



Operations Bridge Suite

Software Version: 2017.04

Installation Guide

Document Release Date: April 2017

Software Release Date: April 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>.

Disclaimer

Certain versions of documents ("Material") accessible here may contain branding from Hewlett-Packard Company (now HP Inc.) and Hewlett Packard Enterprise Company. As of September 1, 2017, the Material is now offered by Micro Focus, a separately owned and operated company. Any reference to the HP and Hewlett Packard Enterprise/HPE marks is historical in nature, and the HP and Hewlett Packard Enterprise/HPE marks are the property of their respective owners.

Contents

Overview	4
Installation	6
Install the ITOM Platform	7
Step 1. Meet the system requirements	8
Step 2. Prepare for the installation	14
Step 3. Download the installation package	16
Step 4 (optional). Install a dedicated NFS server	17
Step 5. Set up one master node	18
Step 6. Set up one or multiple worker nodes	25
Step 7. Verify the installation	27
Install the Operations Bridge Suite	28
Step 1. Import suite images to the local registry	29
Step 2. Activate a suite license	32
Step 3. Run the Suite Installer	34
Step 4. Verify the Suite installation	47
Uninstall the ITOM Platform	49
Troubleshooting	50
Send documentation feedback	57

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 following guide leads you through the installation process of the containerized HPE Operations Bridge Suite. The installation involves multiple steps for each capability that you want to set up.

To learn more about the capabilities of the Operations Bridge Suite, see the *Operations Bridge Suite 2017.04 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.

The high-level process of installing the Operations Bridge Suite is as follows:

1. *Prerequisites.* You must have an active support contract, and obtain an Operations Bridge Suite license on the [HPE Software Entitlement Portal](#).
2. Create a Docker account, and send the Docker ID, your company name, contact information, HPE Passport email address, and SAID to HPE to be able to access the Suite container images. Send your email to one of the addresses below, based on your location:
 - **APJ:** dockersupport.apj@hpe.com
 - **AMERICAS:** dockersupport.ams@hpe.com
 - **EMEA:** dockersupport.emea@hpe.comHPE will enable your Docker ID and send you a confirmation.
3. Prepare one or more systems or VMs for the suite installation.
4. Start the installation as described in "[Installation](#)" on page 6.
 - a. Install a dedicated NFS server, or decide to use the first master node as the NFS server.
 - b. Install the ITOM Platform on the master node.
 - c. Log on to the ITOM Platform.

- d. Install the ITOM Platform on the worker nodes.
- e. Install the Operations Bridge Suite.

The master node coordinates 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.

Installation

Installing, administering, and using the Operations Bridge Suite in a container deployment leverages the container technology based on Docker and Kubernetes. Docker provides a way to run almost any application securely isolated in a container, and Kubernetes automates the deployment, scaling, and management of containerized applications.

Each suite capability is deployed as a containerized application that is integrated with other suite capabilities. To install the containerized suite, first install the ITOM Platform as container management framework, then install the Operations Bridge Suite by using the ITOM Platform's Suite Installer.

The Operations Bridge Suite 2017.04 comprises the following capabilities:

- Operations Bridge Manager (OMi)
- Business Value Dashboard
- Performance Engine
- Operations Bridge Reporter

Follow these instructions to install the ITOM Platform and the Operations Bridge Suite:

["Install the ITOM Platform" on page 7](#)

["Install the Operations Bridge Suite" on page 28](#)

Install the ITOM Platform

The Operations Bridge Suite must be deployed on the ITOM Platform, where you can deploy and administer suites.

Follow the instructions in the following chapters to install the ITOM Platform:

- ["Step 1. Meet the system requirements" on page 8](#)
- ["Step 2. Prepare for the installation" on page 14](#)
- ["Step 3. Download the installation package" on page 16](#)
- ["Step 4 \(optional\). Install a dedicated NFS server" on page 17](#)
- ["Step 5. Set up one master node" on page 18](#)
- ["Step 6. Set up one or multiple worker nodes" on page 25](#)
- ["Step 7. Verify the installation" on page 27](#)

Step 1. Meet the system requirements

The ITOM Platform allows you to deploy a suite in an environment that is comprised of one master node 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 node, and the master node then sends the requests to the worker nodes.

Your environment must meet both ITOM Platform and Operations Bridge Suite requirements, as listed below.

ITOM Platform system requirements

Supported configurations

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.

In a production environment, you can use one master node, multiple worker nodes, and a separate NFS server.

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.

