# ITSM Automation NG Express

Software release version: 2017.07

## Installation Guide

Document release date: July 2017
Product release date: July 2017

Hewlett Packard Enterprise

# Legal Notices

## Warranty

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

The information contained herein is subject to change without notice.

## Restricted rights legend

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

## Copyright Notice



## Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

# Documentation updates

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/.

# Install

The procedure to install ITOM Container Deployment Foundation (CDF) and ITSMA depends on whether your deployment will be on-premises or cloud-based.

## On-premises

Install CDF and ITSMA on premises

## Cloud-based

Install CDF and ITSMA in the cloud

# Install the suite on premises

ITSMA 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 (" ITOM Container Deployment Foundation (CDF)" ), and then install the ITSMA suite from a graphic user interface. The suite components are deployed quickly and integrated seamlessly, requiring little user intervention.

ITSMA can be deployed on on-premises servers (that is, physical or virtual servers) or cloud servers. This section describes the steps for on-premises deployment. For information about cloud deployment, see Install the suite on AWS.

- On-premises deployment supports two deployment modes: fully containerized mode, and mixed mode. For more information, see Deployment modes.
- 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 either during or after the installation. For more information about supported platforms, see Support matrix (on-premises).
- Mixed mode may not require a suite license depending on your actual situation. For more information, see Install an ITSMA suite license.

## Installation procedure

The ITSMA NG Express installation procedure comprises five stages. Click here to view the installation workflow.

The following table lists detailed information about the steps illustrated in the workflow diagram. You can use the checklist in the following table to keep track of completed installation steps. You can also click the link for each step to view detailed instructions.

- In a production environment, using dedicated NFS servers is recommended. Use a master node as the NFS server only in a test environment.
- CDF has two user roles: IT Administrator and Suite Administrator. The out-of-box IT Administrator user account is **admin**/**cloud**.

| Stage | Steps |
| --- | --- |
| Plan | 1.1 Read about the support matrix, sizing recommendations and other information<br><br>1.2 Decide your deployment mode |

| | |
|---|---|
| Prepare | 2.1 Enable your Docker Hub account<br><br>2.2 Prepare your cluster machines to meet the prerequisites<br><br>2.3 Prepare external databases for CDF and ITSMA if you do not want to use their built-in database<br><br>2.4 Set up an Access Server if your organization's network includes a DMZ<br><br>2.5 Download the CDF installation package |
| Install CDF | 3.1 Set up an NFS share for CDF<br><br>3.2 Configure the install.properties file<br><br>3.3 Set up the master node<br><br>3.4 Set up the worker nodes |
| Install ITSMA | 4.1 Download the ITSMA suite images from Docker Hub to CDF<br><br>4.2 Set up three NFS shares for ITSMA<br><br>4.3 Run the Suite Installer<br><br>4.4 Verify the suite installation |
| Perform post-installation configuration | **For fully containerized mode**:<br><br>5.8 Activate an ITSMA suite license if required<br><br>**For mixed mode scenario 1**:<br><br>5.2 Configure the URL of the external Service Manager server and specify an integration account<br><br>5.3 If you are using the Solr search engine and do not want to move to the containerized Smart Analytics (SMA), install a Solr plugin to enable KM search for self-service users using Service Portal<br><br>5.4 If you are using the Solr search engine and want to move to the containerized SMA, purchase an SMA license and migrate from Solr to the containerized SMA<br><br>5.5 If you are using SMA with your external Service Manager, migrate to the containerized SMA<br><br>5.6 (Optional) Switch to the containerized Chat in ITSMA<br><br>5.7 Connect to the containerized Service Portal, which is intended for self-service users<br><br>5.8 Activate an ITSMA suite license if required<br><br>**For mixed mode scenario 2**:<br><br>5.1 Set up an integration between the containerized Service Management and external CMDB<br><br>5.8 Activate an ITSMA suite license if required |

## Plan an on-premises suite deployment

Before your proceed, review the following information to plan your suite deployment:

- Support matrix (on-premises)
- Sizing
- CDF installation configuration

### Support matrix (on-premises)

This section provides support matrix information of ITOM Container Deployment Foundation (CDF) and ITSMA NG Express.

- Supported environments

## 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 comprises one or three master nodes and multiple worker nodes for load balancing purposes. Client requests are sent to the load balancer, and then redirected to the master nodes, and finally to the worker nodes. An NFS server is required to store data for CDF and the suite.

| | |
|---|---|
| Test environment | In a test environment, you can use the following configuration:<br>• One master node (used as a worker node and an NFS server as well)<br>• One worker node |
| Production environment | One or three master nodes, multiple worker nodes, and one dedicated NFS server.<br>For details, see Sizing. |

## Operating systems

The master node, worker nodes, and the NFS server hosts must use the same operating system that is described below.

| | |
|---|---|
| ***Virtual or physical hosts:*** | • 64-bit CentOS 7.2, 7.3<br>• 64-bit RHEL 7.2, 7.3<br>• 64-bit OEL 7.3<br>If you are using CentOS/RHEL/OEL 7.3 with a kernel version of 3.10.0-514.21.2, you need to upgrade to 3.10.0-514.26.2.el7.x86_64 or above to avoid a known issue. For more information, see ITSMA suite issues. |

## Required network identification

Only IPv4 and FQDN are supported.

## Databases

Both ITOM Container Deployment Foundation (CDF) and ITSMA support the use of external databases instead of an internal PostgreSQL database.

> The ITSMA suite has an embedded PostgreSQL database, which is for demontration environments only. For production environments, use external databases.

| | |
|---|---|
| CDF | CDF supports the following external databases:<br>• PostgreSQL 9.5.7 or later<br>• Oracle 12c |
| ITSMA | You have the option to use external databases for the suite components listed below:<br>• **Service Portal**: PostgreSQL 9.5.7 or later for Linux<br>• **Service Management**: Oracle 12c, PostgreSQL 9.5.7 or later for Linux<br>• **CMDB**: Oracle 12c<br>Ensure that the contrib package is installed on the Service Portal PostgreSQL database server. |

### LDAP servers

The internal OpenLDAP server bundled with ITSMA is only for demonstration purposes.

Technically, ITSMA supports all external LDAP servers. The following LDAP servers have been certified by HPE:

- OpenLDAP
- Microsoft Active Directory

### SAML 2.0 identity provider

ITSMA supports SAML SSO using the following identity provider (IdP): Microsoft Active Directory Federation Services (ADFS) 2 or 3.

For more information about, see Configure SAML SSO.

### Mixed mode support

When deployed in mixed mode, ITSMA can be integrated with the following HPE applications:

| | |
|---|---|
| Scenario 1 (using external SM and UCMDB) | • Service Manager (SM) 9.41 or later, 9.5x (9.50 or later)<br>• Universal CMDB (UCMDB) 10.3x (10.30 or later), 10.2x (10.20 or later) |
| Scenario 2 (using external UCMDB only) | • Universal CMDB 10.22, 10.30, 10.31, 10.32, 10.33 |

For additional integration support in mixed mode, see Additional Integration Support.

For more information about mixed mode, see Deployment modes.

