Operations Bridge Suite

Software Version: 2017.08

Administration Guide

Document Release Date: August 2017
Software Release Date: August 2017

Hewlett Packard Enterprise

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/.

This site requires that you register for an HPE Passport and to sign in. To register for an HPE Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** on the HPE Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support site at: https://softwaresupport.hpe.com/.

Most of the support areas require that you register as an HPE Passport user and to sign in. Many also require a support contract. To register for an HPE Passport ID, click **Register** on the HPE Support site or click **Create an Account** on the HPE Passport login page.

To find more information about access levels, go to: https://softwaresupport.hpe.com/web/softwaresupport/access-levels.

HPE Software Solutions Now accesses the Solution and Integration Portal website. This site enables you to explore HPE product solutions to meet your business needs, includes a full list of integrations between HPE products, as well as a listing of ITIL processes. The URL for this website is https://softwaresupport.hpe.com/km/KM01702731.

Contents

Administration

This guide 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" on page 6.

For more information on the suite administration, see "Administer the Operations Bridge Suite" on page 44.

Administer the Container Deployment Foundation

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

Access the Management Portal

To access the Management Portal, do the following:

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.

   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. See the *Operations Bridge Suite Installation Guide*.

Manage users

This section provides information on how to manage a user.

How to display, create, delete, or edit a user

Click **ADMINISTRATION** > **User Management**. The User page opens.

For each user, this page displays the user name, display name, e-mail, and user group.

- **Create a user.** To create a user, click **ADD**. Enter the user name, password, email, group, and (optionally) a display name. Click **SAVE**.
- **Delete a user.** To delete a user, click and select **Delete**.
- **Edit a user.** To edit or view a user information, click and select **View/Edit**.

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.

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. HPE recommends 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.

This section includes the following tasks:

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" on the next page and "Assign labels to nodes" on the next page.
   - Add a node. See "Add a node" on the next page.
   - **REFRESH.** Click to refresh the display.
   - Click the relevant node to see its details. See "View the node details" on the next page.

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

To manage node labels:

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 a node

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

   Enter the host name of the node, the name of a user that can remotely execute commands on the host (typically the root user), and the password of the specified user, and click **ADD** to remotely install the extra node.

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 licenses

The License page enables you to manage your suite licenses.

This section includes the following tasks:

View existing licenses

1. 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 a license

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

2. Click **Archive**.

The licenses are removed from the list of installed licenses in the License Management table and become unavailable for customers to fetch and activate the products.

Restore an archived license

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 a license 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.

