ITSM Automation NG Express

Software release version: 2017.04

Installation Guide

Document release date: April 2017
Product release date: April 2017

Hewlett Packard Enterprise

Please note that this document has been exported from the HPE Software Documentation Portal wiki, which is the primary mode of documentation delivery. For the most current documentation, go to: https://docs.software.hpe.com

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted rights legend

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

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

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.

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE Support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

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 HPESW 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/.

Contents

Contents

This section provides instructions about installing ITOM Container Deployment Foundation (CDF) and ITSMA Suite Express.

ITMSA NG Express leverages container technology from 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. ITSMA NG Express components are deployed as containerized applications that are integrated with each other. What you need to do is to first install a container management framework (referred to as "ITOM CDF" in the documentation), and then install the ITSMA suite from a graphic user interface (the "Suite Installer"). The suite components are deployed quickly and integrated seamlessly, requiring little user intervention.

> Third-party components such as a PostgreSQL database, Tomcat, and OpenLDAP are bundled in the suite images to simplify the deployment process in a test environment; in a production environment, you can configure external Oracle and PostgreSQL databases during installation and configure an external LDAP Server after installation. For more information, see Support matrix.

Installation procedure

The ITSMA NG Express installation procedure includes the following major steps.

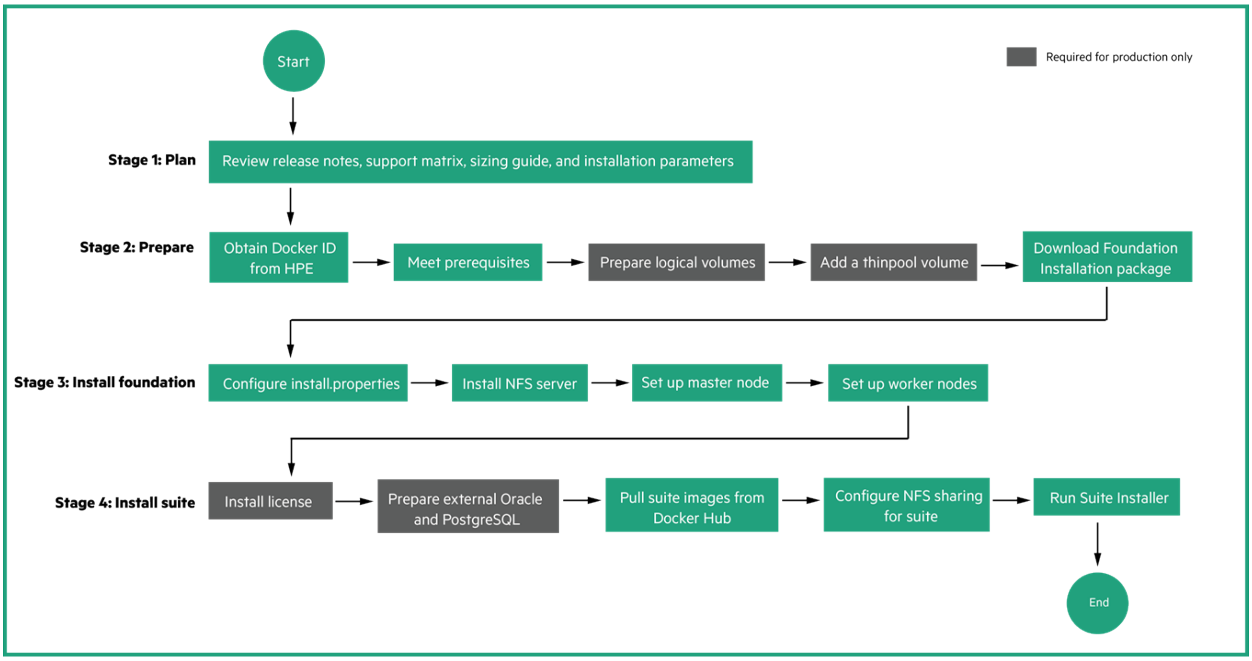
Start
Required for production only
Stage 1: Plan
Review release notes, support matrix, sizing guide, and installation parameters
Stage 2: Prepare
Obtain Docker ID from HPE
Meet prerequisites
Prepare logical volumes
Add a thinpool volume
Download Foundation Installation package
Stage 3: Install foundation
Configure install.properties
Install NFS server
Set up master node
Set up worker nodes
Stage 4: Install suite
Install license
Prepare external Oracle and PostgreSQL
Pull suite images from Docker Hub
Configure NFS sharing for suite
Run Suite Installer
End

Stage	Steps	Notes
Plan	1. Review the Release notes, and Get started sections to obtain a basic understanding of the suite. 2. Read about the hardware requirements, operating systems, sizing, and more to plan your suite deployment. 3. Prepare your installation environment based on the requriements in step 2: - One master node, multiple worker nodes depending on your deployment size, and a separate NFS server - External database (Oracle and PostgreSQL) - External LDAP server (Microsoft Active Directory) 4. Make sure that you have an HPE Passport account (which you can create at https://softwaresupport.hpe.com) so that you can access the HPE Software Entitlement Portal.	• The built-in Postgres database and OpenLDAP server are designed only for demonstration purposes. Configure external ones in production environments.
Prepare	1. Obtain a Docker Hub account from HPE. 2. Prepare your cluster machines to meet the prerequisites. 3. Mount a logical volume on each node for installation. 4. Add a thinpool volume on each node. 5. Download the ITOM CDF installation package from the HPE Software Entitlement Portal.	• You must use a Docker Hub account provided by HPE to download ITSMA suite images from Docker Hub. • Step 3 and 4 are needed for production environments only.

Install ITOM CDF ("Foundation")	1. Configure the install.properties file. 2. Install an NFS server. 3. Set up the master node. 4. Set up the worker nodes.	• In a production environment, you need to copy the downloaded ITOM CDF installation package to all cluster nodes (master and workers). • In a test environment, you can download the installation package to the master node only and add the worker nodes from the ITOM CDF user interface (see Install ITOM CDF on the worker nodes). • The reason why different worker node setup methods are required for production and test environments is that a production environment requires you to add a thinpool device for all nodes, and the thinpool device configuration does not work for the worker nodes if you add them from the ITOM CDF user interface. • The install.properties file defines a bunch of installation parameters. In a production environment, you can configure this file on the master node and copy it to all worker nodes. In a test environment, you only need to configure this file on the master node because you can add the worker nodes from the user interface (which is much simpler). • In a production environment, using a dedicated NFS server is recommended. Use the master node as the NFS server only in a test environment. • ITOM CDF has two user roles: IT Administrator and Suite Administrator. The out-of-box IT Administrator user account is **admin**/**cloud**.

Install ITSMA	1. Install an ITSMA license. 2. Prepare your databases. 3. Import the ITSMA suite images from Docker Hub to the local registry of the master node. 4. Configure NFS sharing for ITSMA. 5. Run the Suite Installer.	License installation • You can skip license installation to use a 30-day trial license. After the installation, you can purchase a perpetual license and replace the trial license. See Replace an ITSMA trial license. • A unique locker code is generated for each ITOM CDF installation and is used for license activation. • Each ITSMA entitlement enables you to activate two sub-licenses: one for the Service Portal capability, and the other for the rest of ITSMA capabilities. • You have two ways to activate your license from the ITOM CDF user interface: install the license files that you received, or enter the activation code that you received from HPE. ITSMA installation • **Important:** During installation, make sure that you skip the LDAP configuration step to use the built-in LDAP server. After installation, you can switch to an external LDAP server. See Configure an external LDAP server. • During installation, you are asked to specify an initial password for the ITSMA administrator user "sysadmin". Be aware that you have no way to reset this password if you forget it.

Post-installation configurations

After installation, you need to perform mandatory and optional configurations to set up the suite, such as external LDAP configuration and master data onboarding. For more information, see Administer the ITSMA suite.

Plan your suite deployment

Before your proceed, review the support matrix information, sizing recommendations, and configuration parameters to plan your suite deployment.

- Support matrix and Sizing recommendations: review the system requirements and sizing recommendations for your suite installation, and prepare your environment accordingly.
- ITOM CDF installation configuration: learn about the mandatory and optional configuration parameters that determine how your suite will be deployed. Decide which parameters you want to configure.

Once you have completed the deployment planning and are ready to deploy the suite, go to Install HPE ITOM CDF.

Support matrix

This section provides support matrix information of ITSMA NG Express.

- Supported environments
- Supported configurations
- Operating systems
- Required Network identification (FQDN)
- Databases
- LDAP
- Browser
- Other requirements

Supported environments

The following environments are supported:

- Physical environment
- Virtual environment (VMware)

Supported configurations

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

> This release of ITSMA NG Express does not support multiple master nodes.

In a test environment, you can use the following configuration:

- One master node (also used as the NFS server)
- One or more worker nodes

In a production environment, you must use the following configuration:

- One master node
- At least three worker nodes
- A separate NFS server

For sizing recommentations, see Sizing recommendations.

Operating systems

The master node, worker nodes, and the NFS server hosts must use the same operating system that is one of the following:

- 64-bit RedHat Enterprise Linux: 7.2, 7.3
- 64-bit CentOS: 7.2, 7.3
- 64-bit Oracle Linux: 7.3

Required Network identification (FQDN)

Only IPv4 is supported.

Databases

You have the option to use the internal database or two external databases:

- Internal: PostgreSQL 9.4 (recommended for test environments)
- External (recommended for production environments):
  - PostgreSQL 9.5.5: for Service Portal
  - Oracle 12c: for Service Management and CMDB

> If you want to use the internal PostgreSQL database in a production environment, see Apply PostgreSQL parameter updates in ITSMA and Bind the internal PostgreSQL database to a dedicated worker node for more information.

LDAP

You have the option to use either the internal or an external LDAP server for the ITSMA suite:

- OpenLDAP: bundled with the suite, which is pre-configured and requires no manual configuration (for test environments only)
- An external LDAP server: Microsoft Active Directory (for production environments)

Browser

Use the the following browsers to access the ITOM CDF and ITMSA:

- Internet Explorer (IE)11
- Latest Firefox
- Latest Chrome

> - For Internet Explorer, you must set **English[en]** in the browser as the language for the English locale. Additionally, there are known issues when accessing ITOM CDF using IE11 (see Known issues, limitations, and workarounds).
> - To access the CMDB Administrator capability UI from the latest Chrome or Firefox, enable JNLP by following instructions here. This is a one-time operation.

Other requirements

Item	Support matrix	Notes
Screen resolution	1600x900, 1280x1024, 1920x1200, or higher	For the client machine running the web browser. The resolutions are applicable to different types of devices, such as laptops, PC monitors, and larger meeting room monitors.
Mobile operating system and browser	• iOS 9.x with Safari • Android 6.x, 5.x with Android browser	For accessing the Mobility capability of the ITSMA suite

Sizing recommendations

When running the Suite Installer to install ITSMA, you need to select a suite size: **Extra Small**, **Small**, or **Medium**. Different hardware configurations are required for different suite sizes.