### Browsers

Use the the following browsers to access the CDF Management Portal and ITSMA:

| Interface | Browser |
|---|---|
| Management Portal | • Microsoft Internet Explorer 11<br>• Google Chrome 48 or later<br>• Mozilla Firefox 45 ESRWindows version<br>• Apple Safari 10.1 |
| Service Portal | • Internet Explorer (IE) 11<br>• Latest Firefox<br>• Latest Chrome<br>• Edge |

| | |
|---|---|
| • Service Management<br>• CMDB<br>• CMDB Browser<br>• Suite Configuration | • Internet Explorer (IE) 11<br>• Latest Firefox<br>• 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 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 | • Android 7.x, 6.x, 5.x with Android browser | For accessing the Mobility capability of the ITSMA suite |

### Language support

| Capabilities/Modules | Supported languages | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | English | Arabic | Brazilian Portuguese | Danish | Dutch | French | German | Greek | Italian | Japanese | Korean | Norwegian | Polish | Turkish | Russian | Simplified Chinese | Spanish |
| Service Management | | | | | | | | | | | | | | | | | |
| CMDB | | | | | | | | | | | | | | | | | |
| • CDF (Management Portal)<br>• Suite Configuration | | | | | | | | | | | | | | | | | |
| Service Portal | | | | | | | | | | | | | | | | | |

- Except for Service Portal, the displayed language for ITSMA capabilities can be switched by changing the language setting in your browser. For Service Portal, the displayed language is determined by the language setting that you specify in your user profile in Service Portal.
- In Service Portal, when you use a language that is supported only by Service Portal (not supported by other ITSMA capabilites), you might see mixed languages or English elements (labels, action labels, drop-down options, and so on) on some forms. This is because such data comes from other ITSMA capabilities. You can install a Service Manager 9.50 language pack in the containerized Service Management by using a workaround that is described in the ***Language support*** section here.
- For suite configuration, when you access the Service Portal theme setting and feature setting pages, you might see a different language than the four supported suite configuration languages. This is because the Service Portal theme setting and feature setting pages are part of the Service Portal application in which the displayed language is specified in the user profile in Service Portal instead of in your browser language setting. To avoid this issue, set the language in Service Portal user profile to the same language that you use in suite configuration.
- If using Firefox or Internet Explorer 11, you must select "<Language>" instead of "<Language> (<Country>)" when you set your browser language.

### Additional integration support

| Product | Versions | Notes |
|---|---|---|

| | | |
|---|---|---|
| Asset Manager (AM) | 9.50 and 9.6x | Supported in mixed mode scenario 1. |
| Universal Discovery (UD) | 10.33 | Supported in fully containerized mode and mixed mode scenario 1. |

# Sizing

## Introduction

This section provides customers with guidelines for implementing ITSMA in the production environment. This section includes hardware sizing recommendations as well as deployment and configuration optimization information. For customers who are going to have a quick look at the system, demo/test environment and corresponding configurations are also provided.

The sizing recommendations are based on substantial load tests implemented in R&D lab which approximately simulate expected user activities and various data volume sizes on backend. In reality, customers may have more complex scenarios which may result in different results. Therefore, we suggest that you review user profiles utilized in this section before diving into sizing recommendations. In addition, a performance test is suggested in your environment to get a more precise performance picture so that you can increase or decrease hardware requirements accordingly.

See the following sections for your ITSMA sizing configurations:

1. Prepare for sizing
2. Set user profile
3. Deployment modes
4. Hardware sizing recommendations
5. Tuning configuration
6. Scaling up and out

## Terminology

| Term | Description |
|---|---|
| Registered users | Users who have credentials (user name and password) to log on to the system, submit/update tickets, order services, and so on. |
| Concurrent users | Online users who have logged on to the system to perform operations and consume system processing resources. They are different from registered users. |
| Throughput | The frequency that end users perform business transactions per hour or per minute. For example, 100 orders per hour, 100 requests per hour, and so on. |
| Workload | User activity imposed on the system. |
| Session length | The duration of the session performing related operations on the system. |
| Pacing time | The time between successive user activities. For example, an end user submits an order each day, considering the 8 working hours a day, the pacing time is 8 hours. |
| User profile | Combination of concurrent users, workload, initial data, user session length, and pacing time. |
| Backend data volume | The active volume of business data in the customer database during daily operations. In general, customers will keep 2 years' business data before archiving. |

## Prepare for sizing

### Introduction

Sizing is the process of determining the underneath hardware resources to support the business requirement targets in production, for both the present and the future. In reality, the most concern on the customer side is the amount of user loads that can be handled by the system. This is true, but still not enough. To answer the question such as:

- What's the hardware resource I should apply to support ***N*** user loads?

There are at least three direct factors customers need to consider:

1. Concurrent users.
2. Throughput.
3. Backend data volume size.

Additional factors that can not be neglected, for example:

1. Business growth in the future 3 to 5 years.

The sizing information in this section is derived from end user and server side perspectives. The end user perspective involves user types, the number of users, their think time on the page, pacing time between their successive activities, and their workload applied on the system. The server side perspective is more focusing on the throughput like transactions per hour and backend data volume size.

No matter you are building up the system from scratch or migrating your system to ITSMA, in order to get more accurate sizing recommendations, you are suggested to follow these steps:

1. Collect user profile information.
2. Determine target deployment mode. Details please refer to Deployment modes.
3. Retrieve sizing recommendations.
4. Consider redundancy for high availability, or design disaster recovery process.
5. Apply tuning configurations that the sizing guide has suggested.

Please note that sizing information provided in this section has included ~30% capacity for the future growth based on lab load tests. For customers who have tremendous different requirements comparing to this section, they may need to apply more hardware resources. Also HPE suggests that you implement individual load and stability tests to make sure that no undersized hardware resource are planned based on the minimum hardware reference.

## Set user profile

### Introduction

This section provides you with the detailed information of user profiles on which sizing recommendations are based. We suggest that you compare it with the user profiles collected in your own environment. More hardware resources are required if you have greater business needs.

ITSMA July release includes five size levels:

1. Demo (or test)
2. Extra small
3. Small
4. Medium
5. Large

You can get a brief outline from the following table.

Table -1: General Concurrent User Size and Initial Data Definition

Edit Document

| Target Size | Service Portal Users | IT Agent Users | Records in SMA - IDOL | CIs and |
|---|---|---|---|---|
| Demo | 5 | 5 | 1K | 1K |
| Extra Small | 50 | 50 | 1M | 2M |
| Small | 200 | 200 | 1M | 2M |
| Medium | 500 | 500 | 2M | 6M |
| Large | 2000 | 1000 | 4M | 25M |

(K=thousand;M=million)

### User profile

User profile information is critical to sizing recommendations, different profiles will produce completely different sizing requirements. It represents real users behaviors and impact on the system during their daily work. To make sure everyone is on the same page, this part describes more details which are utilized by the sizing section.

### *Concurrent users*

Concurrent users are real users who are active on the system and consuming server resources such as memory, CPU, and I/O. For example, at any time on work days, there may be only 2% ~ 5% of registered portal users are logged on to the system. You should plan enough hardware resources to accommodate maximum concurrent users during peak hours.

In general, the sizing section differentiates two types of end users:

1. Service Portal users.
   Session duration of Service Portal users tend to be shorter and they approach the system only a couple of times a day. They view/search catalog items, submit orders or support ticket, they may also trigger chatting with IT agent users in the backend, and so on.
2. IT agent users.
   IT agent users generate much longer session durations on the system, they fulfill portal users' requests, and go through the whole ticket life cycle. They may stay active on the system during most work hours.

### *Session length*

The sizing section uses the following session durations as baseline for hardware resource suggestions.

1. Service Portal users.
   After logging on to the system, each portal user has about 30 minutes session duration on the system. They will produce 1 ~ 2 orders or requests. For example, for extra small size, at any time there will be 50 concurrent portal users active on the system, each of them will complete his/her work in 20 ~ 30 minutes.
2. IT agent users.
   IT agent users have longer session durations, each of them will be active for  about 1.5 hours after logon, during which, 4 ~ 6 tickets will be handled and closed by him/her.

### *Pacing time*

In order to make sure that each size of deployment can support targeted number of concurrent users as claimed, no pacing time is applied during load tests in R&D lab.

### *Think time*

Think time is applied to simulate user's idle time between page operations on web UI. For example, portal users may view catalog details before submitting an order. Or, IT agent users may read descriptions before fulfilling customer's requests.

During sizing tests, portal users' think time is set between 30 seconds and 40 seconds, while for IT agent users, the think time is set between 35 seconds and 45 seconds.

### *Workload*

Following's are benchmark scenarios utilized in R&D lab for sizing recommendations.

#### **Service portal users**

### I - View ticket

Note: The view scenarios will be executed by each XSP user after logon. It is supposed that users have logged on.