View the existing images

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

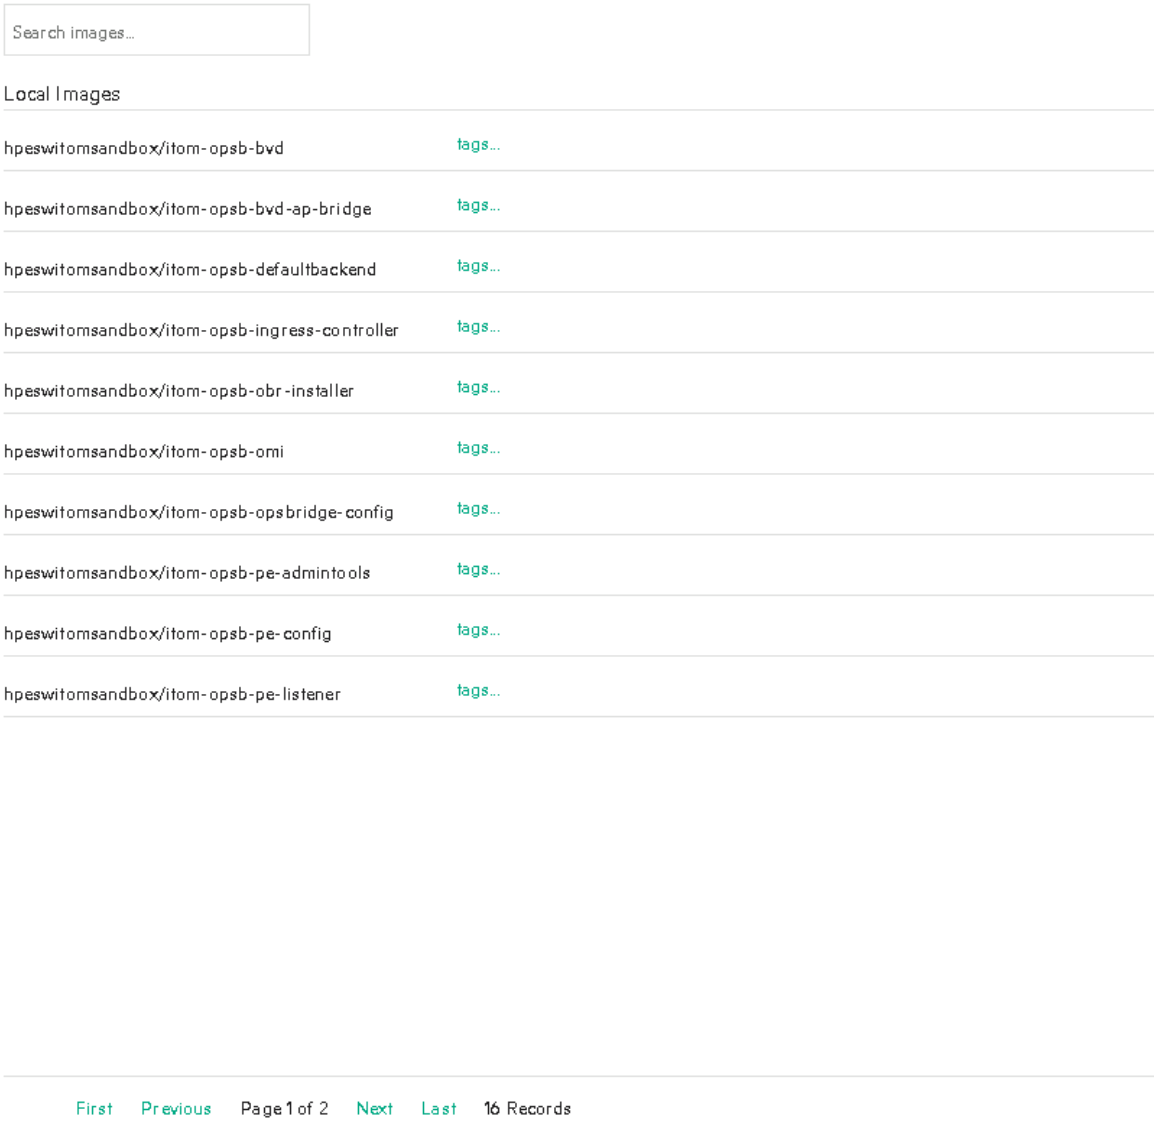

Modify CDF's external database

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

> **Usage 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 <foundation_install_dir>/objectdefs/idm.yaml`

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

Set up LW-SSO

The LWSSO page enables you to set up Lightweight Single Sign-On (LW-SSO) with other products, and configure a customized timeout for the IDM token.

> **Note:** The InitString and Domain of **LWSSO** have default values. You must input the current user's password and then click **Show InitString** to see the InitString's default value. You can also change these default values according to your requirements.

Set up single sign-on with other products

To set up LW-SSO with other products, do the following:

1. Click **ADMINISTRATION** > **LWSSO**.
2. Enter the InitString and Domain and click **UPDATE**.

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

Configure a customized timeout for the IDM token

To configure a customized timeout for the IDM token, do the following:

1. For the **IDM Token timeout configuration**, enter a specific value in minutes and click **UPDATE**. By default, the timeout session is 30 minutes.
2. Restart the IDM pods by running the following commands:

   ```
   cd <foundation_install_dir>/objectdefs

   kubectl delete -f idm.yaml

   kubectl delete -f idm-pg.yaml

   kubectl create -f idm-pg.yaml

   kubectl create -f idm.yaml
   ```

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.

Namespace

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.

1. 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

1. 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.

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.

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

You can select another namespace.

View the deployments

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

The page displays the CPU and memory usage history of the selected namespace during the past 15 minutes, the name of the available deployments, their labels, the number of pods, the creation timestamp of the deployment, and its images.

You can:

- Click a deployment to display its details.

  The details include information about the new replica set, the old replica sets, and the events that

took place.

- Click **Actions** and select **Delete**, to delete the deployment.

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

1. 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.

Replication controllers

The Replication Controllers page provides details about the Replication Controllers.

View the Replication Controllers

1. 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" below.
  - **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**.

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.

Pet Sets

The Pet Sets page provides information about pet sets.

A Pet Set is a Controller that provides a unique identity to its Pods. It provides guarantees about the ordering of deployment and scaling.

View the Pet Sets

1. Click **RESOURCES** > **Workloads** > **Pet Sets** to display the current Pet Sets.
2. Click a Pet Set to view its details.

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

1. 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" on the next page.
   - 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" below.

View a pod's details

1. 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" on the next page.

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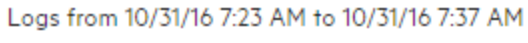 The timestamp of the currently displayed log.
   -  Use the relevant buttons to navigate between logs.

Services and discovery

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" below.

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.

Persistent Volume Claims

The Persistent Volume Claims page displays information about the currently running persistent volumes.

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.

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 a persistent volume claim details

1. 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 see the contents of **itom-vol**, go to the master node (the NFS server) and enter **cd /var/vols/itom/core**. 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.

Configuration

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.

Security

This section is intended for the ITOM Container Deployment Foundation implementers and system administrators who need to implement their ITOM Container Deployment Foundation environment in a secure manner.

This section includes the following information:

Secure Implementation and Deployment

This section provides information on implementing and deploying the ITOM Container Deployment Foundation container-based platform in a secure manner.

Technical system landscape

The ITOM Container Deployment Foundation is a platform that integrates with other Suites. The ITOM Container Deployment Foundation container-based platform is written in Java and JavaScript and Go.

For more information about typical deployment schemes and options, see ITOM Container Deployment Foundation Architecture in the *Quick Access Guide*.

Security in the ITOM Container Deployment Foundation configurations

The ITOM Container Deployment Foundation configurations may be deployed in the following three implementations. See ITOM Container Deployment Foundation Architecture in the *Quick Access Guide*.

- Single mode (one master node).
- 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.

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

External Authentication

With additional configuration, it is possible to supplement or replace the default authentication and authorization provider for the ITOM Container Deployment Foundation by using a variety of industry-standard protocols and tools such as LDAP and Lightweight Single Sign-On. See Configure LDAP or LW-SSO.

Common Security Considerations

The ITOM Container Deployment Foundation can only be deployed on supported operating systems. See Operating System in the *Support Matrix*.

