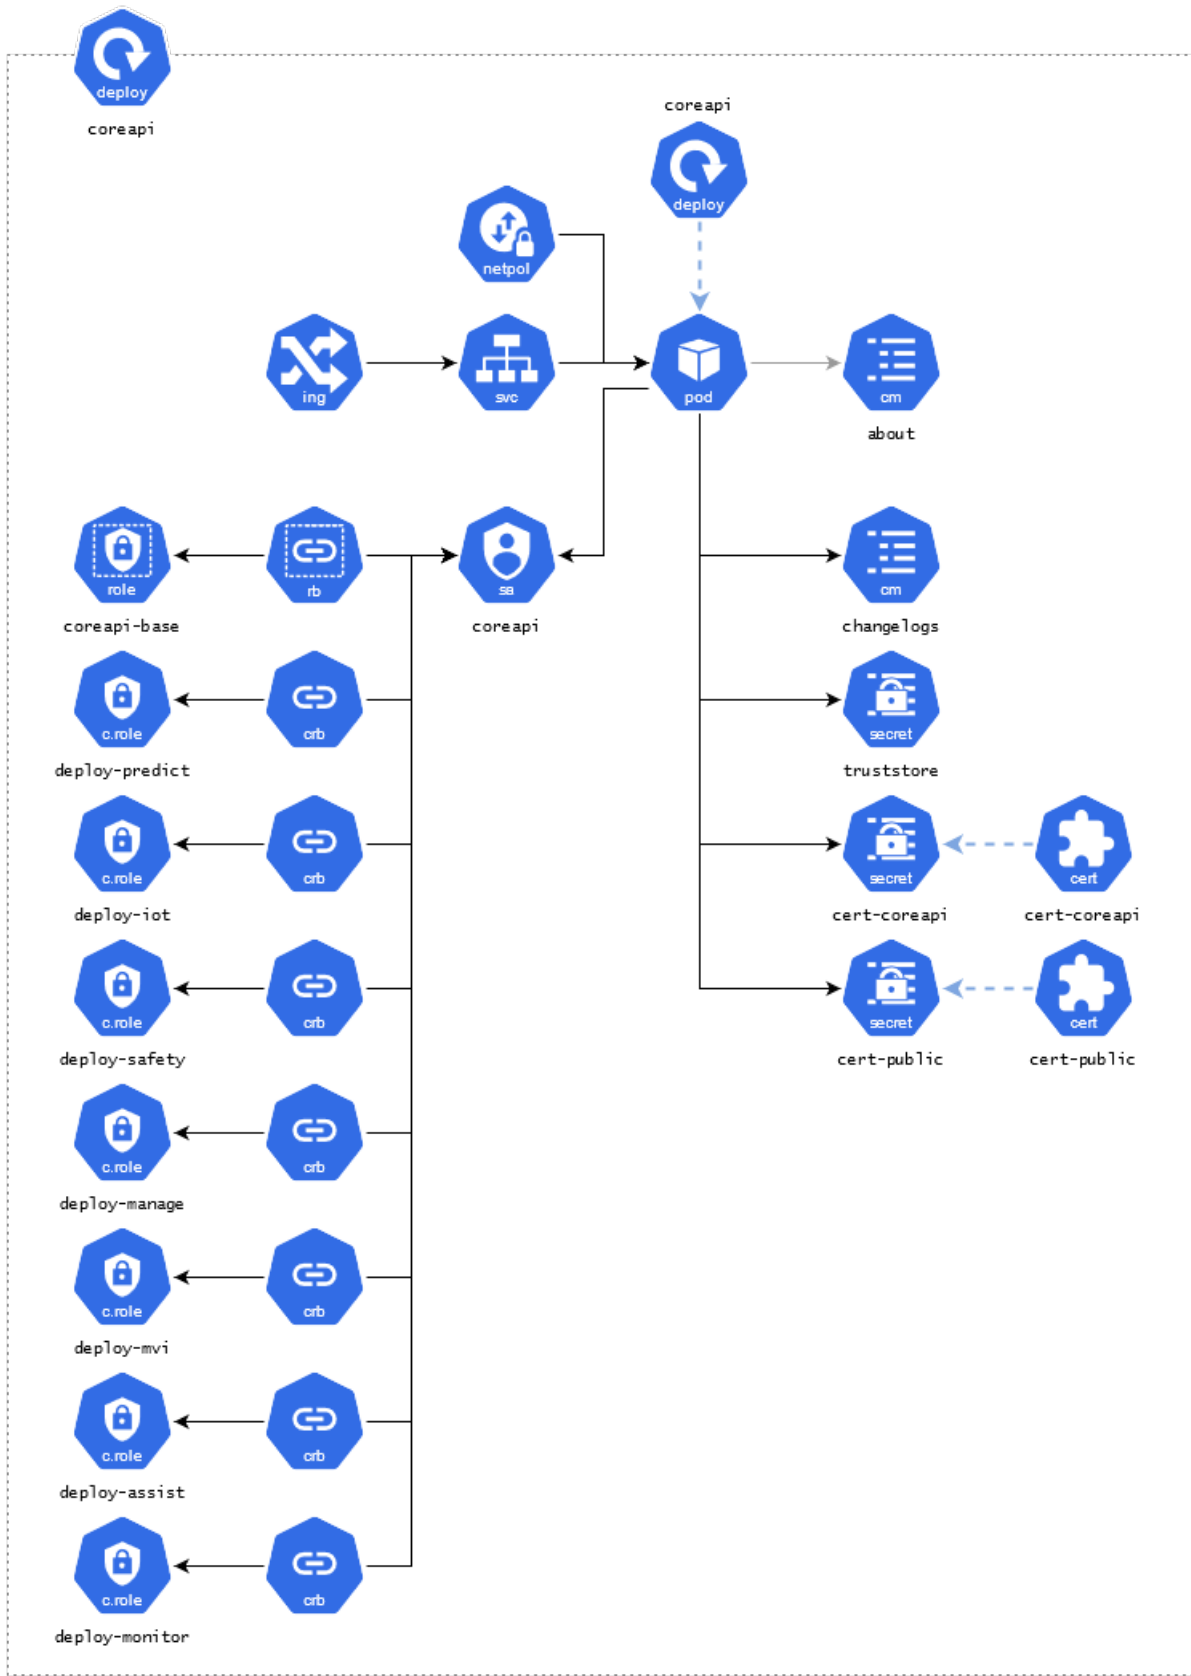
entitymgr-ws

entitymgr-ws watches Workspace.core.mas.ibm.com, manages Workspace creation in Maximo Application Suite.