1. Click top-right notes, click Requests.
2. Choose the first item to view OPEN requests.
3. Go back to the Your Requests page.
4. Click tab - CLOSED.
5. From top-right notes, click Approval.
6. Click top-right to-do items.
7. Click the top-left port logo and go back to the homepage.
8. Choose category - TECHNICAL AND SUPPORT SERVICES, and click it.
9. Click tab - OFFERINGS.
10. Click tab - ARTICLES.
11. Go back to the homepage.
12. In the search bar, type any word like "printer", click search. Repeat 5 times.
13. Go back to the homepage.

### II - New support ticket

1. Open the logon page.
2. Enter the username and password, click Login.
3. Run all steps in I - View scenario.
4. Choose category - APPLICATIONS, and click it.
5. Click tab - OFFERINGS.
6. Type PerformanceSupportItem in the search bar, and click search.
7. Click the next page.
8. Choose one item.
9. Select Urgency - 2-High, click Submit.
10. Click the ticket number to verify that the supplier request ID is returned.
11. Add the first comment under the ticket.
12. Add the second comment under the ticket.
13. Go back to the homepage.
14. Log out.

### III - New order - single item

1. Open the logon page.
2. Type username and password, click Login.
3. Run all steps in I - View scenario.
4. Choose category - PERSONAL PRODUCTIVITY SERVICES, and click it.
5. Click tab - OFFERINGS.
6. Type in "PerformanceTestItem" in the search bar, and search.
7. Click the next page.
8. Add one item to the order.
9. Type in field information, and click Checkout.
10. Click Submit.
11. Click the order number to check the status.
12. Go back to homepage.
13. Log out.

### IV - New order - multiple items

1. Open the logon page.
2. Type in username and password, click login in.
3. Run all steps in I - View scenario.
4. Choose category - PERSONAL PRODUCTIVITY SERVICES, and click it.
5. Click tab - OFFERINGS.
6. Add the first item to order.
7. Type in field information, click ADD TO CART.
8. Click CONTINUE BROWSING.
9. Choose categpry - PERSONAL PRODUCTIVITY SERVICES, and click it.
10. Click tab - OFFERINGS.
11. Type in name in the search bar, and search.
12. Add one item from the search results to order.
13. Type in field information, click ADD TO CART.
14. Click CHECK OUT.
15. Type in field information, and click SUBMIT.
16. Check the order status.
17. Go back to homepage.
18. Log out.

**IT agent users**

## I - Service Desk

1. Open Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expend Service Desk in the navigation panel.
4. Click Create New Interaction.
5. Select category - incident.
6. In field - Contact, type in FAL, select "FALCON, JENNIFER" for Contact in the drop down list.
7. Select Impact - 2 - Site/Dept, urgency - 2 - High, Notify By - E-mail, type in title - this is for itsma-sm-webtier-sd test.
8. Fill in Affected Service - MyD, and select MyDevices from auto-completed list.
9. Type in description like - printer, smart request will return relating ticket automatically. Repeat 5 times.
10. Click Save.
11. In Subcategory, click Fill.
12. Select hardware for subcategory, and hardware failure for area.
13. Click Escalate.
14. Click Create New Incident.
15. Click Save on the incident detail page.
16. Change incident Status to - Resolved, on tab - Proposed Solution, type in Solution - Resolved by user instructions, click Save.
17. In Assignee, click fill.
18. Select falcon from the drop down list.
19. Click Close.
20. Select Closure Code - Solved by User Instruction, and type in comments like - solved by user instruction. Click Finish.
21. Click Yes when prompted whether to close related interaction.
22. Click Cancel to quit the incident detail page.
23. Click Back to cancel previous interaction detail page.
24. Log out from Service Manager.

**II - Incident Managent**

1. Enter the Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expend Incident Management in the navigation panel.
4. Click Create New Incident.
5. Select category - incident.
6. Type in title - this is for itsma-sm-webtier-im test.
7. Type in description like - network, smart request will return relating ticket automatically. Repeat 5 times.
8. Type in Affected Service - Applications, click Save.
9. Click continue to client new incident.
10. Click Cancel to quit the current incident detail page.
11. Click Search Incidents to open the incident search page.
12. Type in Incident ID, and click Search.
13. Change incident status to Work in Progress.
14. In Subcategory, click Fill.
15. Select subcategory - failure, Area - job failed, click Save.
16. Change incident status to Resolved, under tab - Proposed Solution, type in resolved by users. Click Save.
17. Change incident status to Work in Progress, click Save.
18. Under Tab - Tasks, click Link New Task.
19. Select category - Investigation.
20. In task detail page, type in description - this task is for investigation of incident, click Save.
21. Select due date (next day), and click Close Task.
22. Select Completion code - Successful, type in Task Outcome - task is finished successfully, click Finish.
23. Click Cancel to quit the current task detail page.
24. After going back to the incident page, change status to Resolved, click Save.
25. Click Close.
26. Select Closure Code - Solved by Workaround, type in solved by workaround in comments, click Finish.
27. Click Cancel to quit the incident detail page.
28. Click Cancel to quit the incident search page.
29. Log out from Service Manager.

## III - Change Management

1. Enter the Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expend Change Management in the navigation panel.
4. Click Create New Change.
5. Expend Category: Normal Change, expend Subcategory: Major.
6. Click Normal Major RFC.
7. Type in title like - this is for itsma-sm-webtier-cm normal change test. Select Impact - 2 - Site/Dept, Urgency - 2 - High.
8. Select Requested End Date - end of the month, select Reason for Change - User Dissatisfaction, type in Service - My Devices.
9. Type in description - this is for itsma-sm-webtier-cm normal change test. This will trigger SMA request to IDOL. Repeat 5 times.
10. Type in Effect of not Implementing - fail to satisfy customer.
11. Click Save.
12. Click Request Validation.
13. Select Risk Assessment - 2 - Some Risk, set change coordinator and owner to OP000001.
14. Click Validation Accepted.
15. Type in Remediation Plan - this is for remediation plan. Click Request Authorization.
16. In More, click Change Phase.
17. Click Build and Test. Click Yes for the confirmation message.
18. Click Cancel on the change detail page.
19. Click Search Changes.
20. Type in change id just created, click Search.
21. Under tab - Tasks, click the task ID.
22. Change Status to Completed, and fill in other needed fields. Click Save.
23. Click Close to close task.
24. Select Closure Code - 2 - Successful (with problems), type in Closure Comments - this is for test task. Click Finish.
25. Cick Cancel.
26. Type in Impl4mentation Plan - this is for implementation plan, Build and Test Plan - this is for build and test plan, Build and Test Result - this is for build and test result.
27. Fill in Scheduled Implementation Start and End time. Click Request Authorization.
28. In More, click Change Phase.
29. Select Deployment. Click Yes for the confirmation.
30. Type in Actual Implementation Start and End date, type in Implementation Comments - the deployment is finished. Click Request CMDB Update.
31. Click Request PIR.
32. Type in Review Results - this change is finished. Type in Closure Comments - user is satisfied. Click Close.
33. Click Finish.
34. Click Cancel to quit the change detail page.
35. Click Cancel to quit the change search page.
36. Log out from Service Manager.

## IV - Problem Management

1. Open Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expend Problem Management in the navigation panel.
4. Click Create New Problem.
5. Type in Title - this is for itsma-sm-webtier-pm test
6. Type in Description - this is for itsma-sm-webtier-pm test 201707. This will trigger SMA request to IDOL. Repeat 5 times.
7. Type in App in Affected Service field, which will trigger auto complete request.
8. Click Save.
9. Click Continue.
10. Type in Subcategory - hardware, type in Area - hardware failure. Click Close.
11. Select Closure Code - Out of Scope, type in Closure Comments - out of scope, click Finish.
12. Click Back.
13. Click Create New Known Error.
14. Type in Title - this is for itsma-sm-webtier-ke test for record KExxxx.
15. Type in Description - this is for itsma-sm-webtier-ke test for record KExxxx, which will trigger SMA request to IDOL. Repeat 5 times.
16. Fill in Affected Servie - MyDevices, Root cause - not known, workaround - not known. Click Save.
17. Click Close.
18. Fill in Subcategory - hardware, Area - hardware failure. Select Closure Code - Resource Challenge, type in Closure Comments and Solution. Click Finish.
19. Click Back to go back to the homepage.
20. Log out from Service Manager.