HPE recommends to follow vendor-provided best practices and security hardening guides for each of the third-party components used in support of your ITOM Container Deployment Foundation deployment, which includes Docker, Kubernetes, NFS, Vault and Nginx. Below are some resources that can serve as a starting point for researching these recommended security considerations:

Docker Security Tips

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

Kubernetes Security Tips

https://kubernetes.io/docs/tasks/debug-application-cluster/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/

The ITOM Container Deployment Foundation Security Parameters

This section contains reference to some of the ITOM Container Deployment Foundationparameters that are relevant to security.

Secure File Storage

The ITOM Container Deployment Foundation allows users to upload files (suite installation binaries) to the ITOM Container Deployment Foundation Server. All files uploaded to the server must be validated, since they can contain viruses, malicious code, or Trojans.

As a result, it is strongly recommended to implement proper antivirus protection for the file storage.

Installation Security

This section provides information on aspects of the installation security.

Operating Systems

On each node, the SSH server is configured with weak cipher and weak KexAlgorithms by default.

To harden SSH on your operating system, set the values of `KexAlgorithms`, `Ciphers` and `MACs` in file: `/etc/ssh/sshd_config` 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

PostgreSQL

See http://www.openscg.com/postgresql-security-guidelines/ for information about PostgreSQL database security solutions.

Application Server Security Recommendations

Always change default passwords.

Always use the minimal possible permissions when installing and running the ITOM Container Deployment Foundation.

Action	Permissions Needed for User
Installing/Running HPE ITOM Container Deployment Foundation	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

The CDF is designed to be part of a secure architecture, and can meet the challenge of dealing with the security threats to which it could potentially be exposed.

To securely deploy the CDF, HPE recommends to use the TLS/SSL communication protocol.

Replace the ingress service certificate with a custom certificate

To replace the certificate and private key of Ingress Service with a custom certificate and private key, follow the steps below:

1. Generate a certificate and private key for the host on which the Ingress Service is running on. Put it somewhere on the master node.

2. on the master node, delete a secret with the following command:

```
kubectl delete secret nginx-default-secret -n core
```

3. on the master node, recreate the secret with a new certificate and private key:

> **Note:** You must keep the format of the following commands as it is, especially the indented spaces.

```
echo "
apiVersion: v1
kind: Secret
metadata:
  name: nginx-default-secret
  namespace: core
data:
  tls.crt: `base64 <your custom certificate file directory> |tr -d \"\n\"`
  tls.key: `base64 <your custom private key file directory> |tr -d \"\n\"`
" | kubectl create -f -
```

4. on the master node, delete and recreate the ingress service:

```
kubectl delete -f ${K8S_HOME}/objectdefs/nginx-ingress.yaml
kubectl create -f ${K8S_HOME}/objectdefs/nginx-ingress.yaml
```

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

Follow the steps below to replace the `client.crt`, `client.key`, `server.crt`, and `server.key` with custom certificates.

1. Generate new server certificates or client certificates with the following commands:

```
cd ${K8S_HOME}/scripts
./renewCert.sh
```

2. Restart the kubelete service with the following commands:

```
 cd ${K8S_HOME}/bin
./ kube-restart.sh
```

3. Delete three default tokens in the core, default, and suite namespaces with the following commands:

   ```
   kubectl get secrets --all-namespaces
   ```

   ```
   kubectl get delete secret xxxx -n default-token-xxxx
   ```

4. Recreate the yaml files with the following commands:

   ```
   cd ${K8S_HOME}/objectives
   ```

   ```
   kubectl delete -f kube-vault.yaml
   ```

   ```
   kubectl delete -f mng-portal.yaml
   ```

   ```
   kubectl delete -f nginx-ingress.yaml
   ```

   ```
   kubectl create -f kube-vault.yaml
   ```

   ```
   kubectl create -f mng-portal.yaml
   ```

   ```
   kubectl create -f nginx-ingress.yaml
   ```

   ```
   kubectl delete -f ingress yaml
   ```

5. Recreate the suite ingress yaml with the following commands:

   ```
   cd /var/vols/itom/core/suite-install/{suite_ingress _yaml_directory}/objectives
   ```

   ```
   kubectl delete -f xxxx-nginx-ingress.yaml
   ```

   ```
   kubectl create -f xxxx-nginx-ingress.yaml
   ```

FAQ

Question

Do I have to add exceptions to the firewall?

Answer

Browsers access HPE CDF via the HTTPS ports (TCP/5443). End users need to add it to the firewall exception policy.

User Management and Authentication

This section provides information related to user management and authentication.

Authentication Model

The ITOM Container Deployment Foundation supports the following authentication methods:

- Username and password authentication

  In an out-of-the-box default installation, the ITOM Container Deployment Foundation requires users to enter username and password credentials to gain access to the application.

- LDAP authentication

  You can integrate the ITOM Container Deployment Foundation to an LDAP directory service to share contact information across your network.

Authorization

This section provides information related to user authorization in ITOM Container Deployment Foundation.

Authorization Model

Access to the ITOM Container Deployment Foundation resources is authorized based on the user's following settings:

- User name
- Session and inactivity timer timeouts

FAQ

Question

Can the ITOM Container Deployment Foundation inherit users' information and authorization profiles from an external repository, such as LDAP?

Answer

No.

Data Integrity

The database server is used as a simple data store and is responsible for all persistent storage. While the database contains definitions describing business logic, no processing is actually performed in this tier, other than create, read, update, and delete (CRUD) operations in response to requests from the ITOM Container Deployment Foundation. Referential integrity is enforced by the application, thereby protecting transactions. In addition, the database captures a complete audit log of all changes to data.