Suite administration**Core API**

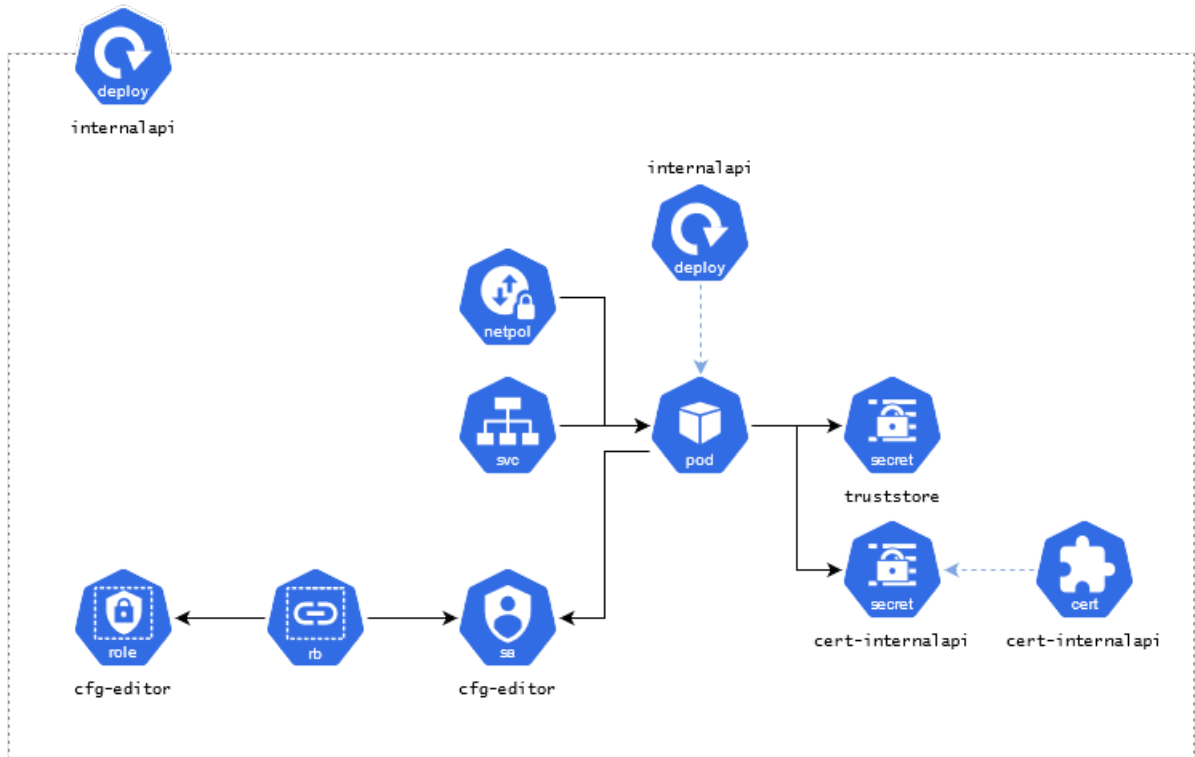
The coreapi deployment provides a RESTful API to support management of the Maximo Application Suite, as an alternative to working directly with Kubernetes resources natively. The API is made

available on the route `https://api.{masdomain}`.



Internal API

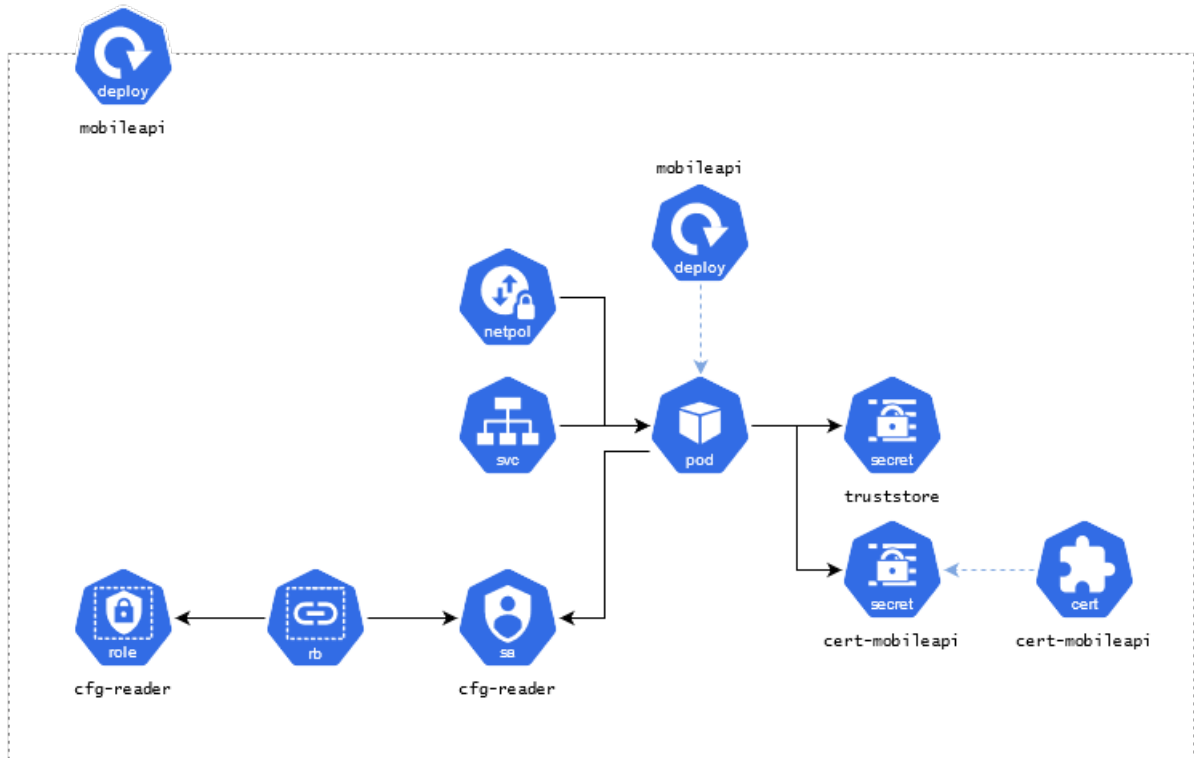
The `internalapi` deployment provides an internal API available to applications in the Maximo Application Suite, for example, user management. `internalapi` is used by internal components only. Application to application communication.



Mobile API

The `mobileapi` deployment provides the backend for the mobile application package API, serving up the navigator application package. Other implementations of the mobile API exist in each application

that supported mobile application packages, Core API controls access to these backend services.



Monitoring agent

The monagent-mas deployment is responsible for tracking the health of the core services. It reports status to the Maximo Application Suite status subresource broken down into three categories:

Mongo configuration status

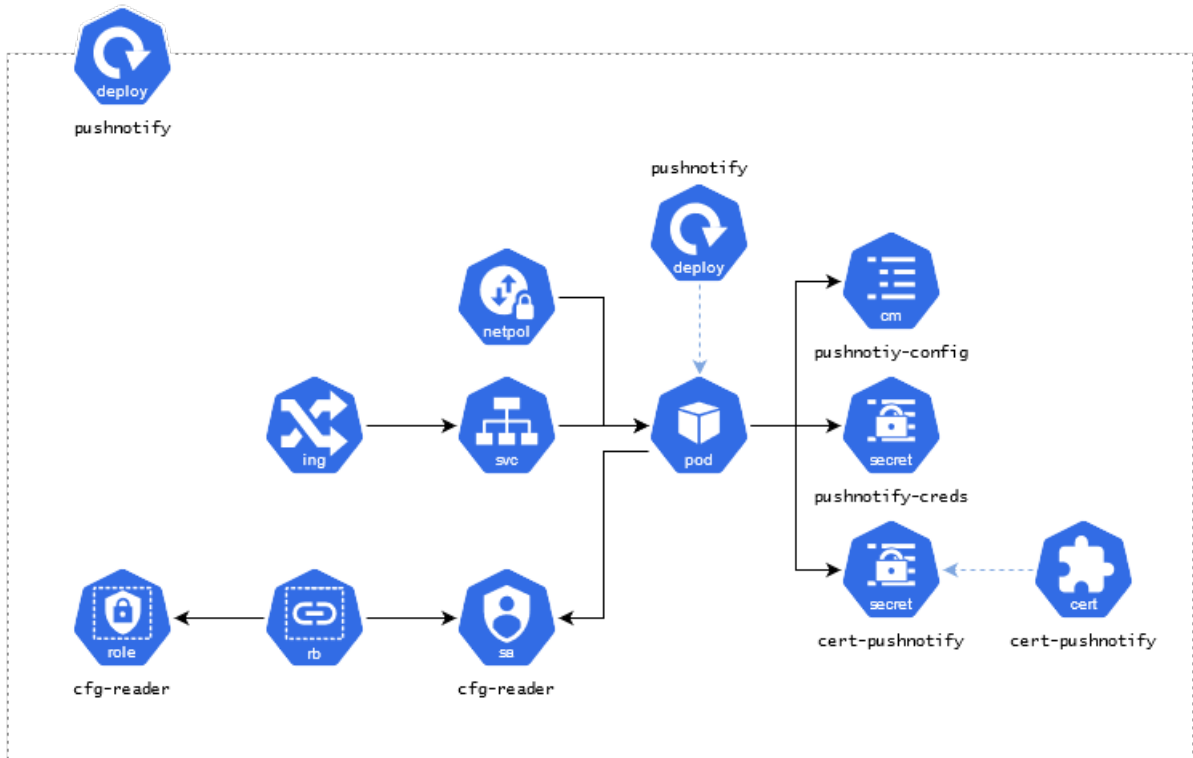
SLS integration status

DRO integration status

PNS integration

The pushnotifications deployment provides support for integration to an external push notification service (PNS). PNS support is an optional extension that is configured by a system administrator, enabling push notification support across Maximo Application Suite applications; it is available at <https://api.{domain}/pushnotification> only if the system scope

PushNotificationCfg resource has been created.



Identity provider

Core IDP

The `coreidp` deployment serves as the identity provider for all applications across Maximo Application Suite.

Core IDP login

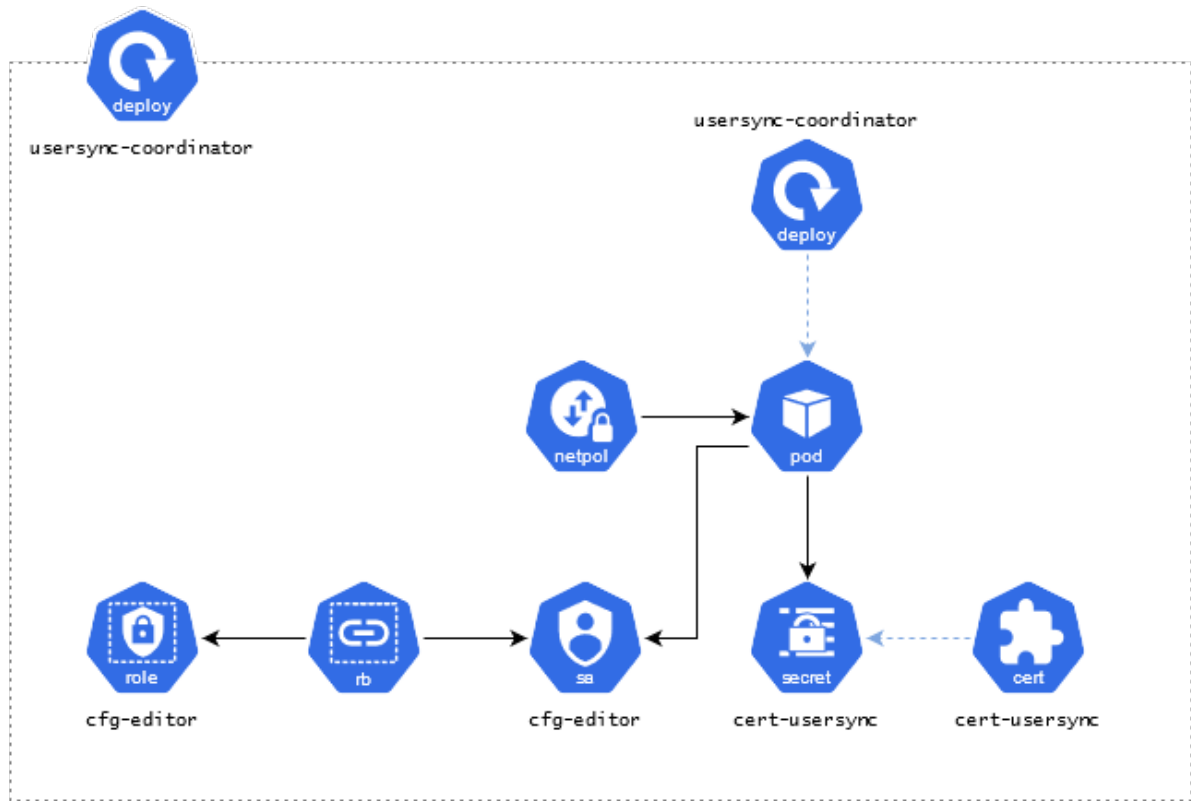
The `coreidp-login` deployment hosts the authentication login screens and superuser login logic for Maximo Application Suite. This is not meant to be a Maximo Application Suite endpoint but used as part of the redirect during Maximo Application Suite authentication flow. The service is available on the route `https://auth.{masdomain}`.

Group sync coordinator

The `groupsync-coordinator` deployment is responsible for coordinating user group synchronization across all installed applications.

User sync coordinator

The `usersync-coordinator` deployment is responsible for coordinating user synchronization across all installed Applications.