**V - Service Catalog**

1. Open the Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expand Service Catalog in the left panel.
4. Click Order from Catalog.
5. Choose and click Personal Productivity Services.
6. Choose and click Hardware Bundles.
7. Choose and click Basic PC Package.
8. Click Add to Cart.
9. Click View Cart/Checkout.
10. Click Submit Request.
11. Type This is for performance test - order from catalog; For Need By - the end of the month; For Urgency - 3 - Average. Click Submit.
12. Click Continue.
13. Click Refresh on the to do queue page.
14. Log out from Service Manager.

**VI - Request Management**

1. Open Service Manager web tier logon page.
2. Type in username and password, click Login.
3. Expand Request Fulfillment in the left panel.
4. Click Create New Request.
5. Go to Category - Generic Request, Subcategory: Hardware, click Order Hardware.
6. Type This is for performance test - request fulfillment in the field Title and Description; For Impact - 3 - Multiple Users; For Urgency - 3 - Average. Click Save. (Make note of request number)
7. Click Cancel to quit the request page.
8. Click Search Requests in the left panel.
9. In Request ID, type in ticket number which you created in step 6, click Search.
10. Update Expected Finish Date - the end of the month, click Save.
11. Click Submit for Approval.
12. Click Cancel to quit the request page.
13. Click Cancel to quit the request search page.
14. Log out from Service Manager.

## Deployment modes

You can deploy ITSMA in containerized mode or mixed mode, and you can use one or three master nodes in your deployment.

### Containerized mode and mixed mode

ITSMA 2017.07 supports two deployment modes:

- **Mixed mode**: provides flexibility to customers who want to leverage containerized services offered by ITSMA such as Service Portal, Smart Analytics and Chat, while still keeping their existing Service Manager or Universal CMDB implementations under classic deployment. Mixed mode supports two scenarios:
    - Scenario 1: ITSMA comprises external Service Manager and UCMDB, and the following containerized services: Service Portal, Smart Analytics, and Chat.
    - Scenario 2: ITSMA comprises external UCMDB and the following containerized services: Service Management, Service Portal, Smart Analytics, and Chat.
- **Fully containerized mode**: ITSMA comprises only containerized services: Service Management, CMDB, Service Portal, Smart Analytics, and Chat.

The following table provides a summary of each mode.

| Deployment mode | Containerized capabilities | External systems used | Note |
|---|---|---|---|
| Fully containerized | • Service Management<br>• CMDB<br>• Chat<br>• Service Portal<br>• Smart Analytics | None | • All of these suite components are pre-integrated and no manual integration setup or migration is needed.<br>• A suite license is required. For more information, see Install an ITSMA suite license.<br>• Chat, Survey, and Smart Analytics are enabled out of the box. |
| Mixed | Scenario 1:<br>• Chat<br>• Service Portal<br>• Smart Analytics<br>Scenario 2:<br>• Service Management<br>• Chat<br>• Service Portal<br>• Smart Analytics | Scenario 1:<br>• HPE Service Manager<br>• HPE Universal CMDB<br>Scenario 2:<br>• HPE Universal CMDB | • Not supported for AWS deployments.<br>• Requires manual post-installation integration setup for the containerized components to work with the external system(s). Manual integration setup and data migration are needed to use Service Portal, Chat, Smart Analytics, and Survey.<br>• A suite license may not be required depending on your actual situation. A Smart Analytics license may be required if you do not already have one. For details, see Install an ITSMA suite license.<br>• Supports HPE Service Manager 9.4x and 9.5x (both codeless and classic/hybrid), HPE Universal CMDB 10.2x and 10.3x (using a generic Service Manager adapter).<br>Since the End User Chat functionality is available only for Service Manager 9.50 or later, to enable Chat for Service Portal users, you must upgrade your external Service Manger system to version 9.50 or later. |

| | | | |
|---|---|---|---|
| | | | • Search engine support: Solr, and IDOL (provided by the containerized Smart Analytics)<br>  • If you are using Solr in your external Service Manager system, you have two options:<br>  • Continue to use Solr: you need to install a Solr plugin. This plugin is required for Service Portal users to search knowledge articles indexed by Solr. However, when a Service Portal user searches catalog items, an IDOL search is performed using the containerized Smart Analytics.<br>  • Move to the containerized Smart Analytics: you need to buy a Smart Analytics license and enable the containerized Smart Analytics in the external Service Manager system.<br>  • If you have already Smart Analytics enabled in your external Service Manager, your IDOL server must have already indexed data. You need to migrate indexed data from the external Smart Analytics to the containerized Smart Analytics. |

### Single-master deployment samples

Demonstrate deployment samples in production are as follows, besides that, a demo/test environment sample is also provided.

For detailed server configurations, please refer to Hardware sizing recommendations.

**NOTES**

1. A master node can also work as a worker node and NFS node only in a demo or test environment. For production, Master, Worker and NFS are located on different servers.
2. Internal LDAP and PostgreSQL database are only suggested for demo/test environment. For production, external LDAP and Oracle or PostgreSQL database servers are utilized.
3. Mixed mode only supports deploying ITSMA on physical, virtual servers.
4. Full containerized mode supports deploying ITSMA on physical, virtual servers and AWS.