> - A CPU load of more than 80% impacts the efficiency of the network transmission significantly inside the platform environment. Make sure that the CPU load is less than 80% by distributing the work load to multiple worker nodes.
> - HPE recommends running Docker with the devciemapper (direct-lvm) storage driver. Make sure that a 60 GB logical device is added to run Docker with the devciemapper (direct-lvm) storage driver. If not, Docker will run with the devicemapper (loop) storage driver. For more information, see Add a logical volume for direct-lvm thinpool (production only).
> - If you want to use the internal PostgreSQL database in a production environment, see Apply PostgreSQL parameter updates in ITSMA and Bind the internal PostgreSQL database to a dedicated worker node for more information.

Extra Small

This deployment size is defined as follows:

- Maximum number of concurrent users in Service Management: 50
- Maximum number of records in Smart Analytics (millions): 1
- Maximum number of CIs and Relationships in CMDB (millions): 1

Hardware requirement:

- **1 master + 2 worker nodes**:
  8 CPU cores / 32 GB RAM / 200 GB disk (at least 60 GB thinpool size) for each node
- **1 separate NFS Server**:
  200GB, RAID 10, IO: 280 MB/s
- **3 additional worker nodes if using built-in PostgreSQL for Service Management, CMDB and Service Portal**:
  4 CPU cores / 16 GB RAM / 200GB disk (at least 60 GB thinpool size) for each node

Small

This deployment size is defined as follows:

- Maximum number of concurrent users in Service Management: 200
- Maximum number of records in Smart Analytics (millions): 1
- Maximum number of CIs and Relationships in CMDB (millions): 1

Hardware requirement:

- **1 master + 3 worker nodes**:
  8 CPU cores / 32 GB RAM / 200GB disk (at least 60GB thinpool size) for each node
- **1 separate NFS Server**:
  200GB, RAID 10, IO: 280 MB/s
- **3 additional worker nodes if using built-in PostgreSQL for Service Management CMDB, and Service Portal**:
  6 CPU cores / 16 GB RAM / 200GB disk (at least 60GB thinpool size) for each node

Medium

This deployment size is defined as follows:

- Maximum number of concurrent users in Service Management: 500
- Maximum number of records in Smart Analytics (millions): 2
- Maximum number of CIs and Relationships in CMDB (millions): 2

Hardware requirement:

- **1 master + 5 worker nodes**:
  8 CPU cores / 32 GB RAM / 200GB disk (at least 60GB thinpool size) for each node
- **1 separate NFS Server**:
  300 GB, RAID 10, IO: 280 MB/s
- **3 additional worker nodes if using built-in PostgreSQL for Service Management and CMDB**:
  8 CPU cores / 32 GB RAM / 200GB disk (at least 60 GB thinpool size) for each node

ITOM CDF installation configuration

ITOM Container Deployment Foundation (CDF) must be running on a Kubernetes (K8S) cluster that comprises a cluster of master and worker nodes. To correctly configure the K8S cluster, you must configure parameters in the **install.properties** file for the nodes. Parameters in this file are described in the following table. Review the parameters and decide on which parameters you need to configure to suit your business needs.

> ITOM CDF supports multiple master nodes; however, the current release of ITSMA supports only one master node. Ignore the parameters in the following table that are applicable only to a cluster configuration that uses multiple master nodes (see the **Notes** column).

Parameter	Description	Notes
MASTER_NODES	Lists the cluster master nodes (IPV4 address or FQDN), separated by a blank and enclosed in double-quotes. If you use more than one master node, you must work with high availability. **Example:** `MASTER_NODES="10.10.10.10 10.10.10.11 10.10.10.12"`	Mandatory
WORKER_NODES	Lists the cluster worker nodes, separated by a blank and enclosed in double-quotes. Suites are run on worker nodes. If you also want to use a master node as a worker node, enter its address IPV4 address or FQDN in *WORKER_NODES*. Typically, a worker node runs the workload when you deploy a suite. By default, when you install a suite, you target a worker node. **Example:** `WORKER_NODES="10.10.10.20 10.10.10.21 10.10.10.22"`	Mandatory

Parameter	Description	Notes
INGRESS_HOST	Defines the IP address (IPV4 address or FQDN) of the node on which you want to start the Ingress Controller. You must use one of the master or worker nodes. 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 of the network (for example, a Help Desk operative on a client machine on another network that needs to access Service Management), 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. **Example:** `INGRESS_HOST=10.10.10.10` (IPV4 address or FQDN of one of the master nodes)	Mandatory
EXTERNAL_ACCESS_HOST	Defines a fully-qualified hostname for external clients to access cluster services. The specified name must resolve the IP address where the ingress is running. 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. **Example:** `EXTERNAL_ACCESS_HOST=myd.XXXX.YYY.net`	Mandatory

Parameter	Description	Notes
NFS_SERVER	Specifies the IPV4 address or FQDN of the NFS server that serves the persistent volumes of the cluster services. 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 ITOM Container Deployment Foundation, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server. **Example**: `NFS_SERVER=16.255.25.255`	Mandatory
CLIENT_CA_FILE	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. 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 installation properties file. **Example:** `CLIENT_CA_FILE=/tmp/ca.crt`	Mandatory **only** for worker nodes
CLIENT_CERT_FILE	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. 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 installation properties file. **Example:** `CLIENT_CERT_FILE=/tmp/client.crt`	Mandatory **only** for worker nodes

Parameter	Description	Notes
CLIENT_KEY_FILE	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. 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 installation properties file. **Example:** `CLIENT_KEY_FILE=/tmp/client.key`	Mandatory **only** for worker nodes

HA_VIRTUAL_IP	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, you do not need to make any configuration. You only need to provide a virtual IP address and make sure the IP address must not be occupied before the installation. The requests to API server should be sent to a VIP for higher availability.
	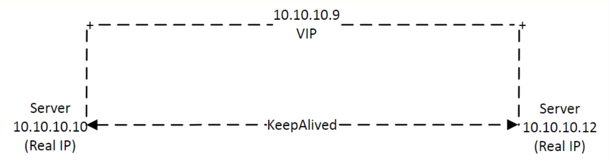 10.10.10.9 VIP
	Server 10.10.10.10 (Real IP)
	KeepAlived
	Server 10.10.10.12 (Real IP)
	Example:
	`MASTER_NODES="10.10.10.10 10.10.10.11 10.10.10.12"`
	`HA_VIRTUAL_IP=10.10.10.9`
		Mandatory **only** if you are using multiple master nodes

Parameter	Description	Notes
HA_NGINX_NODES	Specifies IPV4 address or FQDNof 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 addresses or FQDNs of the master nodes enclosed in double-quotes. **Example:** `HA_NGINX_NODES="10.10.10.10 10.10.10.11"`	Mandatory **only** if you are using multiple master nodes
PEER_CA_FILE	Specifies the CA certificate for TLS authentication. The value of the parameter is the file name of the CA certificate, including the absolute path. **Example:** `PEER_CA_FILE=/tmp/ca/crt`	Mandatory **only** if you are using multiple master nodes
PEER_CERT_FILE	Specifies the certificate for TLS authentication. The value of the parameter is the file name of the certificate, including the absolute path. **Example:** `PEER_CERT_FILE=/tmp/server.crt`	Mandatory **only** if you are using multiple master nodes
PEER_KEY_FILE	Specifies the private key for TLS authentication. The value of the parameter is the file name of the private key, including the absolute path. **Example:** `PEER_KEY_FILE=/tmp/server.key`	Mandatory **only** if you are using multiple master nodes

Parameter	Description	Notes
NFS_FOLDER	Specifies the root folder (fully-qualified directory) for the persistent volume that the NFS server exports. 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 ITOM Container Deployment Foundation, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server. **Example**: `NFS_FOLDER=/var/vols/itom`	Optional
ROOTCA	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. When you install the ITOM Container Deployment Foundation, all communication between the components is secured by using https. Therefore, communications use certificates to maintain security. These certificates can be self-signed or signed with a Certificate Authority provided by the customer. The default value is a self-signed certificate. **Example**: `ROOTCA=/tmp/ca.crt`	Optional

Parameter	Description	Notes
ROOTCAKEY	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. When you install the ITOM Container Deployment Foundation, all communication between the components is secured by using https. Therefore, communications use certificates to maintain security. These certificates can be self-signed or signed with a Certificate Authority provided by the customer. The default value is a self-signed certificate. **Example**: `ROOTCA=/tmp/ca.key`	Optional
NFS_STORAGE_SIZE	Specifies the size of the NFS volume exported by the NFS server. 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 ITOM Container Deployment Foundation, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server. **Example**: `NFS_STORAGE_SIZE=50Gi`	Optional
K8S_HOME	Specifies the installation directory (fully-qualified directory) for the ITOM Container Deployment Foundation binaries. **Example**: `K8S_HOME=/opt/kubernetes`	Optional

Parameter	Description	Notes
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=6443`	Optional
THINPOOL_DEVICE	Specifies the Docker devicemapper storage driver (format: path to a device). To configure the thinpool device, see https://docs.docker.com/engine/userguide/storagedriver/device-mapper-driver/#configure-direct-lvm-mode-for-production. If this parameter is specified, the installation uses the devicemapper(direct-lvm) Docker storage driver. If it is not specified, the installation uses devicemapper (loop). For a production environment, HPE recommends using devicemapper (direct-lvm). **Example**: `THINPOOL_DEVICE= /dev/mapper/docker-thinpool`	Optional

Parameter	Description	Notes
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 download the images from the Internet, for which you need to use a proxy. **Example**: `DOCKER_HTTP_PROXY="<Your Proxy>"`	Optional
DOCKER_HTTPS_PROXY	Specifies the proxy settings for Docker. Configure this parameter if access to Docker Hub or the 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="<Your Proxy>"`	Optional
REGISTRY_ORGNAME	Specifies the name of the organization where suite images are placed (format: a string). The default name is hpeswitomsandbox. **Example**: `REGISTRY_ORGNAME=hpeswitomsandbox`	Optional
FLANNEL_IFACE	Specifies the interface for docker inter-host communication to use (format: a single IPV4 address or interface name). **Example**: `FLANNEL_IFACE=10.10.10.10`	Optional

Prepare for the installation

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

- Obtain a Docker Hub account from HPE
- Meet the prerequisites
- (Optional) Prepare logical volumes for the cluster nodes
- Add a logical volume for direct-lvm thinpool (production only)
- Download and verify the ITOM CDF installation package

Obtain a Docker Hub account from HPE

You must use a Docker Hub account provided by HPE to download (pull) ITSMA suite images from Docker Hub. Based on your region and existing entitlements, contact your licensing team for a Docker ID:

- For Americas region, contact dockersupport.ams@hpe.com.
- For APJ region, contact dockersupport.apj@hpe.com.
- For EMEA region, contact dockersupport.emea@hpe.com.