Catalog management

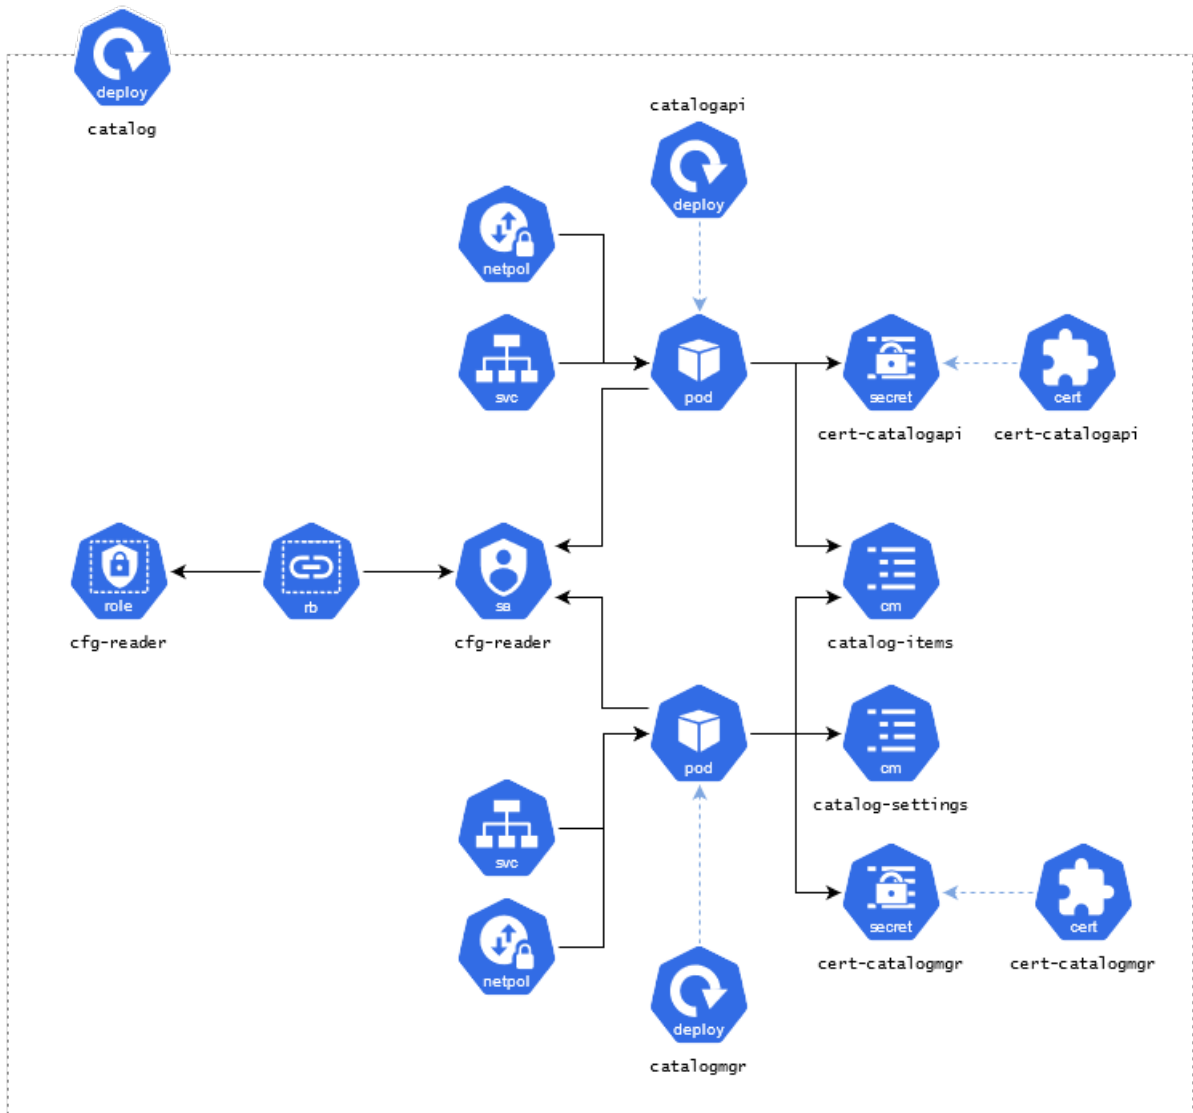
The Maximo Application Suite Catalog acts as the mechanism for customers to discover the resources in Maximo Application Suite that they are interested in, it allows the Maximo Application Suite to abstract the actual resource (that is, implementation) of a capability away from how it is presented as a catalog item.

Catalog API

The `catalogapi` deployment provides read access to the catalog inventory. The catalog API is exposed from the endpoints in [Core API](#) which proxy requests to the internal catalog API service.

Catalog manager

The `catalogmanager` deployment provides internal inventory management and AppPoint reservation APIs.



Console

Admin dashboard

The admin-dashboard deployment provides the administration console available on the route `https://admin.{masdomain}`.

Suite homepage

The homepage deployment provides the primary home screen for the Maximo Application Suite, available on the route `https://home.{masdomain}`.

Application navigator

The navigator deployment serves the application navigator available at `https://{workspace}.home.{masdomain}`.

Licensing and usage data collection

Account AppPoints reporter

The accappoints deployment submits events to IBM Data Reporter Operator (DRO) on an hourly basis. It obtains and converts data from the IBM Suite License Service.

- AppPoint reports are converted to account contractual usage events
- License usage reports are converted to account adoption usage events

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator ”](#) on page 7.

Adoption usage API

The `adoptionusageapi` deployment provides an internal API, which enables applications in the Maximo Application Suite to report metrics that are used to generate AppPoint and License usage reports.

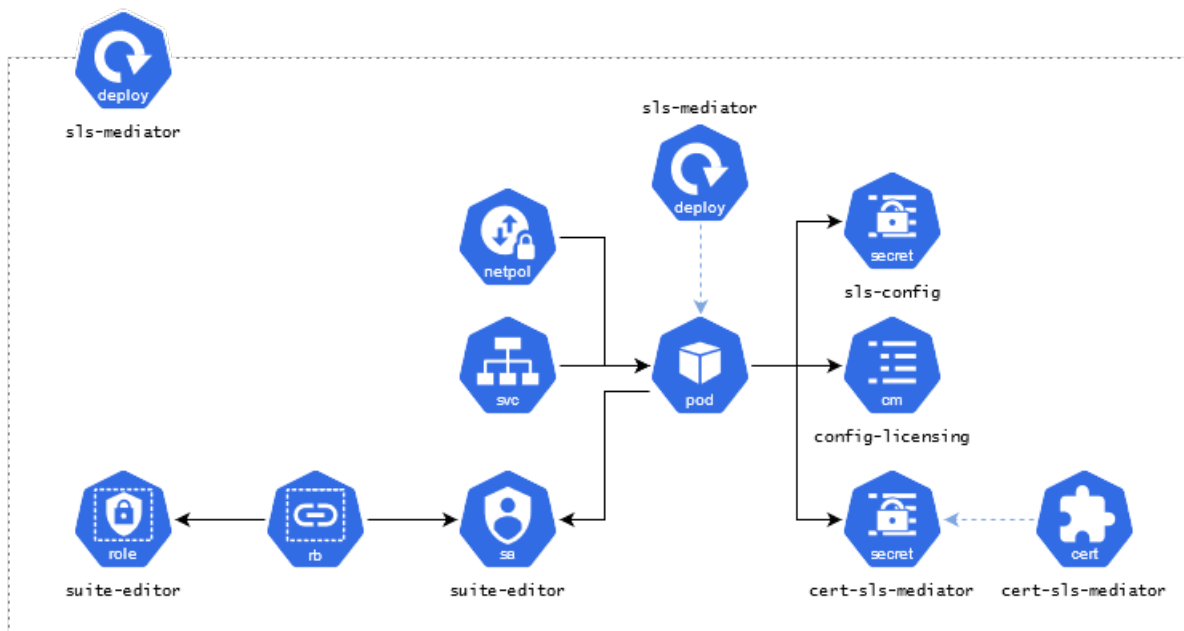
Adoption usage reporter

The `adoptionusage-reporter` deployment pulls the data that is related to adoption of different applications by users. It gathers data in terms of number of users and total AppPoints of these users who login to each of the Maximo Application Suite applications, whenever the users login to these applications.

This application runs as a cronjob and sends this data to Data Reporter Operator (DRO). DRO in turn sends this data to the IBM growth stack to provide IBM an insight into how customers are using the Maximo Application Suite.

