Administer

Operations Bridge Suite 2017.11

MICRO FOCUS

Administer

This section describes administration tasks that you can perform on the Management Portal of the Container Deployment Foundation. On the Management Portal, you can manage the shared services infrastructure and all suite products, including the Operations Bridge Suite deployment and configuration.

To access, open `https://<external_access_host>:5443` in a supported web browser and provide the administrator password.

For more information on the Container Deployment Foundation administration, see Administer the Container Deployment Foundation.

For more information on the suite administration, see Administer the Operations Bridge Suite.

Administer the Container Deployment Foundation

You can perform the following tasks to administer the Container Deployment Foundation:

Access Kubernetes API server with a bearer token

Access the Management Portal

Add or remove machines from a cluster

Administer IdM

Change your password

Customize kubelet parameters

Edit hard eviction thresholds

Manage licenses

Manage nodes

Manage resources

Modify CDF's external database

Monitor infrastructure status

Network and communication

Restart CDF

Security

Set up a host name resolution using a hosts file

View existing images

Access Kubernetes API server with a bearer token

A bearer token file for accessing the Kubernetes API is a csv file with at minimum of 3 columns: token, user name, and user uid. You can add more groups when needed by adding extra columns

and double quoting the group names, for example, **"group1"**. The rows of the csv file list the information of different tokens.
The token authentication is disabled by default. You can enable the token authentication with the following steps.

1. Run the following commands:
2. cd {K8S_HOME}/runconf
   vim kube-apiserver.yaml
3. Add the specified token directory to the --token-auth-file opinion line.
   For example:

```
--token-auth-file=<your token directory>/token
```

4. Restart kubelet with the following commands:

cd {K8S_HOME}/bin
./kube-restart.sh
If you have multiple master nodes, you must use the same bearer token file for every node.
To use the bearer token authentication via an HTTP request, you must pass the value of the bearer token to the HTTP header.
The bearer token must be in character sequence, using no encoding or quoting. For example: a bear token is 31ada4fd-adec-460c-809a-9e56ceb75269. When adding the bearer token to an HTTP header, it should look like this:
Authorization: Bearer 31ada4fd-adec-460c-809a-9e56ceb75269

Access the Management Portal

To access the management portal, do the following:

1. Launch the management portal from a supported web browser:

   https://<*external_hostname*>:5443

   <*external_hostname*> is the fully qualified domain name of the host which you specified in the Connection step during the CDF configuration. Usually, this is the master node's FQDN. As a result, the ITOM Suites login screen should be displayed.
2. Log in to the Management Portal as the admin user.
   Use the password that you specified at initial login.

Add or remove machines from a cluster

You can edit your existing installation by adding or removing machines and by editing the hard eviction thresholds of the worker nodes.

Add more machines to the Kubernetes cluster

To add more machines to the existing Kubernetes cluster, see Add nodes.

Remove machines from the Kubernetes cluster

You can remove machines from the existing Kubernetes cluster as follows:

1. Log on the machine that you want to remove.
2. Go to the installation directory, and run the uninstall command:

```
cd <K8S_home>
./uninstall.sh
```

Administer IdM

The IdM Administration provides the identity management services for CDF. It helps to manage users, groups, roles and single sign-on (SSO) to allow users the access to multiple applications with the same user name and password.

You can access the IdM Administration page by clicking **ADMINISTRATION > IdM Administration.**

Click **System Settings** on the top right to modify the IdM configuration (will apply to all organizations).
To prolong the IdM request token time, and the management session period, set the Request Token Life Time and Access Token Lifetime tag respectively.

See the details about the basic system settings in the table below.

HPSSO	Value	The domain name is listed in the value tag. LW-SSO supports a single domain. All the servers using LW-SSO must have the same domain.
	Initial String	The key for the encryption of LW-SSO. This is the shared secret of all servers protected by LW-SSO and connected to the same authentication point server. The initial string must be the same for all the servers. The minimum length of the initial string is 32 bits.
TOKEN	Encrypted signing key	Keys used to calculate the message digest to validate the message integrity.
	Request Token Life Time	IdM request token life time in minutes.
	Access Token Lifetime	IdM token life time in minutes. Users can change the Access token lifetime to prolong the life time of the management portal.
SMAL	Entity Base URL	The entity ID of IdM's SAML metadata will be based on this URL.
	Keystore Path	Keystore path for SAML and WS-Trust.
	Keystore Default Key Name	Keystore default key name for SAML and WS-Trust.

	Keystore Default Key Password	Keystore default password for SAML and WS-Trust.
	Keystore Password	Keystore password for SAML and WS-Trust.
	Keystore Provider	Keystore provider for SAML and WS-Trust.
	Keystore Type	Keystore type for SAML and WS-Trust.
LDAP	Extended attributes	Properties for the LDAP configuration.
	Nested Group Level	IdM LDAP nested group level

You can click **Advanced** to show the advanced settings.

Add Organization

From **IdM Administration > Add**, you can create an organization.

1. Enter the following information for a new organization:
   Name, Display Name, Integration User and Password
2. Click **Create**.

Delete Organization

From **Organization > Delete Organization**, click **Delete** to delete the organization.

Configure LW-SSO

- Set up single sign-on with other products
- Configure a customized timeout for the IdM token

Note
The Initial String and Creation Domain of LWSSO have their own default values. You can change these default values according to your needs.

To set up single sign-on with other products, follow these steps:

1. Click **ADMINISTRATION** > **IdM Administration > System Settings**.
2. Enter or edit the Initial String and Creation Domain, and then click **Save**.

Tip
You can copy and paste the value of the InitString directly into other products to set up the LW-SSO integration.

To configure a customized timeout for the IdM token, edit the Access Token Lifetime parameter and enter a specific value (expressed in minutes), and then click **Save**. The default session timeout value is 30 minutes.

You must restart the IdM pods after you configure the timeout for the IdM token. To restart the pods, go to **UI Resources > "core" Namespace > Workloads > Pods** and delete the IdM pods by clicking **Actions > Delete**.

Customize password policy

You can customize your password policy for the organizations.

To add a password policy, click **Add Password Policy** from the Password Policy tab. Enter the password policy name, lockout check time, length check time, expiration check time and additional policy checks. Then click **Save**.

To remove the policy, click **Remove**.

Customize the organization

The Customization tab allows you to add or edit the generic key pair parameters for an organization.

Add Groups Into LWSSO Cookie	N	Specify whether to enable the addition of groups into the LW-SSO cookie.
Add Permissions into LWSSO Cookie	N	Specify whether to enable the addition of permissions into the LW-SSO cookie.
Add Roles into LWSSO Cookie	N	Specify whether to enable the addition of roles into the LW-SSO cookie.
Authentication Flow	Y	Specify the authentication flow. For example, seeded, database_user, ldap, ad, jaas, aml, cac, and iwa.
Disclaimer Text	Y	Specify whether the portal displays a disclaimer text.
Featured Category	N	Specify if the category should be featured.
Languages	N	Specify whether the portal should multiple languages.
Order Recipient Enabled	N	Specifiy whether the recipient is in order.
Portal End Date	N	The portal end date.
Portal Enforce End Date	N	Specify whether the portal has the enforce end date.

Portal Footer Message	Y	Specify whether the portal has a footer message.
Portal Legal Notice URL	Y	Specify whether the portal has a legal notice URL.
Portal Show Confirm Dialog	Y	Specify whether the portal shows the confirmation dialog.
Portal Show Legal Notice	Y	Specify whether the portal shows the legal notice.
Portal Show Terms Of Use	Y	Specify whether the portal shows the terms of use.
Portal Terms of Use URL	Y	Specify whether the portal terms use URL.
Portal Title	Y	Specify the portal title.
Portal Welcome Message	Y	Specify the portal welcome message.
Security Level	N	Specify the security level of your metadata.
Theme Name	N	Specify the theme name of your metadata.

