



# Operations Bridge Suite

Software Version: 2017.08

## Installation Guide

Document Release Date: November 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

© 2015 - 2017 Hewlett Packard Enterprise Development LP

### 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

Overview .....	4
Plan your suite deployment .....	7
Prepare for the installation .....	9
Enable your Docker Hub account .....	9
Meet the hardware requirements .....	10
Meet the software requirements .....	14
Unzip the ITOM CDF installation package .....	16
Install the Container Deployment Foundation .....	17
Configure the install.properties file .....	17
Optional. Install an NFS server .....	25
Install ITOM CDF on the (first) master node .....	27
Install ITOM CDF additional master nodes .....	28
Add worker nodes .....	29
Verify the ITOM CDF installation .....	30
Install the Operations Bridge Suite .....	32
Prepare the suite images .....	32
Run the suite installer .....	35
Activate a suite license .....	48
Verify the suite installation .....	49
Edit the installation .....	51
Upgrade .....	53
Reconfigure .....	58
Uninstall .....	59
Troubleshoot .....	60
Send documentation feedback .....	70

## Overview

HPE Operations Bridge Suite helps transform your IT organization from a cost function to a value creator by simplifying and automating IT operations. The suite enables you to sense your environment through automated discovery and monitoring. The activities in your environment can be analyzed to predict and solve critical problems and increase performance.

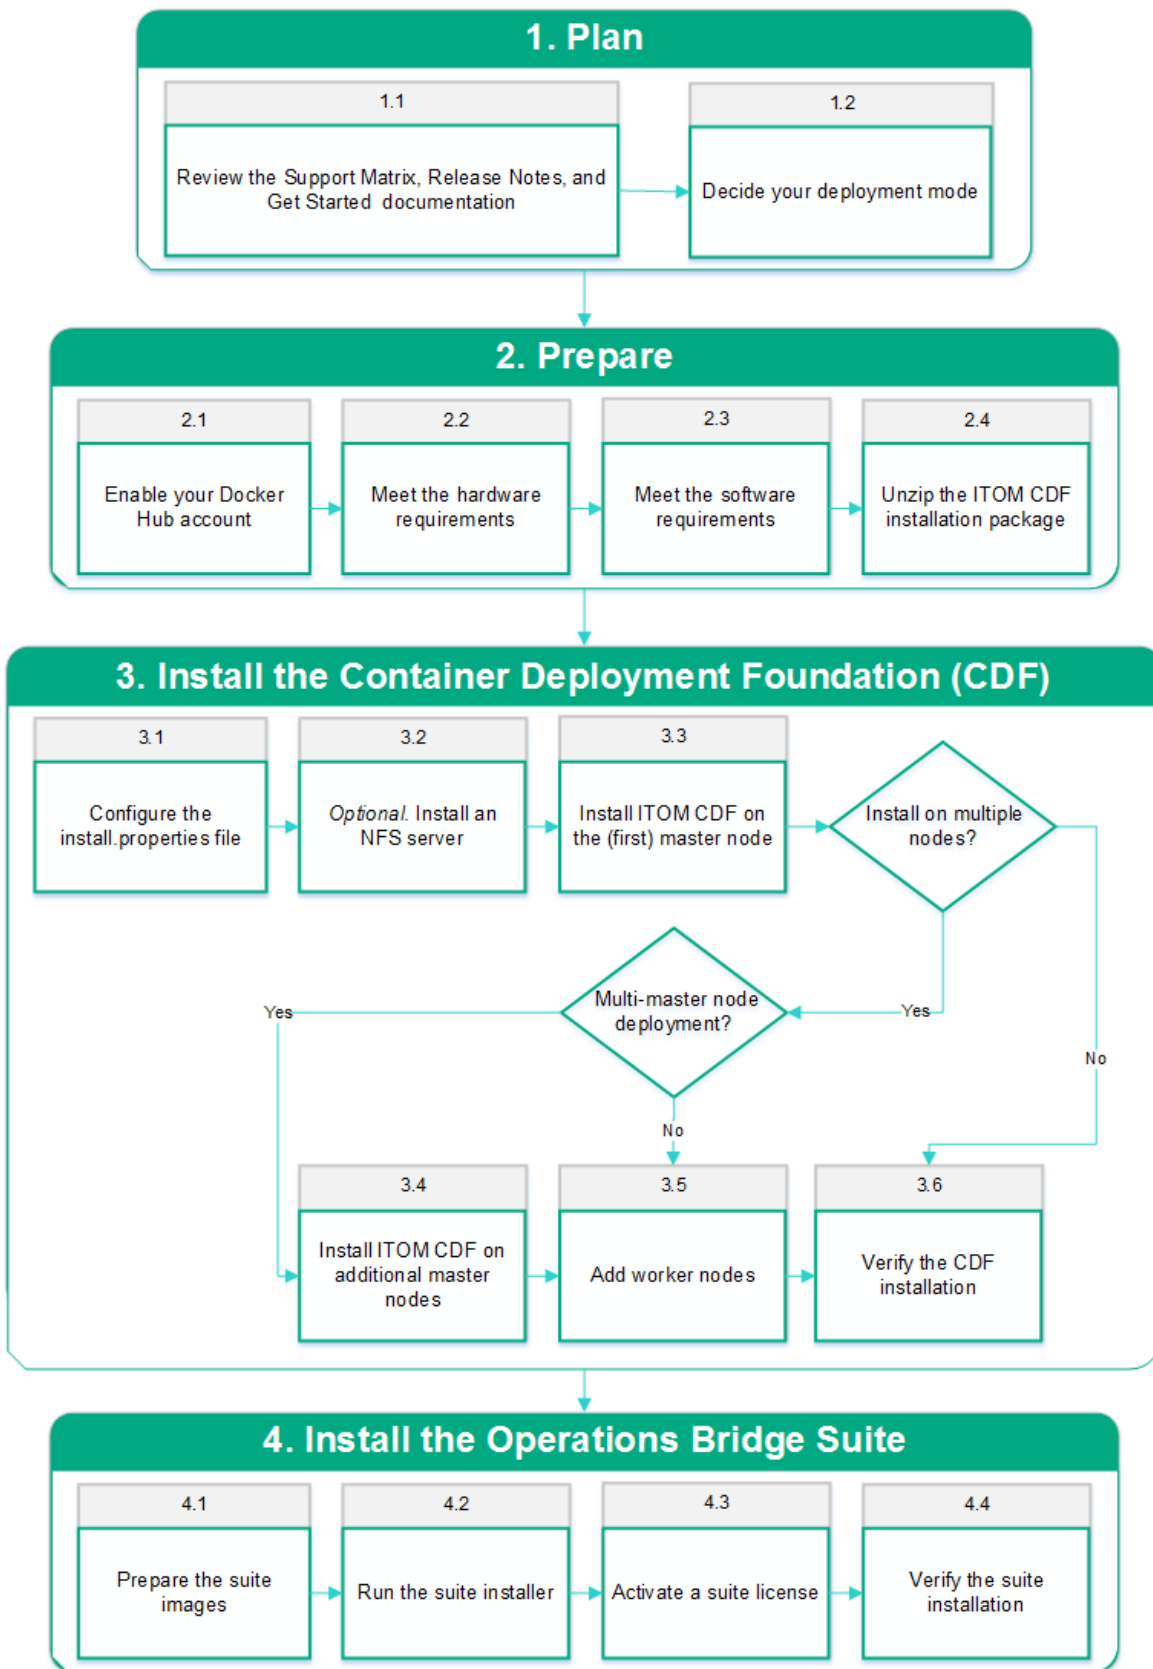
The Operations Bridge Suite is installed in the containerized mode that leverages technology based on Docker and Kubernetes. In this mode, each suite capability is deployed as a containerized application that is integrated with other suite capabilities. You first install a container management framework (referred to as ITOM Container Deployment Foundation (CDF)) and then install the Operations Bridge Suite from a graphic user interface based on this framework. The Operations Bridge Suite capabilities are deployed quickly, requiring little user intervention.

To learn more about the capabilities of the Operations Bridge Suite, see the *Operations Bridge Suite 2017.08 Concepts Guide*.

**Important:** The container-based deployment currently allows you to install the capabilities of the Express and Premium versions of the Operations Bridge Suite. Capabilities of the Ultimate version must be installed separately.

## Installation overview

Your installation steps will vary depending on your deployment mode. Use the following workflow diagram to help you decide the steps to follow depending on your deployment mode:



The following table lists detailed information about the steps illustrated in the workflow diagram. You can click each step to view detailed instructions.

State	Steps
Plan	<p>1.1 <a href="#">"Review the support matrix" on the next page, the Release Notes and the Get started documentation.</a></p> <p>1.2 <a href="#">"Decide your deployment mode" on the next page</a></p>
Prepare	<p>2.1 <a href="#">"Enable your Docker Hub account" on page 9</a></p> <p>2.2 <a href="#">"Meet the hardware requirements" on page 10</a></p> <p>2.3 <a href="#">"Meet the software requirements" on page 14</a></p> <p>2.4 <a href="#">"Unzip the ITOM CDF installation package" on page 16</a></p>
Install ITOM CDF	<p>3.1 <a href="#">"Configure the install.properties file" on page 17</a></p> <p>3.2 <a href="#">"Optional. Install an NFS server" on page 25</a></p> <p>3.3 <a href="#">"Install ITOM CDF on the (first) master node" on page 27</a></p> <p>3.4 <a href="#">"Install ITOM CDF additional master nodes" on page 28</a></p> <p>3.5 <a href="#">"Add worker nodes" on page 29</a></p> <p>3.6 <a href="#">"Verify the ITOM CDF installation" on page 30</a></p>
Install the Operations Bridge Suite	<p>4.1 <a href="#">"Prepare the suite images" on page 32</a></p> <p>4.2 <a href="#">"Run the suite installer" on page 35</a></p> <p>4.3 <a href="#">"Activate a suite license" on page 48</a></p> <p>4.4 <a href="#">"Verify the suite installation" on page 49</a></p>

# Plan your suite deployment

The Container Deployment Foundation allows you to deploy a suite in an environment that is comprised of one or multiple master nodes and multiple worker nodes for load balancing and failover purposes. Client requests are sent to the load balancer, which redirects the requests to the master nodes, and the master nodes then sends the requests to the worker nodes.

