Install

Operations Bridge Suite 2017.11

MICRO FOCUS

Install

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 the 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.

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.

Upgrade overview

You can upgrade your existing 2017.08 deployment to 2017.11 as described in the Upgrade section.

Installation overview

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

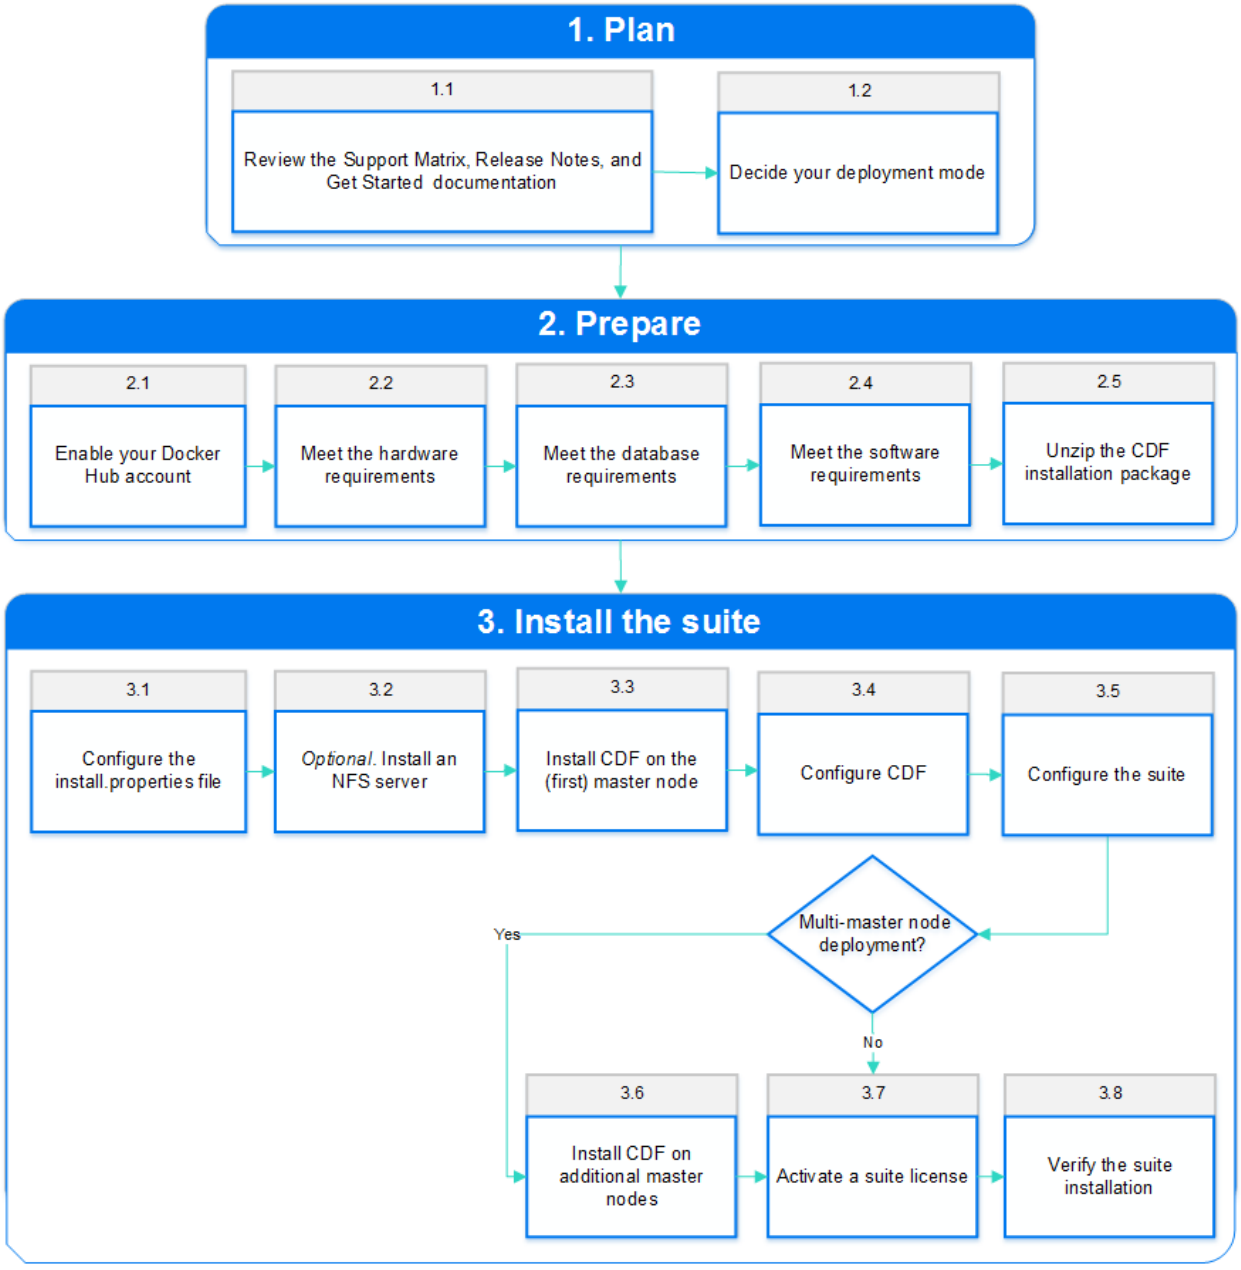

Plan your suite deployment

Go to start of metadata

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 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 and applying additional configurations, you can make your system highly available. For more information, see Configure scaling and high availability.

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.

Prepare for the installation

- Enable your Docker Hub account
- Meet the hardware requirements
- Meet the database requirements
- Meet the software requirements

- Unzip the CDF installation package

Enable your Docker Hub account

You must create a Docker Hub account and then send an email 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 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 customer SAID (must be valid and active)
    - ITOM Suite edition (Operations Bridge Suite)

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

- AMERICAS: dockersupport.ams@hpe.com
- APJ: dockersupport.apj@hpe.com
- EMEA: dockersupport.emea@hpe.com

Your Docker ID will be enabled and you will receive 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, we recommend 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

Review the following hardware requirements based on your deployment:

- If you do not want to enable OMi HA, see Standard setup (no OMi high availability).

- If you want to enable OMi HA, see Standard setup (no OMi high availability) for all capabilities other than OMi, and for OMi, see HA setup (OMi high availability).
- Additionally, the disk space requirements for CDF must be met - see CDF requirements.

Standard setup (no OMi high availability)

Component	RAM	Processors	Disk space
CONTAINER DEPLOYMENT FOUNDATION (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
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) - 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