The data backup procedure is also an integral part of data integrity and while the ITOM Container Deployment Foundation does not provide native backup capabilities, the following guidelines should be considered:

- Database backup is especially important before critical actions such as upgrades.
- Backup files should be stored properly according to the industry best practices to avoid unauthorized access.
- Since database backups can be a resource intensive process, HPE strongly recommends to avoid running backups during peak demand times.

Encryption

This section provides information on data encryption in the ITOM Container Deployment Foundation platform.

TLS/SSL Data Transmission

An IDM server is used for the 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 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.

The ITOM Container Deployment Foundation was configured to use TLS/SSL to transmit data between the server and browsers.

Customers can change the default value of SSL CIPHER through the following steps:

1. On the master node, change the `ssl-ciphers` value in file `$K8S_HOME/objectdefs/nginx-ingress.yaml`.
2. Recreate the ingress container with the commands below:

   ```
   kubectl delete -f $K8S_HOME/objectdefs/nginx-ingress.yaml
   ```

   ```
   kubectl create -f $K8S_HOME/objectdefs/nginx-ingress.yaml
   ```

Encryption of stored database fields

The ITOM Container Deployment Foundation uses proprietary algorithms when encrypting data stored in the database and uses IDM to manage user names and passwords.

Docker logs

This section provides information related to Docker logs.

Log and trace model

Recommendations:

- Pay attention to the log level and do not leave tracing or debug parameters enabled unnecessarily.
- Pay attention to log rotation/switching.

Log rotation

Follow the steps below to configure the maximum log file size and maximum log file number. By default, the maximum log file size is 10 MB, and the maximum number of the log file is 5.

1. Open the `docker` file with the following commands.

   ```
   cd /opt/kubernetes/cfg
   ```

   ```
   vim docker
   ```

2. Change the value of `max-size` and `max-file` in the parameter `DOCKER_LOG_OPTS`.

   For example:

   ```
   DOCKER_LOG_OPTS="--log-driver=json-file --log-opt
   ```

```
labels=io.kubernetes.container.name,io.kubernetes.pod.uid --log-opt max-
size=12m --log-opt max-file=6"
```

3. Restart Docker to enable the changes with the following command:

```
systemctl restart docker
```

> **Note:** The default maximum log size number and maximum log file number is recommended. Do not set a large number for the `max-size` and `max-file`. Too large maximum size and maximum file number may affect the free disk size.

Enable firewall on a running node

Follow the steps below on each running node to enable firewall.

On the NFS server

Run the following commands to enable 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
```

On the running master nodes

For the single-master node deployment

Run the following commands to enable firewall on the running 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 --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNodeIP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --reload
```

For the multiple-master node deployment

Run the following commands to enable firewall on each running 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 --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode1IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode2IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode3IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --reload
```

On the running worker nodes

For the single-master node deployment

Run the following commands to enable firewall on each running 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 --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNodeIP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --reload
```

For the multiple-master node deployment

Run the following commands to enable firewall on each running 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 --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode1IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode2IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source
address="<MasterNode3IP>/32" port protocol="tcp" port="10255" accept'

firewall-cmd --reload
```

Data backup for the single-master cluster

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

For 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.

Network and Communication Security

HPE recommends that you add the iptables rules listed below to the following below tables.

> **Important:** Apart from the listed ports, all other ports should be blocked at the localhost level.

Required ports	Service	Add rules on Server	Direction	Short description
111	NFS	NFS server	Nodes ->NFS Server	NFS server port access by all nodes
2049	NFS	NFS server	Nodes ->NFS Server	NFS server port access by all nodes
2380	Etcd	Master Node	Master<-> Master	Etcd service port for etcd cluster communication

Required ports	Service	Add rules on Server	Direction	Short description
4001	Etcd	Master Node	Nodes -> Master	Etcd service port for connection from client
4194	Kubernetes	All Nodes in Cluster	Localhost only	cAdvisor for local kubelet
5000	Private Registry	All Nodes in Cluster	Localhost only	Registry port for local host
5443	MngPortal	Ingress Node	All -> Ingress Node	The port exposed on ingress node. all clients could access this port
8200	Vault	Master Node	Nodes->Master	Vault port for client connection
8443	kubernetes	Master Node	Nodes->Master	API server port for client connection
10250	Kubernetes	All Nodes in Cluster	Nodes->Nodes	Kubernetes port for internal communication
10251	Kubernetes		Nodes->Nodes	Kubernetes port for internal communication
10252	Kubernetes		Nodes->Nodes	Kubernetes port for internal communication
10255	Kubernetes		Nodes->Nodes	Kubernetes port for internal communication
20048	NFS	NFS server	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
```

Change the host name of the installed cluster node

Follow the steps below to change the host name of the installed cluster node.

Note: For ITOM Container Deployment Foundation, the ingress is in core namespace.