To plan your suite deployment, review the support matrix, the supported configurations, and the ITOM CDF configuration parameters.

## Review the support matrix

1. Download the [Support Matrices for Operations Center products](#).
2. Open SUMA.htm and select **Operations Bridge Suite (container deployment)** from the product list.

The master node and each worker node must run one of the operating systems listed when filtering for the Container Host component.

## Decide your deployment mode

Master nodes coordinate all activity in your cluster, such as scheduling applications, maintaining applications' desired state, scaling applications, and rolling out new updates. Worker nodes run the applications. A node is a VM or a physical computer that serves as a worker machine in a Kubernetes cluster.

A Kubernetes cluster that handles production traffic should have a minimum of one master and three worker nodes. By deploying multiple master and worker nodes, you can make your system highly available. For more information, see the [Kubernetes HA documentation](#).

The Operations Bridge Suite uses NFS to store run time, configuration, and log data. You can use a separate NFS server, or use a master node as NFS server.

Additionally, the suite uses databases for the CDF and for the suite components. You can use embedded databases that run in containers, or you can connect to externally installed databases.

### Single node

In a testing environment, you can use one system as master and worker node (single node deployment) with the system also serving as NFS server.

### **Single master multi-worker deployment**

You can use one master node and multiple worker nodes to have multiple nodes on which you can run the capabilities' workloads on. You can decide if you want to use a separate NFS server, or if you want to use the master node as NFS server.

### **Multi-master multi-worker deployment**

In a production environment, you use multiple master nodes, multiple worker nodes, and (highly recommended) a separate NFS server.

To find out more about how to calculate your minimum system requirements, see ["Meet the hardware requirements" on page 10](#).



## Prepare for the installation

Perform the following steps to prepare for the installation of the ITOM ITOM Container Deployment Foundation (CDF) and the Operations Bridge.

- "Enable your Docker Hub account" below
- "Meet the hardware requirements" on the next page
- "Meet the software requirements" on page 14
- "Unzip the ITOM CDF installation package" on page 16

## Enable your Docker Hub account

You must create a Docker Hub account and then ask HPE to enable your Docker Hub account so that you can download (pull) Operations Bridge Suite images from Docker.

1. Create a Docker account on <https://hub.docker.com>.
2. Log in to <https://hub.docker.com> with your Docker ID.
3. In the top right corner of the page, click Settings and take a screenshot to include your Docker ID and the linked email address.
4. Send the following information together with the screenshot to the HPE software fulfillment and licensing team for your region to enable your Docker account:
  - Your company name
  - Your contact information and HPE Passport email address
  - Your HPE customer SAID (must be valid and active)
  - HPE ITOM Suite edition (Operations Bridge Suite)

Send your email to one of the addresses below, based on your region:

- **AMERICAS:** [dockersupport.ams@hpe.com](mailto:dockersupport.ams@hpe.com)
- **APJ:** [dockersupport.apj@hpe.com](mailto:dockersupport.apj@hpe.com)
- **EMEA:** [dockersupport.emea@hpe.com](mailto:dockersupport.emea@hpe.com)

HPE will enable your Docker ID and send you a confirmation.

# Meet the hardware requirements

To fully prepare your system for the suite installation, review the following hardware requirements.

## Hardware requirements

The minimum hardware requirements for your system depend on the capabilities you decide to install. The total minimum requirements are calculated by summing up the requirements per capability.

The sum of all worker node resources must equal or exceed the total requirements for the capabilities. As a best practice, HPE recommends not to run workloads on the master node.

The required resources for OMi depend on the size of your deployment.

- Small OMi deployment: up to 2000 monitored nodes send events to OMi
- Medium OMi deployment: up to 5000 monitored nodes send events to OMi
- Large OMi deployment: more than 5000 monitored nodes send events to OMi

Component	RAM	Processors	Disk space
<b>CONTAINER DEPLOYMENT FOUNDATION</b> (on the master nodes)			
Container Deployment Foundation	16 GB	8 CPU cores	200 GB
NFS server (if the master is used as NFS server)	-	-	100 GB
<b>CAPABILITIES</b> (on the worker nodes)			
Operations Bridge Manager (OMi) - small deployment	16 GB	4 CPU cores	50 GB
Operations Bridge Manager (OMi) - medium deployment	27 GB	6 CPU cores	75 GB
Operations Bridge Manager (OMi) - large deployment	40 GB	8 CPU cores	100 GB
Business Value Dashboard (BVD)	6 GB	4 CPU cores	30 GB
Performance Engine (PE)	8 GB	4 CPU cores	100 GB
Operations Bridge Reporter (OBR) - small deployment with about 100 nodes			

OBR Server	8 GB	4 CPU cores	150 GB
PostgreSQL	1 GB	1 CPU core	30 GB
Collector	4 GB	1 CPU core	50 GB
Operations Bridge Reporter (OBR) - large deployment with more than 5000 nodes			
OBR Server	16 GB	12 CPU cores	2.5 TB
PostgreSQL	4 GB	2 CPU cores	200 GB
Collector	4 GB	2 CPU cores	30 GB

**Note:** Vertica and Business Objects are not containers, but they require additional resources on a separate system. For more information, see the *OBR Interactive Installation Guide*.

HPE recommends the mount point `/opt/kubernetes` for the master and worker disk space. For the NFS server, the mount point `/var/vols` is recommended if the master node is used as the NFS server.

### Example

You want to install OMi, BVD, and PE. You plan to run a small deployment of OMi on one worker node, and scale out BVD so that you have two BVD deployments. You want to have enough resources for OMi to be moved from one node to another, and also have enough resources to safely take down one of the worker nodes and have the other two worker nodes handle the workload.

So you calculate your minimum requirements per two worker nodes.

Capability	Resources	Scale out multiplier
OMi	16 GB RAM, 4 CPU cores, 50 GB disk space	1
BVD	6 GB RAM, 4 CPU cores, 30 GB disk space	2
PE	8 GB RAM, 4 CPU cores, 100 GB disk space	1
SUM overall	36 GB RAM, 16 CPU cores, 210 GB disk space	
SUM per two worker nodes	18 GB RAM, 8 CPU cores, 105 disk space	

Each of the three worker nodes requires at least 18 GB RAM, 8 CPU cores, and 105 GB disk space.

As the master node is not used as NFS server, it requires at least 16 GB RAM, 8 CPU cores, and 200 GB disk space.

## Database requirements

**Note:** When using an external database, make sure you configure the database to accept remote connections. For external PostgreSQL databases, configure the `pg_hba.conf` file on the PostgreSQL server.

### Suite database requirements

When configuring the Operations Bridge Suite, you can choose between an internal PostgreSQL database or an external PostgreSQL database.

- **Internal PostgreSQL.** There are no specific requirements for the internal PostgreSQL database.
- **External PostgreSQL.** A database for use by the Operations Bridge Suite must already be configured. The name of the database must be `autopassdb`. In addition, the user that accesses the database must have permission to create tables.

For a list of supported PostgreSQL database versions, see the support matrix for the Operations Bridge Suite.

### BVD database requirements

When configuring BVD, you can choose between an external PostgreSQL database and an internal PostgreSQL database.

There are no specific requirements for the internal PostgreSQL database. The database instance is installed and configured in a separate container, and database files are stored on the NFS server.

The requirements for the external PostgreSQL database are as follows:

- **Hardware.** For PostgreSQL hardware requirements, see the PostgreSQL documentation available at:

<http://www.postgresql.org/docs/manuals/>

- **PostgreSQL version.** For a list of supported PostgreSQL database versions, see the support matrix at:

[Support Matrices for Operations Center products](#)

Download and extract the support matrix files, open the document `SUMA.htm` and select **Operations Manager i Business Value Dashboard** from the product list.

- **Installation.** For details on the PostgreSQL software installation, see the installation guide in the documentation for your specific PostgreSQL version.

- **Configuration.** A database for use by BVD must already be configured. The name of the database must not be `postgres`, and the database must use `password` for the authentication, not `MD5`. In addition, the user that accesses the database must have permissions to create tables.
- **Data migration.** If you were using BVD 10.12 or 10.61, specify the external PostgreSQL of your former deployment during the configuration to migrate your data to BVD 10.62 (Operations Bridge Suite 2017.08). The migrated data includes your dashboards, instances, API key, dashboard customizations, CSS customizations, and data integrations.

Do the following to migrate your data to BVD 10.62:

- a. Stop your existing BVD deployment. BVD must no longer be active on the database.
- b. *BVD 10.12 migrations only.* Use a database tool, for example PgAdmin, to open the BVD database.
  - i. Edit the table `bvdLdapServerConfigurations`.
  - ii. Remove the single line that the table contains. This is the LDAP server configuration for 10.12, which is no longer required.  
Do **not** drop the table.
- c. When running the Suite Installer, specify the external PostgreSQL database of your former deployment.
- d. *Optional.* To also migrate your LDAP user permissions and assignments, specify the LDAP server you previously used for BVD during the LDAP configuration. If the same LDAP server is configured, BVD will apply the already configured permissions and role assignments.  
For more information about the LDAP configuration, see the [Operations Bridge Suite Online Help](#).

## PE database requirements

Performance Engine requires an external Vertica database. If your Operations Bridge Suite container deployment includes Performance Engine and Operations Bridge Reporter, the Vertica instance is shared between OBR and PE.

HPE Vertica does not support VMware Vmotion and Logical Volume Manager (LVM) on any system where database files are stored. HPE recommends VMware ESX 5.5 Hypervisor to virtualize the HPE Vertica Analytics Platform, with VMware Tools installed on each virtual machine.

## OBR database requirements

Operations Bridge Reporter requires an external dedicated Vertica database. Vertica is not deployed in a container, but the resources are required for an installation of Vertica on a standalone virtual machine. Use the classic OBR installer and select **Vertica database** to install Vertica on a virtual machine. If

your Operations Bridge Suite container deployment includes the Performance Engine (PE) capability, the Vertica instance can be shared between OBR and PE.