Vertica and Business Objects are not containers, but they require additional resources on a separate system. For more information, see the *Install* section in the OBR Help Center.

CDF requirements

In addition, the directories of the cluster nodes must have sufficient free space as shown in the following table.

Directory	Equivalent directory	Host	Required free disk space	Description
/opt/kubernetes	To specify a customized directory, follow one of the steps below: • • Modify the K8S_HOME parameter in install.properties • Run the following command during the installation: `./ install.sh --k8s-home <your_home_directory>`	Master	20 GB	This directory is for the Kubernetes server, CDF installer and containers.
		Worker	According to the setting of `storage_gb` in suitefeatures.json	
/	N/A	Master	5 GB	If the mount point of / and /var is the same, the required free disk space is 10 GB.
		Worker	5 GB	
/var		Master	5 GB	
		Worker	5 GB	
/tmp	To specify a customized directory, follow one of the steps below: • • Modify the TMP_FOLDER parameter in install.properties	Master	5 GB	This folder is required for the CDF build.
		Worker	5 GB	

	• Run the following command during the installation: `./ install.sh --tmp-folder <your_tmp_directory>`			
/var/opt/kubernetes	Run the following command to specify a customized directory: `./downloadimages.sh --dir <your_directory>`	Master	13 GB	This directory includes the CDF images and all suite images. The subdirectory /offline/suite_images can be removed after uploading the suite images.
	N/A	Worker	N/A	N/A

HA setup (OMi high availability)
OMi's RAM and number of required processors depend on your decision to enable high availability. If you decide to set OMi up to be highly available, refer to these OMi requirements instead of the OMi requirements listed above (the requirements for the other capabilties and CDF remain the same):

CAPABILITIES (on the worker nodes)			
Operations Bridge Manager (OMi) - small deployment	24 GB	5 CPU cores	50 GB
Operations Bridge Manager (OMi) - medium deployment	43 GB	8 CPU cores	75 GB
Operations Bridge Manager (OMi) - large deployment	64 GB	11 CPU cores	100 GB

We recommend 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 non-HA 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.