You can do the following operations to the key pair parameters:

- Add: Click **Add** to add a new key pair parameter.
- Edit: Click **Edit**  to edit a key pair parameter. Enter the value of the Key and click **Save** to save the modification.
- Remove: Click remove  and then confirm to remove one key pair

Customization for Localization

To show the messages in a local language, you can add the language suffix from the table below to the key of Portal Footer Message and Portal Welcome Message. Then add the value in the local language.

Swedish	.sv
Spanish	.ar
Russian	.ru
Japanese	.ja

Italian	.it
German	.de
French	.fr
English(US)	.es
English(UK)	.en
Chinese	.zh

Note
To implement the changes for the language localization, you need to log out of the management portal and then log back into the management portal.

Manage users, groups, and roles

Manage users

This section provides information on how to manage a user.

To manage users, click **ADMINISTRATION** > **IdM Administration,** select the organization name, then click on the Users tab. This page displays the email address and user group for each user.

- To create a user, click **ADD**. Enter user name and password. Click **Add Attributes** to add user attributes. Then click **Save**.

Note
You can choose if you want to enter a password for the user. Users with password are IdM internal users. Users without password are from other authentication flows, such as LDAP, SAML or JAAS. You can add passwords for those users from other authentications to create an internal IdM user with the same user name. To delete an internal user, you can just delete the password.

- To delete a user, click **Remove**, and then click **Remove** again to confirm the deletion.
- To edit or lock a user, click **Edit**.

Change a user's password

To change a user's password, see Change your password.

Manage groups and roles

Adding groups helps to manage the roles and permissions that are assigned to the users. You can do the following:

- Add: Click **Add** to add a new group.

- Edit: Click **Edit**  to edit the settings of a group. Enter a group name, add associated group rules and associated roles of the group and click **Save** to save the modification.
- Remove: Click Remove, and then click confirm to remove a group.

Adding roles to a user helps to manage the permissions assigned to users. You can do the following:

- Add: Click **Add** to add a new role.
- Edit: Click **Edit** to edit a group's settings. Enter a role name, description of the role and the associated permissions. Click **Save** to save the modification.
- Remove: Click **Remove** and then confirm to remove a role.

Change your password

To change your password, follow these steps:

1. Click **ADMINISTRATION** > **IdM Administration**
2. Select the organization name, then click on the users tab and edit your user entry.
3. Click **Reset/Delete Password** to reset the password.
4. Enter a new password, and confirm the new password.
   The password should meet the password policy if you have set one password policy in the IdM Administration.
5. Click **OK**.

Customize kubelet parameters

You can modify the default values of the kubelet parameters and add some customized parameters for kubelet. Follow the steps below to customize the parameters.

1. Log on to any of the cluster node.
2. Edit or add the parameters in the kubelet.service under the /usr/lib/systemd/system directory.
3. Restart the kubelet with the following commands:

```
systemctl daemon-reload
systemctl restart kubelet
```

Edit hard eviction thresholds

CDF uses a hard eviction policy for worker nodes. When a hard eviction threshold is met, Kubernetes ends the pod immediately. The eviction can also delete dead pods, dead containers, and unused images when the disk space reaches the thresholds.

To edit the hard eviction threshold, follow these steps:

1. Log on to the worker node for which you want to edit the eviction threshold.
2. Edit the relevant parameter values in the /usr/lib/systemd/system/kubelet.service file.

For example, you modify the following default thresholds, according to your needs:

```
vim /usr/lib/systemd/system/kubelet.service--
evictionhard=memory.available<500Mi,nodefs.available<5Gi,imagefs.available<5Gi--system-
reserved=memory=1.5Gi
```

3. Run the following commands to enable the new thresholds:
   systemctl daemon-reload
   systemctl restart kubelet

Manage licenses

The License page enables you to manage your suite licenses.
This section includes the following tasks:

- View existing licenses
- Install licenses
- Archive licenses
- Restore archived licenses
- Delete licenses from the License Manager
- View the Licenses Report

View existing licenses

Click **ADMINISTRATION** >**License** > **View Licenses**.

Select the relevant product in **Select Product**. The page displays the license's feature ID and version, product number, capacity, start date, expiry date, the date when it was installed, who installed it, and the Lock Code.

Install licenses

1. Click **ADMINISTRATION** >**License** > **Install Licenses**.
2. Click **Choose file** to select the license file in your local system.
3. Click **Add More Files** to select another license file in your local system.
4. Click **Next**.

The licenses that have been installed are displayed.
You can select the license keys and click **Install Licenses** to install the licenses.

Archive licenses

1. In the **View Licenses** tab, select the unused licenses you want to archive.
2. Click **Archive**.

Restore archived licenses

1. In the **Archived Licenses** tab, select the product whose archived licenses you want to restore.
2. Select the licenses that you want to restore.
3. Click **Restore**.

The licenses are again displayed in the License Management pane and customers can check them out.

Note

If ID locked licenses are auto archived, they cannot be restored unless all the licenses locked to a lock value belonging to same feature are either deleted or archived.

Delete licenses from the License Manager

1. In the **Archived Licenses** tab, select the product whose licenses you want to delete.
2. Select the license to delete.
3. Click **Delete** and confirm the deletion.

View the Licenses Report

Click **ADMINISTRATION** > **LICENSE REPORT**.

The license report page tracks and displays the licenses currently installed and used on the License Manager. It also displays specific check out information about a feature license including the product name and version, the requester ID, and the timestamp of when it was accessed last.

You can export the license report details to Excel.

Manage nodes

The Nodes page provides the CPU and Memory usage history of the selected Namespace, a list of the predefined labels, and the list of nodes of the selected Namespace.

Tip

When the CPU load is over 80%, it significantly impacts the efficiency of network transmission between the base infrastructure environment. We recommend to control the CPU load so it is less than 80% by separating the suite instance into multiple worker nodes: adding more worker nodes and killing the pods on heavy-load nodes and deploying those pods on the newly added worker nodes.

View the existing nodes

1. Click **ADMINISTRATION** > **Nodes**.
2. The area displays the CPU and memory usage of the selected namespace during the past 15 minutes, the list of the node labels, and the status, labels, readiness, and creation timestamp of the nodes corresponding to the selected namespace.

You can do the following:

- Define a set of labels you want to use and then assign them to nodes by dragging them to the node. See Add/delete labels and Assign labels to nodes.
- Add a node. See Add nodes .
- **REFRESH.** Click to refresh the display.

- Click the relevant node to see its details. See View the node details .

Add/delete labels

1. Click **ADMINISTRATION** > **Nodes**.
2. To add a label in the **Predefined Labels** area, enter the **value** and click **[+]**. The label is added to the list.
3. To delete a label: in the **Predefined Labels** area, click **[-]** for the relevant label.

Assign labels to nodes

1. Click **ADMINISTRATION** > **Nodes**.
2. **To assign a label to a node:** drag the relevant label the **Predefined Labels** area to the relevant node in the **Nodes** area.
3. **To create a new label and assign it to a node:** in the relevant node row, click **[+]** below the list of labels, enter the **key** and click **OK**. You do not need to add the **value** of the label.
4. **To unassign a label:** in the **Nodes** area, click **[-]** for the relevant label and node.
5. **To filter the labels:** enter the relevant string or keyword in the Labels box in the table header. The labels with names that include the relevant string are listed.