HPE Vertica does not support VMware Vmotion and Logical Volume Manager (LVM) on any system where database files are stored. HPE recommends VMware ESX 5.5 Hypervisor to virtualize the HPE Vertica Analytics Platform, with VMware Tools installed on each virtual machine.

## Client system requirements

- **Web browser configuration.** The web browser must be configured as follows:
  - The browser must be set to accept third-party cookies and allow session cookies.
  - The browser must be set to enable JavaScript execution.
  - The browser must allow pop-ups from the OMi application.
  - The browser must have Java enabled to run applets.
  - Internet Explorer users must:
    - Configure the caching mechanism to automatically check for newer versions of stored web pages (**Internet options > General > Browsing history > Settings > Temporary Internet Files > Check for newer versions of stored pages: Automatically**).
    - Enable the use of TLS 1.0 or later (**Internet Options > Advanced > Security**)
    - Turn off Compatibility View (in Internet Explorer 11 only)
- **Fonts.** The following fonts must be installed:
  - Arial
  - Meiryo (for Japanese locales)
  - Malgun Gothic or Arial (for Korean locales)
  - SimHei or SimSun (for Simplified Chinese locales)
- **Screen resolution.** 1600x900 or higher (recommended); 1280x1024 (supported).

## Meet the software requirements

The following prerequisites must be met for the installation:

- Make sure that the nodes and NFS server for the installation meet the minimum system requirements. For details, see ["Meet the hardware requirements" on page 10](#).

- The master and worker nodes must have a static IP address.
- The host names of the master and the worker nodes must be DNS resolvable (not only via `/etc/hosts`). Alternatively, it is also possible to resolve the host names / IP addresses via the following local hosts files:

```
/etc/hosts
```

```
/var/vols/itom/core/baseinfra-1.0/kube-dns-hosts/hosts
```

- The `/tmp` directory of the (first) master node must have at least 2.5 GB of space available when adding worker nodes from the management portal.
- If the machine already has Docker or Kubernetes installed, uninstall them.
- Make sure you configured your firewall to allow the necessary ports. For details, see the *Operations Bridge Suite Administration Guide*, or the [Online Help](#).
- The following ports are needed on all nodes during and after the installation, and should not be used by another application: 383, 443, 2380, 4001, 4243, 5000, 5432, 5443, 8080, 8200, 8201, 8443, 10249, 10250, 10251, 10252, 10255, 31387, 31389.

The following ports must be open for system processes: 111 (rpcbind), 2049 (NFS), 20048 (rpc.mountd).

**Note:** The installation script checks and reports if necessary ports are in use.

- Check if you have installed the following rpm packages on all nodes:

```
rpm -qa | grep -E "java-1.8.0-openjdk|libgcrypt|libseccomp|libtool-ltdl|net-  
tools|nfs-utils|systemd-libs|device-mapper-libs|lsof|unzip|chrony|rpcbind"
```

**Note:** `systemd-libs` must be version 219 or higher.

If one or multiple of the packages are not installed, install them using `yum install`:

```
yum install java-1.8.0-openjdk libgcrypt libseccomp libtool-ltdl net-tools nfs-  
utils systemd-libs device-mapper-libs lsof unzip chrony rpcbind
```

If you installed Chrony, run the following commands afterwards:

```
systemctl start chronyd  
systemctl enable chronyd
```

- Remove the shared NFS folder if you have previously installed the Container Deployment Foundation. The default folder is `/var/vols/itom/core`.  
For example: `rm -rf /var/vols/itom/core/*`

Also remove the directory on the NFS server where you stored Operations Bridge suite data, if you previously installed the Operations Bridge Suite, for example:

```
rm -rf /var/vols/itom/opsbridge/*
```

- The NFS server, the master nodes, and the worker nodes must be installed under the same subnet.
- Make sure that the browser cache is cleared.
- The time on all master and worker nodes must be the same. To synchronize the time on your nodes, you can, for example, use NTP or VMWare tools.
- For all processes in your `/etc/` environment, make sure that `https_proxy` and `http_proxy` settings are not set (`unset https_proxy; unset http_proxy`). Alternatively, add the IP address of the master node to the `no_proxy` list for all master and worker nodes.

**Example:** `export no_proxy="localhost,127.0.0.1,<master_node_IP>"`

In a multiple master node deployment, add the IP address you specified in `HA_VIRTUAL_IP` (virtual IP shared by multiple master nodes) to the `no_proxy` list.

Make sure that these settings are consistent on all master and worker nodes.

## Unzip the ITOM CDF installation package

To unzip and move the Container Deployment Foundation installation package, follow these steps:

1. Download the CDF and Suite installation package `CDF1706-00136-15000.zip` from the location that was communicated to you after obtaining your license and sending your Docker ID to HPE. For more information, see ["Overview" on page 4](#).
2. Move or copy the installation package (`CDF1706-00136-15000.zip`) to the master node, then unzip the file and the `HPESW_ITOM_Suite_Platform_<version>.zip` file it contains to a temporary directory.

For example:

```
unzip CDF1706-00136-15000.zip
```

```
unzip HPESW_ITOM_Suite_Platform_2017.08.00200.zip -d ITOM
```

**Note:** In the following Container Deployment Foundation installation steps, the temporary directory `HPESW_ITOM_Suite_Platform_<version>` will be referred to as `<foundation_temp_dir>`.

The CDF installation package includes Docker und Kubernetes binaries.



# Install the Container Deployment Foundation

The Operations Bridge Suite must be deployed on the Container Deployment Foundation (CDF), where you can deploy and administer suites.

Follow the instructions in the following chapters to install the Container Deployment Foundation:

- ["Configure the install.properties file" below](#)
- ["Optional. Install an NFS server" on page 25](#)
- ["Install ITOM CDF on the \(first\) master node" on page 27](#)
- ["Install ITOM CDF additional master nodes" on page 28](#)
- ["Add worker nodes" on page 29](#)
- ["Verify the ITOM CDF installation" on page 30](#)

## Configure the install.properties file

On the (first) master node, go to the `<foundation_temp_dir>` directory, and edit the `install.properties` file by setting the following parameters:

```
* MASTER_NODES="<master node IP address>"
* WORKER_NODES="<worker node 1 IP address> <worker node 2 IP address> <worker node 3 IP address>"
* EXTERNAL_ACCESS_HOST=<master node FQDN>
* NFS_SERVER=<master node IP address>
* DEFAULT_DB_TYPE=<EMBEDDED | EXTERNAL_PG >
* REGISTRY_ORGNAME=hpeswitom
```

Additionally, configure the following proxy settings if access to the Docker Hub or registry requires a proxy to connect to the internet:

```
* DOCKER_HTTP_PROXY="<HTTP proxy URL>:<port>"
* DOCKER_HTTPS_PROXY="<HTTPS proxy URL>:<port>"
* DOCKER_NO_PROXY=<IP address>
```

To reuse the `install.properties` file for additional master and worker node installations, list all IP addresses or FQDNs for all cluster nodes that you are going to install. Make sure that all FQDNs are resolved to the correct IP addresses, not the loop back IP 127.0.0.1.

This configuration uses the master node as the NFS server. If you installed a separate NFS server, configure the NFS server IP in the `NFS_SERVER` parameter. For a full description of the parameters in the `install.properties` file, see ["Parameters in the install.properties file" below](#).

**Caution:** The worker node IP addresses must be separated with a space, and the master node and worker nodes must have a static IP address. Additionally, the `EXTERNAL_ACCESS_HOST` parameter must be set to an FQDN with only lowercase letters.

## Parameters in the install.properties file

The following parameters in the `install.properties` file are required to correctly configure the Kubernetes cluster.

**Note:** The table below lists settings that are only mandatory if you are using multiple master nodes. Also, the file contains settings for Oracle databases. Note that Oracle databases are currently not supported by the Operations Bridge Suite.

Parameter	Description	Notes
MASTER_NODES	<p>Lists the cluster master nodes (IPv4 format), separated by a blank and enclosed in double quotes. If you use more than one master node, you must work with high availability.</p> <p><b>Example:</b></p> <pre>MASTER_NODES="10.10.10.10 10.10.10.11 10.10.10.12"</pre>	Mandatory
WORKER_NODES	<p>Lists the cluster worker nodes, separated by a blank and enclosed in double quotes.</p> <p>If you also want to use a master node as a worker node, enter its address in <code>WORKER_NODES</code>.</p> <p>Typically, a worker node runs the workload when you deploy a suite. By default, when you install a suite, you target a worker node.</p> <p><b>Example:</b></p> <pre>WORKER_NODES="16.255.255.255"</pre>	Mandatory
EXTERNAL_	Defines a fully qualified domain name for external clients to	Mandatory

ACCESS_HOST	<p>access cluster services. The specified name must resolve the IP address where the ingress is running. The host name must be DNS resolvable, not only via /etc/hosts.</p> <p><b>Example:</b></p> <pre>EXTERNAL_ACCESS_HOST=mysd.example.net</pre>	
NFS_SERVER	<p>Specifies the IP (IPv4) address of the NFS server that serves the persistent volumes of the cluster services.</p> <p><b>Example:</b></p> <pre>NFS_SERVER=16.255.25.255</pre>	Mandatory
CLIENT_CA_FILE	<p>Specifies the CA certificate that is used for TLS authentication to the API server. The value is the file name of the CA certificate including the absolute path.</p> <p>When the master node is installed, it will generate a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.</p> <p><b>Example:</b></p> <pre>CLIENT_CA_FILE=/tmp/ca.crt</pre>	Mandatory only for worker nodes
CLIENT_CERT_FILE	<p>Specifies the certificate that is used for TLS authentication to the API server. The value is the file name of the certificate including the absolute path.</p> <p>When the master node is installed, it will generate a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.</p> <p><b>Example:</b></p> <pre>CLIENT_CERT_FILE=/tmp/client.crt</pre>	Mandatory only for worker nodes
CLIENT_KEY_FILE	<p>Specifies the private key that is used for TLS authentication to the API server. The value is the file name of the private key including the absolute path.</p> <p>When the master node is installed, it will generate a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.</p> <p><b>Example:</b></p>	Mandatory only for worker nodes