Meet the 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

BVD requires a database to store information. When configuring BVD, you can either select an internal PostgreSQL to have the database created for you, or specify an existing external 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 2017.11). 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 10.63:

1. Stop your existing BVD deployment. BVD must no longer be active on the database.
2. *BVD 10.12 migrations only*. 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. When running the Suite Installer, specify the external PostgreSQL database of your former deployment.
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.
   For more information about the LDAP configuration, see Configure LDAP authentication.

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.

Vertica does not support VMware Vmotion and Logical Volume Manager (LVM) on any system where database files are stored. We recommend VMware ESX 5.5 Hypervisor to virtualize the 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.

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

OMi database requirements

You can use an internal PostgreSQL database instance, or a remote 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 Help Center > *Develop* > *Prepare the database environment*.

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

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.
- 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 Enable a firewall on a node.
- The following ports are needed on all nodes during and after the installation, and should not be used by another application: 383, 443, 2380, 3000, 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).

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

- The system user needs the UID and GID 1999 so that the PostgreSQL container can access the NFS share.

- 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|httpd-tools"

systemd-libs must be version 219 or higher.

If one or multiple of the packages are not installed, install them by 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 httpd-tools
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 of the system clock on all master nodes, all worker nodes, and on the client systems 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.

- If you install OBR, your environment must fulfill the external SAP BO requirements as described in the OBR Help Center > Install > Installation tasks (interactive Installation Guide).

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).

Unzip the CDF installation package

To unzip and move the CDF installation package, follow these steps:

1. Download the CDF and Suite installation package CDF1710-00292-15001-install.zip from the location that was communicated to you after obtaining your license and sending your Docker ID via email. For more information, see Enable your Docker Hub account.
2. Move or copy the installation package (CDF1710-00292-15001-install.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 CDF1710-00292-15001-install.zip
unzip HPESW_ITOM_Suite_Platform_2017.10.00292.zip -d ITOM
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 suite

The Operations Bridge Suite must be deployed on the Container Deployment Foundation (CDF), where you can deploy and administer suites.
Follow these steps to install CDF and the suite:

- Configure the install.properties file
- Optional. Install an NFS server
- Install CDF on the (first) master node
- Configure CDF
- Configure the suite
- Install CDF on additional master nodes

- Activate a suite license
- Verify the suite installation

Configure the install.properties file

On the (first) master node, go to the `<foundation_temp_dir>` directory, and edit the `install.properties` file. The parameter **HA_VIRTUAL_IP** is mandatory if you want to install multiple master nodes. To specify the **HA_VIRTUAL_IP**, you can configure it in the install.properties file. Alternatively, you can add the **--ha-virtual-ip** command line option when running the **./install** command during the CDF installation.
Other than **HA_VIRTUAL_IP**, all other parameters are optional. You can leave them empty to use the default values.

`HA_VIRTUAL_IP`	Sets up a virtual IP address (single IPv4 address enclosed in double quotes) when setting up multiple master nodes. A virtual IP (VIP) is an IP address that is shared by all members of a HA server pool. The VIP is used for the connection redundancy by providing fail-over for one machine. When a member of the pool goes down, the other pool member takes over the VIP address and responds to requests sent to the VIP. The VIP and each pool member must exist in the same sub-net. Since the VIP does not correspond to an actual physical network interface, users do not need to make any configuration. They only need to provide a virtual IP address and make sure that the IP address is not occupied before the installation. The requests to the API server should be sent to a VIP for higher availability. **Example**: `HA_VIRTUAL_IP=18.16.10.9`	Mandatory only if you are using multiple master nodes
`K8S_HOME`	Specifies the installation directory (fully-qualified directory) for the core platform binaries. **Example**: `K8S_HOME=/opt/kubernetes`	Optional
`KEEPALIVED_NOPREEMPT`	Specifies whether the Keepalived is in preempt mode. When the Keepalived is in preempt mode, VERP will preempt a lower priority machine when a higher priority machine comes online. When the Keepalived is in nopreempt mode, it allows the lower priority machine to maintain the master role even when a	Optional