SLS mediator

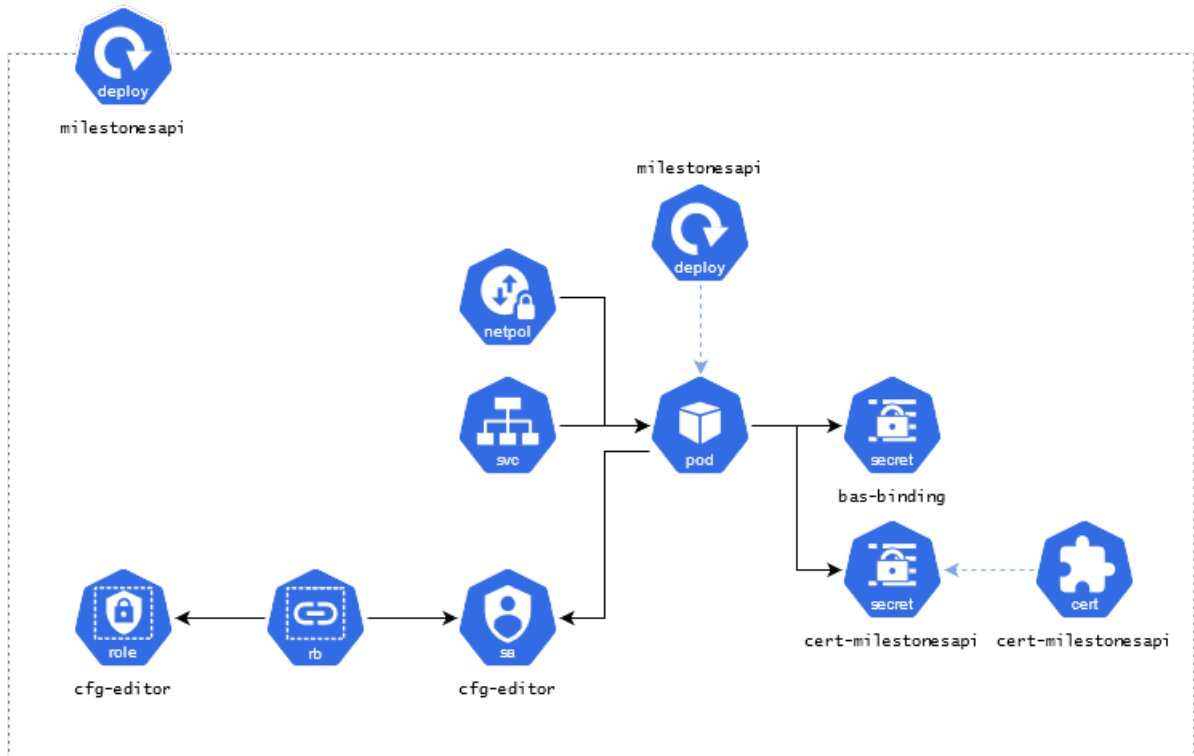
The `licensing-mediator` deployment provides internal APIs that act as the bridge between the Maximo Application Suite installation and the Suite License Service (SLS). It will also periodically run a synchronization process, which ensures licensing information in the user registry and in SLS are in alignment.



Milestones API

The `milestonesapi` deployment is responsible for reporting critical user events, which are known as "milestones", to Data Reporter Operator (DRO). DRO forwards these events into the IBM Growth Stack, which includes tools to help IBM gain insight into customer usage and assist with campaign

administration.



Maximo Application Suite pod details

Pods are used for IBM Maximo Application Suite and IBM Maximo Manage.

Maximo Application Suite core pods

ibm-mas-operator

The operator that manages Maximo Application Suite core. Watches the Maximo Application Suite CR specifications. This operator deploys the following pods:

- ltpakeysgenerator
- coreapi
- internalapi
- admindashboard
- homepage
- navigator
- usersynccoordinator
- groupsynccoordinator
- workspacecoordinator
- catalogmgr
- catalogapi
- mobileapi
- monagentmas

ibm-truststore-mgr-controller-manager

The operator subsystem that is responsible for handling the truststore request and adding the provided truststore in the format that is consumed by the servers.

{instance_name}-accappoints

A reporter pod that pulls data from IBM Suite License Service, converts AppPoints reports to Account Contractual Usage events and license usage reports to Account Adoption Usage events, and pushes the events to Behaviour Analytics Services (BAS) every hour.

For more information, see [“Administering licenses and AppPoints usage”](#) on page 813.

{instance_name}-admin-dashboard

The Maximo Application Suite core admin dashboard UI pod, which is the user interface for Maximo Application Suite administration for system administration work and user management.

The URL is `https://admin.{masdomain}/`

For more information, see [“Administering Maximo Application Suite ”](#) on page 722.

{instance_name}-adoptionusageapi

An API pod that enables the adoption usage process by providing an API, which is started by coreapi that is used to share information about users when they login to any application.

{instance_name}-adoptionusage-reporter

A reporter pod that pulls the data that is related to adoption of different applications by users. It gathers the number of users and total AppPoints of these users who log in to each of the Maximo Application Suite applications, whenever the users log in to these applications. This application runs as a cronjob and sends this data to Data Reporter Operator (DRO). DRO in turn sends this data to IBM growth stack including Segment and Amplitude where different IBM roles like Product management or Operations can understand adoption patterns of different Maximo Application Suite applications by using relevant dashboards.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator ”](#) on page 7.

{instance_name}-catalogapi

A pod that provides read access to the catalog inventory.

For more information, see [“Deploying applications, add-ons and industry solutions”](#) on page 291.

{instance_name}-catalogmgr

A pod that provides the inventory management and AppPoint reservation APIs.

{instance_name}-coreapi

The Maximo Application Suite CoreAPI pod that provides an HTTPS API to support working with Kubernetes resources. Primarily acts as a proxy to the Kubernetes APIs and MongoDB integration.

{instance_name}-coreidp

A Liberty-based server that handles the authentication, access management, and user privileges authorization that is managed by Maximo Application Suite. This pod integrates with internal application by using the OpenID Connect (OIDC) flow and SAML.

For more information, see [“Authentication methods”](#) on page 605.

{instance_name}-coreidp-login

The coreidp login pods that host the login page and superuser login logic for Maximo Application Suite. Not meant to be a Maximo Application Suite endpoint but used as part of the redirect during the Maximo Application Suite authentication flow.

The URL is `https://auth.{masdomain}/login/`

{instance_name}-entitymgr-addons

Add-ons configuration.

For more information, see [“Deploying add-ons”](#) on page 408.

{instance_name}-entitymgr-bascfg

The entity manager that is used to manage DRO integration with Maximo Application Suite. This pod generates internal certificates, performs configuration verification, and deploys all Data Reporter Operator related pods. This pod watches the **BASCfg** CR. This operator deploys the following pods:

- milestonesapi
- adoptionusageapi
- adoptionusagereporter
- accappoints
- adoptionusagemetering

For more information, see [Behavior Analytics Services](#).

{instance_name}-entymgr-coreidp

The entity manager that is used to manage Coreidp integration with Maximo Application Suite. Generates internal certificates, performs validations, and deploys coreidp-login and coreidp pods. This pod watches the Coreidp CR. This operator deploys the oidcclientregcoreidpcoreidplogin pod

{instance_name}-entymgr-idpcfg

The entity manager that is used to manage identity provider (IDP) integration with Maximo Application Suite. Generates internal certificates and performs configuration and verification. This pod watches SAML and LDAP configuration through the IDPCfg CR. This operator deploys the oidcclientregcoreidpcoreidplogin pod.