	<pre>CLIENT_KEY_FILE=/tmp/client.key</pre>	
HA_VIRTUAL_IP	<p>Sets up a virtual IP address (single IPv4 address enclosed in double quotes) when setting up multiple master nodes. The IP address must not be occupied before the installation. The virtual IP, the master node, and the worker nodes must all exist in the same subnet.</p> <p><b>Example:</b></p> <pre>HA_VIRTUAL_IP="18.16.10.9"</pre>	Mandatory only if you are using multiple master nodes
PEER_CA_FILE	<p>Specifies the CA certificate for TLS authentication. The value of the parameter is the file name of the CA certificate, including the absolute path.</p> <p><b>Example:</b></p> <pre>PEER_CA_FILE=/tmp/ca/crt</pre>	Mandatory only if you are using multiple master nodes
PEER_CERT_FILE	<p>Specifies the certificate for TLS authentication. The value of the parameter is the file name of the certificate, including the absolute path.</p> <p><b>Example:</b></p> <pre>PEER_CERT_FILE=/tmp/server.crt</pre>	Mandatory only if you are using multiple master nodes
PEER_KEY_FILE	<p>Specifies the private key for TLS authentication. The value of the parameter is the file name of the private key, including the absolute path.</p> <p><b>Example:</b></p> <pre>PEER_KEY_FILE=/tmp/server.key</pre>	Mandatory only if you are using multiple master nodes
NFS_FOLDER	<p>Specifies the root folder (fully-qualified directory) for the persistent volume that the NFS server exports.</p> <p><b>Note:</b> If a container stops and is restarted, all changes made inside the container are lost. When you install the infrastructure, you must either install an NFS server or use the (first) master node as the NFS server. The NFS server shares out the network volumes.</p> <p><b>Example:</b></p> <pre>NFS_FOLDER=/var/vols/itom/core</pre>	Optional
ROOTCA	<p>Specifies the root or intermediate CA certificate for generating</p>	Optional

	<p>server and client certificates. The value of the parameter is the file name of the CA certificate, including the absolute path.</p> <p>When you install the Container Deployment Foundation, all communication between the components is secured via TLS. Therefore, communications use certificates to maintain security. These certificates can be self-signed or signed with a Certificate Authority. The default value is a self-signed certificate.</p> <p><b>Example:</b></p> <pre>ROOTCA=/tmp/ca.crt</pre>	
ROOTCAKEY	<p>Specifies the CA key for generating server and client certificates. The value of the parameter is the file name of the CA key, including the absolute path.</p> <p>When you install the infrastructure, all communication between the components is secured via TLS. Therefore, communications use certificates to maintain security. These certificates can be self-signed or signed with a Certificate Authority. The default value is a self-signed certificate.</p> <p><b>Example:</b></p> <pre>ROOTCA=/tmp/ca.key</pre>	Optional
NFS_STORAGE_SIZE	<p>Specifies the size of the NFS volume exported by the NFS server.</p> <p><b>Example:</b></p> <pre>NFS_STORAGE_SIZE=50Gi</pre>	Optional
K8S_HOME	<p>Specifies the installation directory (fully-qualified directory) for the core platform binaries.</p> <p><b>Example:</b></p> <pre>K8S_HOME=/opt/kubernetes</pre>	Optional
MASTER_API_PORT	<p>Specifies the HTTP port for the Kubernetes (K8S) API server.</p> <p>If you want to use K8S, you must dock to the K8S API server. The kubectl command line tool communicates with the K8S server.</p> <p><b>Example:</b></p> <pre>MASTER_API_PORT=8080</pre>	Optional
MASTER_API_SSL_PORT	<p>Specifies the HTTPS port for the K8S API server.</p> <p>If you want to use K8S, you must dock to the K8S API server. The</p>	Optional

	<p>kubect1 command line tool communicates with the K8S server.</p> <p><b>Example:</b></p> <pre>MASTER_API_SSL_PORT=8443</pre>	
THINPOOL_DEVICE	<p>Specifies the path to a Docker device mapper storage driver.</p> <p>To configure the thinpool device, see the <a href="#">Docker documentation</a>.</p> <p><b>Note:</b> If this parameter is specified, the installation will use the devicemapper(direct-lvm) Docker storage driver. If it is not specified, the installation will use devicemapper(loop). For production environments, HPE recommends devicemapper(direct-lvm).</p> <p><b>Example:</b></p> <pre>THINPOOL_DEVICE=/dev/mapper/docker-thinpool</pre>	Optional
DEFAULT_DB_TYPE	<p>Specifies the database type that the CDF will use for the service connection.</p> <p>The possible values for DEFAULT_DB_TYPE are EMBEDDED, EXTERNAL_PG, and EXTERNAL_ORA.</p> <ul style="list-style-type: none"> <li>EMBEDDED: CDF will use the embedded, containerized PostgreSQL database for the installation.</li> <li>EXTERNAL_PG: CDF will use an external PostgreSQL database for the installation.</li> </ul> <p>When you decide to use an external PostgreSQL database, you must additionally specify the DEFAULT_DB_HOST, DEFAULT_DB_PORT, and the DEFAULT_DB_NAME or the DEFAULT_DB_CONNECTION_URL properties.</p> <p><b>Example:</b></p> <pre>DEFAULT_DB_TYPE=EMBEDDED</pre>	Mandatory
DEFAULT_DB_HOST	<p>Specifies the database host when choosing the external PostgreSQL database as the DEFAULT_DB_TYPE. You can enter the FQDN or the IP address of the database host.</p> <p><b>Example:</b></p> <pre>DEFAULT_DB_HOST=10.10.10.10</pre>	<p>Mandatory only if you choose EXTERNAL_PG as the DEFAULT_DB_TYPE and the DEFAULT_DB_CONNECTION_URL is not</p>

		specified.
DEFAULT_DB_PORT	<p>Specifies the database port when choosing the external PostgreSQL database as the DEFAULT_DB_TYPE.</p> <p><b>Example:</b></p> <pre>DEFAULT_DB_PORT=5432</pre>	Mandatory only if you choose EXTERNAL_PG as the DEFAULT_DB_TYPE and the DEFAULT_DB_CONNECTION_URL is not specified.
DEFAULT_DB_NAME	<p>Specifies the database name when choosing the external PostgreSQL database as the DEFAULT_DB_TYPE.</p> <p><b>Example:</b></p> <pre>DEFAULT_DB_NAME=CDF_DB</pre>	Mandatory only if you choose EXTERNAL_PG as the DEFAULT_DB_TYPE and the DEFAULT_DB_CONNECTION_URL is not specified.
DEFAULT_DB_CONNECTION_URL	<p>Specifies the database connection URL when choosing the external PostgreSQL database as the DEFAULT_DB_TYPE. This parameter is left empty when the DEFAULT_DB_HOST, DEFAULT_DB_PORT, and DEFAULT_DB_NAME are specified.</p> <p><b>Example:</b></p> <pre>DEFAULT_DB_CONNECTION_URL=jdbc:postgres:thin:@IP:port:db_name</pre>	Mandatory only if you choose EXTERNAL_PG as the DEFAULT_DB_TYPE and the DEFAULT_DB_HOST, DEFAULT_DB_PORT, and DEFAULT_DB_NAME are not specified.
DOCKER_HTTP_PROXY	<p>Specifies the proxy settings for Docker. Configure this parameter if access to the Docker hub or registry requires a proxy (the default value is no proxy). The value of the parameter is any valid HTTP proxy URL.</p> <p>When you install suites and launch containers on Docker inside the K8S cluster, you may need to download the images from the internet, for which you need to use proxies.</p>	Optional

	<b>Example:</b> <pre>DOCKER_HTTP_PROXY="http://web.proxy.host.domain:8080"</pre>	
DOCKER_HTTPS_PROXY	<p>Specifies the proxy settings for Docker. Configure this parameter if access to the Docker hub or registry requires a proxy (the default value is no proxy). The value of the parameter is any valid HTTPS proxy URL.</p> <p>When you install suites and launch containers on Docker inside the K8S cluster, you may need to download the images from the internet, for which you need to use proxies.</p> <b>Example:</b> <pre>DOCKER_HTTPS_PROXY="https://web.proxy.host.domain:8080"</pre>	Optional
DOCKER_NO_PROXY	<p>Specifies the IPv4 IP address or FQDN that does not need the proxy settings for Docker.</p> <b>Example:</b> <pre>DOCKER_NO_PROXY=127.0.0.1</pre>	Optional
REGISTRY_ORGNAME	<p>Specifies the organization name where the suite images are placed. The default name is hpeswitomsandbox.</p> <b>Example:</b> <pre>REGISTRY_ORGNAME=hpeswitom</pre>	Mandatory
FLANNEL_IFACE	<p>Specifies the IPv4 address or the interface name for the Docker inter-host communication to use. This setting is used when the nodes have more than one network adapter.</p> <b>Example:</b> <pre>FLANNEL_IFACE=10.10.10.10</pre>	Mandatory only if you install CDF on a node which has more than one network adapter.
CLOUD_PROVIDER	<p>Specifies the cloud provider when installing the CDF on a cloud server.</p> <b>Example:</b> <pre>CLOUD_PROVIDER=AWS</pre>	Optional
AWS_REGION	<p>Specifies the AWS region to use when choosing AWS as the cloud provider. The default value of this parameter is an empty</p>	Mandatory only if you



	string.  <b>Example:</b>  AWS_REGION=us-east-1	choose AWS as the cloud provider.
--	--	-----------------------------------

## Optional. Install an NFS server

The Container Deployment Foundation requires an NFS server. You can either use the master node as the NFS server or you can set up a separate NFS server. The latter is recommended for production environments.

If you want to use the master node as the NFS server instead, skip this step and go to ["Install ITOM CDF on the \(first\) master node" on page 27](#).

To install a dedicated NFS server, you can use any operating system that provides NFS. Additionally, the NFS server must meet the following hardware requirements: 16 GB RAM, 8 CPU cores, and 100 GB free disk space.

### NFS directories overview

The following NFS directories must be set up during the installation:

Exported NFS file system	Proposed directory	Content
<CDF_core>	/var/vols/itom/core	CDF related configuration and data files.
<opsbridge_config>	/var/vols/itom/conf	Operations Bridge Suite related configuration files.
<opsbridge_data>	/var/vols/itom/data	Operations Bridge Suite related database and runtime files.
<opsbridge_log>	/var/vols/itom/log	Operations Bridge Suite related log files.