	higher priority machine comes back online. By default, the value of *KEEPALIVED_NOPREEMPT* is true. **Example:** `KEEPALIVED_NOPREEMPT=false`	
`MASTER_API_PORT`	Specifies the HTTP port for the Kubernetes (K8S) API server. If you want to use K8S, you must dock to the K8S API server. The `kubectl` command line tool communicates with the K8S server. **Example**: `MASTER_API_PORT=8080`	Optional
`MASTER_API_SSL_PORT`	Specifies the HTTPS port for the K8S API server. If you want to use K8S, you must dock to the K8S API server. The `kubectl` command line tool communicates with the K8S server. **Example**: `MASTER_API_SSL_PORT=8443`	Optional
`THINPOOL_DEVICE`	Specifies the path to a Docker device mapper storage driver. To configure the thinpool device, see the Docker documentation. 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, Micro Focus recommends devicemapper (direct-lvm). **Example**: `THINPOOL_DEVICE=/dev/mapper/docker-thinpool`	Optional
`DOCKER_HTTP_PROXY`	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. When you install suites and launch containers on Docker inside the K8S cluster, you may need to	Optional

	download the images from the internet, for which you need to use proxies. **Example**: `DOCKER_HTTP_PROXY="`http://web.proxy.host.domain:8080`"`	
`DOCKER_HTTPS_PROXY`	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. 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. **Example**: `DOCKER_HTTPS_PROXY="`https://web.proxy.host.domain:8080`"`	Optional
`DOCKER_NO_PROXY`	Specifies the IPv4 address or FQDN that does not need the proxy settings for Docker. **Example**: `DOCKER_NO_PROXY=127.0.0.1`	Optional
`FLANNEL_BACKEND_TYPE`	Specifies the backend type of the flannel network. The acceptable values for this parameter are `host-gw` and `vxlan`. By default, the value is `host-gw`. When setting up CDF cluster nodes in different subnets, set this parameter to `vxlan`. When setting up CDF cluster nodes in the same subnet, set this parameter to `host-gw`. **Example:** `FLANNEL_BACKEND_TYPE=host-gw`	Mandatory only if you install CDF on a node that has more than one network adapter.
`FLANNEL_IFACE`	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. **Example**: FLANNEL_IFACE=10.10.10.10	Mandatory only if you install CDF on a node which has more than one network adapter.

REGISTRY_ORGNAME	Specifies the organization name where the suite images are placed. The default name is `hpeswitomsandbox`. **Example**: `REGISTRY_ORGNAME=hpeswitom`	Optional
CLOUD_PROVIDER	Specifies the cloud provider when installing the CDF on a cloud server. **Example**: `CLOUD_PROVIDER=AWS`	Optional
AWS_REGION	Specifies the AWS region to use when choosing AWS as the cloud provider. The default value of this parameter is an empty string. **Example**: `AWS_REGION=us-east-1`	Mandatory only if you choose AWS as the cloud provider.
AWS_MASTER_NODES	List the Pv4 address or FQDN of cluster master nodes for AWS provider. **Example**: `AWS_MASTER_NODES=10.10.10.10 10.10.10.11 10.10.10.12`	Mandatory only if you install multiple master nodes on AWS.
SYSTEM_USER_ID	Specifies the user ID that is used to start the process in container. By default, the value is 1999. The value of the SYSTEM_USER_ID is between 100000 and 2000000000. This value must be the same value of the NFS folder owner's user ID. **Example**: `SYSTEM_USER_ID=1999`	Optional
SYSTEM_GROUP_ID	Specifies the group ID that is used to start the process in container. By default, the value is 1999. **Example**: `SYSTEM_GROUP_ID=1999`	Optional

`TMP_FOLDER`	Specifies the absolute path of the temporary folder for placing temporary files. By default, it is `/temp`. **Example:** TMP_FOLDER=/tmp	Optional

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.
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.