#### *ITSMA demo/test environment*

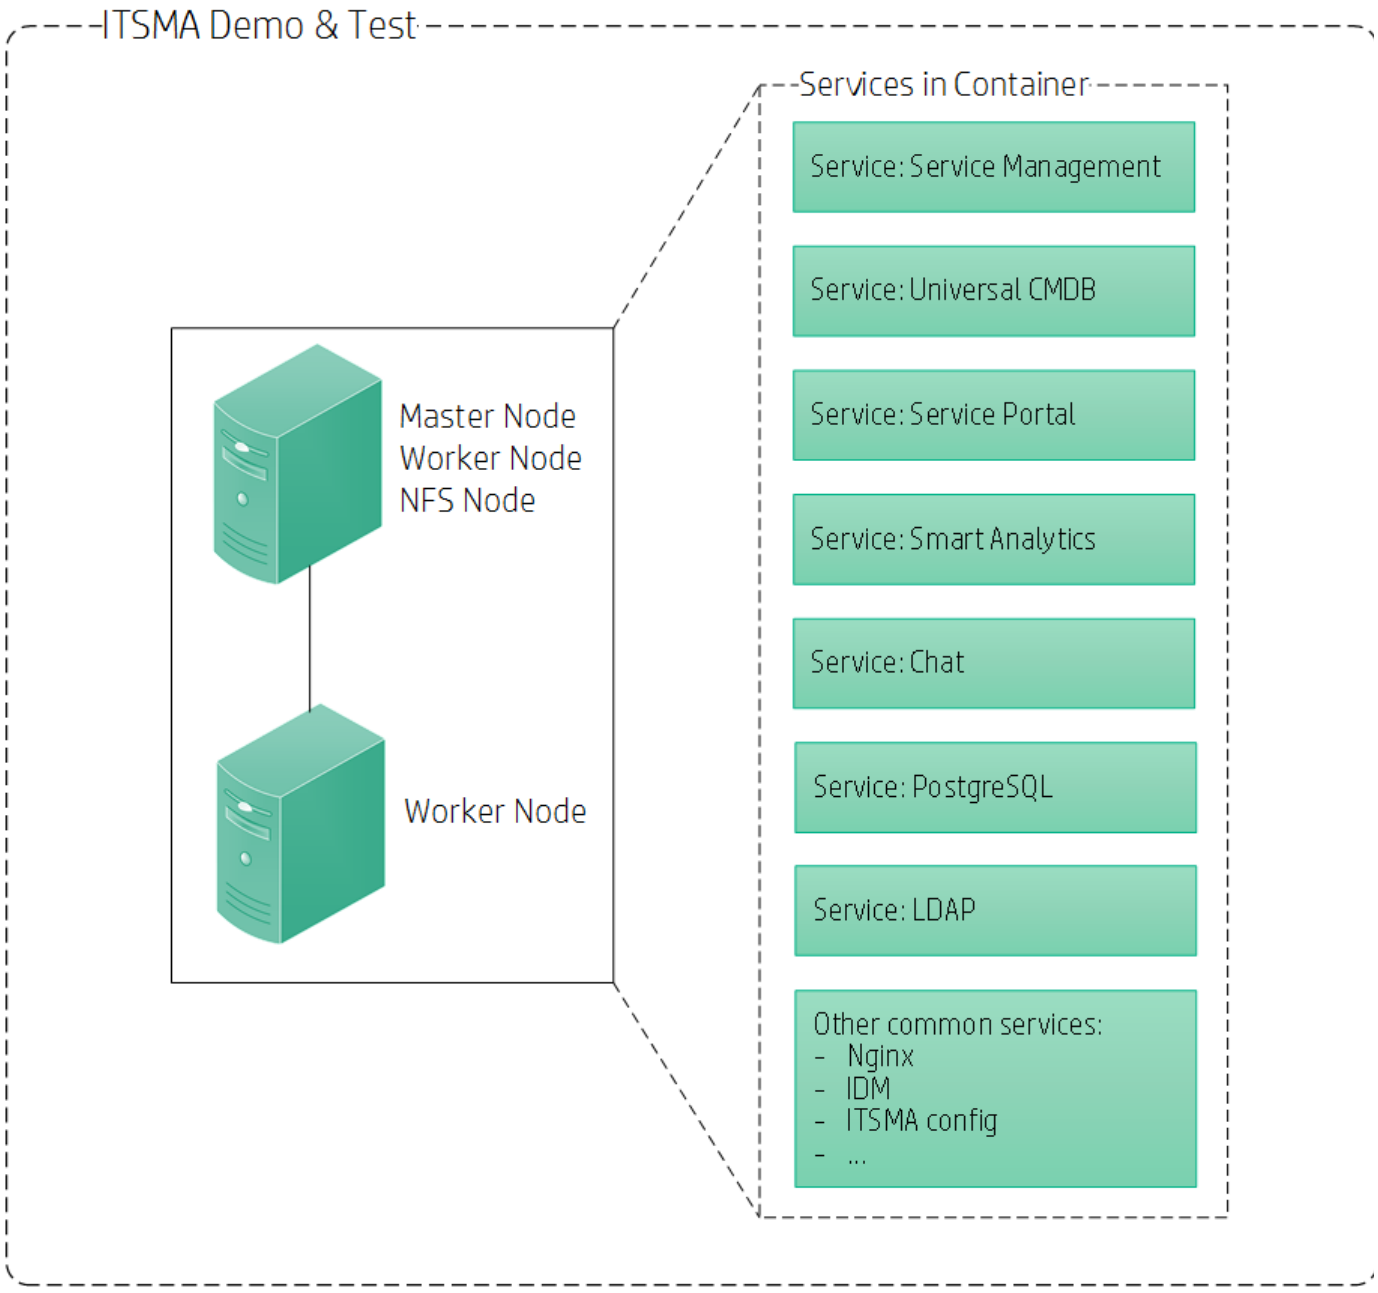


***ITSMA mixed mode scenario 1***

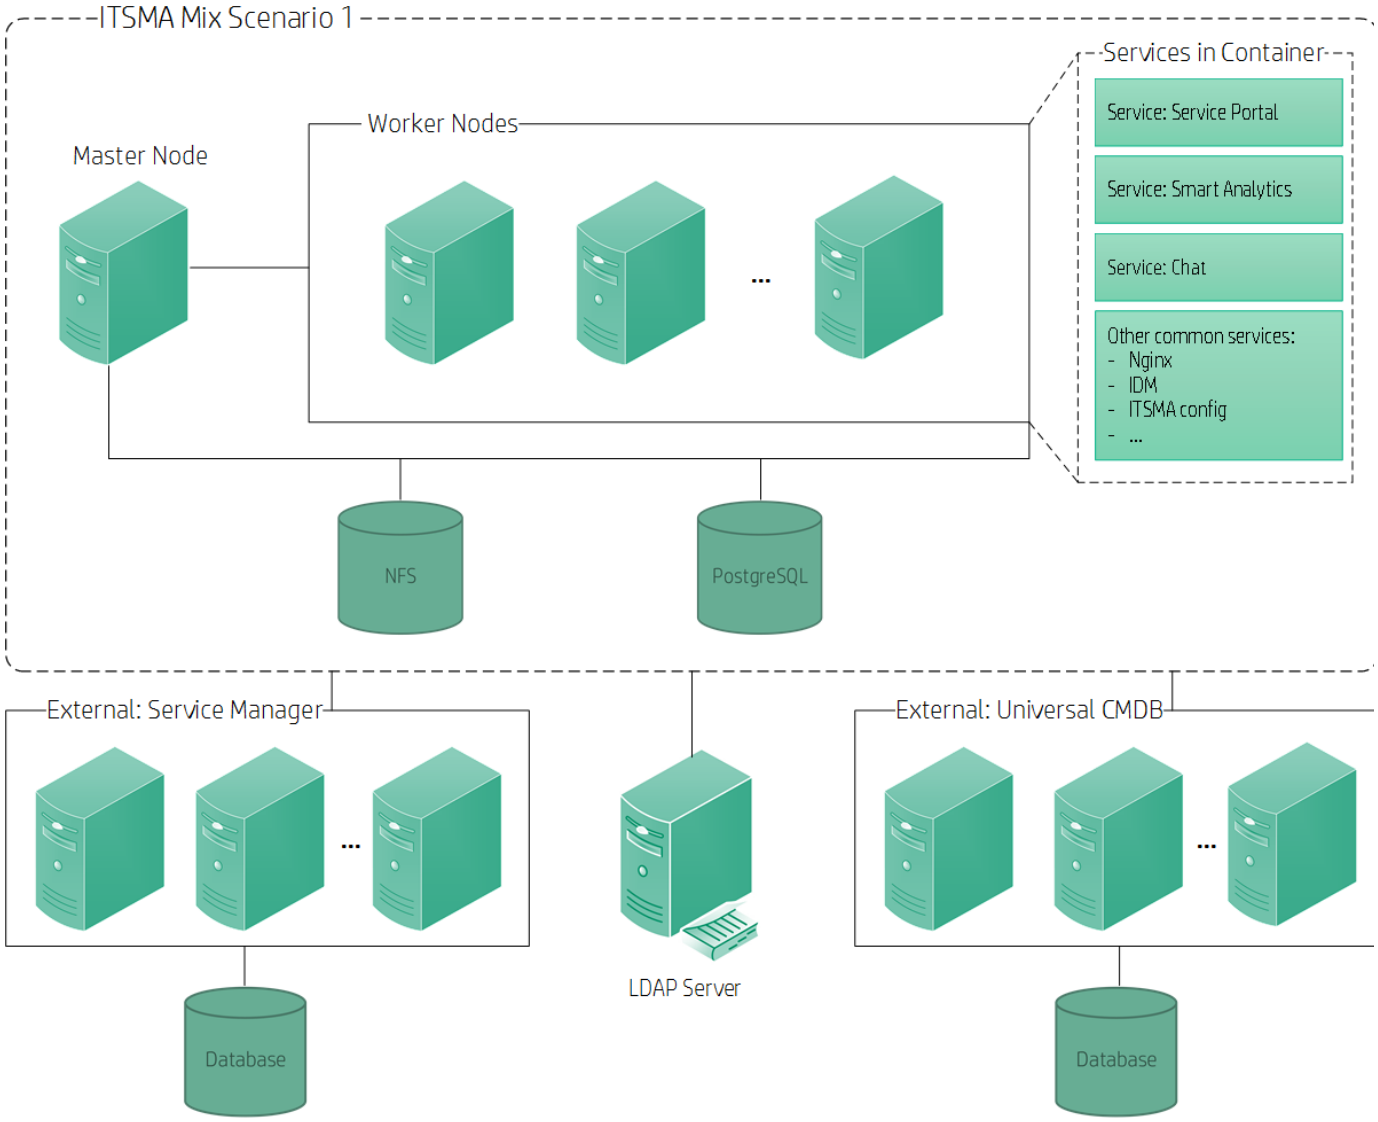


***ITSMA mixed mode scenario 2***

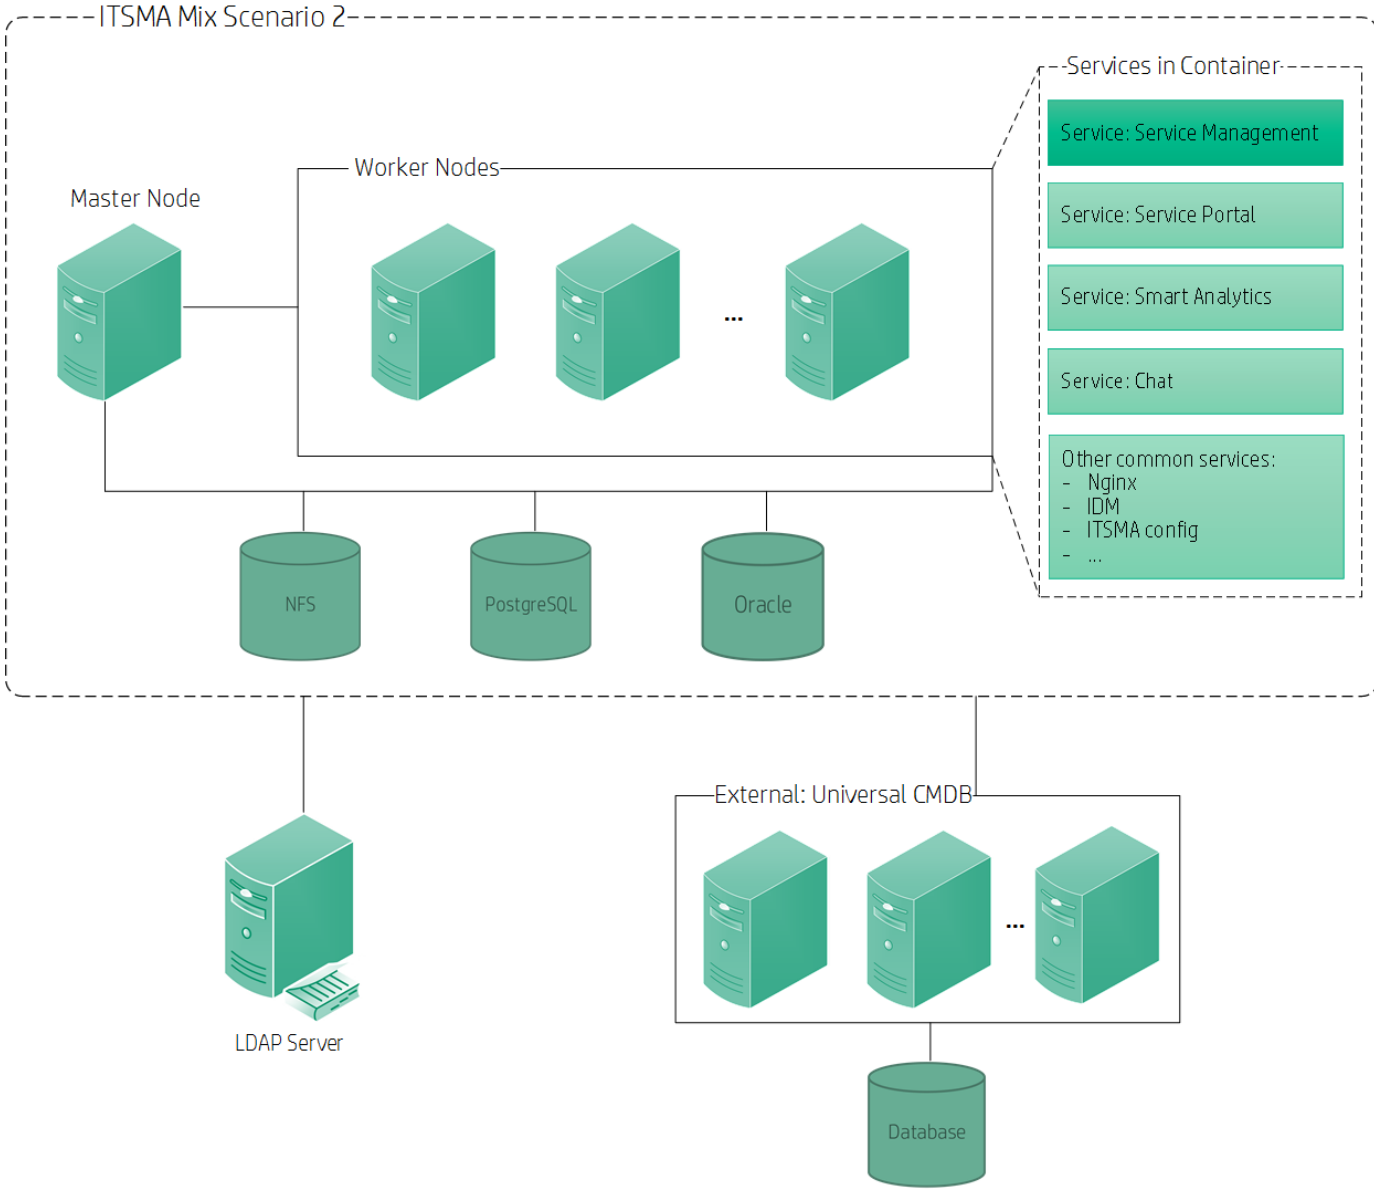


***ITSMA fully containerized mode***

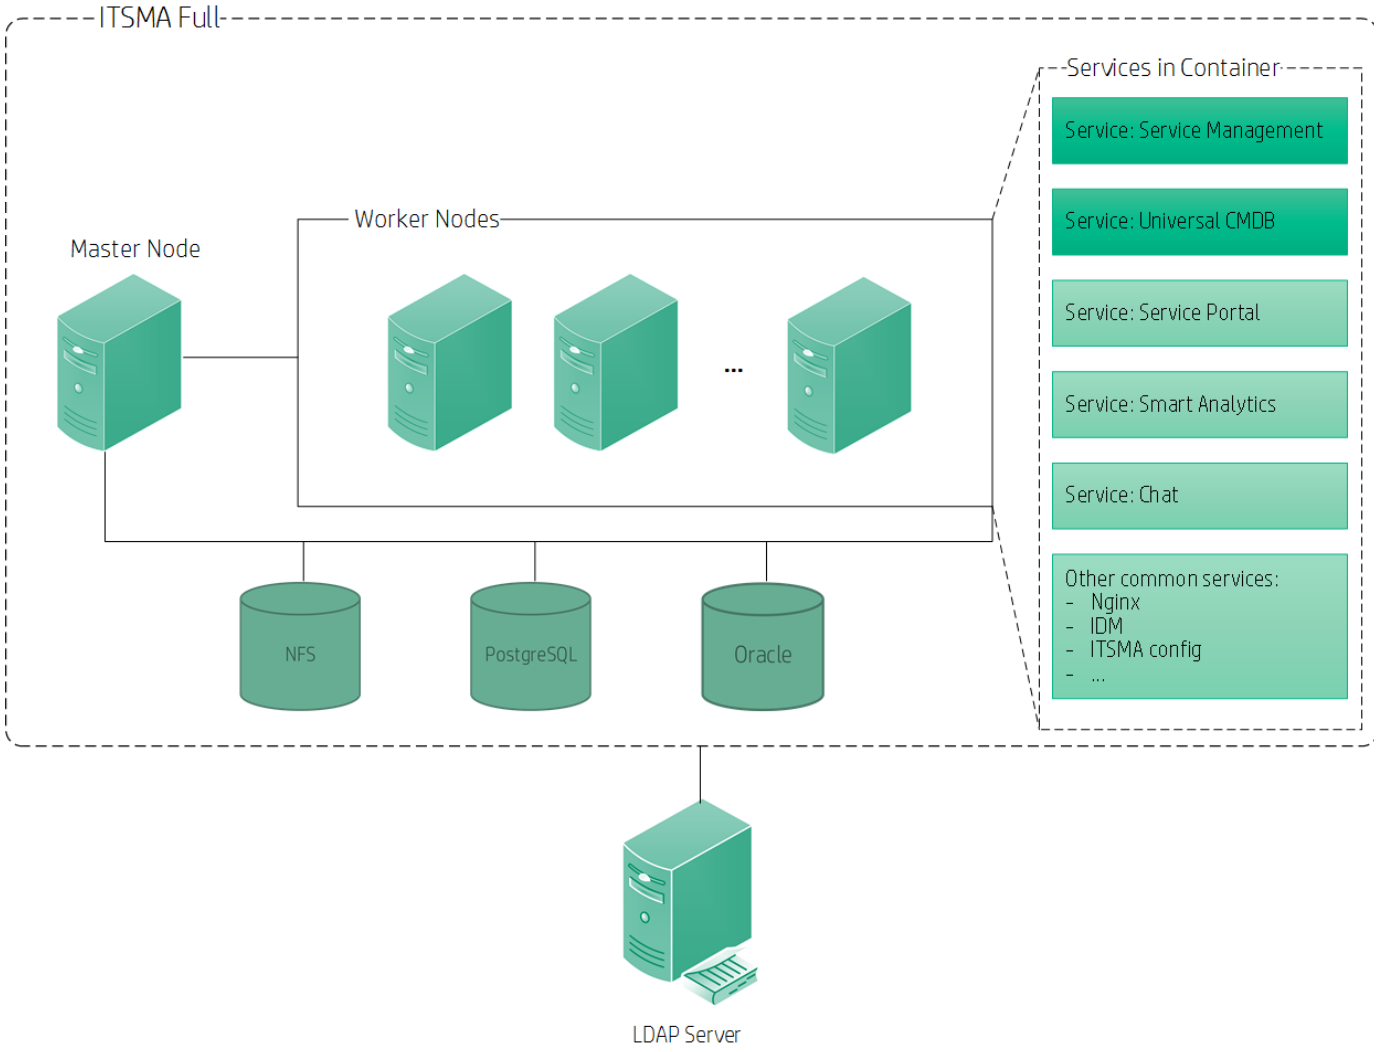


### HA deployments

You can achieve high availability of ITSMA by using our multi-master failover mechanism. On a multi-master setup, a virtual IP address can be defined, assigned and kept alive across the master nodes. When one master goes down, a seconday master takes over automatically.

Be aware of the following:

- Only *three* master nodes are supported.
- The ITSMA multi-master failover mechanism works only when one master goes down. If two masters go down, the ITSMA environment goes down.
- To set up a multi-master environment, you must use the CDF package released with the ITSMA 2017.06.001 patch (see Download the CDF installation package (on-premises)); additionally, if deploying ITSMA on AWS, you must use the AWS package that supports HA (see Download CDF and AWS packages).

## Hardware sizing recommendations



### Introduction

This section demonstrates hardware sizing configurations for different user sizes. It covers mixed scenarios running on physical / virtual servers, or full containerized environment running on physical / virtual servers/ AWS.

After customers adopt new hardware resources, they are suggested to do some tuning mentioned in Tuning configuration for better performance and stability before and after ITSMA installation.

**NOTES**

1. For ITSMA deployed on physical or virtual servers, the underneath storage I/O requirement is suggested to be at least 280 MB/s. For ITSMA on AWS, please refer to the sizing details table.
2. Customers can switch to SSD for storage purpose, which provides higher I/O speed and bandwidth. This is typical for large size deployment.
3. For Processor type, Intel Xeon E5, E7 or peers are suggested. For Processor speed, Demo/test environment is suggested to have a processor frequency higher than 1.9 GHz, while for production environment, the processor frequency is suggested to be at least 2.3 GHz. Higher speed will certainly bring in performance improvements.
4. The suggested hardware resources are dedicated for ITSMA without sharing with other product lines.