Meet the prerequisites

Before you proceed to the ITOM Container Deployment Foundation (CDF) installation, make sure your environment meets the following prerequisites:

- You have prepared the needed number of machines and the appropriate physical volume, and the nodes and NFS server for the installation meet the minimum system requirements. For details, see Plan your suite deployment.
- There is sufficient disk space allocated for the ITOM CDF installation. See (Optional) Prepare logical volumes for the cluster nodes.
- You have set the performance rule in the physical server BIOS to "High Performance".
- The master node has a fully qualified domain name (FQDN) in lowercase.
- The NFS server, master node, and work nodes are installed in the same subnet, and have a static IP address.
- You know the **root** user password. This is because you will need to install the suite as **root** or a user with sudo access.
- The machine does not have Docker or Kubernetes installed (if it does, you need to remove them first).
- Your existing firewall is disabled on the host. If not, you can disable it by running the following commands:
  systemctl stop firewalld
  systemctl disable firewalld
- The following ports, which are needed during the installation, are not in use on the host VM: 111, 2049, 2380, 4001, 4194, 5000, 5443, 8080, 8200, 8443, 10250, 10251, 10252, 10255, 20048.

- Any existing shared NFS folder is removed if you installed ITOM CDF previously. The default folder is /var/vols/itom/core.
  For example: **rm -rf /var/vols/itom/core**
- Network Time Protocol (NTP) is configured on all cluster hosts by using Chrony. Chrony is installed by default on some versions of CentOS. However, if Chrony is not installed or running on your system, use the following commands to install and enable Chrony and verify it is running:
  # yum install chrony
  # systemctl start chronyd
  # systemctl enable chronyd
  # chronyc tracking

  > You can use other tools to synchronize system time, for example, ntp.

- The **/tmp** folder of the target system has enough free space (at least 2.5G), which is required when adding worker nodes from the Management Portal.
- The following packages are installed on the relevant hosts (the master and worker nodes and NFS server). You can run the following command to install the packages: **yum install [package name]**.

Package	Install on
device-mapper-libs	master, workers
java-1.8.0-openjdk	master only
libgcrypt	master, workers
libseccomp	master, workers
libtool-ltdl	master, workers
lsof	master, workers
net-tools	master, workers
nfs-utils	master, workers
rpcbind	master, workers, NFS server
systemd-libs (version >= 219)	master, workers
unzip	master, workers

- The IP address of the master node is added to the no_proxy lists for the master node and worker nodes.

(Optional) Prepare logical volumes for the cluster nodes

Follow the steps below on each cluster node to ensure that you have enough logical volumes for the ITOM Container Deployment Foundation (CDF) installation.

> This task is required for production environments only. You can choose any volume group name, logical volume name, and disk location address for your installation according to your system.

1. Prepare a physical disk for the ITOM CDF cluster nodes. The physical volume of your system must meet the system requirements (see Support matrix for the supported operating systems).
2. Create a volume group by running the following command:
   # vgcreate [volume group name] [logical volume name]
   For example:
   # vgcreate core-platform /dev/sdb
3. Create a logical volume for the ITOM CDF installation by running the following command:
   # lvcreate -l 100%FREE -n [logical volume name] [volume group name]
   For example, utilize 100% of the volume group:
   # lvcreate -l 100%FREE -n mylv core-platform
4. Activate the volume group by running the following command:
   # vgchange -ay [volume group name]
   For example:
   # vgchange -ay core-platform
5. Format the file system by running the following command:
   # mkfs.ext3 [logical volume path]
   For example:
   # mkfs.ext3 /dev/core-platform/mylv
6. Mount the volumes under the folder in which you will install ITOM CDF by running the following command:
   # mount [logical volume path] [platform installation folder]
   For example:
   # mount /dev/core-platform/mylv /opt/coreplatform

   > [platform installation folder] must be the same as the **K8S_HOME** parameter value in the install.properties file (default value: **/opt/coreplatform**). For more information, see Configure the install.properties file.

Add a logical volume for direct-lvm thinpool (production only)

In a production environment only, you need to prepare a data volume for direct-lvm thinpool on docker. This is a one-time effort. We recommend that you create a VM template so that you can always reuse the same configuration for future installations.

> You need to perform this step on all nodes in the cluster (that is, the master node and worker nodes). Additionally, after the configuration, be sure to configure the **THINPOOL_DEVICE** parameter in the install.properties file (see Configure the install.properties file).

This step includes the following tasks:

- Task 1: Add a volume
- Task 2: Prepare the volume
- Task 3: Specify the THINPOOL_DEVICE parameter in install.properties

Task 1: Add a volume

The steps below assume that you are using an ESX/ESXi server for VM management.

1. Type 'fdisk -l' to check the partitions on your disk.

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *        2048     1026047      512000   83  Linux
/dev/sda2         1026048   125829119    62401536   8e  Linux LVM
```

2. Add a volume to the host server and restart the server. Usually, THINPOOL for each standard ITSMA host server (master or worker) requires a disk space of 45 to 60 GB.
   The following screenshot shows an example using vSphere.

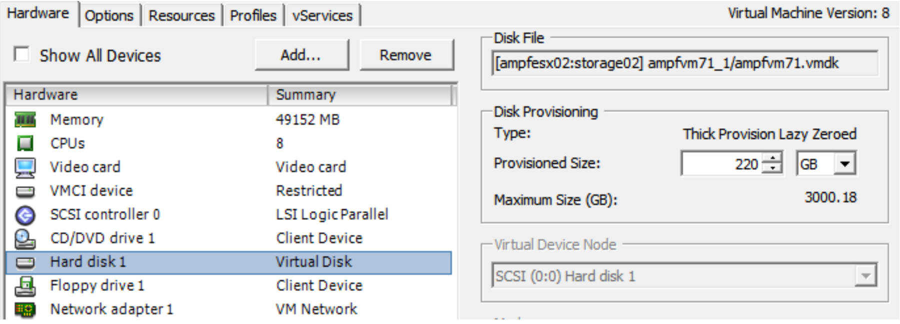

Hardware	Summary
Memory	49152 MB
CPUs	8
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
CD/DVD drive 1	Client Device
Hard disk 1	Virtual Disk
Floppy drive 1	Client Device
Network adapter 1	VM Network

Disk File: [ampfesx02:storage02] ampfvm71_1/ampfvm71.vmdk

Disk Provisioning
Type: Thick Provision Lazy Zeroed
Provisioned Size: 220 GB
Maximum Size (GB): 3000.18

Virtual Device Node: SCSI (0:0) Hard disk 1

3. Type 'fdisk -l' to check the partitions on your disk again. You should see more disk size is available.
4. Use fdisk to allocate the volume that you added:
   a. Type 'fdisk /dev/sda'.
   b. Type 'n' to create a new partition.
   c. Type 'p' to use the primary.
   d. Choose the newly added volume.
   e. Type 'w' to save and exit.
   f. Type 'fdisk -l' to check your partitions.

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *        2048     1026047      512000   83  Linux
/dev/sda2         1026048   125829119    62401536   8e  Linux LVM
/dev/sda3       251658240   461373439   104857600   83  Linux
```

5. Reboot the server to make the new volume accessible to the OS.
6. Make sure that there is no docker service started.

Task 2: Prepare the volume

Prepare the new volume by following the instructions below.

> For more information, visit the following links:
>
> https://docs.docker.com/engine/userguide/storagedriver/device-mapper-driver/#device-mapper-and-docker-performance
> https://github.com/docker/docker/issues/21701

```
# install lvm2
yum install lvm2
# ideally, perform these tasks before you start docker for the first time OR make sure docker is
stopped; all containers, images and data will be lost during this process
# This guide assume '/dev/sda3' is your new device
# create a physical volume (replace /dev/sda3 with your block device)
pvcreate /dev/sda3
# create a volume group named 'docker' (replace /dev/sda3 with your block device)
vgcreate docker /dev/sda3
# create a thin pool named 'thinpool'; in this example, the data LV is 95% of the 'docker' volume group
size (leaving free space allows for auto expanding of either the data or metadata if space is runs low
as a temporary stopgap)
lvcreate --wipesignatures y -n thinpool docker -l 95%VG
lvcreate --wipesignatures y -n thinpoolmeta docker -l 1%VG
# convert the pool to a thin-pool
lvconvert -y --zero n -c 512K --thinpool docker/thinpool --poolmetadata docker/thinpoolmeta
#  configure autoextension of thin pools via a lvm profile
vi /etc/lvm/profile/docker-thinpool.profile
# specify the value for 'thin_pool_autoextend_threshold' (where the number is the % of space used
before lvm attempts to autoextend the available space; 100 = disabled)
    thin_pool_autoextend_threshold = 80
# modify the autoextend percentage for when thin pool autoextension occurs (where the number is the %
of space to increase the thin pool; 100 = disabled)
    thin_pool_autoextend_percent = 20
# example /etc/lvm/profile/docker-thinpool.profile:
activation {
    thin_pool_autoextend_threshold=80
    thin_pool_autoextend_percent=20
}
# apply the lvm profile
lvchange --metadataprofile docker-thinpool docker/thinpool
# verified the lv is monitored
lvs -o+seg_monitor

# if docker was previously started, clear your graph driver directory
rm -rf /var/lib/docker/*

# make sure to monitor your thin pool and volume group free space!  it will auto-extend but the volume
group can still fill up
```

```
# monitor logical volumes
lvs
lvs -a (to see the data and metadata sizes)
# monitor volume group free space
vgs
# logs can show the auto-extension of the thin pool when it hits the threshold
journalctl -fu dm-event.service
```

Task 3: Specify the THINPOOL_DEVICE parameter in install.properties

For production environment, specify the THINPOOL_DEVICE parameter in the install. properties file before you set up a master node and worker nodes. For more information, see Configure the install.properties file.

THINPOOL_DEVICE=/dev/mapper/docker-thinpool

Download and verify the ITOM CDF installation package

Once your environment is ready (see Meet the prerequisites), you can download the ITOM Container Deployment Foundation (CDF) installation package to the master node and verify the package.

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

1. Download the installation package (CDF1704-15000.zip) to the master node from the HPE Software Entitlement Portal.
2. Unzip the file to a temporary directory. For example:
   unzip CDF1704-15000.zip -d ITOM

   > - The installation package is signed.
   > - An ITOM CDF license is included in each suite license. You only need an ITSMA suite license to install ITOM CDF and the ITSMA suite.

   The CDF1704-15000.zip file that you downloaded includes the following files:
   - HPESW_ITOM_Suite_Platform_YYYY.MM.nnnnn.zip
   - HPESW_ITOM_Suite_Platform_YYYY.MM.nnnnn.zip.HPc
   The **HPESW_ITOM_Suite_Platform_YYYY.MM.nnnnn.zip** file includes the files and directories listed in the following table.