Hardware requirements

Depending on the capabilities you decide to install, your system must meet different minimum hardware requirements. The total minimum requirements are calculated by summing up the requirements per capability.

The sum of all worker node resources must match the totaling 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
ITOM PLATFORM (on the master node)			
ITOM Platform	16 GB	8 CPU cores	200 GB
NFS server (if the master is used as NFS server)	-	-	100 GB
CAPABILITIES (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)			
OBR Server	8 GB	8 CPU cores	100 GB
PostgreSQL	4 GB	4 CPU cores	20 GB
Collector	8 GB	4 CPU cores	50 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.

Operations Bridge Suite system requirements

Depending on the capabilities you want to install, your system must meet different database and client system requirements. Review the requirements listed below to make sure your container is correctly set up to host the suite.

Database requirements

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 not be `postgres`. In addition, the user that accesses the database must have permissions to create tables.

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

OMi database requirements

- **Internal PostgreSQL.** With an internal PostgreSQL database, the database instance is installed and configured in a separate container, and database files are stored on the NFS server.
- **Remote PostgreSQL.** If you use a remote database instance, OMi can configure it for you or you can configure it directly in the database management system (for example, if your organization does not allow the usage of administrator credentials during setup).

For detailed database requirements and instructions on creating database instances manually, see the *OMi Database Guide*.

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.

Caution: Make sure you configure the `pg_hba.conf` file on the PostgreSQL server to accept remote connections. Otherwise, the connection to the PostgreSQL database cannot be established.

- **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, specify the external PostgreSQL of your 10.12 deployment during the configuration to migrate your 10.12 data to the BVD container deployment.

The migrated data includes your dashboards, instances, API key, dashboard customizations, CSS customizations, and data integrations.

Before you specify the details of your 10.12 database in the Suite Installer, perform the following steps:

- a. Stop BVD 10.12. BVD must no longer be active on the database.
- b. 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.

To also migrate your LDAP user permissions and assignments, specify the LDAP server you used for BVD 10.12 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 Administration Guide*.

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 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.
- 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).

Step 2. Prepare for the installation

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 ["Step 1. Meet the system requirements" on page 8](#).
- The master node and the 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`).
- Use the root user for the installation.
- If the machine already has Docker or Kubernetes installed, uninstall them.
- Disable your existing firewall on all nodes by running the following commands:

```
systemctl stop firewalld
systemctl disable firewalld
```

The firewall has to stay disabled also after the installation.

- The following ports are needed on all nodes during and after the installation, and should not be used by another application: 2380, 4001, 4194, 5000, 5443, 8080, 8200, 8443, 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: Java is only required on the master node. `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 ITOM Platform. 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 node, and the worker nodes should be installed under the same subnet.
- Make sure that the browser cache is cleared.
- The time on the master and all worker nodes should be the same. To synchronize the time on your nodes, you can, for example, use NTP or VMWare tools.

Step 3. Download the installation package

To download and verify the ITOM Platform installation package, follow these steps:

1. Download the ITOM Platform and Suite installation package from the [OMi download area](#) on SSO.
2. Move or copy the installation package (CDF1704-15000.zip) to the master node, then unzip the file and the HPESW_ITOM_Suite_Platform_2017.03.<version>.zip file it contains to a temporary directory.

For example:

```
unzip CDF1704-15000.zip
```

```
unzip HPESW_ITOM_Suite_Platform_2017.03.00135.zip -d ITOM
```

Note: In the following ITOM Platform installation steps, the temporary directory HPESW_ITOM_Suite_Platform_2017.03.<version> will be referred to as *<platform_temp_dir>*.

The ITOM Platform installation package includes Docker and Kubernetes binaries.

Step 4 (optional). Install a dedicated NFS server

The ITOM Platform 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 ["Step 5. Set up one master node" on page 18.](#)

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.

Follow the steps below for the installation:

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

```
mkdir -p /var/vols/itom/core  
chown -R 1999:1999 /var/vols/itom/core
```

Note: You can expose a differently named folder. Name the exposed folder when installing the ITOM Platform.

3. Create a directory to store the suite data, and adapt the directory permissions:

```
mkdir -p /var/vols/itom/<opsbridge_directory>  
chown 1999:1999 /var/vols/itom/<opsbridge_directory>
```

Replace `<opsbridge_directory>` with a directory name you choose, for example `opsbridge`.