Add nodes

1. Click **ADMINISTRATION** > **Nodes**.
2. In the Nodes area, click **+ ADD**.

   Enter the following information:

   - The node's host name
   - Your username
   - Your password
   - The *THINPOOL_DEVICE* parameter specifies the path to the Docker devicemapper storage driver.
   - The *FLANNEL_IFACE* parameter specifies the interface for Docker inter-host communication as a single IPv4 address or interface name. This parameter is used when the nodes have more than one network adapter so that Flannel can set up the correct routing table entries.

Tip

You can add multiple nodes simultaneously with **+ ADD**:

- Enter multiple host names or IP addresses separated by a space.
- Enter the user name and password.

The added nodes share the same user name and password. The installation of each node runs in parallel.

View the node details

1. Click **ADMINISTRATION** > **Nodes**.
2. In the Nodes area, select a node name from nodes list.

The page displays the CPU and memory usage history of the selected node for the past 15 minutes.

The **Details** area displays details about the selected node as well as system information.

The **Allocated resources** area displays the minimum CPU requests, CPU limits, memory requests, and memory limits for the container as well as the percentage of <what is in use>/<what is available>. By default, pods run with unbounded CPU and memory limits. The format is: <what is in use>/<what is available>.

The **Conditions** area displays the type, status, last heartbeat and transaction time, reason, and message.

The **Pods** area displays the CPU and memory usage history of the pod for the past 15 minutes, the name of the pod, the status, number of restarts in the cycle, the amount of time passed since the pod has been created, the cluster IP, as well as the CPU and memory usage of the pod.

You can do the following:

- Click a Pod name to open the Workloads - Pods page for the pod.
- Click the menu icon to review the pod log.
- Click **Actions** and select **Delete** to delete the pod.

The **Events** area displays the message, source, sub-object, count, first seen, and last seen information.

Manage resources

The **Resources** menu enables you to deploy containerized applications to a Kubernetes cluster, troubleshoot them, and manage the cluster and its resources itself. You can use it to get an overview of applications running on the cluster, as well as for creating or modifying individual Kubernetes resources and workloads, such as Daemon sets, Pet sets, Replica sets, Jobs, Replication controllers and corresponding Services, or Pods.

It also provides information on the state of Pods, Replication controllers, etc. and on any errors that might have occurred. You can inspect and manage the Kubernetes resources, as well as your

deployed containerized applications. You can also change the number of replicated Pods, delete Pods, and deploy new applications using a deploy wizard.

For more information, see the following topics:

- Manage namespaces
- Manage workloads
- Manage services and ingresses
- View persistent volume claims
- Manage secrets and config maps

Manage namespaces

This section provides details about the selected Namespace.

Kubernetes supports multiple virtual clusters backed by the same physical cluster. These virtual clusters are called namespaces.

Select the namespace

You select a namespace to filter the information in the pages of the UI and display only the items related to the namespace.
Click **RESOURCES** > **Namespace** and select the relevant namespace.
Resources will be displayed filtered by the specific namespace.
The page shows the CPU and memory usage history for the selected namespace, for the past 15 minutes, the name of the namespace, its labels, pods, the timestamp of the creation of the namespace and its images.
Click the relevant namespace to display more details.

View the namespace details

Click **RESOURCES** > **Namespace** and select the relevant namespace. You can also click **Workloads** > **Namespaces**, and click the relevant namespace.
The page shows details about the namespace and details about the events occurring in the core such as messages, source, count, first seen and last seen.

View persistent volume claims

A persistent volume claim is bound to a persistent volume. The claim is subsequently used inside a container volume specification. This provides volume technology abstraction for the suite deployment, as suites request size and access type rather than a certain specific storage provider. Each suite will have at least one persistent volume but may have more depending on the suite.

A volume is a directory, possibly with some data in it, which is accessible to the containers in a pod.

View the Persistent Volume Claims

Click **RESOURCES** > **Persistent Volume Claims**.

The page displays the name of the persistent volume, the volume it belongs to, the labels, and the timestamp of the creation of the persistent volume.

Each suite will have at least one persistent volume but may have more depending on the suite.

You can click the relevant volume to display its details.

View the details of Persistent Volume Claims

Click **RESOURCES** > **Persistent Volume Claims**, and then click the relevant Persistent Volume Claims. The page that opens displays detailed information about the persistent volume claim.

Tip
To view the contents of **itom-vol**, go to the master node (the NFS server) and enter **cd /var/vols/itom/**. It contains the **baseinfra-<version-number>** and the **suite-install** subdirectories.

Enter **ls -R baseinfra-<version-number>**; this shows the **PrivateRegistry**.

Enter **ls -R suite-install/**; this shows information about the containers that includes the configuration information to deploy the supported suites.

Manage secrets and config maps

Click **RESOURCES** > **Configuration** to display information about Secrets and Config Maps.

Secrets

The Secrets page provides information about secrets that are currently running.

A secret stores sensitive data, such as authentication tokens, which can be made available to containers upon request.

View the Secrets

Click **RESOURCES** > **Configuration** > **Secrets** .

The page displays the list of secrets and their age.

You can click the relevant secret to display its details. The page displays the details of the selected secret and its data.

Config Maps

The Config Maps page provides information about the config maps that are currently running.

The ConfigMap API resource holds key-value pairs of configuration data that can be consumed in pods or used to store configuration data for system components such as controllers. ConfigMap is similar to Secrets, but designed to more conveniently support working with strings that do not contain sensitive information.

View the Config Maps

Click **RESOURCES** > **Configuration** > **Config Maps**.
The page displays the names of the configuration map and its labels, and the amount of time passed since the configuration map was created.

- Click and select **Delete** to delete the config map.
- Click the relevant config map to display its details. The page displays the selected config map details, and its related data.

Manage services and ingresses

Click **RESOURCES** > **Services** and discovery to display information about services and Ingress.

Services

The Services page provides information about services.

A service defines a set of pods and a means by which to access them, such as single stable IP address and corresponding DNS name (such as a web service or API server) that directs and load balances traffic to the set of pods that it covers.

View services

Click **RESOURCES** > **Services and Discovery** > **Services**.

The page displays the names of the services attached to the selected namespace, the labels assigned to the service, the IP of the related cluster, and the internal and external endpoints.

- Click **Actions** and select **Delete** to delete the service.
- Click the relevant service to display its details. See View a service's details.

View a service's details

Click **RESOURCES** > **Services and Discovery** > **Services**, and then click the relevant Service.

The page displays details about the service and the connection as well as information about the related pods.

Ingress

An Ingress is a collection of rules that allow inbound connections to reach the cluster services.

It can be configured, for example, to give services externally reachable URLs, load balance traffic, terminate SSL, or offer name based virtual hosting. Users request ingress by POSTing the Ingress resource to the API server. An Ingress controller is responsible for fulfilling the Ingress, usually with a load balancer, though it may also configure your edge router or additional frontends to help handle the traffic in an HA manner.

View ingress

Click **RESOURCES** > **Services and discovery** > **Ingress**.

The page displays the names of the ingresses attached to the selected namespace, the labels assigned to the ingress, the IP of the related cluster, and the internal and external endpoints.

Click an ingress to view its details.

Manage workloads

This section displays information about Namespaces, Deployments, Replica Sets, Replication Controllers, Daemon Sets, Jobs, Pods, filtered by the selected namespace.

Click **RESOURCES** > **Workloads**.