Name	Description	Type
bin	The bin directory includes: • All the runtime files that are core of the container platform: docker runtime binaries (**docker**, **docker-containerd**, **docker-containerd-ctr**, **docker-container-shim**, **dockerd**, **docker-proxy**, **docker-runc**), the binary to access the distributed configuration database (etcdctl), the runtime to interact with Kubernetes(**kubectl**). • The scripts used to check the ITOM Container Deployment Platform**(kube-restart.sh**, **kube-start.sh**, **kube-stop.sh**). • The script to check that everything is running **(kube-status.sh**). • The script used during installation to create the configuration for Docker (**mk-docker-opts.sh**) and **vault** that is used for security purposes to store sensitive information and to generate and manage certificates for the ITOM Container Deployment Platform and the suite deployment.	Directory
cfg	The initial user and role information that will be seeded into IDM to create user accounts (single sign on).	Directory
images	All the core platform images and share services images	Directory
install	The binary that needs to be run to install ITOM Container Deployment Platform	File
install.properties	The properties file used to configure the installation.	File
jar		Directory
manifests	The manifests contain YAML files that describe how to deploy a container The image for Kubernetes.	Directory
objectdefs		Directory
rpm		Directory
scripts		Directory

Name	Description	Type
tools		Directory
uninstall.sh	Use to uninstall ITOM Container Deployment Platform	File
version.txt		File
zip		Directory

3. Verify and decrypt the installation package before installing:
   a. Download the HPE BinaryChecker from the same location as the ITOM CDF package, and then unzip the file. For example:
   unzip Release-BinaryChecker_12_7_2015.zip -d BinaryChecker
   b. Run the BinaryChecker under the proper directory to verify the signed ITOM CDF installation package: **hpbinarychecker --input <signedbinaryfile.HPc>**.
   Example:
   hpbinarychecker --input testbinary.HPc (or)
   hpbinarychecker --input testbinary.HPc –verbose

   > The binary checker commands are different for 32-bit and 64-bit systems. The HPESW_ITOM_Suite_Platform_YYYY.MM.nnnnn.zip and HPESW_ITOM_Suite_Platform_YYYY.MM.nnnnn.zip .HPc file must be placed under the same directory before running the command to verify and decrypt the signed ITOM CDF installation package.

   c. If the binary check fails, download the ITOM CDF package again, and then use the BinaryChecker to verify the newly downloaded package.
   Next, install ITOM CDF on one master node. For details, see Install ITOM CDF on the master node.

Install HPE ITOM CDF

The ITSMA suite must be deployed on ITOM Container Deployment Foundation (CDF), which provides a graphic user interface for suite administrators to deploy and administer suites. Before you can deploy the ITSMA suite, you must install ITOM CDF first.

Follow these steps to install ITOM CDF:

- Configure the install.properties file
- Install an NFS server
- Install ITOM CDF on the master node
- Install ITOM CDF on the worker nodes

If want to uninstall ITOM CDF after the installation, see Uninstall HPE ITOM CDF.

Configure the install.properties file

To correctly configure the Kubernetes cluster, you must configure the install.properties file of the master node. Once you have set up the properties file and installed ITOM Container Deployment Foundation (CDF) on the master node, you can reuse the file to install all worker nodes.

1. Make sure that you have already downloaded and unzipped the ITOM CDF installation package to a temporary directory on the master node. For details, see Download and verify the ITOM CDF installation package.
2. On the master node, go to the <ITOM CDF folder> directory, and then edit the `install.properties` file. At a minimum, you must specify the following parameters:
   * MASTER_NODES="<master node FQDN or IP address>"
   * WORKER_NODES="<worker node 1 FQDN or IP address> <worker node 2 FQDN or IP address> ...<worker node n FQDN or IP address>"
   * INGRESS_HOST=<master node IP address or FQDN>
   * EXTERNAL_ACCESS_HOST=<master node FQDN>
   * NFS_SERVER=<NFS Server IP address or FQDN>
   * NFS_FOLDER=/var/vols/itom/core
   * K8S_HOME=/opt/coreplatform
   * THINPOOL_DEVICE=/dev/mapper/docker-thinpool

Be sure to set the **EXTERNAL_ACCESS_HOST** parameter to the fully qualified domain name (FQDN) with only lowercase letters (do not use the IP address).
Be sure to set the **WORKER_NODES** parameter to the FQDNs or IP addresses of the worker nodes, separated by a space and enclosed in double-quotes. However, in a test environment, you can alternatively set the **WORKER_NODES** parameter to the master node FQDN or IP address to first install ITOM CDF on the master node only and then add worker nodes from the ITOM Platfrom UI (the "Management Portal"). See Install ITOM CDF on the worker nodes).
The default value of the **NFS_FOLDER** parameter is used here in the example above. You can use this default value or change it to another value; however, be aware that you must use this **NFS_FOLDER** parameter value when setting up the NFS server (see Install an NFS server).
The value of the **K8S_HOME** parameter (default value: **/opt/coreplatform**) must be the same as the **[Platform installation folder]** value that you used in (Optional) Prepare logical volumes for the cluster nodes.
The **THINPOOL_DEVICE** specifies the Docker devicemapper storage driver, which is required for a production environment. See Add a logical volume for direct-lvm thinpool (production only).

When you set FQDNs for the cluster nodes in the `install.properties` file, make sure that the FQDNs are resolved to correct IP addresses, not the loop back IP 127.0.0.1.

For a full description of the parameters in the `install.properties` file, see ITOM CDF installation configuration.

Install an NFS server

A Network File System (**NFS**) server allows you to mount your local file systems over a network so that remote hosts can interact with them as they are mounted locally on the same system. To install ITOM Container Deployment Foundation (CDF), you must install an NFS server first.

In a test environment, you can set up an NFS server on the master node; in a production environment, you are required to set up a dedicated NFS server.

To install an NFS server, follow these steps:

1. Log in to the NFS server host as root.
2. Make sure that the **rpcbind** package is installed on the host (see Meet the prerequisites). If the package is not already installed, run the following command to install it:
   yum install rpcbind
3. Install the NFS server by running the following command:
   yum install -y nfs-utils
4. Create an NFS folder for the ITOM CDF installation by running the following commands:

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

> The NFS directory that you create in this step must be the same as the **NFS_FOLDER** parameter value in the install.properties file (see Configure the install.properties file). Here, the default value **/var/vols/itom/core** is used as an example.

5. In the **/etc/exports** file, expose the NFS server that you created by adding the following line:
   ```
   <NFS folder> *(rw,sync,anonuid=1999,anongid=1999,all_squash)
   ```
   For example:
   ```
   /var/vols/itom/core *(rw,sync,anonuid=1999,anongid=1999,all_squash)
   ```
6. Run the following command:
   ```
   exportfs -ra
   ```
7. Run the following commands to enable the rpcbind and nfs services:
   ```
   systemctl restart rpcbind
   systemctl enable rpcbind
   systemctl restart nfs-server
   systemctl enable nfs-server
   ```

Now, go to Install ITOM CDF on the master node.

Install ITOM CDF on the master node

ITOM Container Deployment Foundation (CDF) must be deployed on a master node and three or more worker nodes. You can copy the ITOM CDF installation package to all nodes and run the installation script on each node in parallel; however, we recommend installing ITOM CDF on the master node first, and then adding the worker nodes from the ITOM CDF user interface (referred to as "the Management Portal "in the documentation).

To install ITOM CDF on the master node, follow these steps:

1. Make sure that you have already downloaded and unzipped the ITOM CDF installation package to a temporary directory on the master node. For details, see Download and verify the ITOM CDF installation package.
2. Make sure that you have configured the `install.properties` file (see Configure the install.properties file).
3. Go to the <ITOM CDF folder> directory, and then run either of the following commands:
   ./install (as the root user)
   sudo ./install (as a non-root user)
   Wait until the installation on the master node is complete. The following messages indicate a successful installation:
   ```
   Adding label role=loadbalancer for node xx.xxx.xxx.xxx ...
   Successfully added the node label.
   Successfully completed configuring the HPE ITOM Core Platform on this server!
   ```

> To see the installation log, run the following command:
> vi /tmp/install-<timestamp>.log

4. (Optional) Run the following commands to see what is installed:
   cd /opt/kubernetes
   ls -l
   For information about what is installed, see Additional information.
5. Run the following commands to make sure that all pods are running:
   kubectl get nodes
   kubectl get pods --all-namespaces
6. Launch the ITOM CDF user interface (referred to as "the Management Portal") from a browser:
   https://*<EXTERNAL_ACCESS_HOST FQDN>*:5443
   Since the current release of the ITSMA suite supports only one master node,
   the EXTERNAL_ACCESS_HOST parameter is set to master node. You must use the master node FQDN in this URL. Do not use the IP address.

   > ITOM CDF provides one user account with full administration privileges: **admin**/ **cloud**. You will be prompted to change this password at initial login.

7. Log in with **admin**/**cloud**. After login, you are required to change the password.
8. Follow the on-screen instructions to change the password.
   Now, you are ready to add three or more worker nodes from the Management Portal. See Install ITOM CDF on the worker nodes.
   If you want to uninstall ITOM CDF, see Uninstall HPE ITOM CDF.

Additional information

The ITOM CDF installer installs the following items:

- Base installation files
- Docker
- Certificates
- etcd
- flannel
- Internal network
- Vault
- Images
- Configuration for K8S
- Persistent volumes
- All the base ITOM CDF services (such as the postgresql for IdM, and the management portal)
- More SSL certificates for Nginx used for proxying requests to the ITOM CDF

The files and directories that were installed are described in the following table.

Name	Description	Type
bin	The bin directory includes: - All the runtime files that are core of the container platform: docker runtime binaries (**docker**, **docker-containerd**, **docker-containerd-ctr**, **docker-container-shim**, **dockerd**, **docker-proxy**, **docker-runc**), the binary to access the distributed configuration database (etcdctl), the runtime to interact with Kubernetes(**kubectl**). - The scripts used to check the ITOM Container Deployment Platform**(kube-restart.sh**, **kube-start.sh**, **kube-stop.sh**). - The script to check that everything is running **(kube-status.sh**). - The script used during installation to create the configuration for Docker (**mk-docker-opts.sh**) and **vault** that is used for security purposes to store sensitive information and to generate and manage certificates for the ITOM Container Deployment Platform and the suite deployment.	Directory
cfg	The cfg directory includes the Docker configuration. It includes docker and docker-bootstrap and idm. There are two Docker daemons running on each node. Only Docker is physically running on the host and everything else is containerized. So services or programs that you would typically run directly on the host, are now also run inside a container: docker-bootstrap instance. It runs etcd and flannel. - To see what is running inside docker, run this command: **docker ps**. Kubernetes is actually running inside Docker. - To see what is running inside the bootstrap docker, run this command: **docker - H unix:///var/run/docker-bootstrap.sock ps**. It runs flannel, vault, and etcd, which are containerized. Docker provides an abstraction layer from the host. - To see what is running inside the bootstrap-docker, which is a separate instance, you need to pass the socket of bootstrap-docker. Run this command: **ps -cf|grep dockerd**. K8S is actually running inside Docker. There are two K8S running: docker and bootstrap-docker that run on two different sockets.	Directory
data	Data that is generated by K8S and is the runtime data for K8S. To see what is in the data directory, run this command: **ls data/***.	Directory