5. For demo environment, the master node will also work as the worker node and the NFS server. In addition, you do not need to prepare an external database server since internal databases are utilized.

## Sizing recommendations

The following are sizing recommendations for one-premises and cloud-based deployments.

**HIgh availability**
Be aware that these recommendations do not take high availability into account and therefore use only one master node. In a production environment, we recommended using three master nodes to achieve high availability supported by the multi-master failover mechanism of ITSMA.

### *Sizing for on-premises deployment*

ITSMA on-premises deployment includes hardware resources for ITSMA servers (master and worker nodes) and external database/NFS servers.

**Hardware recommendation for ITSMA servers (master and worker nodes)**

Edit Document

| Installation Mode | Target Size | Hardware Configuration | |
|---|---|---|---|
| | | Usage | Number of machines |
| Mixed mode (Scenario 1) | Demo | ITSMA master node | 1 |
| | | ITSMA worker node | 1 |
| | Extra Small | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Small | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Medium | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Large | ITSMA master node | 1 |
| | | ITSMA worker node | 5 |
| Mixed mode (Scenario 2) | Demo | ITSMA master node | 1 |
| | | ITSMA worker node | 1 |
| | Extra Small | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Small | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Medium | ITSMA master node | 1 |
| | | ITSMA worker node | 6 |
| | | ITSMA master node | 1 |

| | | | |
|---|---|---|---|
| | Large | ITSMA worker node | 10 |
| Fully containerized mode | Demo | ITSMA master node | 1 |
| | | ITSMA worker node | 1 |
| | Extra Small | ITSMA master node | 1 |
| | | ITSMA worker node | 3 |
| | Small | ITSMA master node | 1 |
| | | ITSMA worker node | 4 |
| | Medium | ITSMA master node | 1 |
| | | ITSMA worker node | 6 |
| | Large | ITSMA master node | 1 |
| | | ITSMA worker node | 11 |

**Hardware recommendation for ITSMA external database and NFS servers**

Edit Document

| **Installation Mode** | **Target Size** | **Hardware Configuration** | |
|---|---|---|---|
| | | **Usage** | **Nu** |