4. Configure the NFS sharing of the ITOM Platform and suite data directories:

```
echo "/var/vols/itom/core *(rw, sync, anonuid=1999, anongid=1999, all_squash)" >> /etc/exports  
  
echo "/var/vols/itom/<opsbridge_directory> *(rw, sync, anonuid=1999, anongid=1999, all_squash)" >> /etc/exports
```

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

```
exportfs -ra
```

Tip: Run `exportfs` to check what has been exported.

Step 5. Set up one master node

The following steps describe how to install the ITOM Platform on a master node.

1. Make sure you have already downloaded the installation package to a temporary directory on the master node. For details, see ["Step 3. Download the installation package" on page 16.](#)
2. Unzip the zip file.
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 master node as the NFS server.

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

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

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

```
<platform_temp_dir>/scripts/setupNFS.sh /var/vols/itom/<opsbridge_directory>
```

Replace *<opsbridge_directory>* with a directory name you choose, for example *opsbridge*.

4. On the master node, go to the *<platform_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>"
* INGRESS_HOST=<master node IP address>
* EXTERNAL_ACCESS_HOST=<master node FQDN>
* NFS_SERVER=<master node IP address>
* REGISTRY_ORGNAME=hpeswitom
```

Additionally, configure the proxy settings if you want to pull the Docker images from the master node and you need a proxy to connect to the internet.

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" on the next page.](#)

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.

5. On the master node, access the `<platform_temp_dir>` directory, and run the following command:

```
./install
```

Wait until the installation on the master node is complete.

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

In the following installation steps, the directory containing the installed ITOM Platform files (`/opt/kubernetes` by default) will be referred to as `<platform_install_dir>`.

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. Note that the Operations Bridge Suite 2017.04 does not support multiple master node setups.

Parameter	Description	Notes
MASTER_NODES	<p>Lists the cluster master nodes (IPv4 format), separated by a blank and enclosed in double quotes.</p> <p>Example:</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>Example:</p> <pre>WORKER_NODES="16.255.255.255"</pre>	Mandatory

INGRESS_HOST	<p>Defines the IP address (a single IPV4 address) of the node on which you want to start the Ingress Controller. You must use one of the master or worker nodes.</p> <p>Everything that runs on a cluster is actually on a private network, which is not externally accessible. If you want any suite functionality to be available from outside the network (for example, a Help Desk operative on client machine on another network that needs to access Service manager), you must provide an ingress into the cluster to be able to access the functionality. This is done by configuring the INGRESS_HOST and EXTERNAL_ACCESS_HOST parameters.</p> <p>Example:</p> <pre>INGRESS_IP=10.10.10.10 (IP address of one of the master nodes)</pre>	Mandatory
EXTERNAL_ACCESS_HOST	<p>Defines a fully qualified domain name for external clients to 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>Example:</p> <pre>EXTERNAL_ACCESS_HOST=myd.XXXX.YYY.net</pre>	Mandatory
NFS_SERVER	<p>Specifies the IP (IPv4) address of the NFS server that serves the persistent volumes of the cluster services.</p> <p>Example:</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>Example:</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</p>	Mandatory only for worker nodes

	<p>certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the <code>install.properties</code> file.</p> <p>Example:</p> <pre>CLIENT_CERT_FILE=/tmp/client.crt</pre>	
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 <code>install.properties</code> file.</p> <p>Example:</p> <pre>CLIENT_KEY_FILE=/tmp/client.key</pre>	Mandatory only for worker nodes
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>Example:</p> <pre>HA_VIRTUAL_IP="18.16.10.9"</pre>	Mandatory only if you are using multiple master nodes
HA_NGINX_NODES	<p>Specifies the IP addresses of the two master nodes that will run Nginx and keepalived for the API server Ingress load balancing. The value of the parameter is a space-delimited list of the two IPv4 address of the master nodes enclosed in double quotes.</p> <p>Example:</p> <pre>HA_NGINX_NODES="10.10.10.10 10.10.10.11"</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>Example:</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>	Mandatory only if you are using

	<p>Example:</p> <pre>PEER_CERT_FILE=/tmp/server.crt</pre>	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>Example:</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>Note: If a container stops and is restarted, all changes made inside the container are lost. If you want to save information such as configuration files, any other files, or databases, they must be located outside the container in a persistent volume provided by a Network File System (NFS). When you install the infrastructure, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server.</p> <p>Example:</p> <pre>NFS_FOLDER=/var/vols/itom</pre>	Optional
ROOTCA	<p>Specifies the root or intermediate CA certificate for generating 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 ITOM Platform, 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>Example:</p> <pre>ROOTCA=/tmp/ca.crt</pre>	Optional
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-</p>	Optional

	<p>signed certificate.</p> <p>Example:</p> <pre>ROOTCA=/tmp/ca.key</pre>	
NFS_STORAGE_SIZE	<p>Specifies the size of the NFS volume exported by the NFS server.</p> <p>Example:</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>Example:</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 <code>kubectl</code> command line tool communicates with the K8S server.</p> <p>Example:</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 <code>kubectl</code> command line tool communicates with the K8S server.</p> <p>Example:</p> <pre>MASTER_API_SSL_PORT=6443</pre>	Optional
THINPOOL_DEVICE	<p>Specifies the path to a Docker device mapper storage driver.</p> <p>To configure the thinpool device, see the Docker documentation.</p> <p>Note: If this parameter is specified, the installation will use the <code>devicemapper(direct-lvm)</code> Docker storage driver. If it is not specified, the installation will use <code>devicemapper(loop)</code>.</p> <p>Example:</p> <pre>THINPOOL_DEVICE=/dev/mapper/docker-thinpool</pre>	Optional
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</p>	Optional

	<p>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> <p>Example:</p> <pre>DOCKER_HTTP_PROXY="http://web.proxy.host.domain:8080"</pre>	
<p>DOCKER_HTTPS_PROXY</p>	<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> <p>Example:</p> <pre>DOCKER_HTTPS_PROXY="https://web.proxy.host.domain:8080"</pre>	<p>Optional</p>
<p>REGISTRY_ORGNAME</p>	<p>Specifies the organization name where the suite images are placed. The default name is hpeswitomsandbox.</p> <p>Example:</p> <pre>REGISTRY_ORGNAME=hpeswitomsandbox</pre>	<p>Optional</p>
<p>FLANNEL_IFACE</p>	<p>Specify the IPv4 address or the interface name for the Docker inter-host communication to use.</p> <p>Example:</p> <pre>FLANNEL_IFACE=10.10.10.10</pre>	<p>Optional</p>

Step 6. Set up one or multiple worker nodes

Add worker nodes via the ITOM Platform as follows:

Note: You can add nodes manually instead. For details, see ["How to add worker nodes manually" below](#).

1. Launch the ITOM Platform 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 ITOM Platform 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 of the node, the name of the root user, and the password of the specified user, and click **ADD** to remotely install the extra node.

How to add worker nodes manually

As an alternative to using the ITOM Platform GUI, you can also add worker nodes manually as follows:

Note: This step is *not* required if you already added nodes using the ITOM Platform GUI.

1. Make sure you have already downloaded the installation package to a temporary directory on all worker nodes. For details, see ["Step 3. Download the 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 <platform_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 master node to all worker nodes into the `<platform_temp_dir>` directory.
6. On each worker node, open the `install.properties` file under the `<platform_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.

Step 7. Verify the installation

Once the ITOM Platform 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 ITOM Platform 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 ITOM Platform 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 ["Troubleshooting" on page 50](#).

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 28](#).

Install the Operations Bridge Suite

Once the ITOM Platform 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 ITOM Platform.

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

- ["Step 1. Import suite images to the local registry" on page 29](#)
- ["Step 2. Activate a suite license" on page 32](#)
- ["Step 3. Run the Suite Installer" on page 34](#)
- ["Step 4. Verify the Suite installation" on page 47](#)

Step 1. Import suite images to the local registry

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

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

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

```
cd <platform_install_dir>/scripts  
./downloadimages.sh -o hpeswitom/
```

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

Suite	OpsBridge
User name and password	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 2017.04

2. On the 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 100GB.

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 repo:

```
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 master node to a machine on which you have internet access, for example into the /tmp directory:

```
/var/vols/itom/core/suite-install/suite_feature/opsbridge/2017.04/opsbridge_
suitefeatures.2017.04.json
```

```
<platform_install_dir>/scripts/downloadimages.sh
```

```
<platform_install_dir>/bin/jq
```

- b. Skip this step if you have already installed jq. If not, move jq to /usr/local/bin/ by using the following commands:

```
chmod 777 jq
```

```
mv jq /usr/local/bin
```

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:

Suite	OpsBridge
User name and password	Enter your Docker Hub account credentials.
Suite version	Enter the suite version you want to install, for example 2017.04

5. Copy all files from /tmp/image_tars to your 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 <platform_install_dir>/scripts
```