Name	Description	Type
images	All the core platform images that have been imported locally. To see what is in the data directory, run this command: **ls images**.	Directory
log	The logs of some of the components that are currently running. - To see what is in the log directory, run this command: **ls log**. - To do a recursive log, run this command: **ls -R log**. All the components put their running information in the logs.	Directory
install-YYYY.MM.nnnnn.log		File
jar		Directory
manifests	The manifests contain YAML files that describe how to deploy a container. It contains YAML files that have to run on every node. They are K8S components: - **kube-apiserver.yaml**, which is for the K8S API server. - **kube-controller-manager.yaml**, which controls access to the K8S server. - **kube-proxy.yaml**, which contains proxy connections. - **kube-scheduler.yaml**, which schedules on what node to execute a container. - **kube-registry-proxy.yaml**, which starts the kube registry proxy container.	Directory
objectdefs	The objectdefs contains more YAML files for autopass, idm, persistent volumes, registry proxies, vault, management portal, Nginx controller, and the suite installer.	Directory
rpm	The rpm is an installable package used to enable the installation of an NFS server. The NFS utility helps sharing data via a networked volume.	Directory
runconf	The runconf is a transient directory used during the installation.	Directory
ssl	The ssl contains all the certificates and the keys that have been generated by the running ITOM Container Deployment Platform.	Directory
scripts		Directory

Name	Description	Type
uninstall.sh	The uninstall script To uninstall ITOM Container Deployment Platform run `./uninstall.sh`. The uninstall process stops containers and removes them, removes daemons, and more. You need to reboot the server afterwards.	File
tools	The support toolset for troubleshooting. For more information, see Support toolset.	Directory
version.txt		File
zip	The zip directory includes a subset of files used to install a new cluster node from the Management Portal (using the **Add Node** functionality).	Directory

Install ITOM CDF on the worker nodes

Once you have installed ITOM Container Deployment Foundation (CDF) on the master node, you are ready to install ITOM CDF on the worker nodes.

- In a production environment, you must run the installation script on each worker node to install ITOM CDF.
- In a test environment, you can directly add the worker nodes from the ITOM CDF user interface (the "Management Portal").

The reason why different worker node setup methods are required for production and test environments is that a production environment requires you to add a thinpool device for all nodes, and the thinpool device configuration does not work for the worker nodes if you add them from the ITOM CDF user interface.

Set up the worker nodes by running the installation script

To do this, follow these steps:

1. Copy the ITOM CDF installation package from the master node to a temporary directory on the worker nodes.
2. Unzip the zip file on all worker nodes.
3. Copy the `install.properties` file from the master node to all worker nodes to overwrite the default copy under the `<ITOM CDF folder>` directory.

4. On the master node, go to the `/opt/kubernetes/ssl` directory, and copy `ca.crt`, `client.crt`, and `client.key` to each worker node. You can copy the files to any directory on each worker node (for example: the `/tmp` directory).
5. On each worker node, open the install.properties file under the <ITOM CDF folder> 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
   ```
6. On each worker node, run the following command under the `<ITOM CDF folder>` directory:
   ```
   ./install
   ```

   > You can run the installation script on the worker nodes in parallel.

7. Log in to the ITOM CDF user interface ("Management Portal") as the **admin** user, and then click **ADMINISTRATION** > **Nodes** to verify that the worker nodes are successfully set up.

Add the worker nodes from the Management Portal

To add worker nodes from the Management Portal, follow these steps:

1. Log in to the Management Portal as the **admin** user.
2. Click **ADMINISTRATION** > **Nodes**.
3. In the Nodes area, click **+ ADD**.
4. Enter the following information for a worker node and then click ADD to remotely install the worker node:
   - the name of the node
   - the name and password of a user that can remotely execute commands on the host - typically the **root** user
5. Repeat the steps to add the rest of the worker nodes.
   Alternatively, enter the information for multiple nodes and add them simultaneously with the **+ ADD** button.
6. Go to the Nodes area, click **Refresh**, and check that the worker nodes were successfully added (which is indicated by a status with a tickle icon).
   You may need to wait for a while (for example, ten minutes) before you can see the newly added worker nodes.

Next, proceed to the suite installation (see Install the ITSMA suite).

> If you want to uninstall ITOM CDF, see Uninstall HPE ITOM CDF.

Uninstall HPE ITOM CDF

If you want to uninstall ITOM Container Deployment Foundation (CDF), uninstall ITOM CDF from each node in the cluster (master and workers) by running the **uninstall.sh** script, as described in the following.

(Optional) Back up the image tars

Before you uninstall ITOM CDF, you can back up image tars from the local private registry to a remote registry.

1. Go to the directory where the **local_backup.sh** file is located: **<installation folder>/scripts**.
2. Move jq file to /usr/local/bin/ using the following commands:
   chmod 777 jq
   mv jq /usr/local/bin
   If you have installed jq, skip this step.
3. Run the following command: **chmod 777 local_backup.sh**.
   Ensure that the script file format is Unix (not dos).
4. Run: ./**local_backup.sh <registryHost>**
   For example:
   ./**local_backup.sh 16.255.255.255:5000**
   The tar files are saved in image_tars/xxx.tar.

Uninstall the ITOM CDF

1. Run **uninstall.sh**. The uninstall process stops containers and removes them, removes daemons, and more.
2. Reboot the server.

Install the ITSMA suite

Once ITOM Container Deployment Foundation (CDF) is installed and an NFS server is set up, you are ready to perform the following steps to install the ITSMA suite.

- Install an ITSMA suite license (in a test environment, you can skip this step to use a trial license; in a production environment, you can also skip this step at this point and then replace the trial license with a perpetual one after the installation)
- Prepare your databases (in a test evironment, you can skip this step to use the built-in PostgreSQL database)
- Import ITSMA suite images from Docker Hub to the local registry
- Configure NFS sharing for ITSMA
- Run the Suite Installer

If want to uninstall the ITSMA suite after the installation, see Uninstall the ITSMA suite.

Install an ITSMA suite license

A license for ITOM Container Deployment Foundation (CDF) is included in each suite license. Each ITSMA suite license includes two sub-licenses: one of which contains a license key for Service Portal and the other contains license keys for the rest of the ITSMA components.

- If you skip this step, when you run the Suite Installer to install ITSMA, a 30-day trial license will be used. You can use a trial license in a test environment; in a production environment, you still have the option to skip license installation at this point and install a perpetual license after the installation (see Replace an ITSMA trial license).
- You can activate your ITSMA license by directly uploading the license files or by entering an activation code.

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

- Step 1: Obtain the locker code from the Management Portal
- Step 2: Obtain an ITSMA license or activation code
- Step 3: Install the ITSMA license

Step 1: Obtain the locker code from the Management Portal

For each ITOM CDF deployment, a unique locker code is generated. You must use this locker code to redeem a suite license.

To obtain the locker code for your ITOM CDF deployment, follow these steps:

1. Log in to Management Portal as the **admin** user (using the new password that you specified after you logged in with the initial password "cloud"):
   https://*<master node FQDN>*:5443
2. Click **ADMINISTRATION** > **License**.

3. Click **Install Licenses**. The screen displays an auto-generated Lock Code.
4. Make a note of the locker code.

Step 2: Obtain an ITSMA license or activation code

Once you have got your locker code, you can use it to obtain an ITSMA suite license or activation code. You can install the license by using either the license files or the activation code.

1. Go to the HPE Software Entitlement Portal.
2. Sign in by using your HPE passport.
3. Locate and click the license link for ITSMA Express.
4. In the **Locking Information** field, enter the locker code that you previously obtained from the Management Portal.
5. Provide information for the other fields as necessary.
6. Click **Next**. The Activation Result page is displayed.
7. Download the license files to your local drive or save the activation code generated for your license in a safe place.

Step 3: Install the ITSMA license

To install the license, you can either upload the two license files or enter the Activation Code generated for your license.

To install the license, follow these steps:

1. Return to the Management Portal.
2. Click **ADMINISTRATION** > **License**.
3. Click **Install Licenses**.
4. Use one of the following methods:
   Method 1:
   a. Use the **Choose File** and **Add More Files buttons** to upload the license files that your downloaded, and then click **Next**.
      The license details are displayed.
   b. Select all licenses that are listed, and then click **Install Licenses**.

   Method 2:
   a. Enter your activation code, and then click **Next**.
   b. Select an environment type.
   c. In the **Quantity to Activate** field, enter a desired number.
   d. Click **Next**.
5. When the installation is complete, click **View Licenses** to view the installed licenses.
6. Click **LICENSE REPORT** to view information about the suite components that consume the licenses.

Prepare your databases

In ITSMA, the Service Management, CMDB, and Service Portal capabilities require a database. You are recommended to configure external databases in a production environment:

- Prepare an external PostgreSQL database for Service Portal.
- Prepare an external Oracle database for Service Management and CMDB.

> The built-in PostgreSQL database in ITSMA is recommended for test environments. However, if you want to use it in a production environment, see Sizing recommendations, Apply PostgreSQL parameter updates in ITSMA, and Bind the internal PostgreSQL database to a dedicated worker node for more information about how to use it.

Prepare an Oracle database for Service Management and CMDB

Before you proceed to the installation, prepare an Oracle database. Later, you will be asked to enter the database connection settings when running the Suite Installer.

To prepare your RDBMS, follow these steps:

1. Create an Oracle database.
   For the supported Oracle database version, see Support matrix. For details about how to create an Oracle database, see the Oracle documentation.
2. Create a login ID and password for Service Management to connect to your Oracle server.
   When you log on to Service Management, it creates a table in the default table space defined for that login ID. The login ID must have the following privileges:
   - Connect
   - Create, Alter, Drop a table
   - Create, Alter, Drop an index
   - Select on v_$parameter
   - Alter Session Privileges

   You can provide these privileges to an Oracle user by using the following oracle statements:
   create user <smadmin> identified by <smadmin> default tablespace <users> quota unlimited on <users>;
   grant connect, resource, select on v_$parameter to <smadmin>;
3. Create a login ID and password for CMDB to connect to your Oracle server.
   The database administrator should create an Oracle schema user with the following database permissions required by CMDB:
   Roles: Connect
   Privileges:
   - CREATE TABLE
   - CREATE VIEW
   - CREATE SEQUENCE
   - CREATE TRIGGER
   - CREATE PROCEDURE
   - UNLIMITED TABLESPACE
   - ALTER USER ${user} DEFAULT ROLE ALL
   - CREATE TYPE
   - EXECUTE ON DBMS_LOB

- EXECUTE ON DBMS_STATS

> The last two permissions (EXECUTE ON DBMS_LOB and EXECUTE ON DBMS_STATS) are granted by default.

Prepare a PostgreSQL database for Service Portal

> - The ITSMA sutie does not support an external PostgreSQL database that has SSL enabled.
> - If your ITSMA suite installation fails for some reasons, before you can reinstall the suite using an external PostgreSQL database, you must first remove the mounted PostgreSQL data in the **/var/vols/itom/itsma/itsma-<namespace>/db/propel** directory on the master node. If you fail to do this, data will not be created in the external PostgreSQL database and therefore the suite installation will fail.

To do this, follow these steps:

1. Create a PostgreSQL database.
   For the supported PostgreSQL database version, see Support matrix. For details about how to create a PostgreSQL database, see the PostgreSQL documentation.
2. Determine a login username and password for Service Portal to connect to the database.
   You will need to configure this login account when configuring the connection settings of this external PostgreSQL database when installing ITSMA (see Run the Suite Installer).
   You have two options:
   - Use the "postgres" user (recommended).
   - Use another user. If you choose this option, do the following:
     a. Create a login user that has the same name as the database and has the create database/extension/user/schema privileges.
     b. Set the database owner to this user.

3. Configure the pg_hba.conf file to allow the following users to connect to their corresponding database:

User	Database
analytics	analytics
bpmdb	bpmuser
catalog	catalog
dashboarddb	dashboard

notificationdb	notification
jumpstart	jumpstart
sxdb	sxuser

For example, add the following lines to the pg_hba.conf file:
host analytics analytics 0.0.0.0/0 md5
host bpmdb bpmuser 0.0.0.0/0 md5
host catalog catalog 0.0.0.0/0 md5
host dashboarddb dashboard 0.0.0.0/0 md5
host notificationdb notification 0.0.0.0/0 md5
host jumpstart jumpstart 0.0.0.0/0 md5
host sxdb sxuser 0.0.0.0/0 md5

4. Restart the PostgreSQL database.

Import ITSMA suite images from Docker Hub to the local registry

Before you can install the ITSMA suite from the ITOM Container Deployment Foundation (CDF) user interface, you must first download the ITSMA suite images from Docker Hub and then import the images to the local registry of the master node.

- Before you proceed, make sure that you have obtained a Docker Hub account from HPE, which is required to pull the suite images. For more information, see Obtain a Docker Hub account from HPE.
- In the following steps, **$K8S_HOME** represents the installation root directory that is configured in the **K8S_HOME** parameter in the install.properties file. See Configure the install.properties file.

To do this, perform the following tasks:

- Task 1: Download the images from Docker Hub
- Task 2: Import the images to the local registry
- Task 3: Verify the images in ITOM CDF

Task 1: Download the images from Docker Hub

The steps vary depending on whether your master node has access to Docker Hub.

If the master node can access Docker Hub:

1. On the master node, run the following command to make sure you can pull images from Docker Hub:

```
docker pull hello-world
```

2. If you can pull the hello-world image, perform the following steps on the master node:
   a. Ensure that the master node is connected to the Internet.

b. Go to the directory in which the **downloadimages.sh** is located: **$K8S_HOME/scripts**.
c. Run the following command: **./downloadimages.sh -o hpeswitom -s itsma -u <username> -p <password>**.

> Use the Docker Hub credentials (username and password) that you obtained from HPE (for details, see Install). You can also run the script without any parameters and enter required parameter values when prompted. To see help information about the parameters, run the **./downloadimages.sh --help** command.

d. Select the ITSMA suite version by entering **2017.04** or **latest**.
The script starts the downloading process. When the script has finished execution, the following message is displayed: **Successfully downloaded the ITSMA suite version: 2017.04 ...**
You should be able to see the tar files of the images in the suite images tar directory (default: **/var/opt/kubernetes/offline/suite_images**).

3. If you cannot pull the hello-world image, perform the following steps on the master node:
   a. Run the following command to enable the service:
   systemctl enable docker.service
   b. Run the following command to create a directory:
   mkdir -p /usr/lib/systemd/system/docker.service.d
   c. Configure a proxy:
   ```
   cat << EOF > /usr/lib/systemd/system/docker.service.d/http_proxy.conf
   [Service]
   Environment="HTTP_PROXY=<Your Proxy>" "HTTPS_PROXY=<Your Proxy>"
   EOF
   ```
   d. If Docker-Content-Trust is turned on, configure your proxy in the **http_proxy.conf** file:
   ```
   export http_proxy= <Your Proxy>
   export https_proxy=<Your Proxy>
   ```
   e. Run the following command to reload the configuration:
   systemctl daemon-reload
   f. Run the following command to restart Docker:
   service docker restart
   g. Make sure that you can pull images from Docker Hub by running the following command:
   docker pull hello-world
   h. Run the following commands to execute the **downloadimages.sh** script:
   ```
   cd $K8S_HOME/scripts
   ./downloadimages.sh -o hpeswitom -s itsma -u <username> -p <password>
   ```
   Where: <username> and <password> are the Doc Hub credentials that you obtained from HPE. For details, see Install.
   i. Select the ITSMA suite version by entering **2017.04** or **latest**.
   The script starts the downloading process. When the script has finished execution, the following message is displayed:
   Successfully downloaded the ITSMA suite version: 2017.04 ...

You should be able to see the tar files of the images in the suite images tar directory (default: **/var/opt/kubernetes/offline/suite_images**).

If the master node cannot access Docker Hub:

If your master node cannot access Docker Hub, you can first download the images to a machine that can access Docker Hub, and then copy the images to the master node.

> Make sure that you run the downloadimages.sh and up loadimages.sh scripts on the same operating system. You may fail to import some images to the local registry if you run the scripts on different operating systems.

1. Find another machine that can access Docker Hub, and perform the following steps on this machine:
    a. Check that the machine has at least 100GB free disk space.
    b. Run the following command to make sure that your operating system is 64-bit and the Linux kernel version is 3.10 or higher:
    ```
    uname –r
    ```
2. Install Docker on this machine by performing the following steps (for more information, refer to the Docker installation documentation):
    a. Configure a yum proxy in the yum.conf file:
    Run the following command:
    ```
    vi /etc/yum.conf
    ```
    Add the following line to this file:
    ```
    proxy=<Your Proxy>
    ```
    b. Run the following command to list the package version in the system:
    ```
    yum list
    ```
    c. Add a yum repository:
    ```
    cat << EOF > /etc/yum.repos.d/docker.repo
    [dockerrepo]
    name=Docker Repository
    baseurl=https://yum.dockerproject.org/repo/main/centos/7/
    enabled=1
    gpgcheck=1
    gpgkey=https://yum.dockerproject.org/gpg
    EOF
    ```
    d. Run the following command to update the source information:
    ```
    yum update --skip-broken -y
    ```
    e. Run the following command to install Docker:
    ```
    yum install -y docker-engine
    ```
    f. Run the following command to enable the docker service:
    ```
    systemctl enable docker.service
    ```
    g. Configure a proxy so that you can download the suite images:
    ```
    cat << EOF > /usr/lib/systemd/system/docker.service.d/http_proxy.conf
    [Service]
    Environment="HTTP_PROXY=<Your Proxy>" "HTTPS_PROXY=<Your Proxy>"
    EOF
    ```

h. If you have Docker-Content-Trust turned on, run the following commands:
export http_proxy= <*Your Proxy*>
export https_proxy=<*Your Proxy*>

i. Run the following command to reload the configuration:
systemctl daemon-reload

j. Run the following command to restart Docker:
service docker restart

> This process may take several minutes.

3. Download the suite images:
   a. Copy the downloadimages.sh script (located in the $K8S_HOME/scripts directory), and the jq file (located in the $K8S_HOME/bin directory) from your master node to the current machine (which can pull images from Docker Hub).
   b. Move the jq file to **/usr/local/bin/** using the following commands:

   > If you have already jq installed, skip this step.

   chmod 777 jq
   mv jq /usr/local/bin

   c. Run the following commands to execute the **downloadimages.sh** script:
   cd $K8S_HOME/scripts
   ./downloadimages.sh -o hpeswitom -s itsma -u <username> -p <password>
   Where: <username> and <password> are the Doc Hub credentials that you obtained from HPE.
   d. Select the ITSMA suite version by entering **2017.04** or **latest**.
   The script starts the downloading process. When the script has finished execution, the following message is displayed:
   Successfully downloaded the ITSMA suite version: 2017.04 ...
   You should be able to see the tar files of the images in the suite images tar directory (default: **/var/opt/kubernetes/offline/suite_images**).
4. Copy the files downloaded in the previous step to the following directory on the master node: **/var/opt/kubernetes/offline/suite_images**.

Task 2: Import the images to the local registry

1. Log on to the master node as the root user.
2. Go to the directory in which the **uploadimages.sh** file is located: **$K8S_HOME/Scripts**.
3. Execute the **uploadimages.sh** script:
   ./uploadimages.sh -s itsma -d /var/opt/kubernetes/offline/suite_images

> You can also run the script without any parameters and enter required parameter values when prompted. To see help information about the parameters, run the **./uploadimages.sh --help** command.

The script starts the uploading process. Wait until the script has finished uploading the images.

Task 3: Verify the images in ITOM CDF

1. Log in to ITOM CDF as the **admin** user.
2. Click **ADMINISTRATION** > **Local Registry**.

Now, you are ready to run the Suite Installer to install the ITSMA suite. See Run the Suite Installer.

Configure NFS sharing for ITSMA

When installing ITOM Container Deployment Foundation (CDF), you set up an NFS server (see Install an NFS server). Before you launch the Suite Installer to install ITSMA, you need to configure a directory on the same NFS server to store ITSMA suite data and share this directory using NFS.
To configure NFS sharing for ITSMA, follow these steps:

1. Log on to the NFS server.
2. Create a directory to store suite information:
   mkdir -p <ITSMA NFS shared folder>

   > This directory will be shared using NFS, to make it available to the suite installation. Later, you will need to specify this directory when running the Suite Installer. See Run the Suite Installer.

   For example:
   mkdir -p /var/vols/itom/itsma/itsma-itsma1

   > You are recommended to use this structure for the directory: **/var/vols/itom/itsma/itsma-itsma<n>** (n=1, 2...) so that you can easily identify the ITSMA NFS directory path.

3. Configure the NFS sharing folder in the /etc/exports file by adding the following line:
   <ITSMA NFS shared folder> *(rw,sync,anonuid=1999,anongid=1999,all_squash)
   For example, add the following line:
   /var/vols/itom/itsma/itsma-itsma1 *(rw,sync,anonuid=1999,anongid=1999,all_squash)
4. Grant the "itsma" user the right permissions:
   sudo exportfs -ra
   groupadd -g 1999 itsma
   useradd -g 1999 -u 1999 itsma
   chown -R itsma:itsma /var/vols/itom/itsma/<ITSMA NFS shared folder>
   For example:

```
sudo exportfs -ra
groupadd -g 1999 itsma
useradd -g 1999 -u 1999 itsma
chown -R itsma:itsma /var/vols/itom/itsma/itsma-itsma1
```

Run the Suite Installer

Once the ITOM Container Deployment Foundation (CDF) UI (referred to as "the Management Portal") is available, you are ready to install the ITSMA suite.

- The ITSMA suite is bundled with an OpenLDAP server and a PostgreSQL database, which you can use in a test environment. In a production environment, you need to configure an external LDAP server (after suite installation), and two external databases (one Oracle database for Service Management and CMDB, and one PostgreSQL database for Service Portal).
- Before you proceed, make sure that the ports to be used by the suite are not in use on the master node. See ITSMA node ports.
- During the suite installation, do not click any browser buttons (such as **Back** or **Refresh**) or any other menu options on the left-side navigation pane; otherwise you will be forced to exit the installation wizard and have no way to return to the wizard.
- You can only install one instance of ITSMA in the Management Portal. During installation, the Suite Installer automatically assigns a namespace for ITSMA. The namespace is "itsma1" for the first-time installation. If you uninstall the first instance and then reinstall ITSMA, the namespace becomes "itsma2", and so on.

The installation wizard will walk you through a process that includes the following major steps:

- Step 1: Select your suite size (Extra Small, Small, or Medium)
- Step 2: Configure external databases (one Oracle and one Postgres)
- Step 3: Specify an initial password for the **sysadmin** user (the suite administrator)
- Step 4: Configure an LDAP server
- Step 5: Install ITSMA

To install the ITSMA suite, follow these steps:

1. Log in to the ITOM CDF UI ("Management Portal") as **admin**:
   https://*<master node FQDN>*:5443
   The initial password for **admin** is **cloud**.

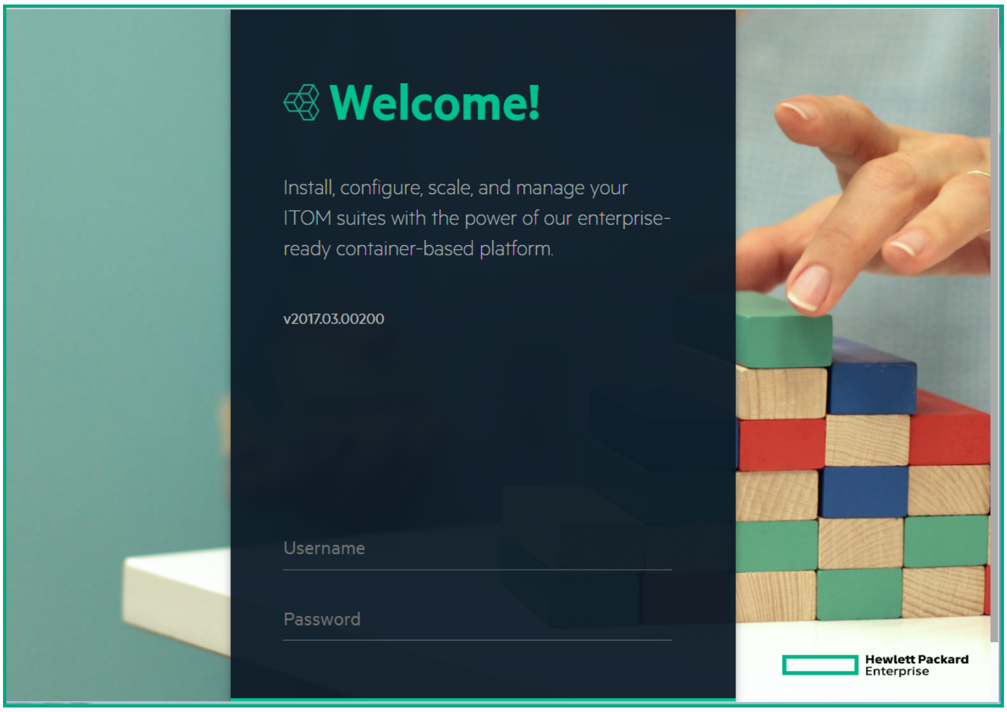

Welcome!

Install, configure, scale, and manage your ITOM suites with the power of our enterprise-ready container-based platform.

v2017.03.00200

Username

Password

2. Start the Suite Installer.

> ITOM CDF supports only one instance of ITSMA. If there is already an instance of ITSMA installed, the **Installation** menu option is not displayed.

   a. On the left-side navigator, expand the **SUITE** node.
   b. Click **Installation**.

3. Review and accept the end user license agreement and HPE privacy policy, and then click **Next**.
4. Select the ITSMA suite.
   a. On the suite selection page, click **ITSMA**, and then click **Next**.
   b. On the summary page that lists the capabilities that will be installed, click **Next**.

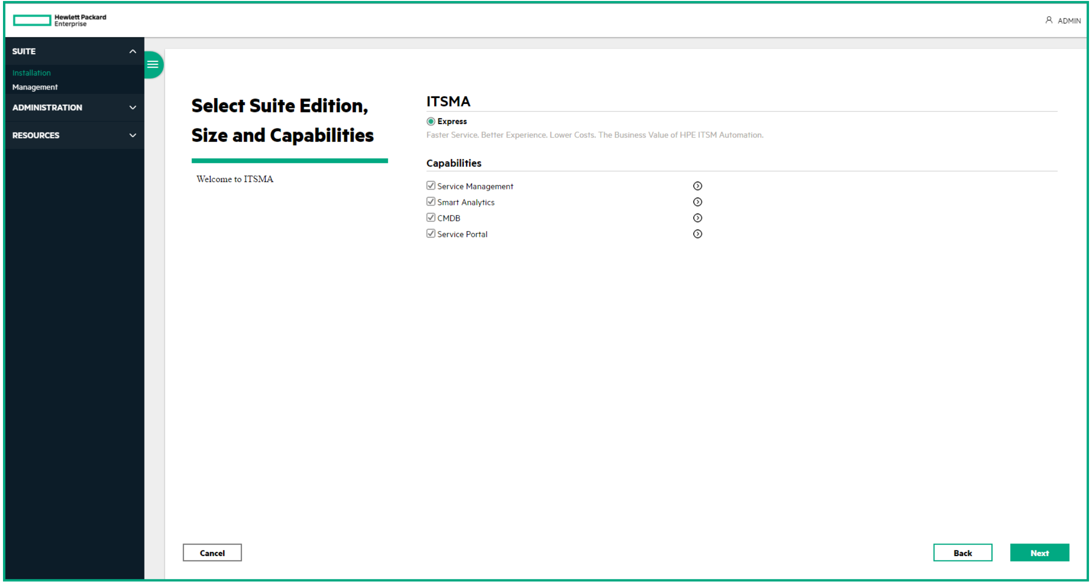

Select Suite Edition, Size and Capabilities

Welcome to ITSMA

ITSMA

● Express

Faster Service. Better Experience. Lower Costs. The Business Value of HPE ITSM Automation.

Capabilities

☑ Service Management
☑ Smart Analytics
☑ CMDB
☑ Service Portal

5. On the **Configure the Suite Storage** page, provide ITSMA NFS sharing configuration information:

   > You need to provide information according to the ITSMA NFS sharing configuration that you completed previously (see Configure NFS sharing for ITSMA). Do not include a trailing slash when entering the ITSMA NFS shared directory (for example: **/var/vols/itom/itsma/itsma-itsma1**); otherwise an error occurs.

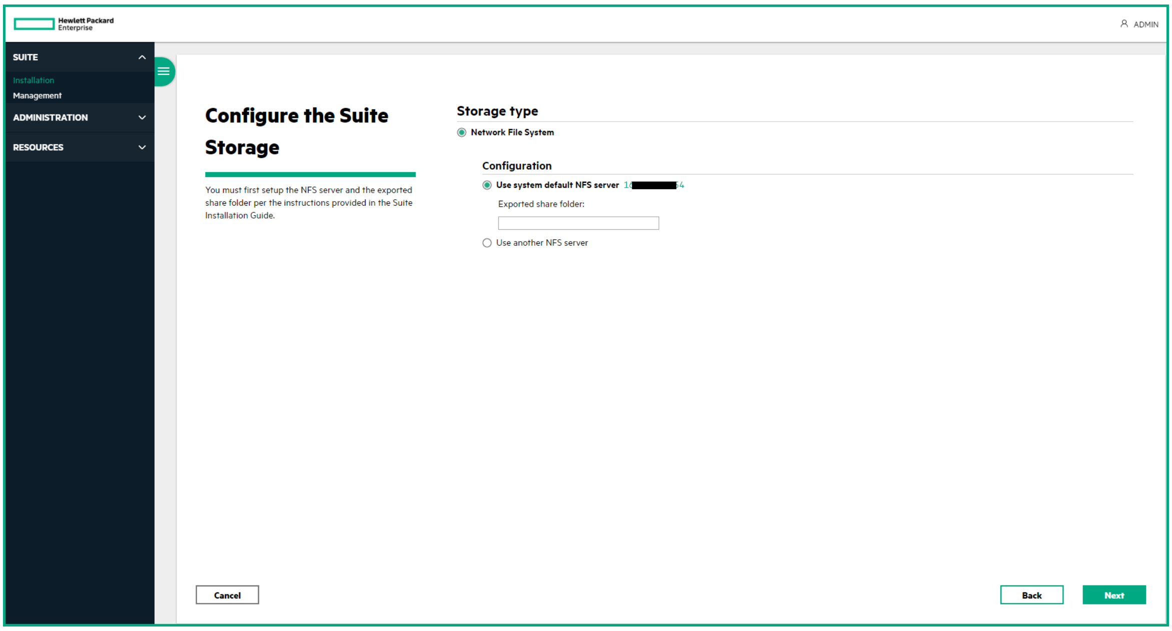

Configure the Suite Storage

You must first setup the NFS server and the exported share folder per the instructions provided in the Suite Installation Guide.

Storage type

● Network File System

Configuration

● Use system default NFS server [redacted]

Exported share folder:

○ Use another NFS server

   - If you used the master node as the NFS server, select the **Use system default NFS server** option, and then enter the ITSMA NFS shared folder created previously.
   - If you used a dedicated NFS server, select the **Use another NFS server** option, and then enter the NFS server host name or IP address and the ITSMA NFS shared folder created previously.

6. The Suite Installer starts loading the suite images, which you have imported to the local registry (see Import ITSMA suite images from Docker Hub to the local registry). This process may take a while. Wait for the image loading to complete.
7. On the **Select Suite Size** screen, select a deployment size depending on your number of concurrent users and hardware configuration: **Extra Small**, **Small**, or **Medium**. For more information, see Support matrix.
8. On the database configuration page, switch on the **External** button to configure your Oracle and Postgres database connection settings.

   > In a test environment, you can skip this step to use a built-in Postgres database by not switching on the **External** button. If you choose to use the built-in database in a production environment, see Sizing recommendations, Apply PostgreSQL parameter updates in ITSMA, and Bind the internal PostgreSQL database to a dedicated worker node for information about additional configurations required.

   Service Management and CMDB:

Setting	Description
Oracle Server Host Name or IP Address	The fully-qualified domain name or IP address of the external Oracle database server
Oracle Server Port	The communications port of the external Oracle database server
SID/ Service Name	Oracle system identifier (SID) or Service Name of the Oracle database: • SID: a name that identifies a specific instance of a database. For any database, there is at least one instance referencing the database. • Service Name: A logical representation of a database, which is the way a database is presented to clients. A database can be presented as multiple services and a service can be implemented as multiple database instances. The service name is a string that is the global database name, that is, a name comprised of the database name and domain name, entered during installation or database creation. If you are not sure what the global database name is, then you can obtain it from the value of the SERVICE_NAMES parameter in the initialization parameter file. For CMDB, only Service Name is supported.