| Mixed mode (Scenario 1) | Extra-small | ITSMA NFS node | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Small | ITSMA NFS node | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Medium | ITSMA NFS node | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Large | ITSMA NFS node | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| Mixed mode (Scenario 2) | Extra-small | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Small | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Medium | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Large | ITSMA NFS node | 1 |
| | | ITSMA SM database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| Fully containerized mode | Extra-small | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Small | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | Medium | ITSMA NFS node | 1 |
| | | ITSMA database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | | ITSMA NFS node | 1 |

| | | | |
|---|---|---|---|
| | Large | ITSMA SM database - Oracle | 1 |
| | | ITSMA database - PostgreSQL | 1 |
| | | ITSMA uCMDB database - Oracle | 1 |

***Sizing for cloud-based deployment (AWS)***

Edit Document

| Installation Mode | Target Size | Hardware Configuration | |
|---|---|---|---|
| | | Usage | Number of mach |
| Fully containerized mode | Demo | ITSMA master + worker node + NFS | 3 |
| | Extra Small | ITSMA master node | 1 |
| | | ITSMA worker node | 4 |
| | | ITSMA NFS node | 1 |
| | | External: ITSMA database - Oracle | 1 |
| | | External: ITSMA database - PostgreSQL | 1 |
| | Small | ITSMA master node | 1 |
| | | ITSMA worker node | 4 |
| | | ITSMA NFS node | 1 |
| | | External: ITSMA database - Oracle | 1 |
| | | External: ITSMA database - PostgreSQL | 1 |
| | Medium | ITSMA master node | 1 |
| | | ITSMA worker node | 6 |
| | | ITSMA NFS node | 1 |
| | | External: ITSMA database - Oracle | 1 |
| | | External: ITSMA database - PostgreSQL | 1 |
| | Large | ITSMA master node | 1 |
| | | ITSMA worker node | 11 |
| | | ITSMA NFS node | 1 |
| | | External: ITSMA database - Oracle | 1 |
| | | External: ITSMA database - PostgreSQL | 1 |

* The external database server hardware configurations listed in the table are based on our testing using AWS RDS. If your database application is deployed on EC2 with EBS rather than with RDS, the external database hardware configurations for Small and Medium deployments can be reduced by half.