The page displays all the resources filtered by the selected namespace:

- The CPU and memory usage of the selected namespace during the past 15 minutes.
- The list of replication controllers linked to the selected namespace.
- The list of pods linked to the selected namespace.

For more information, see the following topics:

- Manage deployments
- Manage replica sets
- Manage replication controllers
- View daemon sets
- Manage pods

View daemon sets

The Daemon Sets page provides information about the Daemon Sets for the selected Namespace.

A Daemon Set ensures that all (or some) nodes run a copy of a pod. As nodes are added to the cluster, pods are added to them. As nodes are removed from the cluster, those pods are garbage collected. Deleting a Daemon Set will clean up the pods it created.

View the daemon sets

1. Click **RESOURCES** > **Workloads** > **Daemon Sets** to display the current daemon sets.
2. Click the relevant daemon set to view its details.

Manage deployments

You create and manage sets of replicated containers (actually, replicated pods) using deployments. A deployment provides declarative updates for pods and replica sets (the next-generation replication controller).

A deployment ensures that a specified number of pod “replicas” are running at any one time. If too many replicas are running, the deployment will kill some. If too few replicas are running, the pod will start more.

View the deployments

Click **RESOURCES** > **Workloads** > **Deployments**.
The page displays the CPU and memory usage history of the selected namespace over the past 15 minutes, the name of the available deployments, their labels, the number of pods, the creation timestamp of the deployment, and the deployment's images.

View deployment details

Click **RESOURCES** > **Workloads > Deployments**, and then click the relevant deployment.

The page displays the CPU and memory usage history of the selected deployment over the past 15 minutes, and details about the selected deployment, including details about the new replica set, the old replica sets, and events.

Manage pods

The Pods page provides information about the pods that are currently running or that have been running for the past 15 minutes. You can also access details about a specific pod as well as its log. By default, pods run with unbounded CPU and memory limits. This means that any pod in the system will be able to consume as much CPU and memory on the node that executes the pod. You may want to impose restrictions on the amount of resources a single pod in the system may consume for a variety of reasons.

View the pods

Click **RESOURCES> Workloads** > **Pods**.
The page displays the CPU and memory usage history of the namespace the pod belongs to, status, number of restarts during the lifecycle of the pod, the amount of time passed since the creation of the pod, the IP address of the pod, the CPU and memory usage of the pod itself in the last 15 minutes.

- Click  to display the log of the pod. See View log.
- Click **Actions** and select to delete the pod or to view and edit its YAML.
- Click the pod itself to display its details. See View a pod's details.

View a pod's details

Click **RESOURCES** > **Workloads** > **Pods**, and then click the relevant Pod.

The page displays the CPU and memory usage history of the pod in the last 15 minutes, the pod details, and the network details. To display the log of the pod, see View log.

Also included is information about the pod containers such as the name, image, environment variables, commands, arguments, and more.

View log

1. Click **RESOURCES** > **Workloads** > **Pods**.
2. Click the relevant pod.
3. Click  in the Pod page or **View logs** in the Pod Details page or click **View logs** in the Container area. The page displays the information for the pod.

You can use the following tools:

-  Toggles to change the size of the font used in the log.
-  Toggles to change the colors of the log: white characters on black background or black characters on white background.
- 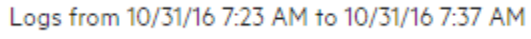

The timestamp of the currently displayed log.
-  Use the relevant buttons to navigate between logs.

Manage replica sets

Replica Sets are the next-generation Replication Controller. The only difference between a Replica Set and a Replication Controller is the selector support. Replica Sets support the new set-based selector requirements whereas a Replication Controller only supports equality-based selector requirements.

This section displays information about replica sets of the selected namespace.

View replica sets

Click **RESOURCES** > **Workloads** > **Replica Sets**.
The page shows the CPU and memory usage history of the selected namespace during the past 15 minutes, the name of the available replica sets for the selected namespace, its labels, pods, images and creation timestamp.
You can click **Actions** and select **Delete** to delete a replica set.

View a replica set's details

1. Click **RESOURCES** > **Workloads > Replica Sets**.

2. Click the relevant replica set.

The page shows details about the selected replica set, the services, pods, and events related to the replica set.

Manage replication controllers

The Replication Controllers page provides details about the Replication Controllers.

View the Replication Controllers

Click **RESOURCES** > **Workloads** > **Replication Controllers** to display the current Replication Controllers.
The page displays the CPU and memory usage of the selected namespace during the past 15 minutes, the list of replication controllers with their name, labels, pods, age, and images of the replication controllers associated with the selected namespace.
You can do the following:

- Click the relevant replication controller to view its details.

The details display the CPU and memory usage history of the selected replication controller for the past 15 minutes, and the services provided by the selected replication controller.

- Click **Actions** and select:
  - **View details.** You can also click the relevant replication controller.
  - **Scale.** See Scale the number of pods linked to the replication controller.
  - **Delete.** The replication controller is deleted.

Scale the number of pods linked to the replication controller

1. Click **RESOURCES** > **Workloads** > **Replication Controllers**.
2. Click and select **Scale**. Enter the relevant number of pods and click **OK**.

Modify CDF's external database

You can modify CDF's external database configuration with the following command: `<foundation_install_dir>/bin/updateExternalIdmDbInfo`

Example

```
updateExternalIdmDbInfo <-t|--dbtype <DB type>> <-u|--user <username>> <-H|--
host <DB host>> <-p|--port <DB port>> <-d|--dbname <DB name>>

updateExternalIdmDbInfo <-t|--dbtype <DB type>> <-u|--user <username>> <-U|--
url <DB connection URL>>

-u|--user External database username.

-H|--host External database host.

-p|--port External database port.

-d|--dbname External database name.
```

```
-U|--url External database connection URL.

-t|--dbtype External database type, optional choices are
("EMBEDDED","EXTERNAL_PG","EXTERNAL_ORA") . The database type must be
capitalized.

-h|--help Show help.
```

When you modified an external default database configuration, you must recreate the IDM pod with the following commands:

```
kubectl delete -f <K8S_home>/objectdefs/idm.yaml

kubectl create -f <K8S_home>/objectdefs/idm.yaml
```

Monitor infrastructure status

You can monitor the infrastructure status of your namespaces, nodes, and persistent volumes.

To access, click **ADMINISTRATION** > **Admin**.

The Admin page displays:

- **Namespaces.** The list of the current default namespaces as well as the namespaces for the suites. Every suite on the same Kubernetes cluster is deployed in a different namespace.
- **Nodes.** The composition of the Kubernetes cluster in terms of servers on which the cluster were installed (master and worker nodes, the physical servers or the VMs).
- **Persistent volumes.** The persistent volume configuration for one or more suites. These volumes contain the data that needs to live outside of the containers.

Network and communication security

We recommend that you add the iptables rules listed below.

Important
Apart from the listed ports on the specific hosts, all other ports should be blocked at the local host level.