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.

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.

Depending on the capabilities you want to deploy, you might have to configure additional volumes later on.

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*>

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
```

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 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.

2. Unzip the zip file on the master node.

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 command 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.

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

./install -m opsbridge-2017.11-metadata.tar.gz
Enter the administrator's password and wait until the installation on the first master node is complete. You can check the installation log at /opt/kubernetes/install-<date><time>.log

5. If you chose to use an external database in the install.properties file, enter the database user name, password (and schema for EXTERNAL_ORA) when prompted.

Configure CDF

After installing CDF on the (first) master node, configure CDF on the management portal as follows:
1 - Select
2 - Download
3 - Prepare
4 - Configure/Deploy

1 - Select

Start the configuration by selecting the suite, capabilities, deployment size, and so on.

1. Log on to the management portal by using the admin credentials you specified earlier:

http://<hostname>:3000
The <hostname> is the FQDN of your (first) master node.

2. Select a suite version from the drop-down list and click **Next**.

3. Agree to the license terms and click **Next**.

4. Select the suite capabilities you want to install and click **Next**.

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.
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.
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.
Operations Bridge Reporter. Operations Bridge Reporter (OBR) is a solution based on the Big Data technology 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, Operations Bridge Reporter also collects and collates specific information on the relationships between the IT elements and the business services.

5. Specify the database that CDF will use for the service connection and run-time data of the suite.

When you choose PostgreSQL or Oracle as the database, make sure that the database is empty.

6. Select a deployment size according to your environment and click **Next**.
7. *Optional*. Specify the external host name and port that will be used to access the management portal. By default, the external host name and port are those of the first installed master node. Click **Use custom certificates** to upload custom certificates for the ingress. Click **Next**.
8. *Optional*. Click **Make master highly available** if you want to add one or two additional master nodes. For each additional master node, enter the host name or IPv4 address, the user name, password, and (optionally) the ThinPool Device path and Flannel IFace if you have multiple active network interfaces.
9. You can click **Allow suite workload to be deployed on the master node** (not recommended) to deploy the worker load on the master node. In that case, you do not have to add worker nodes.
   For each worker node, click **Add** and specify the type, the host name or IP address, the user name, the password, and (optionally) the ThinPool Device path and Flannel IFace if you have multiple active network interfaces. Click **Save**, and then **Next** once you have added all worker nodes.
10. Specify the FQDN of the dedicated NFS server or your (first) master node. Then specify the mount points / exported paths for the NFS volumes db-volume, log-volume, and conf-volume. Additionally, if you install OMi, configure the mount points required for the StatefulSet pods omi-volume-0 and omi-volume-1.
    Click **Validate** for each file system, then click **Next** to continue.
11. Click **Yes** to go to the next configuration section. You cannot go back to the Select phase after clicking Yes.

2 – Download

You must download the suite images from Docker Hub and import the images to the local registry of the (first) master node. If your first master node has internet accConfigure the suiteess, follow the next steps. If it does not have internet access, see Download suite images to another machine.

1. Select Offline Download and click **Next**.
2. Read the instructions for the offline download and click **Next**.
3. Click **Get script package** to download a zip file, and click **Next**.
4. Copy the downloaded zip file to the suite image directory. By default, the suite image directory is **/var/opt/kubernetes/offline/suite_images.**
5. Download the suite images as follows:
   a. Extract the zip package to a local directory.
   b. Navigate to the **offline_image_download** directory and run the following command:

./downloadimages.sh
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.
Suite version	Enter the suite version you want to install, for example 2017.11

6. Import the suite images as follows:

    a. Navigate to the **<*K8S_home*>/scripts/** directory (by default, this is **/opt/kubernetes/scripts**).

    b. Run the following command:

./uploadimages.sh

    c. Wait until the images are successfully imported.

7. Go back to the management portal, and click **Next**. Check how many suite images were imported. You can click **Check Again** to see the details of the updated images. Then click **Next** to go to the next configuration section. You cannot go back to the Download phase after clicking Next.

3 - Prepare

The prepare phase makes your deployment ready for configuration.

1. Review the list of master and worker nodes that are being deployed. If there is an error displaying that the deployment of a node failed, click **Retry** to deploy the node again. Then click **Next**.
2. On the Deployment of Infrastructure Services page, click **Next**.