### Installation

Follow the steps below for the installation:

1. Install the NFS server: `yum install -y nfs-utils`
2. Create a directory to store the CDF data, and adapt the directory permissions:

```
mkdir -p <CDF_core>
```

```
chown -R 1999:1999 <CDF_core>
```

**Note:** If you expose a folder that is not named `core`, specify the exposed folder when installing CDF.

3. Create directories to store the suite configuration data, the database data, and the log data, and adapt the directory permissions:

```
mkdir -p <opsbridge_config>
```

```
mkdir -p <opsbridge_data>
```

```
mkdir -p <opsbridge_log>
```

```
chown 1999:1999 <opsbridge_config>
```

```
chown 1999:1999 <opsbridge_data>
```

```
chown 1999:1999 <opsbridge_log>
```

4. Configure the NFS sharing of the CDF and suite directories:

```
echo "<CDF_core>*(rw, sync, anonuid=1999, anongid=1999, all_squash)" >>  
/etc/exports
```

```
echo "<opsbridge_config> *(rw, sync, anonuid=1999, anongid=1999, all_squash)" >>  
/etc/exports
```

```
echo "<opsbridge_data> *(rw, sync, anonuid=1999, anongid=1999, all_squash)" >>  
/etc/exports
```

```
echo "<opsbridge_log> *(rw, sync, anonuid=1999, anongid=1999, all_squash)" >>  
/etc/exports
```

5. Restart the NFS service to activate the directory sharing:

```
exportfs -ra
```

**Tip:** Do the following to check what has been exported:

1. Disable the firewall on the NFS server.
2. Restart the NFS service by running the following commands:

```
systemctl restart rpcbind  
systemctl enable rpcbind  
systemctl restart nfs-server
```

```
systemctl enable nfs-server
```

3. Run `exportfs`

## Install ITOM CDF on the (first) master node

The following steps describe how to install the Container Deployment Foundation on the (first) master node.

1. Make sure you have already downloaded the installation package to a temporary directory on all master nodes. For details, see ["Unzip the ITOM CDF installation package" on page 16](#).
2. Unzip the zip file on the master node.

**Note:** In the following installation steps, the directory containing the installed Container Deployment Foundation files (`/opt/kubernetes` by default) will be referred to as `<foundation_install_dir>`.

3. *Skip this step if you use a dedicated NFS server.* If you did not install a dedicated NFS server, you must set up the first master node as the NFS server.

- a. On the master node, run the following command to set up the core NFS share:

```
<foundation_temp_dir>/scripts/setupNFS.sh
```

- b. Then run the following commands to set up the Operations Bridge NFS share:

```
<foundation_temp_dir>/scripts/setupNFS.sh <opsbridge_config>
```

```
<foundation_temp_dir>/scripts/setupNFS.sh <opsbridge_data>
```

```
<foundation_temp_dir>/scripts/setupNFS.sh <opsbridge_log>
```

Replace `<opsbridge_config>`, `<opsbridge_data>`, and `<opsbridge_log>` with the directory names of your choice, located at `/var/vols/itom/`. For more information, see ["NFS directories overview" on page 25](#).

4. On the (first) master node, access the `<foundation_temp_dir>` directory, and run the following command:

```
./install
```

Wait until the installation on the first master node is complete.

**Tip:** You can check the installation log at `/opt/kubernetes/install-<date><time>.log`

5. If you chose to use an external PostgreSQL database in the `install.properties` file, enter the database user name and password for `EXTERNAL_PG` when prompted.

## Install ITOM CDF additional master nodes

The following steps describe how to install the Container Deployment Foundation on additional master nodes.

1. Make sure you have already downloaded the installation package to a temporary directory on all additional master nodes. For details, see ["Unzip the ITOM CDF installation package" on page 16](#).
2. Unzip the zip file on every master node.
3. On each additional master node, run the following command to initialize the environment variables:

```
cd <foundation_install_dir>
```

```
source /etc/profile
```

4. Copy the server certificate files (`ca.crt`, `server.crt`, and `server.key`) from the `<K8S_home>/ssl` directory of the first master node to any local directory on each additional master node (for example: the `/tmp` directory).

The default `<K8S_home>` directory is `/opt/kubernetes/ssl`.

5. Copy the `install.properties` file from the first installed master node to each additional master node into the `<foundation_install_dir>` directory.
6. Edit the `install.properties` file on each additional master node as follows:

```
PEER_CA_FILE=/tmp/ca.crt
```

```
PEER_CERT_FILE=/tmp/server.crt
```

```
PEER_KEY_FILE=/tmp/server.key
```

Replace `/tmp` with the temporary directory you used earlier.

7. On each additional master node, run the following command:

```
./install
```

**Note:** You must configure the master nodes one after another.

# Add worker nodes

Add worker nodes via the Management Portal as follows:

**Note:** You can add nodes manually instead. For details, see ["How to add worker nodes by using the install.properties file" below](#).

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 with the user name **admin**, and the password **cloud**.
3. The password must be changed after you logged in for the first time. Follow the instructions to change the password.
4. Go to **ADMINISTRATION > Nodes**.
5. In the Nodes area, click **+ ADD**.

Enter the host name or IP address of the node, the name of the root user, and the password of the specified user. Optionally, specify the `THINPOOL_DEVICE` and/or the `FLANNEL_IFACE`.

`THINPOOL_DEVICE`: specifies the path to the Docker devicemapper storage driver.

`FLANNEL_IFACE`: specifies the interface for the Docker inter-host communication as a single IPv4 address or interface name. This setting is used when the worker nodes have more than one network adapter.

Click **ADD** to remotely install the extra node.

## How to add worker nodes by using the install.properties file

As an alternative to using the Management Portal GUI, you can also add worker nodes by using the `install.properties` file as follows:

**Note:** This step is *not* required if you already added nodes using the Management Portal.

1. Make sure you have already downloaded the installation package to a temporary directory on all worker nodes. For details, see ["Unzip the ITOM CDF installation package" on page 16](#).
2. Unzip the zip file on all nodes.
3. On each worker node, run the following command to initialize the environment variables:

```
cd <foundation_temp_dir>
```

```
source /etc/profile
```

4. Copy the client certificate files (ca.crt, client.crt, and client.key) from the <K8S\_home>/ssl directory of the first master node to any local directory on each worker node (for example: the /tmp directory).

The default <K8S\_home> directory is /opt/kubernetes/ssl.

5. Copy the install.properties file from the first master node to all worker nodes into the <foundation\_temp\_dir> directory.
6. On each worker node, open the install.properties file under the <foundation\_temp\_dir> directory, and set the following parameters to the corresponding file paths (for example /tmp):

```
CLIENT_CA_FILE=/tmp/ca.crt
```

```
CLIENT_CERT_FILE=/tmp/client.crt
```

```
CLIENT_KEY_FILE=/tmp/client.key
```

7. On each worker node, run the following command:

```
./install
```

**Tip:** You can run the installation script on the worker nodes in parallel.

## Verify the ITOM CDF installation

Once the Container Deployment Foundation installation is complete, verify the installation as follows:

**Tip:** You can check the installation log at /opt/kubernetes/install-**<date><time>**.log

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. If that is not the case, see ["Troubleshoot" on page 60](#).

2. Log in with the user name **admin**, and the password **cloud**.
3. The password must be changed after you logged in for the first time. Follow the instructions to change the password.

**Note:** If you want to change the password again later, you can click the ADMIN avatar on the upper right corner of the screen and then select **Change Password**.

Now you can proceed with installing the Operations Bridge Suite; see ["Install the Operations Bridge Suite" on page 32](#).

# Install the Operations Bridge Suite

Once the Container Deployment Foundation is installed, you are ready to install the Operations Bridge Suite.

You need to import suite images to the local registry, activate an Operations Bridge Suite license, and run the Suite Installer from the Management Portal.

Follow the instructions in the following chapters to install and configure the Operations Bridge Suite:

- ["Prepare the suite images" below](#)
- ["Activate a suite license" on page 48](#)
- ["Run the suite installer" on page 35](#)
- ["Verify the suite installation" on page 49](#)

## Prepare the suite images

Before you can install the Operations Bridge Suite, you must import the suite images to the local registry of your first master node.

If your first master node has internet access, follow the next steps. If it does not have internet access, see ["Download suite images to another machine" on the next page](#).

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

```
cd <foundation_install_dir>/scripts  
./downloadimages.sh
```

**Important:** Make sure you are running the `downloadimages` and `uploadimages` scripts from the `<foundation_install_dir>` directory (by default, that is `/opt/kubernetes`).

This script starts the installation and pulls images. You are prompted for the following information:

<b>Suite</b>	OpsBridge
<b>User name and password</b>	Enter your Docker Hub account credentials. If the master node does not have an internet connection, press Ctrl+C,



and continue with the steps described in ["Download suite images to another machine" below](#).

**Suite version**

Enter the suite version you want to install, for example 2017.08

2. On the first master node, run the following command to upload the downloaded images into the local registry:

```
./uploadimages.sh
```

When prompted for the Suite, enter OpsBridge.

## Download suite images to another machine

If the master node does not have an internet connection and cannot access Docker Hub, you must manually export and import the images to the local registry of your master node.

**Note:** This step is *not* required if you ran the `downloadimages.sh` script on a master node with internet connection.

To do this, perform the following tasks:

1. Find another machine that can access Docker Hub, and get your current kernel version:

```
uname -r
```

Make sure that your operating system is 64-bit, the Linux kernel version is 3.10 or higher, and the free disk space is about 100 GB.

2. On the machine that can access Docker Hub, install Docker. For more information, see the [Docker installation documentation](#).