1. Generate a new certificate for the new host name.
   - You can use a certificate signed by your enterprise RootCA.
   - Alternatively, you can generate a self-signed certificate by following the steps below:

     Log on to one of the master nodes and run the following commands:

     ```
     cd $K8S_HOME/ssl

     openssl genrsa -out ${hostName}.key 4096

     openssl req -new -key ${hostName}.key -subj "/CN=${hostName}" -out
     ${hostName}.csr

     openssl x509 -sha256 -req -in ${hostName}.csr -CA "${cacert}" -CAkey
     "${cakey}" -CAcreateserial -out ${hostName}.crt -days 365

     chmod 0400 ${hostName}.crt ${hostName}.key

     rm -f ${hostName}.csr
     ```

2. Follow the commands below to replace the defect nginx secret.

   ```
   kubectl delete secret nginx-default-secret -n core

   echo "

   apiVersion: v1

   kind: Secret

   metadata:

   name: nginx-default-secret

   namespace: core

   data:

   tls.crt: `base64 -w 0 $K8S_HOME/ssl/${hostName}.crt`
   ```

```
tls.key: `base64 -w 0 $K8S_HOME/ssl/${hostName}.key`
"| kubectl create -f -
```

3. Follow the commands below to restart the ingress container.

```
kubectl delete -f $K8S_HOME/objectdefs/nginx-ingress.yaml
kubectl create -f $K8S_HOME/objectdefs/nginx-ingress.yaml
```

4. On all master nodes, replace the `hostname` in the suite.yaml, mng-portal.yaml, and idm.yaml under the `$K8S_HOME/objectdefs` directory with the following commands.

```
sed -i -e "s/old_hostname/new_hostname/g" $K8S_HOME/objectdefs/idm.yaml
sed -i -e "s/old_hostname/new_hostname/g" $K8S_HOME/objectdefs/suite.yaml
sed -i -e "s/old_hostname/new_hostname/g" $K8S_HOME/objectdefs/mng-portal.yaml
```

5. On one of the master nodes, restart the idm, suite, and mng-portal with the following commands.

```
kubectl delete -f $K8S_HOME/objectdefs/idm.yaml
kubectl delete -f $K8S_HOME/objectdefs/suite.yaml
kubectl delete -f $K8S_HOME/objectdefs/mng-portal.yaml
kubectl create -f $K8S_HOME/objectdefs/idm.yaml
kubectl create -f $K8S_HOME/objectdefs/suite.yaml
kubectl create -f $K8S_HOME/objectdefs/mng-portal.yaml
```

Customize the parameters for kubelet

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
```

Restart the ITOM Container Deployment Foundation

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
   ```

Administer the Operations Bridge Suite

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

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 Foundation installation. 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.

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.

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.

Configure high availability

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

> **Important:** Only by using *all* of these techniques, your system will be highly available.

- **Highly available redundant storage**. Use a redundant NFS server for your container deployment.
- **Highly available database instances**. Use a redundant external database.
- **Kubernetescluster with multiple master nodes**. See "High availability of the Kubernetes cluster" below.
- **Operations Bridgeservices 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. Scaling is currently only available for BVD. See "Scale a BVD deployment horizontally" on the previous page and the example in Hardware requirements.
- **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" on page 46.

Other Operations Bridge capabilities

Apart from BVD, no other capability supports 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. For Performance Engine and Operations Bridge Reporter, this can take several minutes. For Operations Manager i, this can take more time, depending on the amount of events and CIs.

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 Foundation installation. Usually, this is the master node's FQDN.

2. Log in as the admin user.
3. Go to **ADMINISTRATION**> **LDAP**. In the Organization List, click **Provider**.

4. Click **ADD CONFIGURATION** to enter a valid LDAP configuration. For details on what to enter for each LDAP setting, see "LDAP settings" below.
5. Click **SAVE**.
6. 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 Foundation installation. 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.

LDAP settings

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

Setting	Description
LDAP Server Information	
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.
Connection Security	Select **Connection Security: SSL** 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.
User Authentication	
User Search Base	Parameters to indicate which attributes are to be included in the user search. **Example**: `CN=Users`
User Name	Name of field that contains the user name. **Example**: `sAMAccountName`
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))`.
Follow Referral	Select to follow LDAP referrals to another server that offers the requested information.
Search Subtree	Select to search the subtree below the base DN (including the base DN level).
User Attributes	
Common Name	Common 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`
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`

User Group	
Group Membership	List of comma-separated LDAP attributes to find groups in a user profile. **Example**: `member,uniqueMember`
Group Name	LDAP name used to identify objects of the type group. **Example**: `cn`
Group Search Filter	LDAP pattern to use when searching for a group account. **Example**: `(objectclass=group)`

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.

For examples specific to OMi and OBR, see "Example: Access OMi command line interfaces" below and "Example: Access OBR command line interfaces: " below.

Example: Access OMi command line interfaces

1. Find the container that contains the CLI.

   ```
   [root@master]# kubectl get pods -n opsbridge1 -o name | grep omi
   pod/<container_id>
   ```

2. Start the shell.

   ```
   [root@master]# kubectl exec -ti <container_id> -n opsbridge1 -c omi bash
   omi:/ #
   ```

3. 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:

   ```
   [root@master]# kubectl get pods -n opsbridge1 -o name | grep obr-server
   pod/<container_id>
   ```

2. Start the shell.

   ```
   [root@master]# kubectl exec -ti <container_id> -n opsbridge1 -c obr-server bash
   obr-server:/ #
   ```

3. 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 delivers additional information and advanced configuration possibilities.
To access the RTSM 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.

Integration

You can integrate the capabilities of the Operations Bridge Suite to fully leverage the suite's benefits. See the following sections for information about OMi, BVD, PE, or OBR integrations.

OMi integrations

Integrating Operations Bridge Manager (OMi) with other software products is a great way to extend your IT management capability. You can integrate OMi with every component of the Operations Bridge Suite Premium:

Integration	Documentation
OMi - BVD	See the Integrate section in the OMi Online Help or the BVD Online Help.
OMi - PE	See "PE integrations" on the next page.
OMi - OBR	See the OBR Integration Guide.
OMi - OA	See "Integrate Operations Agent with OMi" below.