{instance_name}-entymgr-jdbccfg

The entity manager that is used to manage JDBC integration with Maximo Application Suite and perform configuration validation. This EntityMgr watches the JDBCCfg CR.

For more information, see [“Updates to JDBC configuration and other data resources are not applied in the Maximo Manage application”](#) on page 903.

{instance_name}-entymgr-kafkacfg

The entity manager that is used to manage Kafka integration with Maximo Application Suite and perform configuration validation. This EntityMgr watches the KafkaCfg CR.

For more information, see [Apache Kafka](#).

{instance_name}-entymgr-mongocfg

The entity manager that is used to manage Mongo integration with Maximo Application Suite and perform configuration validation. This EntityMgr watches the MongoCfg CR.

For more information, see [MongoDB](#).

{instance_name}-entymgr-objectstorage

The entity manager that is used to manage ObjectStorage integration with Maximo Application Suite and perform configuration validation. This EntityMgr watches the ObjectStorageCfg CR.

{instance_name}-entymgr-pushnotificationcfg

The entity manager that is used to manage PushNotification integration with Maximo Application Suite and perform configuration validation. This EntityMgr watches the PushNotificationCfg CR. This operator deploys the pushnotificationsservice pod.

For more information, see [Push notifications](#).

{instance_name}-entymgr-scimcfg

The entity manager that is used to manage the System for Cross-domain Identity Management (SCIM) integration with Maximo Application Suite and perform configuration validation and resources creation, such as scimsync-agent job and scimsync liberty pod. This EntityMgr watches the SCIMCfg CR. This operator deploys the scimsyncagentscimsync pod.

{instance_name}-entymgr-slscfg

The entity manager that is used to manage IBM Suite License Service (SLS) integration with Maximo Application Suite and perform configuration validation and resources creation, such as licensing-mediator pod. This pod is also responsible for registering the SLS client in the SLS server. This pod watches the SLSCfg CR. This operator deploys the licensing mediator pod.

For more information, see [“IBM Suite License Service”](#) on page 213.

{instance_name}-entymgr-smtpcfg

The entity manager that is used to manage SMTP integration with Maximo Application Suite and perform configuration validation. This pod watches the SMTPCfg CR.

For more information, see [“Simple Mail Transfer Protocol configuration” on page 634.](#)

{instance_name}-entymgr-watsonstudiocfg

The entity manager that is used to manage IBM Watson Studio integration with Maximo Application Suite and perform configuration validation. This pod watches the WatsonStudio CR.

For more information, see [IBM Watson Studio.](#)

{instance_name}-entymgr-ws

The entity manager that is used to manage workspace creation in Maximo Application Suite. This EntityMgr watches the Workspace CR.

{instance_name}-groupsync-coordinator

The group sync manager pod that coordinates the group sync activities between Maximo Application Suite and applications.

{instance_name}-homepage

The pod for the Maximo Application Suite core homepage UI. The main page when you are not running in a workspace-based endpoint.

The URL is `https://home.{masdomain}/`

{instance_name}-internalapi

The internal API pod that provides the internal version of the Maximo Application Suite Administrative API, for example, user management. Used only by internal components for application-to-application communication.

{instance_name}-licensing-mediator

Proxy server for that SLS service. Acts as an integration point between Maximo Application Suite and SLS.

{instance_name}-milestonesapi

The milestones API pod that is responsible for reporting critical user events, which are known as milestones, to Data Reporter Operator.

{instance_name}-mobileapi

The mobile API pods that provide the mobile application package API and serve the navigator application package. Integrates with coreapi.

{instance_name}-monagent-mas

A pod that is responsible for monitoring the health of general components in Maximo Application Suite and reporting to the Suite CR.

{instance_name}-navigator

A workspace-based UI that acts as a home page where users can access any application.

{instance_name}-usersync-coordinator

The user sync manager pod that coordinates the user sync activities between Maximo Application Suite and Applications.

For more information, see [“Administering users and user access in Maximo Application Suite in 9.0 and earlier” on page 796.](#)

workspace-coordinator ({instance_name}-workspace-coordinator)

The workspace sync manager pod that coordinates the workspace sync activities between Maximo Application Suite and applications.

Maximo Manage pods

ibm-mas-manage-operator

The overall main operator that is responsible for creating the following suboperators and the artifacts that are in the Manage application level.

- Workspace operator (entitymgr-ws)
- Application status operator (entitymgr-appstatus)
- BDI operator (entitymgr-bdi)
- User sync agent (usersyncagent)
- Group sync agent (groupsyncagent)

Uses the ManageApp CR.

ibm-truststore-mgr-controller-manager

The truststore manager that creates and updates various truststores that are consumed by Manage servers. Truststore manager is created when the Manage operator is installed.

admin-build-config (admin-build-config-#-build)

The Red Hat OpenShift BuildConfig that generates the image for the manage-maxinst pod. Can have the following statuses:

- Completed if the image is successfully built.
- Error if a problem occurred during the build process.

It is created by the ManageWorkspace operator during the reconciling of the ManageWorkspace.

Note: If any issues occur when the image is pulled, the log includes the details of downloading base images/customization archive.

ui/cron/mea/report-build-config-#-build

The job pod Red Hat OpenShift BuildConfig that generates the build liberty server images for the ii, cron, mea, and report server bundles. Can have the following statuses:

- Completed if the image is successfully built.
- Error if a problem occurred during the build process.

ManageWorkspace operator during reconciling of the ManageWorkspace.

Note: Usually the issue is captured by the admin-build-config pod unless it is a network issue.

{instance_name}-entitymgr-appstatus

The operator that tracks the status conditions for all the custom resources (CRs) that are managed by the Workspace operator. including the following CRs:

- ManageWorkspace
- ManageBuild
- ManageDeployment
- ManageServerBundle

{instance_name}-entitymgr-bdi

The operator that is responsible for deploying the necessary Kubernetes resources to run the build data interpreter (BDI) application, for example, Deployment, Service, NetworkPolicy, and Configmaps. Uses the BuildDataInterpreter CR that is created by the Manage workspace operator. The BDI is installed when Maximo Aviation or Maximo Asset Configuration Manager is installed as a component of Maximo Manage. A separate pod that runs independently of other Manage pods. Validates and reports configuration status of configuration-managed assets in Manage.

{instance_name}-entitymgr-ws

The operator that is responsible for creating, managing, upgrading, and deleting the Manage system in the workspace.

- Maximo Manage Database maxinst/updateDB
- The deployment of the server bundles MXServer running inside a Liberty container.
- Creates the Maximo Manage Build CR that builds the Admin Image Maxinst/Updatedb and the Websphere Liberty Profile images ServerBundle.

{instance_name}-groupsyncagent

Synchronize the user group into Maximo Manage from Maximo Application Suite. The valid Manage user will be added to the group by using the Maximo integration framework. The Maximo Enterprise Adapter or all server bundle is required for group sync tasks.

For more information, see [Security and user management](#).

{instance_name}-{workspace_name}-all/cron/report/mea/ui

Specifications that describe a Manage server bundle. Represents a set of Manage Liberty servers that load balance for a type of work load. The all server bundles is the Manage server bundle, which includes the following types:

cron

Components that are required for cron tasks.

mea

The enterprise web services API.

Note: This server bundle type is required for user/group synchronization unless the administrator selects the all server bundle for the Manage deployment.

report

Components for reports.

ui

The user interface components.