- a. Configure a yum proxy: `vi /etc/yum.conf`
- b. Add the following line: `proxy=<your_proxy>`
- c. List the package version in the system: `yum list`
- d. Add the yum repository:

```
cat << EOF > /etc/yum.repos.d/docker.repo
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7/
enabled=1
gpgcheck=1
```

```
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- e. Update the source information: `yum update --skip-broken -y`
- f. Install Docker: `yum install -y docker-engine`
- g. Enable the service: `systemctl enable docker.service`
- h. Configure a proxy so you can download the official images:

```
mkdir -p /usr/lib/systemd/system/docker.service.d/
cat << EOF > /usr/lib/systemd/system/docker.service.d/http_proxy.conf
[Service]
    Environment="HTTP_PROXY=http://<web-proxy-host>:<port>/" "HTTPS_
    PROXY=http://<web-proxy-host>:<port>/"
EOF
```

Replace `<web-proxy-host>` and `<port>` with your proxy settings.

- i. Reload the configuration: `systemctl daemon-reload`
- j. Restart docker: `service docker restart`

The restart may take several minutes.

### 3. Export the images:

- a. Copy the following files from your first master node to a machine on which you have internet access, for example into the `/tmp` directory:

```
<foundation_install_dir>/scripts/downloadimages.sh
<foundation_install_dir>/bin/jq
<foundation_install_dir>/scripts/image-list.json
<foundation_install_dir>/scripts/deployments.json
```

- 4. Access the directory into which you copied the files, for example `/tmp`, and run the following command:

```
./downloadimages.sh -o hpeswitom/
```

You are prompted for the following information:

<b>Suite</b>	OpsBridge
<b>User name and password</b>	Enter your Docker Hub account credentials.
<b>Suite version</b>	Enter the suite version you want to install, for example 2017.08

5. Copy all files from `/var/opt/kubernetes/offline/suite_images/` on the machine with internet access to your first master node into the directory `/var/opt/kubernetes/offline/suite_images/`.
6. On the master node, run the following commands to upload the downloaded images into the local registry:

```
cd <foundation_install_dir>/scripts  
./uploadimages.sh
```

When prompted for the Suite, enter `OpsBridge`.

**Note:** Alternatively, you can run `./uploadimages.sh -d <directory>` to upload the images from any temporary directory you specify.

7. *Optional.* You can verify that the images are listed in the local registry by accessing the Management Portal as the admin user, and checking the list images in **ADMINISTRATION > Local Registry**.

## Run the suite installer

**Important:** During the suite installation, do not use any browser buttons (such as Back or Refresh) on the current installation wizard page; otherwise, unexpected errors might occur.

To install the suite, 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 to the Management Portal as the admin user. Use the password you specified after your initial login.
3. On the left side navigation, expand the **SUITE** menu and click **Installation**. Agree to the license agreement and, optionally, the privacy policy, and click **Next**.
4. Select **Operations Bridge** and the latest version of the suite, and click **Next**.
5. Select the edition of the suite (Express or Premium) and the capabilities of the Operations Bridge

Suite that you want to install.

When installing the Operations Bridge Suite Premium, you can select the following capabilities:

- **Operations Bridge Manager.** Operations Bridge Manager (OMi) provides the ability to sense, analyze and adapt to manage IT services that support digital business. With advanced event correlation, log intelligence, predictive analytics and automation you can remediate issues across all your technologies to prioritize business targets.
  - **Business Value Dashboard.** Business Value Dashboard (BVD) brings your data to life. Use BVD to create custom, flexible dashboards that visualize information in an appealing way and that can be accessed anywhere, anytime, from any device. Incorporate your own graphics, add color to identify status, and receive real-time updates—so you always understand the value driven by your IT environment.
  - **Performance Engine.** The Performance Engine (PE) is an add-on component of Operations Bridge Manager (OMi) that provides streaming of custom metrics and system metrics in a large scale environment.
  - **Operations Bridge Reporter.** Operations Bridge Reporter (OBR) is a solution based on the Big Data technology HPE Vertica and has been built to specifically address the challenges of reporting in dynamic IT environments. In addition to consolidating performance data and metrics from multiple domain-focused collectors, HPE Operations Bridge Reporter also collects and collates specific information on the relationships between the IT elements and the business services.
6. In the **Configure the Suite Storage** section, specify the IP address or FQDN of the dedicated NFS server or your (first) master node. Then specify the mount points for the NFS volumes `db-volume`, `log-volume`, and `conf-volume`.

Click the **Edit** button in the Config column to configure each volume mount point. Specify the FQDN or IP address of the NFS server (when using the master node, this is the master node's IP address). Select the mount path from the drop-down list and click **Apply**.

After you have configured all volumes, click **Next**. The configuration wizard is displayed.

- 
7. Configure the suite defaults. The Suite Defaults configuration defines general settings that all capabilities of the suite share.

## Suite Defaults > Configuration Type

Select the configuration type of the suite.

### Custom configuration

*Default.* Displays the complete configuration wizard. You can specify custom values for all capabilities.

### Express configuration

Uses default values for some of the settings, to speed up the configuration process. When this option is chosen, the suite by default uses an internal PostgreSQL database, a TLS certificate automatically generated by the Management Portal, a 60-day evaluation license, and the same password for the administrator and the PostgreSQL database user. The Management Pack for Infrastructure is installed automatically.

## Suite Defaults > Login

Define the default administrative user credentials for all capabilities.

If you chose the express configuration, this will be the global password.

If you chose the custom configuration, you can later specify individual user credentials for the different capabilities.

### Login

The login name is `admin`.

### Password

Specify a password for the administrator user. You can change this password again after the installation.

**Note:** The password must consist of eight characters or more, and contain at least one upper-case letter, one lower-case letter, one digit, and one special character.

## Suite Defaults > Database

Configure the default database for the Operations Bridge Suite.

If you chose the express configuration, this will be the database for all capabilities.

If you chose the custom configuration, you can later specify individual databases for the different capabilities.

### Database type

You can select one of the following database types: Internal PostgreSQL, External

PostgreSQL.

**Host**

*External database only.* The name of the host machine on which the database is installed.

**Port Number**

*External database only.* The database listening port.

Default: 5432

**Database user**

The name of a user with administrative permissions on the specified database.

**Password**

The password of the specified user.

## Suite Defaults > Connection

Specify your load balancer information. The load balancer is used to access the different user interfaces of the Operations Bridge Suite capabilities.

**External Hostname**

The external hostname of the load balancer.

In a single server environment, enter the FQDN you specified as `EXTERNAL_ACCESS_HOST` in the `install.properties` file. This hostname must be resolvable via the DNS server, not only via `/etc/hosts`.

**Port Number**

The port of the load balancer.

Default: 443

**SSL Certificate File**

Click **Use the certificate generated by the Management Portal** to use the automatically generated certificate file.

Click **Upload certificate** to browse your files and select the load balancer's server and CA root certificate files. The Operations Bridge Suite supports server certificates in P12 and PFX format, and CA certificates in PEM format.

- 
8. *Optional.* Configure Operations Bridge Manager (OMi).

## OMi > Login

Define the administrative user credentials for OMi. You can later change the password in your account settings.

### **Use suite default administrative user account**

Select to use the administrative user account credentials that you specified during the suite configuration.

### **Custom credentials**

Select to specify custom credentials for OMi.

The administrative user name is `admin`, the password can be changed in the OMi user interface at a later time.

The JMX password is used by the OMi administrator for all JMX consoles (user name: `admin`) and for the RTSM JMX console (user name: `sysadmin`).

## OMi > Database

Configure a database to store all OMi related information. You can choose to use the database specified for the Operations Bridge Suite, use an OMi specific internal database, create a new database, or you can connect to an already existing database.

### **Use suite default database settings**

Select to use the database that you specified during the suite configuration. You can specify the names of the Management, RTSM, and Event schemas.

### **Custom database settings for Operations Bridge Manager**

Select to create a new database for this OMi instance or connect to an existing database.

If you decide to use a remote database instance, you can preconfigure it or OMi can configure it for you. For detailed information on deploying the database servers in your system for use with OMi, and creating the databases manually, see the *OMi Database Guide*.

If you decide to use an internal PostgreSQL database instance, OMi installs and configures the instance for you.

### **Database Type**

Select the appropriate database type: Internal PostgreSQL or External PostgreSQL.

**Host**

The name of the host machine on which the database is installed. Alternatively, you can also specify the IP address of the host machine.

**Port Number**

The database listening port.

Default: 5432 (Postgres)

**Login**

The name of a user with administrative permissions on the specified database.

**Password**

The password of the specified user.

**Use TLS**

*Optional.* Click Use TLS to encrypt the communication with the database.

The server must be running with TLS communication enabled and it must provide a certificate for use by OMi.

**Management Schema**

For storage of system-wide and management-related metadata.

**Event Schema**

For storage of events and related data, such as annotations, as well as for storage of configuration data, such as event correlation rules.

**RTSM Schema**

For storage of RTSM data. The RTSM (Run-time Service Model) is OMi's embedded CMDB, which acts as the central repository for configuration information that is collected and updated from the various OMi data collection processes.

## OMi > Server Deployment

Define the size of your OMi deployment. The number of monitored nodes determines the amount of memory on your system.

**Number of monitored nodes**



Select the appropriate number of monitored nodes that send events to OMi: up to 2000, up to 5000, or more than 5000.

This includes all nodes that are present as CIs and that send events to OMi (for example, nodes connected to Operations Manager, nodes directly connected to OMi, and target connectors).

## OMi > Management Packs

Select the HPE OMi Management Packs to install in your OMi environment.

You can choose not to install dependent management packs. However, if you do so, the functional scope of the selected management packs might be reduced.

Management packs provide add-on content on top of OMi. They deliver automatic and end-to-end monitoring solutions of infrastructure and applications. Management packs enable users to monitor, detect, troubleshoot, and remediate issues in the IT domain. They increase the productivity of users by optimizing and automating various tasks, and reduce the mean time to resolve (MTTR) incidents.

Management packs discover application domains and proactively monitor the domains for availability and performance issues. They include, for example, management templates, aspects, policy templates, performances graphs, troubleshooting tools, auto remediation flows, and topology-based event correlation (TBEC) rules.

To install management packs after the first configuration, run `opr-mp-installer.sh`. For details about how to run OMi command-line tools from within the container, see the *Operations Bridge Suite Administration Guide*.

`opr-mp-installer` by default installs management packs from the `/opt/HP/BSM/opr/mgmtpacks` directory inside the OMi container. In this directory, you can find all management packs that can be selected during the suite installation.

Once installed, management packs cannot be removed.

**Note:** To update a management pack to a later version than the one included with OMi, download its installation package from the [HPE Marketplace](#) website and install the management pack manually. You can also install additional management packs that are not bundled with OMi.

To install the downloaded management pack, put the management pack zip file into a location that is accessible to the OMi container, then specify this location when executing the `opr-mp-installer` script using the `-i <input_path>` option.

For example, a suitable location would be a `mgmtpacks` directory in the `./omi/var/opt/OV/shared/server/conf/` subfolder on the `<opsbridge_config>` volume of the NFS share. You could then execute the `opr-mp-installer` tool as follows:

```
opr-mp-installer -install <mp_name> -i  
/var/opt/OV/shared/server/conf/mgmtpacks
```

---

9. *Optional*. Configure Business Value Dashboard (BVD).

## BVD > Login

Define the administrative user credentials for BVD. One built-in super-admin user is defined for every installation of BVD. You can later change the password in your account settings.

### Use suite default administrative user account


Select to use the administrative user account credentials that you specified during the suite configuration.

### Custom credentials

Select to specify custom credentials for BVD.

#### Name

Login name of the built-in BVD super-admin.

The built-in super-admin is not listed among the users in user management. If you have logged in as the super-admin, you can change the user's information, including password and contact information, in the **My Account** page in the  **Personal User Settings** menu.

Default: admin

#### Password

Password of the built-in super-admin.

BVD enforces a strong password policy. The password must be at least eight characters long, and meet at least two of the following requirements: one upper-case letter, one digit, and one special character. Special characters should be ASCII characters only.

## BVD > Database

Configure a database to store all BVD related information. You can choose to use the database specified for the Operations Bridge Suite, create a new, embedded database, or you can connect to an already existing database.

### Use suite default database

Select to use the database that you specified during the suite configuration. If you chose an external database, enter a database name.

### Custom database for BVD related data

Select to specify an existing, already configured database for BVD.

To migrate data from a previous BVD 10.12 installation, make sure you performed the migration steps described in the BVD database requirements. Then you can proceed with specifying the external PostgreSQL database that you used for your 10.12 deployment.

**Note:** Before connecting to an external PostgreSQL database, make sure the database is installed as required by BVD.

### Database type

Choose the type of database to be used.

External PostgreSQL: for use with an external PostgreSQL database.

Internal PostgreSQL: for use with the embedded PostgreSQL database.

### Host

The name of the host machine on which PostgreSQL is installed.

Default: localhost

### Port

The PostgreSQL listening port.

Default: 5432

### Database

The name of the PostgreSQL database.

### Login

The name of a user with administrative permissions on the PostgreSQL database.

Default: dbadmin

### **Password**

The password of the BVD administrative user to access the PostgreSQL database.

## **BVD > Security**

Configure security settings for BVD.

### **Allow to embed BVD in iframes**

Determines if BVD can be embedded into other web pages as a iframe. If checked, the browser allows framing from other domains.

Be aware that this might enable an attacker to perform cross-site scripting attacks against BVD.

## **BVD > Aging**

Configure the controller process that scans the database configuration.

By default, up to 500 data records per data channel are stored in the database. You can modify the default and adjust additional data aging settings.

### **Data Records**

Purge old data records based on their age. The **Maximum Age** is the time period (in days) during which data records are kept in the database. Records older than the configured time period are automatically deleted by the aging process.

The value must be an integer greater than 0.

Default: 10 days

### **Data Channel Statistics**

Time period (in days) during which a data channel is available in the list of data channels in the widget properties. If a data channel does not receive any data during the configured time period and the data channel is not associated with a widget, it is deleted from the data store. If the data channel is associated with a widget, the channel is not deleted even if the data last received for the channel is older than the configured time period.

The value must be an integer greater than 0.

Default: 1 day

---

10. *Optional.* Configure the Performance Engine.

PE > Login

**Password**

Password for the Performance Engine. The password must be at least sixteen characters long, and contain at least one lower-case letter, one upper-case letter, one digit, and one special character.

PE > Vertica Database

***Optional.* Configure Vertica Database**

Select to configure a Vertica database for storing and retrieving historical performance data. When installing the Performance Engine without this option, the embedded data store of the Performance Engine allows you to retrieve data only for a limited time period. By additionally configuring a Vertica database, you can access data that has been collected for a longer time period.

For information about how to install Vertica, see the *OBR Interactive Installation Guide*.

**Vertica hostname**

The hostname of your Vertica database (if your Vertica instance is not shared).

**Port**

The Vertica listening port. Default: 5433 (if your Vertica instance is not shared)

**Database name**

The name of the Vertica database.

**Database user name**

The name of a user with administrative permissions on the Vertica database.

**Database password**

The password of the administrative user to access the Vertica database.

PE > Connection

**Enable https**

Select to enable secure connection between OMi and Performance Engine.

### Server Certificate

Run the following command in OMi to download the certificate:

```
/opt/OV/bin/ovcm -issue -file <FILE_NAME.crt> -name <PE Node> -pass  
<Password> -ca
```

Click **Choose File** to browse to the location where you downloaded the certificate.

### Certificate Password

Enter the password that you specified when downloading the OMi server certificate.

---

11. *Optional.* Configure Operations Bridge Reporter.

## OBR > Login

Define the administrative user credentials for Operations Bridge Reporter. You can later change the password in your account settings. For more information, see the *Operations Bridge Reporter Administration Guide*.

### Use suite default administrative user account

Select to use the administrative user account credentials that you specified during the suite configuration.

### Custom credentials

Select to specify custom credentials for OBR.

## OBR > Time Zone Selection

Select the time zone in which you want the Operations Bridge Reporter to operate. The time zone that you select applies to the OBR system and reports. However, the run-time information for processes like collection and work flow streams is always based on local time zone irrespective of this selection.

### GMT

Select to use the Greenwich Mean Time (GMT).

### Local

Select to use the time zone of your local system.

## OBR > Vertica Database

Configure a Vertica database to store performance data.

For information about how to install Vertica, see the *OBR Interactive Installation Guide*.

### **Vertica hostname**

The hostname of your Vertica database (if your Vertica instance is not shared).

### **Port**

The Vertica listening port. Default: 5433 (if your Vertica instance is not shared)

### **Database name**

The name of the Vertica database.

### **Database user name**

The name of a user with administrative permissions on the Vertica database.

### **Database password**

The password of the administrative user to access the Vertica database.

## OBR > Management Database

OBR comes with a management database that stores the OBR configuration and run-time data. Create a new user account for the management database administrator to access this database.

The management database refers to the Online Transaction Processing (OLTP) store used by HPE OBR to store its run-time data such as data process job stream status, runtime information for individual steps, and data source information.

### **Database Admin (DBA)**

The password of the database administrator. The login name is postgres.

### **Database User**

The password of the management database user. The login name is pmdb\_admin.

## OBR > Reporting Platform

OBR uses SAP BusinessObjects for report generation. The Operations Bridge Reporter includes the SAP BusinessObjects BI launch pad portal that enables you to view the generated reports.

### Business Objects hostname

The hostname of the system that hosts the BusinessObjects BI platform.

**Note:** After the OBR container is deployed, you must configure OBR to collect data from the data sources. For more information on configuring OBR, see the *Operations Bridge Reporter Configuration Guide*.

12. On the Configuration Complete page, click **Next** to start the installation.

**Caution:** Do not refresh the page during the installation; otherwise, you will quit the installation and log out of the Management Portal.

Wait until the installation is complete.

# Activate a suite license

**Tip:** In a testing environment, you can skip this step and use a 60-day trial license for the suite. The trial license is used automatically if you do not install a perpetual license.

The suite license contains keys for the ITOM Container Deployment Foundation as well as all capabilities of the suite. Therefore, the suite license is the only license you need to install the suite.

To activate a license for the suite, perform the following steps.

## 1 — Activate a suite license

1. Go to the [HPE Software Entitlement Portal](#).
2. Obtain an Operations Bridge Suite license.
3. Activate the license. Enter any valid IP address in the **Locking Information** field — this must not be the IP address of your master or worker nodes.
4. Download the license file to your local drive.




## 2 — Install the suite license

To install the suite license on the Container Deployment Foundation, 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.

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.

## Verify the suite installation

Once the Suite installation is complete, verify the installation as follows:

1. On the master node, run the following command:

```
kubectl get ns
```

The namespace of your Suite deployment should appear in the list.

2. Continue to run the following command:

```
kubectl get pods --namespace <namespace>
```

All container processes are displayed with the status **Running** and the **READY** column must show that all processes are ready (for example 2/2, not 1/2).

Alternatively, you can also verify the status of the pods via the Management Portal:

- a. Launch the Management Portal and log on as administrative user.
- b. Access **RESOURCES** and select the namespace of the Operations Bridge Suite.
- c. Click **Workloads > Pods**. All pods must have the status **Running** or **Succeeded**.

**Important:** After all pods have the status `Running`, it might take 20 to 45 minutes until you can launch your capabilities.

3. *Optional.* Launch your installed capabilities:

**OMi:** `https://<external_access_host>/omi` or `https://<external_hostname>:<port>/omi`

**BVD:** `https://<external_access_host>/bvd`

**OBR:** `https://<external_access_host>/OBRAApp`

`<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.

`<external_hostname>` and `<port>` are the external hostname and port of the load balancer that was specified in the Connection step of the suite configuration.

4. *Optional.* If you installed OMi, you can check the status of your OMi deployment with `serverStatus.jsp`:

`https://<external_access_host>/topaz/serverStatus.jsp`

You can configure additional settings, like scaling and LDAP. For more information, see the *Operations Bridge Suite Administration Guide* or the [Online Help](#).

# Edit the installation

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 the Operations Bridge Suite Administration Guide or the [online help](#).

## Remove machines to 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 <foundation_temp_dir>
./uninstall.sh
```

## Edit hard eviction thresholds

The ITOM Container Deployment Foundation adheres to Kubernetes' hard eviction policy for the existing worker nodes. Once a hard eviction threshold is met, Kubernetes will kill the pod immediately. The eviction can also delete dead pods, dead containers, and unused images when the disk space meets the thresholds. For more details, see the [Kubernetes documentation](#).