Most major integrations between OMi and other HPE products are described in *the OMi Integrations Guide*.

For a complete list of available product integrations, see the Integrations Catalog on HPE Software Support.

Integrate Operations Agent with OMi

If Operations Agent is installed on one of the master or worker nodes, additional steps are required to integrate Operations Agent with OMi .

> **Caution:** If the agent is installed on an external node instead, these steps are **not** required.

Configure the agent to use a different port than the one used by default, and configure OMi to use this non-default port:

1. Make sure that your system has a virtual external hostname (for example `kubecluster.example.com`), which is different than the physical node name (for example `dock.example.com`).
   To do this, add an entry to the DNS server for the virtual hostname of the suite installation. This

hostname must be resolvable for all agents that are installed on one of the Kubernetes nodes.

2. Configure the agent to use a different server port (for example 384):

   ```
   ovconfchg -ns bbc.cb -set SERVER_PORT 384
   ```

3. Configure Operations Agent to use the virtual external hostname (`kubecluster.example.com`) as the management server name:

   ```
   opcactivate.sh -s <virtual_external_hostname>
   ```

4. Configure OMi to connect to the master node with the specified port (for example 384):

   ```
   ovconfchg -ns bbc.cb.ports -set PORTS dock.example.com:384
   ```

BVD integrations

Business Value Dashboard (BVD) can be integrated out-of-the-box with OMi and Operations Connector (OpsCx). You can also create your own integrations for any data source by writing an adapter for BVD. For more information, see the BVD Help Center.

PE integrations

Performance Engine (PE) can be integrated with OMi.

After installing Performance Engine, you must configure the infrastructure settings to integrate Performance Engine with OMi.

1. In OMi, go to **Administration** > **Setup and Maintenance** > **Infrastructure Settings**.
2. Select the **Applications** context.
3. Select the **Performance Engine** from the drop-down list.
4. In the Performance Engine Node Infrastructure Setting, click  **Edit**.

   You can use this parameter to configure the Performance Engine Node details from which OMi Performance Dashboard must request data. The value is required in the format `http://<external_access_host>:<port>`.

   The default port is `31387`.
5. Click **Save**. Alternatively, click **Restore Default** > **Save**, to reset the default values.
6. In the Performance Engine Node password, click  **Edit** and specify the password that you set

during the Performance Engine configuration.

7. Click **Save**.

OBR integrations

Operations Bridge Reporter can be integrated with OMi, and with other HPE products. For a complete list of available product integrations, see the Integrations Catalog on HPE Software Support.

The OBR - OMi integration is documented in the *OBR Integration Guide*.

Troubleshoot

This section provides information that can help you troubleshoot problems you may encounter when installing and using the ITOM Container Deployment Foundation and the Operations Bridge Suite.

Manual verification commands

The following commands can be used to troubleshoot the ITOM Container Deployment Foundation and the Operations Bridge Suite container deployment, for example to list namespaces and services.

`/opt/kubernetes/bin/kube-status.sh`
Displays the status of the K8S cluster.

`/opt/kubernetes/bin/kube-stop.sh`
Stops the K8S cluster.

`/opt/kubernetes/bin/kube-restart.sh`
Restarts the K8S cluster.

`/opt/kubernetes/bin/kube-start.sh`
Starts the K8S cluster.

`kubectl`
The command to interact with Kubernetes (K8S).

> **Tip:** To shorten the `kubectl` command, run the following command:
>
> `ln -s /usr/bin/kubectl /usr/bin/kl`
>
> This enables you to type `kl` instead of `kubectl`.

`kubectl cluster-info`
Summarizes information about some of the services that are running on the cluster, including Kubernetes master, KubeDNS for service discovery, and the endpoints of the KubeRegistry (if you are running a registry).

`kubectl get nodes`

Lists all nodes in the cluster.

`kubectl describe nodes <node_IP>`
Provides more specific information about the node, such as labels, events, capacity, CPU, memory, the maximum number of pods that the node can support, system information on the node, external IP address, the pods that are running, the list of namespaces, and resources.

`kubectl get pods`
Lists all pods in the default namespace (used to separate the Container Deployment Foundation services from the deployed suites).

`kubectl get pods -n=<namespace>`
Lists all the pods that are running on the specified namespace.

For example, run `kubectl get pods -n=opsbridge1` to get a list of the pods running in the namespace `opsbridge1`.

`kubectl get pods --all-namespaces`
Lists all the pods that are currently running in the cluster.

`kubectl describe pod <pod_name> -n=<namespace>`
Displays details about a specified pod in a specified namespace, such as the image it is running, the port it is exposing, and the command (/hyperkube) that is running inside the container itself with their options, volumes, and more.

`kubectl exec <pod_name> -c <container> -n <namespace>`
Executes a command in the specified container. If no container is specified, the first container in the pod is selected.

> **Example**: `kubectl exec omi-1949254658-p3ipj -c omi -n opsbridge1 bash -ti`
>
> Executes a bash shell in the OMi container with the pod name `omi-1949254658-p3ipj` and the namespace `opsbridge1`. By executing a bash shell in the OMi container, you can call CLIs from inside the container. For more information, see the *Operations Bridge Suite Administration Guide*.

`kubectl get services --all-namespaces`
Displays all the services running in the cluster.

`kubectl logs <pod_name> -n=<namespace>`
Displays the log output for the specified pod.

Log files

To troubleshoot your issue, you can review the following log files.