NFS Server	111	NFS	Nodes ->NFS Server	NFS server port access by all nodes

NFS Server	2049	NFS	Nodes ->NFS Server	NFS server port access by all nodes
Master Node	2380	Etcd	Master<-> Master	Etcd service port for etcd cluster communication
Master Node	4001	Etcd	Nodes -> Master	Etcd service port for connection from client
All Nodes in Cluster	4194	Kubernetes	Localhost only	Cadvisor for local kubelet
All Nodes in Cluster	5000	Private Registry	Localhost only	Registry port for local host
Ingress Node	5443	MngPortal	All -> Ingress Node	The port exposed on ingress node. all clients could access this port
Master Node	8200	Vault	Nodes->Master	Vault port for client connection
Master Node	8201	Vault	Nodes->Master	Vault port for peer member connection
Master Node	8443	kubernetes	Nodes->Master	API server port for client connection
All Nodes in Cluster	10250	Kubernetes	Nodes->Nodes	Kubernete port for internal communication
All Nodes in Cluster	10251	Kubernetes	Nodes->Nodes	Kubernete port for internal communication
All Nodes in Cluster	10252	Kubernetes	Nodes->Nodes	Kubernete port for internal communication
All Nodes in Cluster	10255	Kubernetes	Master ->Nodes	Kubernete port for internal communication
NFS Server	20048	NFS	Nodes ->NFS Server	NFS server port access by all nodes

Example

The cluster is installed on `10.10.10.10, 10.10.10.11, 10.10.10.12.` The Master Node on: `10.10.10.10`

To add an iptable rules to port 8443 on the master node do the following:

```
iptables -I INPUT 1 -p tcp -m tcp -s 0.0.0.0/0 --dport 8443 -j DROP

iptables -I INPUT 1 -p tcp -s 127.0.0.1 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.10 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.11 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.12 --dport 8443 -j ACCEPT
```

Restart CDF

Follow the steps below to stop the ITOM Container Deployment Foundation:

1. On each master node, run the following commands:

   cd $K8S_HOME/bin
   ./kube-stop.sh

2. On each worker node, run the following commands:

   cd $K8S_HOME/bin
   ./kube-stop.sh

Follow the steps below to restart the ITOM Container Deployment Foundation:

1. On each master node, run the following commands:

cd $K8S_HOME/bin
./kube-start.sh

2. On each worker node, run the following commands:

cd $K8S_HOME/bin
./kube-start.sh

Security

This section is intended for Suite Management Portal container-based platform implementers and system administrators who need to implement their Suite Management Portal environment in a secure manner.

For more information, see these topics:

- Installation security recommendations
- Network and communication
- Authorization
- Data integrity
- Encryption
- Logs
- Enable a firewall on a node
- Back up data for a single-master cluster

Technical system landscape

ITOM Container Deployment Foundation (CDF) is a container that integrates with other suites. CDF is written in Java, JavaScript, and Go.

Security in CDF configurations

CDF configurations may be deployed in the following three modes:

- Single node mode
- Distributed mode 1 (one master node and multiple worker nodes)
- Distributed mode 2 (multiple master nodes and multiple worker nodes)

All of these implementations share the same basic out-of-the-box security configuration options:

- In an out-of-the-box installation, Transport Layer Security/Secure Socket Layer (TLS/SSL) security is enabled between the browser and the CDF server by default.
- In an out-of-the-box installation, CDF requires users to enter username and password credentials to gain access to the application.

External authentication

Though CDF cannot inherit users' information and authorization profiles from an external repository, suite users can use the industry-standard protocols and tools provided by identification management (IDM) integrated into CDF to get the users' information and authentication profiles. For example, suite users can configure LDAP or Single Sign-On provided by IDM to get external authentication profiles.

Common security considerations

CDF can only be deployed on supported operating systems.

We recommend that you follow vendor-provided best practices and security hardening guides for each of the third-party components in your CDF deployment. This includes Docker, Kubernetes, Vault, Nginx, and NFS. The following resources serve as a starting point for researching these recommended security considerations:

Docker Security Tips

https://www.docker.com/docker-security

Kubernetes Security Tips

http://kubernetes.io/docs/troubleshooting/

Vault Security Tips

https://www.hashicorp.com/security.html

Nginx Security Tips

http://nginx.org/en/security_advisories.html

NFS Security Tips

http://www.cert.org/historical/advisories/

Authorization

This section provides information related to user authorization in Suite Management Portal platform.

Authorization Model

Access to the Suite Management Portal platform resources is authorized based on the user's following settings:

- User name
- Session & Inactivity timer timeouts

Note
CDF cannot inherit users' information and authorization profiles from an external repository, such as LDAP.

Back up data for a single-master cluster

To back up the data in the `data` directory for the single-master cluster, use the `etcdctl backup` command.

Example

```
etcdctl backup \
```

```
--data-dir %data_dir% \
--backup-dir %backup_data_dir%
```

You can also use the `etcdctl backup` command to back up all the exported folders in the NFS server too.

The `etcdctl backup` command will rewrite some of metadate contained in the backup (specifically, the node ID and cluster ID), which means that the node will lose its former identity.

Note
In order to recreate a cluster from the backup, you will need to start a new, single-node cluster. The metadata is rewritten to prevent the new node from inadvertently being joined onto an existing cluster.

Data integrity

The data backup procedure is also an integral part of data integrity. As CDF does not provide native backup capabilities, please consider the following guidelines:

- Database backup is especially important before critical actions such as upgrades.
- Backup files should be stored according to industry best practices to avoid unauthorized access.
- As database backup can be a resource intensive process, we strongly recommend that you avoid running backup operations during peak demand times.

Enable a firewall on a node

Enable a firewall on the NFS server

Run the following commands to enable a firewall on the NFS server:

```
systemctl start firewalld;systemctl enable firewalld
firewall-cmd --permanent --add-port=111/udp
firewall-cmd --permanent --add-port=111/tcp
firewall-cmd --permanent --add-port=2049/tcp
firewall-cmd --permanent --add-port=20048/tcp
firewall-cmd --reload
```

Enable a firewall on the master nodes

Run the following commands to enable a firewall on each master node:

```
systemctl start firewalld; systemctl enable firewalld
firewall-cmd --permanent --add-port=4001/tcp
```

```
firewall-cmd --permanent --add-port=2380/tcp

firewall-cmd --permanent --add-port=8200/tcp

firewall-cmd --permanent --add-port=8201/tcp

firewall-cmd --permanent --add-port=8443/tcp

firewall-cmd --permanent --add-port=10250/tcp

firewall-cmd --permanent--direct --add-rule ipv4 filter FORWARD 1 -o docker0
-j ACCEPT -m comment --comment "docker subnet"

firewall-cmd --permanent --direct --add-rule ipv4 filter INPUT 1 -i docker0 -
j ACCEPT -m comment --comment 'kube-proxy redirects'

firewall-cmd --reload
```

Enable a firewall on the worker nodes

Run the following commands to enable a firewall on each worker node:

```
systemctl start firewalld; systemctl enable firewalld

firewall-cmd --permanent --add-port=10250/tcp

firewall-cmd --permanent--direct --add-rule ipv4 filter FORWARD 1 -o docker0
-j ACCEPT -m comment --comment "docker subnet"

firewall-cmd --permanent --direct --add-rule ipv4 filter INPUT 1 -i docker0 -
j ACCEPT -m comment --comment 'kube-proxy redirects'

firewall-cmd --reload
```

Encryption

This section provides information on data encryption in the Suite Management Portal platform.

TLS/SSL data transmission

An IDM server is used for authentication. The IDM server is monitored by a single center policy server, and consists of a user repository, a policy store, and a web server agent installed over each of the capability's web servers that communicates with the policy server. The IDM server controls users' access to various organizational resources, protecting confidential personal and business information from unauthorized users.

For optimal security, we recommend that you either configure a TLS connection between the suite and the IdM server, or have the suite server and the IDM servers on the same secure internal network segment. Authentication is performed by the IDM server, and authorization is handled by the capabilities.

ITOM Container Deployment Foundation (CDF) uses TLS/SSL to transmit data between the server and browsers.