For more information, see [Server bundle overview](#).

{instance_name}-{workspace_name}-manage-maxinst

The Manage admin server, which is used for administrative tasks such as running the maxinst/updated B and installing languages. Contains the full SMP folder, which provides access to all binary files, dbc files, and admin tools utilities.

For more information, see [Administrative tasks for Maximo Application Suite](#).

{instance_name}-{workspace_name}-jmsserver-0

The Manage server bundle for stand-alone JMS components. For more information, see [Configuring a JMS server for Maximo Manage](#).

Only one JMS server is allowed. For more information, see [Scaling JMS servers](#)

{instance_name}-monitoragent

The Manage Monitoring agent that collects statistics about application usage and performance. Administrators can configure Prometheus service monitoring and then use visualization software, such as Grafana, to view the usage information that the agent collects.

For more information, see [Monitoring agent for Maximo Manage](#).

{instance_name}-usersyncagent

Synchronizes users from the Maximo Application Suite user registry to the user registry in Maximo Manage. Any action, such as create, update, or delete, for a Maximo Application Suite user who has access to Maximo Manage will be synchronized into Maximo Manage through the Maximo integration framework. The Maximo Enterprise Adapter or all server bundle is required to perform the user sync tasks.

For more information, see [User synchronization](#)

Enabling access for identity provider administration by using APIs

You can enable user access for identity provider (IdP) management in nonproduction instances by using the IDP_ADMIN API.

About this task

A user with IdP management access can configure LDAP and SAML authentication, SMTP, and user registry synchronization. This user can also customize the user interface and manage certificates. IdP management access is a subset of system configuration with fewer access privileges. For example, this access type does not include the ability to deploy or activate applications or configure database connections.

Procedure

1. Create an API key by using REST API so that you can authenticate Maximo Application Suite APIs. For more information, see [Maximo Application Suite APIs 8.11](#) or [Maximo Application Suite APIs 9.0](#)
2. Enable IDP_ADMIN access for the user by issuing a POST request for the /v3/users API.
3. Set up the access for local Identity Provider (IdP) by issuing a PUT request for the /v3/users/{userId}/idps/{idpId} API.

The following example shows the payload for the user creation where the permission key for idpAdmin is set as true

```
{
  "emails": [
    {
      "value": "test122@ibm.com",
      "type": "Work",
      "primary": true
    }
  ],
  "phoneNumbers": [],
  "addresses": [],
  "givenName": "",
  "familyName": "",
  "title": "",
  "displayName": "",
  "permissions": {
    "systemAdmin": false,
    "userAdmin": false,
    "apikeyAdmin": false,
    "idpAdmin": true
  },
  "entitlement": {
    "application": "NONE",
    "admin": "NONE",
    "alwaysReserveLicense": false
  },
  "username": "test122",
  "id": "test122",
  "owner": "local"
}
```

Results

After you enabled the access permissions, you can configure LDAP and SAML authentication, SMTP, and user registry synchronization. You also have the permissions to customize the user interface and manage certificates.

Glossary

air gap

A Red Hat OpenShift cluster that does not have internet connectivity. See also [“Disconnected” on page 926](#).

Ansible

An IT automation engine that automates cloud provisioning, configuration management, application deployment, intra-service orchestration, and many other IT needs.

For more information, see [Ansible](#).

Ansible collection

A distribution format for Ansible content that can include playbooks, roles, modules, and plug-ins. As modules move from the core Ansible repository into collections, the module documentation moves to the collections pages. You can install and use collections by using [Ansible Galaxy](#).

For more information, see [Using Ansible collections](#).

Ansible playbook

A configuration management system for defining and managing the tasks that are required to deploy complex applications.

For more information, see [Ansible playbooks](#).

Ansible role

An automation mechanism that automatically loads Ansible artifacts, such as related variables, files, tasks, and handlers. Grouping content in roles facilitates reuse and sharing.

For more information, see [Ansible roles](#).

Apache Kafka

A publish/subscribe distributed messaging system to manage messages that are sent to and received from external interfaces.

For more information, see [Apache Kafka](#).

Apache Spark

With IBM Analytics for Apache Spark for IBM Cloud, you can run jobs on an Apache Spark cluster.

- Run Jupyter Notebook and jobs from other tools in IBM Watson Studio analytics projects by selecting an Apache Spark environment runtime. Install this service either before or after you install the Watson Studio service.
- Run Spark SQL or jobs for data transformation, data science, or machine learning by using Spark job APIs. The Spark job APIs do not require the Watson Studio service.

application scope

The configuration that is set for and used by a single application, for example, the JDBC connection that is used by the application.

Amazon Web Services CLI

The command-line interface that is used for managing Amazon Web Services (AWS) services.

Amazon Simple Email Service (SES)

An email service that enables developers to send mail from within any application. You can configure Amazon SES for use cases that include transactional, marketing, or mass email communications. For more information, see [Amazon Simple Email Service](#).

bastion node

A specialized host that is exposed on a public network and configured to securely provide services to the network-isolated Red Hat OpenShift cluster, for example, a local image registry mirror.

Catalog source

A repository of CSV files, custom resource definitions, and packages that define an application.

certificate authority

A trusted third-party organization or company that issues digital certificates. The certificate authority (CA) typically verifies the identity of the individuals who are granted the unique certificate.

CloudFormation

A service in Amazon Web Services (AWS) that is used to model, provision, and manage AWS resources by using infrastructure as code. For more information, see [Amazon Web Services CloudFormation](#).

cloud service provider

A third-party company that offers a cloud-based platform, infrastructure, application, or storage services.

cluster

A group of nodes or machines in the Kubernetes infrastructure, which is composed of primary and secondary nodes.

common core services

A set of shared components that are used by multiple services in Cloud Pak for Data. For more information, see [Shared cluster components](#).

compute node

A Red Hat OpenShift Container Platform node on which IBM Maximo Application Suite application, and dependency workloads are scheduled. Also known as a *worker node*.

container

The basic unit of Red Hat OpenShift Container Platform applications.

control plane node

The Red Hat OpenShift Container Platform node on which key control plane services are scheduled. Also known as a *master node*.

custom resource

An endpoint in the Kubernetes API that stores a collection of API objects of a certain kind.

Db2 Warehouse

An analytics data warehouse that features in-memory data processing and in-database analytics. It is client-managed and optimized for fast and flexible deployment, with automated scaling that supports analytics workloads.

deployments

The process of installing and configuring a software application and all its components. The `Deployment` and `DeploymentConfig` API objects in Red Hat OpenShift Container Platform provide two similar but different methods for fine-grained management over common user applications.

Disconnected

A Red Hat OpenShift cluster without internet connectivity. See also [“air gap” on page 923](#).

Docker

An open platform that developers and system administrators can use to build, ship, and run distributed applications.

dynamic catalog

An curated operator catalog that is continuously updated. If you use the dynamic catalog, you always have access to the latest operator updates.

The IBM Maximo Application Suite team takes a snapshot of the online IBM operator catalog and tests compatibility of all dependent IBM operators with supported releases of IBM Maximo Application Suite, which allows the team to intercept any breaking changes before they reach your cluster. No updates are made to this catalog without extensive testing with all in-support version of Maximo Application Suite.

Elastic Compute Cloud (EC2)

A web service that is provided by Amazon Web Services to offer various types of computing capacity on the cloud.