Installation

```
/opt/kubernetes/install-<date><time>.log
```

NFS share

- /var/vols/itom/log/omi/opt/HP/BSM/log/topaz_all.log
- /var/vols/itom/log/omi/opt/HP/BSM/log/jboss7_boot.log
- /var/vols/itom/log/omi/opt/HP/BSM/log/supervisor/nanny_all.log
- /var/vols/itom/log/opsbridge-opsbridge/pe/logs

Login

```
/var/vols/itom/log/omi/opt/HP/BSM/log/jboss/login.log
```

OBR

Configuration

```
<NFS_conf_directory>/OBR/reporting/... (OBR server)
<NFS_conf_directory>/OBR/reporting-collector/... (OBR reporting collector)
<NFS_conf_directory>/OBR/reporting-content/... (OBR content pack artifacts)
```

Logs

```
<NFS_log_directory>/OBR/reporting/... (OBR server)
<NFS_log_directory>/OBR/reporting-collector/... (OBR reporting collector)
```

Data

```
<NFS_data_directory>/OBR/reporting/... (OBR server)
<NFS_data_directory>/OBR/reporting-collector/... (OBR reporting collector)
<NFS_data_directory>/OBR/MgmtDB/... (OBR PostgreSQL instance)
```

Support toolset

The support toolset helps to collect information about Docker, Kubernetes, suites, commands, directories, and files as listed below:

- Docker: containers, inspect, docker service systemd logs
- Kubernetes: nodes, pods, namespaces, images, containers, cluster-info, describe, logs
- Suite: suite-db dump, suite data, modules, product deployments, features

- Commands defined by users
- Directories and files defined by users

You can view the summary information on a console. For the detailed output information, you can view them in an encrypted tar file.

Use the toolset

Run the following commands to use the toolset:

1. `cd <K8S_HOME>/tools/support-tool`
2. `# ./support-dump [ -c <dump_filename_with_path> ] [-u <username> [-p <password>]] [-P <package_password>]`
3. Unpack the dumpfile:

   `# dd if=xxxx.des3 |openssl des3 -d -k <package_password>|tar zxf -`

> **Example**
>
> - Create a dump file with the default file name in the default directory.
>
> `# ./support-dump`
>
> - Create an example dump file `dump.des3` in the directory `/var/test`.
>
> `# ./support-dump -c /var/test/dump.des3`
>
> - Create a dump file with the user name `admin` and the password `123456`. Additionally, specify the package password `abcdef`.
>
> `# ./support-dump -u admin -p 123456 -P abcdef`

Configuration file

The support toolset provides a configuration file with some predefined `[commands]`, `[files]`, and `[dirs]` to specify your deployment's information. You can also define your own commands, files, and directories in the configuration file. Alternatively, create other configuration files in the same directory. The default configuration file is `conf/supportdump.config`.

- The outputs of the same command will be saved into one file. For example, the all the outputs of the `cat` command will be saved in the `cat.out` file.
- All directories, files, and outputs of commands will be stored in the `<local_ip>-<NodeType>/os` directory.

- Wildcards can be used in file and directory names. For example `/etc/sysconfig/network-scripts/ifcfg-*`
- Single environment variables are supported.
- One or multiple files (separated by spaces) following a directory will be excluded from the support toolset collection.

> **Example** :
>
> `<K8S_HOME>/cfg *_User.json`
>
> The support toolset collects all files and directories located in `<K8S_HOME>/cfg` except the `* _ User.json` file(s).

Dump file

The default support dump file is called `dmp/support_data_YYYYMMDD-hhmmss.des3`. The dump file contains the `support_data_YYYYMMDD-hhmmss.log` of the running support toolset and the `ITOM_Core_Platform` directory for the dump files. The table below shows the dump files in the `ITOM_Core_Platform` directory.

Common problems and limitations

You may encounter the following problems and limitations when installing or administering the Container Deployment Foundation and the Operations Bridge Suite.

Management Portal is not accessible

Description

After the installation of the Container Deployment Foundation, the Management Portal cannot be accessed at https://`<external_access_host>`:5443.

Possible solutions

- Make sure you entered the correct URL and port.
- Make sure you can access the host: `ping <external_access_host>`
- Check your browser's proxy settings.
- Check the installation logs in `/opt/kubernetes/install-<timestamp>.log`.
- Empty the NFS folder and then reinstall the Container Deployment Foundation.

- See also "Management Portal is not accessible: nginx controller is Pending" below, "Management Portal is not accessible: Gateway time out" on the next page and "Login to Management Portal is not possible: IDM service is not ready yet" on the next page.

Management Portal is not accessible: nginx controller is Pending

Description

After the installation of the Container Deployment Foundation, the Management Portal cannot be accessed at https://<*external_access_host*>:5443.
When running `kubectl get pods --all-namespaces`, the nginx ingress controller status is Pending.

Cause and solution

The map hash bucket size might be too small. Check if that is the case by running the following commands:

```
kubectl describe nginx-ingress-controller-u69gg
```

```
kubectl logs nginx-ingress-controller-u69gg
```

If an error is displayed similar to `nginx: [emerg] could not build map_hash`, increase the `map_hash_bucket_size` as follows:

1. Access the file `/opt/kubernetes/objectdefs/nginx-ingress.yaml`
2. Locate the specified `map_hash_bucket_size` (32 by default) and increase it, for example to 128
3. Run the following commands to recreate the `nginx-ingress.yaml` file:

   ```
   kubectl delete –f /opt/kubernetes/objectdefs/nginx-ingress.yaml
   ```

   ```
   kubectl create –f /opt/kubernetes/objectdefs/nginx-ingress.yaml
   ```

4. *Optional*. If you get a warning about failed scheduling, the scheduling constraints could not be fulfilled. Execute the following command to fix this:

   ```
   kubectl label nodes role=loadbalancer –all
   ```

   The nginx pod container should then be started automatically.