```
./uploadimages.sh
```

When prompted for the Suite, enter OpsBridge.

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

Step 2. 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 ITOM Platform license is included in the suite licenses. You only need a suite license to install a suite. The suite license contains license keys for all capabilities of 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 ITOM Platform, do the following:

1. Launch the ITOM Platform 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 ITOM Platform installation. Usually, this is the master node's FQDN.

2. Log in as the admin user.
3. Click **ADMINISTRATION > 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.

Step 3. 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 ITOM Platform 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 ITOM Platform installation. Usually, this is the master node's FQDN.

2. Log in to the ITOM Platform 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. Select **Use system default NFS share** if you are using the master node as NFS server, and **Use other NFS server** if you are using your own NFS server. The exported share is the directory you created to store the suite data: `/var/vols/itom/<opsbridge_directory>`

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 ITOM Platform, a 60-day evaluation license, and the same password for the administrator and the PostgreSQL database user.

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.

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 ITOM Platform** 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 OMi administrator 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, that 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 reduce.

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`. 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 `<OMi_HOME>/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 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 "[Database requirements](#)" on [page 10](#). 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.

Scanning Interval

Time interval (in minutes) at which the aging process scans the database to identify and automatically delete data records.

The value must be an integer greater than 0.

Default: 30 minutes

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.

-
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 ITOM Platform.

Wait until the installation is complete.

Step 4. 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**.

Alternatively, you can also verify the status of the pods via the ITOM Platform:

- a. Launch the ITOM Platform 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**.

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 ITOM Platform 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 the `serverStatus.jsp` tool in OMi:

```
/opt/HP/BSM/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](#).

Uninstall the ITOM Platform

You can backup image tars from local private registry to remote registry before you uninstall the ITOM Platform.

Optional. Back up the image tars

1. Go to directory where the `local_backup.sh` file is located: `<platform_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 ITOM Platform

1. Go to the `<platform_install_dir>` directory, and run `uninstall.sh`

The uninstallation process stops the containers and removes them.

2. Reboot the server.

Troubleshooting

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

- ["Manual verification commands" below](#)
- ["Log files" on the next page](#)
- ["Common problems and limitations" on page 52](#)

Manual verification commands

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

```
./kube-status.sh
```

Displays the status of the K8S cluster.

```
./kube-stop.sh
```

Stops the K8S cluster.

```
./kube-restart.sh
```

Restarts the K8S cluster.

```
./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/k1
```

This enables you to type `k1` 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 ITOM Platform 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> --namespace=<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> --namespace=<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/opsbridge/*<namespace>*/omi/opt/HP/BSM/log/topaz_all.log
- /var/vols/itom/opsbridge/*<namespace>*/omi/opt/HP/BSM/log/jboss7_boot.log
- /var/vols/itom/opsbridge/*<namespace>*/omi/opt/HP/BSM/log/supervisor/nanny_all.log
- /var/vols/itom/opsbridge/opsbridge-opsbridge/pe/logs

Login

/var/vols/itom/opsbridge/*<namespace>*/omi/opt/HP/BSM/log/jboss/login.log

OBR

- /var/vols/itom/opsbridge/*<namespace>*/obr/obr-server/adminServer/logs (Administration UI)
- /var/vols/itom/opsbridge/*<namespace>*/obr/obr-server/Flink/log (Flink)
- /var/vols/itom/opsbridge/*<namespace>*/obr/obr-server/log (OBR platform)
- /var/vols/itom/opsbridge/*<namespace>*/obr/obr-server/data (OBR data)
- /var/vols/itom/opsbridge/*<namespace>*/obr/obr-server/config (OBR configuration)

Common problems and limitations

You may encounter the following problems and limitations when installing or administering the ITOM Platform and the Operations Bridge Suite.

ITOM Platform is not accessible

Description

After the installation of the ITOM Platform, the ITOM Platform 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 ITOM Platform.
- See also ["ITOM Platform is not accessible: nginx controller is Pending"](#) below, ["Troubleshooting on page 50"](#) and ["Login to ITOM Platform is not possible: IDM service is not ready yet"](#) on the next page.

ITOM Platform is not accessible: nginx controller is Pending

Description

After the installation of the ITOM Platform, the ITOM Platform 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`

ITOM Platform is not accessible: Gateway time out

Description

After the installation of the ITOM Platform, the ITOM Platform 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 ITOM Platform is not possible: IDM service is not ready yet

Description

After the installation of the ITOM Platform, the ITOM Platform 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
```

"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 ITOM Platform, 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 ITOM Platform, 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 ITOM Platform 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 ITOM Platform.

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

Description

After the installation of the ITOM Platform, 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.

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 2017.04)

Just add your feedback to the email and click send.

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

We appreciate your feedback!