3. Review the list of core foundation services that are being deployed, and click **Next**.
4. Click **Next** to go to the next configuration section. You cannot go back to the Prepare phase after clicking Next.

4 - Configure / Deploy

In the last phase, you configure the suite. For details, see Configure the suite.

Configure the suite

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

In the Configure/Deploy configuration section, you configure and deploy the suite:

1. 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.NoteThe 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. This hostname must be resolvable via the DNS server, not only via /etc/hosts.
Port number: The port of the load balancer. Default: 443
TLS 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.

2. *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 defaults: Select to use the administrative user account credentials that you specified during the suite configuration.
Custom configuration: 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 Help Center > *Develop* > *Prepare the database environment*.
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. If you configure OMi with an external PostgreSQL, you have to create an additional database user with "create" permissions
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)LoginThe 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 > High Availability
Decide if you want to enable high availability for OMi by using two worker nodes for automatic failover.

Enable high availability for OMi: Tick the checkbox to enable HA for OMi. Note that you can should only set up high availability if you have at least two worker nodes and two master nodes. For more information on high availability, see Configure scaling and high availability.
Note
It is currently not possible to unconfigure high availability for OMi once it has been enabled and deployed.

OMi > Management Packs
Select the 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 Access Command Line Interfaces.
opr-mp-installer by default installs management packs from the /opt/HP/BSM/opr/mgmtpacksdirectory 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.

To update a management pack to a later version than the one included with OMi, download its installation package from the 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

3. *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 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 former deployment.
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

4. *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

***Optional*. 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.

5. *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.
***Optional*. Enable TLS**: If TLS is enabled on Vertica, click this checkbox to enable TLS communication between the OBR server and Vertica. You will have to upload the CA certificate file from the Vertica system and specify a truststore password. If TLS is not enabled on Vertica, uncheck this option and proceed.
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 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.NoteAfter 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*.

6. 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.

Install CDF on additional master nodes

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

1. On the first installed master node, go to the <*K8S_home*> directory.

The default <*K8S_home*> directory is /opt/kubernetes.

2. Open **start_lb.sh** with a text editor.
3. Locate the parameters "PEER1_IP" and "PEER2_IP" in the **start_lb.sh** file.
4. Replace "PEER1_IP" and "PEER2_IP" with the IPv4 addresses of the second and third master nodes.
5. Run the following command:

./start_lb.sh

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 CDF 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 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, do the following:

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

https://<*external_hostname*>:5433
<*external_hostname*> is the fully qualified domain name of the host which you specified in the Connection step during the CDF configuration. Usually, this is the master node's FQDN.

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**.

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_hostname*>/omi or https://<*external_hostname*>:<*port*>/omi
BVD: https://<*external_hostname*>/bvd
OBR: https://<*external_hostname*>/OBRApp
<*external_hostname*> is the fully qualified domain name of the host which you specified in the Connection step during the CDF configuration. Usually, this is the master node's FQDN.
If you specified a load balancer in the Connection step of the suite configuration, <*external_hostname*> and <port> are the external hostname and port of the load balancer.

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

https://<*external_hostname*>/topaz/serverStatus.jsp

5. *Optional*. If you installed OBR, you can access the BO UI as described in the OBR Help Center > Use > Log on > Logging On.

Reconfigure

You can reconfigure OMi, BVD, PE, and the suite defaults 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 master node.

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 your capabilities as required. For information about the available settings, see Configure the suite.

   Note the following:
   - You can add Management Packs, but you cannot remove them.
   - You cannot change the OMi high availability setting.
   - For BVD, specify your former database in order to migrate your previous BVD data automatically.
   - If you reconfigure PE, ensure that you integrate PE with OMi again. For more information about the integration, see PE integrations.

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 the following command:

./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 servers.
4. *Optional*. In multi-master node deployments without DNS, delete the hosts files in the /etc directory.