User Name and **Password** (for Service Management)	The user name and password that Service Management uses to connect to the external Oracle database This user must have the privileges required by Service Management, as described in Prepare your databases.
User Name and **Password** (for CMDB Administrator)	The user name and password that CMDB Administrator uses to connect to the external Oracle database This user must have the privileges required by CMDB, as described in Prepare your databases.

Service Portal:

Setting	Description
PostgreSQL Server Host Name or IP Address	The fully-qualified domain name or IP address of the external PostgreSQL database server
PostgreSQL Server Port	The communications port of the external PostgreSQL database server
User Name and **Password**	The user name and password that Service Portal uses to connect to the external PostgreSQL database server This user must have the privileges required by Service Portal, as described in Prepare your databases.

When the configuration is complete, click **Test Connection** to make sure you can successfully connect to the databases and then click **Next**.

9. Configure an initial password for the suite administrator, and then click **Next**.
   The ITSMA suite provides a seeded user named **sysadmin**, which has full administrator privileges for ITSMA. This user account is stored in the database for the internal HPE Identity Manager (IdM) server. After the suite installation is compete, you can use the **sysadmin** user name and the specified initial password to log in to ITSMA (see Log in to ITSMA). You can also change the initial password after the initial login (see Change the suite administrator password).
10. Skip the LDAP configuration to use the internal LDAP server, and then click **Next**. A warning is displayed, asking you to confirm if you want to start the installation process.

> During the suite installation, do not switch on the **External** button to set up an external LDAP server.
> To configure an external LDAP server, first use the default internal LDAP during the suite installation, and then switch to external LDAP after you install the suite. In a production environment, you must configure an external LDAP server, because the internal one is for demonstration purposes only. For instructions, see Configure an external LDAP server.

11. Click **Next**. A warning is displayed, asking you to confirm if you want to start the installation process.
12. Click **YES** to start the installation. The installation progress is displayed on the screen.

> Once the installation process has started, even if you accidentally exit the wizard, for example, because you accidentally refreshes the browser, the installation process still continues running in the backend.

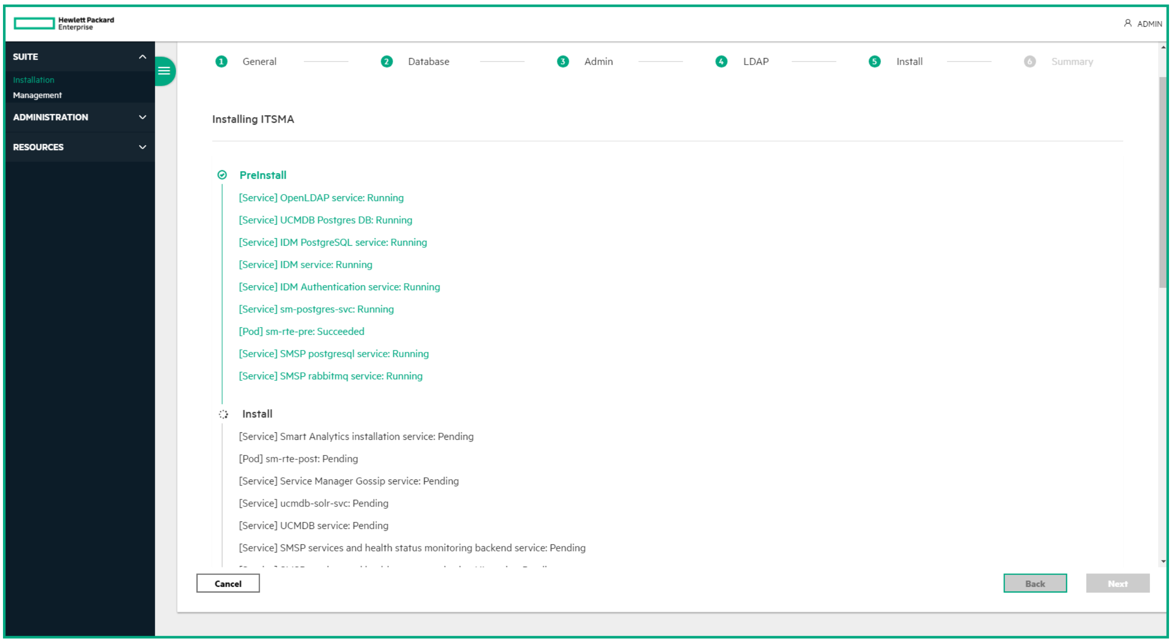

Installing ITSMA

PreInstall

[Service] OpenLDAP service: Running

[Service] UCMDB Postgres DB: Running

[Service] IDM PostgreSQL service: Running

[Service] IDM service: Running

[Service] IDM Authentication service: Running

[Service] sm-postgres-svc: Running

[Pod] sm-rte-pre: Succeeded

[Service] SMSP postgresql service: Running

[Service] SMSP rabbitmq service: Running

Install

[Service] Smart Analytics installation service: Pending

[Pod] sm-rte-post: Pending

[Service] Service Manager Gossip service: Pending

[Service] ucmdb-solr-svc: Pending

[Service] UCMDB service: Pending

[Service] SMSP services and health status monitoring backend service: Pending

13. Wait until the following Congratulations page is displayed. The ITSMA suite has been successfully installed.

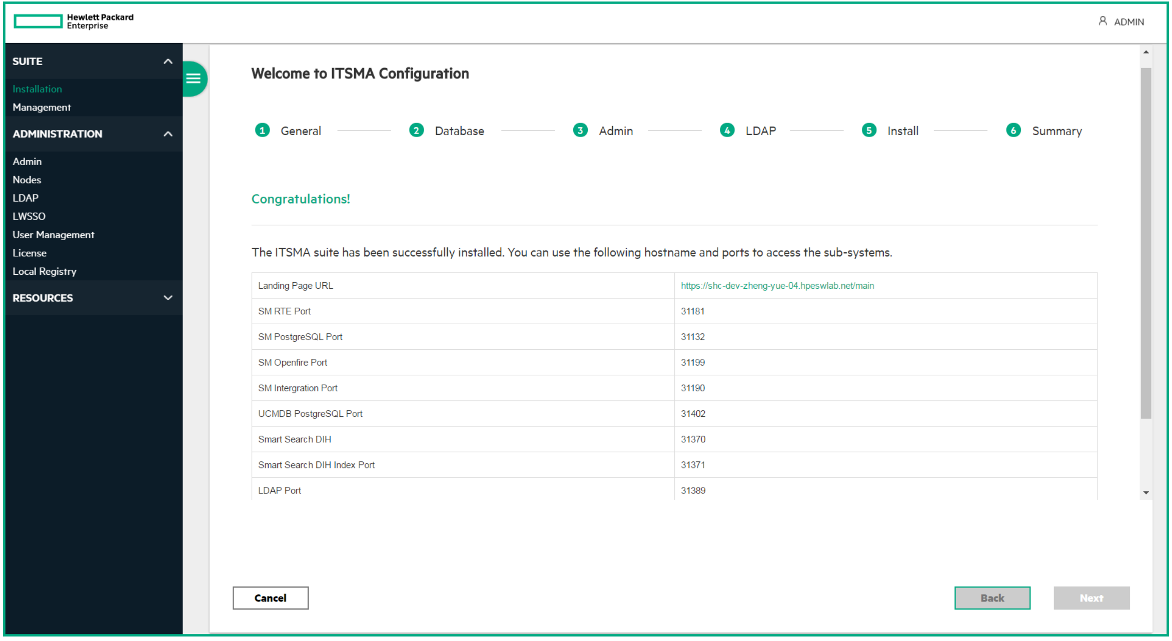

Welcome to ITSMA Configuration

Congratulations!

The ITSMA suite has been successfully installed. You can use the following hostname and ports to access the sub-systems.

Landing Page URL	https://shc-dev-zheng-yue-04.hpeswlab.net/main
SM RTE Port	31181
SM PostgreSQL Port	31132
SM Openfire Port	31199
SM Intergration Port	31190
UCMDB PostgreSQL Port	31402
Smart Search DIH	31370
Smart Search DIH Index Port	31371
LDAP Port	31389

14. Verify the installation on the master node.
    a. Log in to the master node as the root user.
    b. Run the following command:

       kubectl get ns

       Find the namespace for ITSMA from the list. For the first installation of ITSMA, ITOM CDF assigns a namespace of "itsma1". If you uninstall ITSMA and then reinstall ITSMA, ITOM CDF assigns "itsma2" as the new namespace, and so on.
    c. Run the following command to make sure that all pods are in Running state.

       kubectl get pods --namespace <namespace>

       For example: **kubectl get pods --namespace itsma1**

> If any pods are not running, see Troubleshoot the ITSMA suite.

15. Launch the ITSMA Suite Portal using one of the following URLs:
    https://<master node FQDN>/main (for the suite administrator "sysadmin")
    https://<master node FQDN>/ess (for end users)
    A login page is displayed.
16. Make sure that you can log in to ITSMA. For details, see Log in to ITSMA.

> After the installation, the ITSMA suite needs a one-time data preparation that may take half an hour to one hour. During this data preparation period, some capabilities such as CMDB and Service Portal are not accessible.

Next, you can:

- Log in to ITSMA as **sysadmin** to configure additional settings of the suite. For more information, see Log in to ITSMA and ITSMA suite administration.
- Perform post-installation configurations. See Set up ITSMA.
- Uninstall the suite if needed. See Uninstall the ITSMA suite.

1. Skip the LDAP configuration to use the internal LDAP server, and then click **Next**. A warning is displayed, asking you to confirm if you want to start the installation process.

Uninstall the ITSMA suite

You can uninstall ITSMA from the Management Portal user interface. To do this, follow these steps:

1. Log on to the Management Portal as the **admin** user:
   https://<master node FQDN>:5443
2. Click **Suite** > **Management**.
3. Click the action icon for the ITSMA suite instance, and then select **Uninstall**.

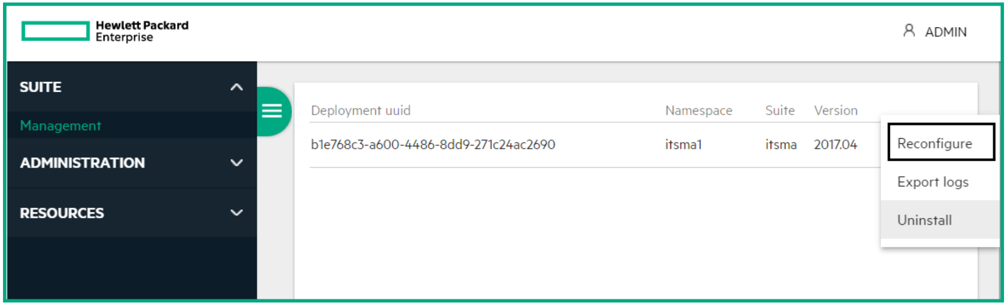

Deployment uuid	Namespace	Suite	Version
b1e768c3-a600-4486-8dd9-271c24ac2690	itsma1	itsma	2017.04

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 ITSMA documents (ITSMA 201704)

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-ITSM@hpe.com.

We appreciate your feedback!