To change the default value of the SSL cipher, follow these steps:

1. On the master node, change the ssl-ciphers value in the $K8S_HOME/objectdefs/nginx-ingress.yaml file.
2. Run the following commands to recreate the ingress container:
   kubectl delete -f $K8S_HOME/objectdefs/nginx-ingress.yaml
   kubectl create -f $K8S_HOME/objectdefs/nginx-ingress.yaml

Encryption of stored database fields

CDF uses proprietary algorithms to encrypt data that is stored in the database, and uses the Identity Manager (IDM) to manage user passwords.

Installation security recommendations

This section provides information on aspects of installation security.

Harden SSH on the operating system

By default, the SSH server is configured with a weak cipher and a weak KexAlgorithms on each node. To harden the SSH server, set the values of **KexAlgorithms**, **Ciphers** and **MACs** in the /etc/ssh/sshd_config file as follows:

- **KexAlgorithms** ecdh-sha2-nistp521,ecdh-sha2-nistp384,ecdh-sha2-nistp256
- **Ciphers** aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr
- **MACs** hmac-sha2-256

Database security recommendations

Refer to the following website for information about PostgreSQL database security solutions:

http://www.openscg.com/postgresql-security-guidelines/

Application server security recommendations

- Always change the default passwords.
- Always use the minimal possible permissions when installing and running CDF (You must install and run root permissions using the sudo command).

Network and communication

This section provides information on network and communication security.

Secure topology

ITOM Container Deployment Foundation (CDF) is designed to be part of a secure architecture and to deal with the security threats to which it could potentially be exposed.

To securely deploy the CDF, we recommends that you use the TLS/ SSL communication protocol.

Import custom certificates during installation

You can specify certificates for ingress service during the CDF installation. On the Connection page, select the your server certificate and root certificate, and click **Upload**.

Update certificates

From the management portal, **ADMINISTRATION** > **Certificate**, select certificates and the key files. Click **Update** to use the selected certificates and keys.

Renew the client.crt, client.key, server.crt, and server.key certificates

You cannot replace these certificate files with your own. When these certificates expire, you must renew them.

To renew the certificates, follow these steps:

1. Generate new server certificates or client certificates with the following commands on each master node:

cd $<*K8S_HOME*>/scripts
./renewCert.sh

2. You will be asked to enter password as root when renewing the certificates on the first master node. You can use the same password for other additional nodes by pressing **enter** for the password.

3. Restart the kubelete service to use the new certificates by entering **yes** for the kubelet restart question on the terminal. Alternatively, you can restart the kubelet service manually later.

Note When some nodes failed to renew certificates, follow the steps below to update the certificates on the nodes that failed certificate renewal:

1. Generate new certificates with the following commands:

cd $<*K8S_HOME*>/scripts
./renewCert.sh
Or copy the certificates from **<*K8S_HOME*>/ssl/new-certs/** to **<*K8S_HOME*>/ssl** manually.

2. Restart the kubelete service manually with the following commands

cd $<*K8S_HOME*>/bin
./ kube-restart.sh

Security recommendations

We recommend that you add the following iptables rules.

Caution Apart from the listed ports, all other ports should be blocked at the localhost level.

All Nodes in Cluster	10250	Kubernetes	Nodes -> Nodes	Kubernete port for internal communication
All Nodes in Cluster	10251	Kubernetes	Nodes -> Nodes	Kubernete port for internal communication
All Nodes in Cluster	10252	Kubernetes	Nodes -> Nodes	Kubernete port for internal communication
All Nodes in Cluster	10255	Kubernetes	Nodes -> Nodes	Kubernete port for internal communication
Ingress Node	5443	MngPortal	All -> Ingress Node	The port exposed on ingress node. All clients could access this port
Master Node	2380	Etcd	Master <-> Master	Etcd service port for etcd cluster communication
Master Node	4001	Etcd	Nodes -> Master	Etcd service port for connection from client
Master Node	8200	Vault	Nodes->Master	Vault port for client connection
Master Node	8201	Vault	Nodes->Master	Vault port for peer member connection
Master Node	8443	Kubernetes	Nodes -> Master	API server port for client connection
NFS server	111	NFS	Nodes -> NFS Server	NFS server port access by all nodes
NFS server	2049	NFS	Nodes -> NFS Server	NFS server port access by all nodes
NFS server	20048	NFS	Nodes -> NFS Server	NFS server port access by all nodes

Example:

Assume that the cluster is installed on 10.10.10.10, 10.10.10.11, 10.10.10.12, and the master node is installed on: 10.10.10.10. In this example. to add iptable rules to port 8443 on the master node, you run the following commands:

iptables -I INPUT 1 -p tcp -m tcp -s 0.0.0.0/0 --dport 8443 -j DROP

iptables -I INPUT 1 -p tcp -s 127.0.0.1 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.10 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.11 --dport 8443 -j ACCEPT

iptables -I INPUT 1 -p tcp -s 10.10.10.12 --dport 8443 -j ACCEPT

Set up a host name resolution using a hosts file

When you want to access a machine within or outside the cluster from a pod of the node cluster via a FQDN, follow the steps below to set up a host name resolution using a hosts file.

1. Navigate to the CDF NFS volume directory with the following command:
   cd /var/vols/itom/core/baseinfra-1.0
2. Go to the DNS host folder with the following command:
   cd kube-dns-hosts
3. Create a hosts file with the following command:
   vi hosts
4. Add the host name and IP address of the machine to the file. For example:
   example.me.com 16.155.198.76
5. Save the hosts file.

Now you can access the machine by using the FQDN.

View existing images

You can view the existing images in the local registry. Click **ADMINISTRATION > Local Registry**. The following page is displayed.

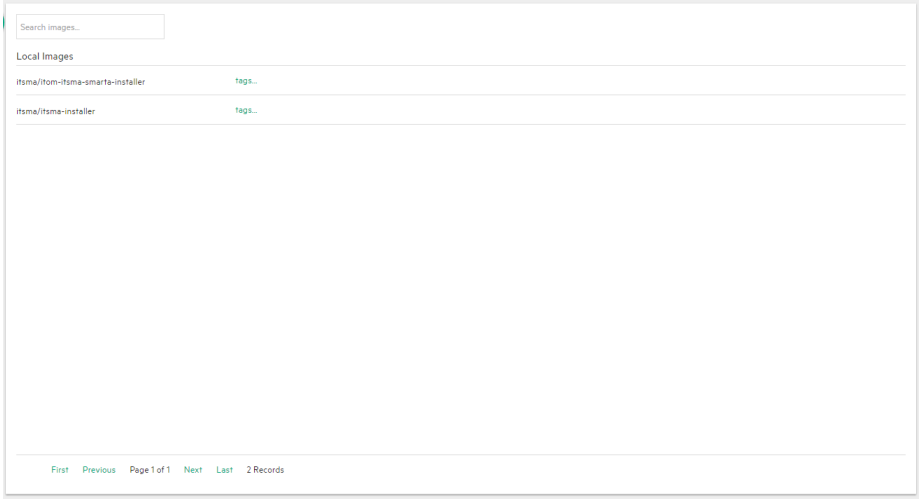

Administer the Operations Bridge Suite

You can perform the following tasks to administer the Operations Bridge Suite in a container deployment:

Access Command Line Interfaces

Access the RTSM JMX Console

Configure LDAP authentication

Configure scaling and high availability

Replace the suite trial license

Access Command-Line Interfaces

OMi and OBR provide several command line interfaces that are useful for automation and troubleshooting. To access the CLIs from within the Operations Bridge Suite container environment, the basic workflow is as follows:

1. Find the container that contains the CLI.

[root@master]# kubectl get pods --all-namespaces | grep <omi|obr-server>

2. Start the shell.
   [root@master]# kubectl exec -ti <*pod_id*> bash -c <omi|obr-server> -n <*namespace*>

For example: kubectl exec -ti omi-2246081285-8u1e0 bash -c omi -n opsbridge1

3. Execute the CLI.

Example: Access OMi command-line interfaces

1. Find the container that contains the CLI.

2. [root@master]# kubectl get pods -n opsbridge1 -o name | grep omi

pod/<*container_id*>

3. Start the shell.

4. [root@master]# kubectl exec -ti <*container_id*> -n opsbridge1 -c omi bash

omi:/ #

5. Execute the CLI, in this example, opr-node:

omi:/ # /opt/HP/BSM/opr/bin/opr-node -username admin -list_nodes –all

Example: Access OBR command-line interfaces

1. Find the container that contains the CLI:

2. [root@master]# kubectl get pods -n opsbridge1 -o name | grep obr-server

pod/<*container_id*>

3. Start the shell.

4. [root@master]# kubectl exec -ti <*container_id*> -n opsbridge1 -c obr-server bash

obr-server:/ #

5. Execute the CLI, in this example, abcMonitor:

obr-server:/ # abcMonitor -streamdef

Access the RTSM JMX console

The RTSM in OMi provides a JMX console that gives additional information and advanced configuration possibilities.
To access the JMX console from your container deployment, open the following URL from a supported web browser:

`https://<external_access_host>/jmx-console`

<*external_access_host*> is the fully qualified domain name of the host which you specified as EXTERNAL_ACCESS_HOST in the install.properties file during the Container Deployment

Foundation installation. Usually, this is the master node's FQDN.
The login user is *sysadmin*.

Configure LDAP authentication

With the default single sign-on authentication strategy for the Operations Bridge Suite, users are authenticated to all installed capabilities with the same credentials. User names and passwords are stored and verified by a central server so that a user needs only one account to access all capabilities.

A suite-specific Identity Management (IDM) server is used for the authentication. The IDM server is monitored by a single central policy server and consists of a user repository, a policy store, and a web server agent installed over each of the capability's web servers communicating with the policy server. The IDM server controls users' access to various organizational resources, protecting confidential personal and business information from unauthorized users.

For optimal security, HPE recommends to either configure a TLS connection between the suite and the IDM server, or have the suite server and the IDM servers on the same secure internal network segment. Authentication is performed by the IDM server, and authorization is handled by the capabilities.

Additionally, you can configure LDAP authentication for BVD. Automatic user creation from LDAP servers simplifies the user management process for administrators as authentication is performed through the LDAP server.

You can use an external LDAP server to store user information (user names and passwords) for authentication purposes, instead of using the internal IDM service. You can manually create BVD users and LDAP users, and use LDAP servers to automatically create LDAP users in BVD.

Note
LDAP should be configured *after* the installation of the Operations Bridge Suite.

How to configure LDAP

1. Launch the Management Portal from a supported web browser:

https://<*external_access_host*>:5443
<*external_access_host*> is the fully qualified domain name of the host which you specified as EXTERNAL_ACCESS_HOST in the install.properties file during the Container Deployment Foundationinstallation. Usually, this is the master node's FQDN.

2. Log in as the admin user.
3. Go to **ADMINISTRATION**> **IdM Administration**. In the Organization List, click **Operations Bridge**.
4. In the Provider Detail, click **Authentication** and then **Add**.
5. Select **LDAP** in the drop-down list, and click **Create**.
6. Enter a valid LDAP configuration. For details on what to enter for each LDAP setting, see LDAP settings.

7. Click **Save**.
8. Log on to your capabilities via LDAP:

OMi: https://<*external_access_host*>/omi
BVD: https://<*external_access_host*>/bvd
<*external_access_host*> is the fully qualified domain name of the host which you specified as EXTERNAL_ACCESS_HOST in the install.properties file during the Container Deployment Foundationinstallation. Usually, this is the master node's FQDN.
Note that additional steps are necessary in OMi, BVD, and OBR to set up LDAP group mappings and permissions. For details, see the Help Centers for each capability.

LDAP settings

The LDAP settings contain parameters for the LDAP server configuration, LDAP attributes, and user login information.

LDAP Server Settings	
Display Name	Name of the LDAP configuration. This name cannot be changed when you reconfigure the settings.
Hostname	Fully-qualified domain name or IP address of the LDAP server. **Example**: `192.0.2.24`
Port	Port of the LDAP server. LDAP servers typically use port 389 or secure port 636.
SSL Connection	Select **SSL Connection** if an LDAPS URL is specified.
Base DN	The Distinguished Name (DN) of the LDAP entity from which you want to start your user search. **Example**: `CN=Users,DC=omi,DC=example,DC=com`
User ID (Full DN)	The Distinguished Name (DN) of a user with search privileges on the LDAP directory server. **Example**: `CN=Administrator,CN=Users,DC=example,DC=com`
Password	Password of the specified user ID.
LDAP Attributes	
Full Name	Full name to be included in the user search. **Example**: `cn`

User Email	Property that contains the user's email address (specific to the selected LDAP vendor, for example MS Active Directory). **Example**: `mail`
Group Membership	List of comma-separated LDAP attributes to find groups in a user profile. **Example**: `member,uniqueMember`
Manager Identifier	Any attribute (for example DN or CN) of the user who is the user's manager. **Example**: `manager`
Manager Identifier Value	The value of the identifier. For example, if you specified the DN in the Manager Identifier field, enter `dn`.
User Avatar	Attribute for the user avatar image. You must specify an LDAP record property name that exists on the LDAP server. **Example**: `cn`
Priority	Specifies the priority of the domain controller. The priority determines the order in which clients contact a domain controller.
Referral Search	Select to follow LDAP referrals to another server that offers the requested information.
User Login Settings	
User Name Attributes	Name of field that contains the user name. **Example**: `sAMAccountName`
User Searchbase	Parameters to indicate which attributes are to be included in the user search. **Example**: `CN=Users`
User Search Filter	LDAP pattern to use when searching for a user account. **Example**: `(sAMAccountName={0})` The user search filter must include the pattern `{0}`, which is replaced with the user name entered on login. IDM does not support LDAP multiple search filter components like `(&(sAMAccountName={{username}})(objectclass=user))`.

Search Subtree	Select to search the subtree below the base DN (including the base DN level).
Group Settings	
Group Search Base	Parameters to indicate which attributes are to be included in the group search. **Example**: ou=Groups,dc=example,dc=net
Group Search Filter	LDAP pattern to use when searching for a group account. **Example**: `cn={0}`

Configure scaling and high availability

You can improve your system availability and reliability by scaling your suite resources as required. You can scale single nodes, as well as multiple nodes.

By managing your resources, you can scale your system out or in. For example, by increasing the number of pod replicas on a deployment, the load of the deployment can be automatically distributed across all pods.

A high availability configuration offers continuous service despite power outages, machine downtime, and heavy load.

Configure high availability

To achieve a high availability of your system, use the following techniques:

Important
Micro Focus strongly recommends to use *all* of the following techniques. That is the only way how you can make your system highly available.

- **Highly available redundant storage**. Use a redundant NFS server for your container deployment.
- **Highly available database instances**. Use a redundant external database.
- Kubernetes **cluster with multiple master nodes**. See High availability of the Kubernetes cluster.
- Operations Bridge **services on multiple worker nodes**. To achieve a high availability of your Operations Bridge capabilities, deploy the capabilities on multiple worker nodes. By scaling your deployment horizontally, worker nodes can take over loads from failed worker nodes. See High availability of Suite capabilities.
- **Keep-alive monitoring of** Kubernetes. Kubernetes automatically detects failures of single services and complete nodes and restarts pods on other nodes.