5. After the OMi configuration, you must repeat steps 2 and 3 for the OMi nginx controller located at `/var/vols/itom/core/suite-install/opsbridge/output/suite-ingress-controller-configmap.yaml`

Management Portal is not accessible: Gateway time out

Description

After the installation of the Container Deployment Foundation, the Management Portal cannot be accessed at https://*<external_access_host>*:5443. The Docker daemon cannot be started, and displays the error message `Gateway time out` when logging into IDM.

Cause and solution

Kubernetes might not be running. Run the following commands to start Kubernetes:

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

Login to Management Portal is not possible: IDM service is not ready yet

Description

After the installation of the Container Deployment Foundation, the Management Portal cannot be accessed at https://*<external_access_host>*:5443. The login failure error `The IDM service is not ready yet` is displayed, and the pods `autopass-lm`, `idm`, and `suite-installer` all have the status `CrashLoopBackOff`.

Solution

1. Run the following command:

   ```
   kubectl delete -f autopass-lm.yaml; kubectl delete -f autopass-pg.yaml; kubectl delete -f idm.yaml; kubectl delete -f idm-pg.yaml; kubectl delete -f suite.yaml
   ```

2. Delete the subfolders located in the NFS subdirectories `<NFS_HOME>/baseinfra-1.0/autopass_db`, `<NFS_HOME>/baseinfra-1.0/idm_db`, and `<NFS_HOME>/baseinfra-1.0/suite_db`.
3. Run the following command:

   ```
   kubectl create -f idm-pg.yaml; kubectl create -f idm.yaml; kubectl create -f autopass-pg.yaml; kubectl create -f autopass-lm.yaml; kubectl create -f suite.yaml
   ```

Reboot does not work. Pods are in status CrashLoopBackOff.

Description

After attempting to reboot, the pods have the status CrashLoopBackOff.

Cause and solution

This is related to the vault-renewal container, which does not get a valid token. You have to delete the failed pods. Once the pods are deleted, they are recreated automatically and should run without error.

You can get the status of all pods with the following command:

```
kubectl get pods --all-namespaces
```

First delete all failed database related pods (`suite-db`, `idm-postgresql`, `postgresql-aplm`). Next, delete all failed pods within the namespace `core`. After that delete all failed pods within the namespace `opsbridge`, starting with `postgres`, `ucmdb`, `omi`, `redis`, `bvd`, `obr-server`, `obr-rc`).

Use the following command to delete the failed pods within the namespaces specified above:

```
kubectl delete pod <pod_name> --namespace <pod_namespace>
```

"502 Bad Gateway" error when attempting to launch OMi

Description

After the installation of the Operations Bridge Suite, a 502 Bad Gateway error is displayed when trying to access OMi.

Cause and solution

The 502 error is displayed because OMi is not yet up and running. Depending on the host machine, it might take up to one hour for OMi to start after the initial configuration.

No server connection: invalid character "{" in host name

Description

A connection to the server could not be established. The log displays that the invalid character "{" is used in the host name.

Cause and solution

The firewall might still be enabled on the NFS server. Make sure that the firewall is disabled.

Pod is in ImagePullBackOff or ErrImagePull status: Image not found

Description

After the installation of the Container Deployment Foundation, one of the pods has the status `ImagePullBackOff` or `ErrImagePull`. When running the command `kubectl describe pod <pod_name> -n <namespace>`, the following error message is displayed:

```
Image <image_name> not found
```

Cause and solution

Make sure the images are pushed into the private docker registry. To confirm, run the following command:

```
docker pull <image_name>
```

Pod is in ImagePullBackOff or ErrImagePull status: Error while pulling image

Description

After the installation of the Container Deployment Foundation, one of the pods has the status `ImagePullBackOff` or `ErrImagePull`. When running the command `kubectl describe pod <pod_name> -n <namespace>`, the following error message is displayed:

```
Error while pulling image: Get http://localhost:5000/v1/repositories/xxx: dial tcp [::1]:5000: getsockopt: connection refused
```

Cause and solution

To resolve this issue, delete the Docker registry and the registry proxy pods, and then restart them.

Worker node installation fails with a Flannel related error

Description

Setting up one or multiple worker nodes fails during the Container Deployment Foundation installation due to an error related to Flannel.

Cause and solution

To troubleshoot and resolve this issue, do the following:

- Double check if the FQDN is resolved to the correct IP address on the master node.
- On the master node, run `kube-restart.sh`
- Reinstall the worker node from the Management Portal.

"503 nginx error" when attempting to run the Suite Installer

Description

After the installation of the Container Deployment Foundation, a 503 Nginx error is displayed when trying to access the Suite Installer.

Cause and solution

This error might be displayed because the time on the master and worker nodes is different. To resolve this issue, synchronize the time on your nodes by using, for example, NTP or VMWare tools.

Worker node does not start

Description

Due to missing disk space, the worker node does not start.

Cause and solution

To solve this problem, make sure that the `/` and `/var` directories have at least 5 GB free disk space.

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Administration Guide (Operations Bridge Suite 2017.08)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to ovdoc-asm@hpe.com.

We appreciate your feedback!