Git

An open source program for source control management.

GNU bash

A shell that was created by the GNU Project. Also known as the *Bourne Again SHell*.

GPU

A specialized processor designed to rapidly manipulate and alter memory to accelerate the creation of images in a frame buffer intended for output to a display.

IBM App Connect

An integration solution that you use to connect your applications.

IBM Cloud Internet Services

A collection of services that provide reliability, performance, and security for Internet-facing applications, websites, and services by using Cloudflare. It includes Domain Name Service (DNS), Global Load Balancer

(GLB), Distributed Denial of Service (DDoS) protection, Web Application Firewall (WAF), Transport Layer Security (TLS), rate limiting, smart routing, and caching.

IBM Entitled Registry

A registry that contains images for the IBM Maximo Application Suite image, its applications, and components. The IBM Entitled Registry also contains images for other IBM products, such as IBM Cloud Pak for Data.

IBM operator catalog

A catalog of product offerings in the form of a catalog index image. To display the IBM offerings in the Red Hat OpenShift operator catalog, you must enable the IBM operator catalog image on your Red Hat OpenShift cluster by deploying a **CatalogSource** resource.

IBM Suite License Service

A token-based licensing system that uses MongoDB as the data store.

IBM Data Reporter Operator

An operator that accepts events and transforms them into reports that are submitted to the Data Service of the IBM Metrics Operator.

Note: Starting in IBM Maximo Application Suite 9.0, 8.11.7, and 8.10.10, the User Data Services (UDS) is deprecated and replaced with IBM Data Reporter Operator (DRO).

For more information, see [“Data Reporter Operator ” on page 7](#).

IBM Cloud Pak for Data

A platform that integrates software components for data analysis and organization connecting siloed data distributed across a hybrid cloud landscape. Deployment options include an on-premises software version that is built on the Red Hat OpenShift Container Platform or a fully managed version that is built on IBM Cloud .

IBM Cloud Pak for Data foundational services

Common services , such as the IBM Certificate Manager, that are used by Maximo Application Suite and its dependencies.

IBM Watson Discovery

An AI-powered intelligent search and text-analytics platform that helps you find valuable information that is buried in your enterprise data.

IBM Watson Machine Learning

A full range of tools and services so that you can build, train, and deploy machine learning models.

IBM Watson OpenScale

An enterprise-grade environment for AI applications that provides your enterprise visibility into how your AI is built and used Its open platform enables businesses to operate and automate AI at scale with transparent, explainable outcomes that are free from harmful bias and drift.

IBM Watson Studio

An environment and tools for work on data to solve your business problems.

IBM Cloud Pak for Business Automation

IBM Cloud Pak for Business Automation assembles certified software from the [IBM Automation Platform for Digital Business](#) on multiple cloud infrastructures. A private cloud vendor can be used as an enabling layer with a user interface and command line to limit access to members of an enterprise and partner networks.

image IDs

A Secure Hash Algorithm (SHA) code that can be used to pull an image. An SHA image ID cannot change. A specific SHA identifier always references the same container image content, for example, `docker.io/openshift/jenkins-2-centos7@sha256:ab312bda324`.

image registry

A content server that can store and serve container images.

image repository

A collection of related container images and tags that identify them.

images

A binary file that includes everything that is needed to run a single container, including the metadata that describes its needs and capabilities. Containers in Red Hat OpenShift Container Platform are based on container images that are formatted for OCI or Docker.

image tags

A label that is applied to a container image in a repository to identify a specific image as distinct from other images in an image stream. Typically, the tag represents a version number.

jq

A command line JSON processor that filters, maps, and transforms JSON data.

Kubernetes

An open source orchestration tool for containers.

Let's Encrypt

An automated and open certificate authority from the Internet Security Research Group (ISRG).

MongoDB

A document database that is designed for ease of development and scaling. Maximo Application Suite uses MongoDB for local user management and data dictionary.

For more information, see [MongoDB](#).

namespace

A virtual cluster within a Kubernetes cluster that can be used to organize and divide resources across multiple users.

See also [“project” on page 929](#)

nodes

A virtual or bare-metal machine in a IBM Cloud Kubernetes Service cluster.

For more information, see [Overview of nodes](#).

Object storage

An approach to addressing and manipulating data storage as discrete units, called objects. Objects are kept inside a single repository and are not nested as files inside a folder inside other folders. Also known as *object-based storage*.

operator

A component of Red Hat OpenShift Container Platform that is used for packaging, deploying, and managing services on the control plane.

For more information, see [Operators overview](#).

pod

A group of containers that are running on a Kubernetes cluster. A pod is a runnable unit of work, which can be either a stand-alone application or a microservice. For more information, see [Using pods](#).

project

A mechanism in Kubernetes for isolating groups of resources in a single Kubernetes cluster. Used interchangeably with *namespace*. See also [“namespace” on page 929](#)

Red Hat OpenShift Container Platform

A platform for developing and running containerized applications.

route

A mechanism that exposes services by assigning hostnames. For more information, see [Route configuration](#).

Simple Storage Service (S3)

An object storage service that is offered by Amazon Web Services.

security context constraint

A mechanism that allows Red Hat OpenShift administrators to control pod permissions. For more information, see [Managing security context constraints](#).

Service Binding Operator

An operator that manages the data plane for applications and backing services. Service Binding Operator reads data that is made available by the control plane of backing services and provides the data to applications based on rules that are provided by the service binding resource.

static catalog

A fixed reference point that does not change, which allows for reproducible installations.

To receive security updates and fixes, you must periodically update the static catalog versions that you installed on the cluster. After you update the catalogs, all operators that you installed from the catalog are automatically updated to the newer version.

Strimzi

Strimzi provides a way to run an Apache Kafka cluster on IBM Cloud Kubernetes Service in various deployment configurations.

For more information, see [Strimzi](#)

suite instance

An installation of IBM Maximo Application Suite that is referenced by a unique name or ID in the Red Hat OpenShift cluster.

For example, in the `mas-myinstance-core` namespace or project that is created to install IBM Maximo Application Suite, `myinstance` is the instance name or ID. You might have multiple instances in the same Red Hat OpenShift cluster that are referenced by different unique names or IDs.

system scope

A configuration that is set for and can be used across the whole suite. For example, a JDBC configuration can be used by all applications in the IBM Maximo Application Suite.

Terraform

An open-source infrastructure as code software tool that enables you to create, change, and improve infrastructure. For more information, see [Terraform](#).

VMware vSphere

A virtualization platform that is used for server virtualization.

worker nodes

A group of nodes or machines in the Kubernetes infrastructure.

Workspace

An aggregation of namespaces.

workspace application scope

A configuration that is set for and used by a single application in the default workspace., for example, the JDBC connection that is used by the application in the default workspace.

workspace scope

The configuration that is set for and used in the default workspace, for example, the JDBC connection that can be used by all applications in the suite.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation*

North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Red Hat® and OpenShift® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Privacy Statement

IBM Software products, including software as service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user's name, user name, password, or other personally identifiable information for purposes of session management, authentication, single sign-on configuration or other usage tracking or functional purposes. These cookies can be disabled, but disabling them will also likely eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek

your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see the IBM Privacy Statement at <http://www.ibm.com/privacy> in the section entitled “Cookies and Similar Technologies”.

Index

A

accessibility [128](#)