High availability of the Kubernetes cluster

When installing multiple master and worker nodes, set the `EXTERNAL_ACCESS_HOST` to an FQDN which is resolved to the `HA_VIRTUAL_IP`. This way, you make sure you can access the CDF with the FQDN defined in `EXTERNAL_ACCESS_HOST`.

By specifying these properties, an ingress instance and `keepalived` are launched on each master node. `keepalived` binds the Virtual IP to a master node. The node with the Virtual IP is in master mode while the other nodes are in standby mode. If the master node with the Virtual IP is down, the Virtual IP is bound to another master node.

High availability of Suite capabilities

BVD

BVD pods are highly available by default, even if you have not scaled them out. However, if one of the worker nodes fails, restarting the pods will cause a short downtime of your system.
To ensure a constant system availability even if one of the worker nodes fails, you can increase the number of BVD receiver (`bvd-receiver-deployment`) and web server (`bvd-www-deployment`) pod replicas.
A short downtime will then only occur if the BVD Redis pod is affected by the crash of a worker node or the Redis process. In this case, Kubernetes restarts the BVD Redis pod. This takes usually less than a minute.
Any data that was sent during the downtime of the BVD Redis pod will be buffered by the BVD receivers, so no data is lost.

For more information, see Scale a BVD deployment horizontally.

OMi

For OMi to be highly available, the following prerequisites must be fulfilled:

- You must have two master nodes and at least two worker nodes that have enough resources to host OMi.
- When running the suite installer, tick the checkbox **Enable high availability for Operations Manager i**.

Note
It is currently not possible to unconfigure high availability for OMi once it has been enabled and deployed.

If high availability for OMi is enabled, two OMi pods are created by the OMi StatefulSet controller which manages the deployment and scaling of the pods. Each OMi pod starts on a different worker node and in sequential order. This means the second replica is not scheduled until the first has completely started. The created OMi pods are single server deployments with the names omi-0 and omi-1. After the initial startup, omi-0 is configured as primary data processing server (DPS) and omi-1 as backup DPS. omi-0 is the active DPS and omi-1 is passive. If, at any point, the active DPS fails, the passive DPS automatically takes over with a small downtime. Any data that was sent during the downtime will be buffered and processed as soon as the healthy DPS has taken over as active DPS, so no data is lost.

For more details regarding server roles, management, configuration troubleshooting and limitations for high availability within the OMi capability, see the OMi Help Center.

Notes and Limitations

- **Node failures**. If a node that hosts one of the OMi pods becomes unreachable because it is down for scheduled maintenance or becomes partitioned from the network, the corresponding OMi pods goes into an "Unknown" or "Terminating" state. The pod is not rescheduled to a different node in the Kubernetes cluster unless the node object is manually deleted, deleted via the Node Controller, forcefully deleted, or shut down by kubelet to remove the pod's entry from the API server.
  This limitation exists because Kubernetes enforces the applications running in a StatefulSet to have a stable network identity and storage. In general, Kubernetes tries to avoid the creation of multiple instances of the same pod to prevent data corruption.
- **Node maintenance**. OMi tolerates node maintenance with minimal downtime on event processing if the active DPS is hosted on the node that is going on maintenance. Zero downtime is tolerated if the passive DPS is hosted on that node.
  The steps to perform node maintenance are described in the Kubernetes documentation.
- **Planned/unplanned node downtime**. On some cloud providers, you can specify the automatic deletion of node objects after an infringement of SLAs/SLOs is detected, meaning if the host is unresponsive for 10 minutes. In that case, Kubernetes automatically starts the process of rescheduling the StatefulSet pods in nodes with sufficient resources. However, when running on-prem, we recommend that you delete the node object from Kubernetes as you power it down, or have a reconciliation loop keeping the Kubernetes idea of nodes in synchronization with the actual nodes available.
- **Node restart**. Depending on the resource availability, Kubernetes will reschedule the OMi pod on a different node or on the same node as soon as the node starts (when the node is in state "Ready").
- **Suite uninstallation**. Uninstalling the Operations Bridge Suite, as well as deleting and/or scaling down the OMi StatefulSet will not delete the contents of the volumes associated with the corresponding OMi pods. This is how Kubernetes ensures data safety, which is generally more valuable than an automatic purge of the releated StatefulSet resources.

Other Operations Bridge capabilities

Operations Bridge Reporter and Performance Engine do not support running pods multiple times on different worker nodes. Therefore, when one of the worker nodes fails, all corresponding pods have to be restarted. The downtime of the capabilities depends on the restart time of those pods. This can take several minutes.

Scale a BVD deployment horizontally

To ensure a high availability of BVD even if one of the worker nodes fails, you can increase the number of BVD receiver (`bvd-receiver-deployment`) and/or web server (`bvd-www-deployment`) pod replicas.

Tip

- Scale out the BVD receiver load if you send a lot of data to BVD. By scaling out, the number of data samples that can be processed by BVD increases.
- Scale out the BVD web server load if BVD is accessed by multiple people at the same time. By scaling out, a higher number of users will be able to access BVD concurrently.

1. Launch the Management Portal from a supported web browser:
   https://<external_access_host>:5443
   <external_access_host> is the fully qualified domain name of the host which you specified as EXTERNAL_ACCESS_HOST in the install.properties file during the Container Deployment Foundation installation. Usually, this is the master node's FQDN.
2. Log on as the admin user.
3. Go to **Resources**, and click **All namespaces**. In the drop-down list, select the namespace for the Operations Bridge Suite that was assigned during the installation (for example `opsbridge`).
4. Go to **Workloads** > **Deployments**. You can either scale out the receiver load (`bvd-receiver-deployment`), or the web server load (`bvd-www-deployment`).

   Caution Do not scale out any of the of the other BVD pods.

5. For the deployment you want to scale out, click **Actions** and select **View/edit YAML**. The **Edit a Replica Set** dialog box opens.
6. Edit the line `spec replicas : 1`. Increase the number of pod replicas as required.
7. Click **Update**.
8. Wait until the deployment is updated. This might take a few minutes.
9. *Optional*. You can verify that the deployment has been updated correctly:
   a. Refresh the **Deployments** page. For the deployment you selected, the number of pods should have increased (for example `2/2` instead of `1/1`).
   b. Go to **Workloads** > **Replica Sets**, and verify that the number of pod replicas for the deployment has increased as specified. The age displays for how long the pods have been running.

Replace the suite trial license

If you do not provide a perpetual license prior to the suite installation, the built-in 60-day trial license (InstantOn) is used.

If later you purchase a perpetual license after the installation, you can replace the trial license with the perpetual license. To do this, follow these steps:

1. Launch the Management Portal from a supported web browser:

https://<*external_access_host*>:5443

<*external_access_host*> is the fully qualified domain name of the host which you specified as EXTERNAL_ACCESS_HOST in the install.properties file during the Container Deployment Foundationinstallation. Usually, this is the master node's FQDN.

2. Log in as the admin user.

3. Click **SUITE > Management**. For your suite deployment, click **Actions** and select **License**.
4. Click **Install Licenses**.
5. Click **Choose File** to browse to the license file on your local drive, then click **Next**.

The license details are displayed.

6. Select all listed licenses and click **Install Licenses**.
7. *Optional*. When the installation is complete, click **View Licenses** to view the installed licenses.