To edit the hard eviction threshold for a worker node, do the following:

1. Log on the worker node for which you want to edit the eviction threshold.
2. Access the file `/usr/lib/systemd/system/kubelet.service`, and edit the parameter values as required. For example, you can edit the following values:

### Examples:

```
vim /usr/lib/systemd/system/kubelet.service
--eviction-
```

```
hard=memory.available<500Mi,nodefs.available<5Gi,imagefs.available<5Gi  
--system-reserved=memory=1.5Gi
```

3. To enable the new thresholds, run the following commands:

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

# Upgrade

The following upgrades are available:

- **"2017.04 > 2017.08 (BVD only)" below.** You can upgrade BVD 10.61 (2017.04 suite installation) to BVD 10.62 (2017.08 suite installation).
- **"BVD 10.12 > 2017.08" below.** You can upgrade BVD 10.12 or 10.61 (classic deployment) to BVD 10.62 (2017.08 suite installation)

---

## 2017.04 > 2017.08 (BVD only)

You can upgrade BVD as part of your existing, containerized 2017.04 suite installation to Operations Bridge Suite 2017.08. Note that an upgrade is only possible if you used an external PostgreSQL for your previous BVD version.

Do the following to upgrade BVD from 2017.04 to 2017.08:

**Note:** At the moment, only BVD can be upgraded to 2017.08. For every other component of the Operations Bridge Suite 2017.04, a fresh installation is required.

1. Uninstall the Operations Bridge Suite 2017.04. For details, see ["Uninstall" on page 59](#).
2. Upgrade CDF from 2017.03 to 2017.06. For details, see ["Upgrade CDF" on the next page](#).
3. Remove the following files from the CDF core NFS directory that you configured during the CDF installation (the proposed directory is `/var/vols/itom/core`):

```
rm /var/vols/itom/core/omi-persistent-volume.yaml
```

```
rm /var/vols/itom/core/namespace.yaml
```

4. Install and configure the Operations Bridge Suite 2017.08. When configuring the database connection for BVD, specify the external PostgreSQL database of your former deployment. For details, see ["Install the Operations Bridge Suite" on page 32](#).

---

## BVD 10.12 > 2017.08

You can migrate your BVD 10.12 data to Operations Bridge Suite 2017.08. Note that an upgrade is only possible if you used an external PostgreSQL for your previous BVD version.

1. Stop your existing BVD deployment. BVD must no longer be active on the database.
  2. Use a database tool, for example PgAdmin, to open the BVD database.
    - a. Edit the table `bvdLdapServerConfigurations`.
    - b. Remove the single line that the table contains. This is the LDAP server configuration for 10.12, which is no longer required.

Do **not** drop the table.
  3. Install and configure the Operations Bridge Suite 2017.08. When configuring the database connection for BVD, specify the external PostgreSQL database of your former deployment. For details, see ["Install the Operations Bridge Suite" on page 32](#).
  4. *Optional.* To also migrate your LDAP user permissions and assignments, specify the LDAP server you previously used for BVD during the LDAP configuration. If the same LDAP server is configured, BVD will apply the already configured permissions and role assignments.
- 

## Upgrade CDF

CDF can be upgraded from version 2017.03 to 2017.06. The upgrade will update all components such as Docker, Kubernetes, Heapster, Vault, Etcd, Flannel and the CDF core to the 2017.06.001 version.

**Note:** We recommend that you back up the following before the upgrade:

- The entire NFS folder in which the CDF data is stored.
- The etcd data.

To upgrade CDF, do the following:

1. Make sure all CDF nodes are running.
2. Run the command: `upgrade.sh -g` on any one of the master nodes to generate the `CDF_upgrade_parameters.txt` under the current directory.

Enter the fully qualified domain name (FQDN) of the host name for external users to access the clusters.

For a single master node, the specified FQDN must be resolved to the master node IP address. For multiple master nodes, the specified FQDN must be resolved to the IP address you specified in the `HA_VIRTUAL_IP` setting.

**Note:** This step will try to remove resources from the yaml file. If you receive an error when

deleting the resources, remove the resources manually before the upgrade. Otherwise, the upgrade may fail due to the existing resources.

3. Stop all the master nodes one after another with the following commands and note which master node you stopped first:

```
docker -H unix:///var/run/docker-bootstrap.sock stop etcd_container
```

```
${K8S_HOME}/bin/kube-stop.sh
```

4. Stop all the worker nodes one after another with the following command:

```
${K8S_HOME}/bin/kube-stop.sh
```

5. *Optional.* If the machine has multiple network interfaces, the user can set the value of FLANNEL\_IFACE in the CDF\_upgrade\_parameters.txt file to specify the interface to be used for docker inter-host communication before the upgrade.
6. Copy the CDF\_upgrade\_parameters.txt to all nodes.

7. Upgrade the master node that was stopped first with the command:

```
upgrade.sh -u /<Parameter file path>/CDF_upgrade_parameters.txt
```

8. Upgrade the rest of the master nodes one after the other (no preferred order) with the following command:

```
upgrade.sh -u /<Parameter file path>/CDF_upgrade_parameters.txt
```

9. Upgrade the worker nodes one after the other (no preferred order) with the following command:

```
upgrade.sh -u /<Parameter file path>/CDF_upgrade_parameters.txt
```

## Retry the CDF upgrade

You may have to retry the upgrade of one or multiple nodes if an error message popped up about the upgrade on the node having failed.

To retry the CDF upgrade on that failed node, do the following:

1. Run the following command: `source /etc/profile`.
2. *Optional.* Run the following command: `upgrade.sh -d /<Parameter file path>/CDF_upgrade_parameters.txt` only when the failed master node is not the first stopped master. This command helps to remove the etcd member from the etcd cluster.
3. Check if the backup-complete file is available under the `/<backup directory>/CDF_201703_`

backup directory.

- If the backup-complete file does not exist, the CDF backup step has failed. Back up the CDF again with the following steps:
    - i. Delete the `<backup directory>/CDF_201703_backup` folder.
    - ii. Run the following command: `upgrade.sh -u <Parameter file path>/CDF_upgrade_parameters.txt`
  - If the backup-complete exists, the CDF backup has been completed.
4. Check the status of the kubelet service with the command: `systemctl status kubelet`
    - If the kubelet service is not active, delete the `kubelet.service` file under the `/usr/lib/systemd/system` directory.
    - If the kubelet service is active, run the command `systemctl stop kubelet` to stop the kubelet service. Then delete the `kubelet.service` file under the `/usr/lib/systemd/system` directory.
  5. Check the status of the docker service with the following command: `systemctl status docker`
    - If the docker service is not active, delete the `docker.service` file under the `/usr/lib/systemd/system` directory.
    - If the docker service is active, run the command `systemctl stop docker` to stop the docker service. Then delete the `docker.service` file under the `/usr/lib/systemd/system` directory.
  6. Check the status of the docker-bootstrap service with the command: `systemctl status docker-bootstrap`.
    - If the docker-bootstrap is not active, delete the `docker-bootstrap.service` file under the `/usr/lib/systemd/system` directory.
    - If the docker-bootstrap is active, run the command `systemctl stop docker-bootstrap` and then delete the `docker-bootstrap.service` file under the `/usr/lib/systemd/system` directory.
  7. Unmount the mounted data with the following commands:
 

```
for data in $(mount | grep "${K8S_HOME}/data/" | cut -d" " -f3 | sort -r);do
umount -f -l $data; done

for data in $(mount | grep "/usr/lib/kubelet" | cut -d" " -f3 | sort -r);do
umount -f -l $data; done
```
  8. Reboot the machine on which you are retrying the upgrade.
  9. Delete the `${K8S_HOME}` directory with the following command:



```
rm -rf ${K8S_HOME}
```

10. Rollback the \${K8S\_HOME} directory with the following command:

```
mv /<backup directory>/CDF_201703_backup ${K8S_HOME}
```

11. Delete the backup-complete file under the \${K8S\_HOME} directory.
12. Recover the docker.service and the docker-bootstrap.service files with the following commands:

```
mv ${K8S_HOME}/docker.service /usr/lib/systemd/system/
```

```
mv ${K8S_HOME}/docker-bootstrap.service /usr/lib/systemd/system/
```

13. *Optional.* Restore the data on the NFS server manually before retrying the upgrade only for the first stopped master node that fails to upgrade.
14. Retry the upgrade with the following command:

```
upgrade.sh -u /<Parameter file path>/CDF_upgrade_parameters.txt
```


# Reconfigure

You can reconfigure the BVD container and the suite defaults in the Operations Bridge Suite via the Management Portal.

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**.
4. For the suite deployment that you want to reconfigure, click  **Actions** next to the suite name and select **Reconfigure**.
5. Follow the instructions of the Suite Installer to reconfigure BVD and the suite defaults as required. For information about the available settings, see ["Run the suite installer" on page 35](#).

**Note:** If you decide to use a different database, your previous BVD data is not migrated automatically.

# Uninstall

You can back up image tars from your local private registry to a remote registry before you uninstall the Container Deployment Foundation.

## Optional. Back up the image tars

1. Go to directory where the `local_backup.sh` file is located: `<foundation_install_dir>/script`.


2. Run: `./local_backup.sh localhost:5000`

For example:

```
./local_backup.sh localhost:5000
```

The tar files are saved in `image_tars/xxx.tar`.

## Uninstall the Operations Bridge Suite

1. In the Management Portal, click **SUITE > Management**.
2. For your current Operations Bridge Suite installation, click  **Actions > Uninstall**.
3. Click **UNINSTALL** to uninstall the suite.

## Uninstall the Container Deployment Foundation

1. On the worker nodes, go to the `<foundation_install_dir>` directory, and run `uninstall.sh`.
2. Once CDF is uninstalled on all worker nodes, go to the `<foundation_install_dir>` directory on the master nodes, and run `uninstall.sh`.

The uninstallation process stops and removes the containers, daemons, and so on.

3. Reboot the server.

# 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" below](#)
- ["Log files" on the next page](#)
- ["Support toolset" on page 62](#)
- ["Common problems and limitations" on page 64](#)

## 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, ucldb, 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 Installation 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](mailto:ovdoc-asm@hpe.com).

We appreciate your feedback!