## Tuning configuration



### Introduction

This page presents all tuning configurations R&D have implemented for optimal performance while avoiding any useless overhead.

## Tuning before ITSMA installation

### *Task 1: Create direct-lvm thinpool*

With devicemapper storage driver, Docker provides two kinds of modes:

1. loop - lvm mode.
2. direct - lvm mode.

For production mode, direct-lvm must be applied for better performance and stability. Follow the steps in Task 2: Add or extend a logical volume for a direct-lvm thin pool to create direct-lvm on the master and worker nodes.

More reference can be found at https://docs.docker.com/engine/userguide/storagedriver/device-mapper-driver/.

### *Task 2: Update OS TCP socket read/write buffer size and port range*

Check OS settings on the TCP socket read/write buffer with the sysctl utility, and in most situations, customers need to increase them.

```
net.core.wmem_max=4194304
net.core.rmem_max=4194304
net.ipv4.tcp_wmem=4096 87380 4194304
net.ipv4.tcp_rmem=4096 87380 4194304
net.ipv4.ip_local_port_range = 1024 65535
```

### *Task 3: Increase OS resource limit*

Double-check OS resource limit with "ulimit -a", increase the number of "open files" and "max user processes".

```
* hard nofile 1000000
* soft nofile 1000000
root hard nofile 1000000
root soft nofile 1000000
* soft nproc 1000000
* hard nproc 1000000
```

### *Task 4: Tune database parameters*

The following table lists the basic PostgreSQL performance tuning parameters for your reference. Make sure that you have tuned shared memory-related kernel parameters and tuned semaphores to accommodate PostgreSQL database parameters changes. For details, refer to https://www.postgresql.org/docs/9.5/static/kernel-resources.html.

| Parameter | Demo | Extra Small | Small | Medium | Large |
|---|---|---|---|---|---|
| effective_cache_size | N/A | 4 GB | 4 GB | 8 GB | 12 GB |
| maintenance_work_mem | N/A | 128 MB | 256 MB | 1 GB | 2 GB |
| max_connections | N/A | 300 | 500 | 1500 | 3000 |
| shared_buffers | N/A | 1 GB | 2 GB | 4 GB | 6 GB |
| work_mem | N/A | 4 MB | 8 MB | 16 MB | 16 MB |

### *Task 5: Enable hot add RAM and CPU when utilizing virtual servers*

Hot add RAM and CPU is very useful and necessary for customers who are utilizing virtual servers, it enables users to extend servers without shutdown time. Therefore, we strongly suggest customers to enable this kind of feature.

The following example shows how to enable this feature when customers are utilizing VMware.

1. Log on to data center through vSphere client.
2. Shut down / power off the servers.
3. Right-click the server name, and click Edit Setting.
4. Go to tab - Options.
5. Under Advanced, select Memory/CPU Hotplug, enable hog add RAM and CPU.

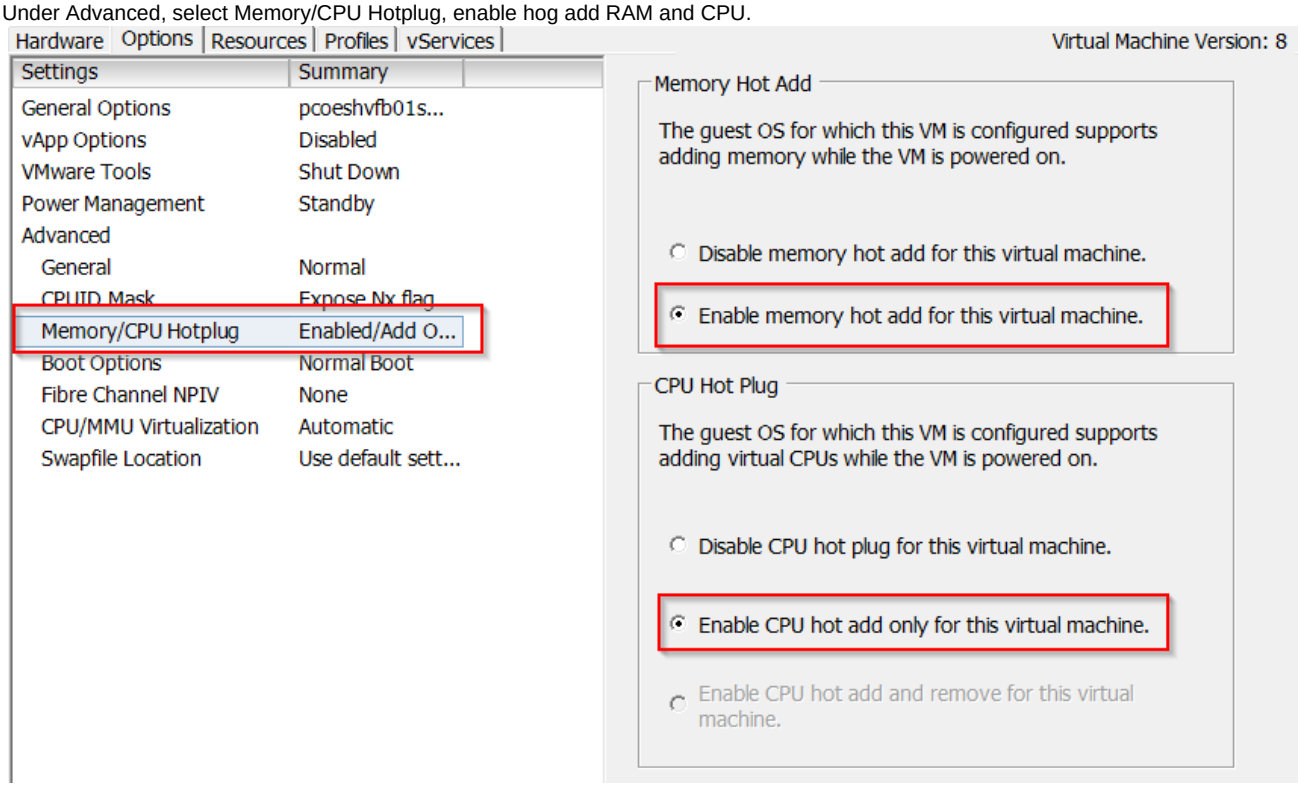


6. Save the changes.

### ITSMA installation

After you perform tuning before ITSMA installation, you can start to install ITSMA according to your corporation sizing. See the ITSMA installation section for instructions.

We recommand that you choose the sizing configuration that is closest to your corporation sizing need. For example, if your cooperate sizing requries 600 current users, you can install ITSMA with the Medium sizing profile, and then scale out the system accordingly. See Scaling up and out for how to scale out your system.

### Tuning after ITSMA installation

#### *Task 1: Browser SSL certificate*

For optimal browser performance, Certification Authority signing SSL certificate must be trusted. After doing this, a green lock should appear in the browser address box, and browser caching resources will work as expected. This is the most important especially when web users are working under some network latency.

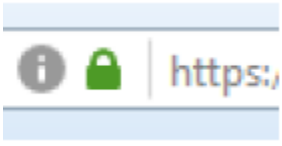


#### *Task 2: Enable TCP socket reuse / recycle*

To avoid lots of "TIME_WAIT" TCP socket connections above medium size deployment which may produce 502/504 errors to the front end, customers may need to further tune TCP socket behaviors from the OS level.

First, the safest way is to try to reuse "TIME_WAIT" socket by applying the following changes.

```
net.ipv4.tcp_tw_reuse=1
```

Alternatively, as an aggressive method, customers can enable faster TCP socket recycling by applying the following changes.

```
net.ipv4.tcp_tw_recycle=1
```

### *Task 3: Increase ITSMA IDM JVM heap size*

For better throughput and performance, customers can increase JVM heap size limit (by default, 1300 MB is provided from OOB) on IDM pods, this is especially useful for medium and bigger sizes.

1. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Go to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
3. Find pod with the name **idm**, and click View/edit YAML.

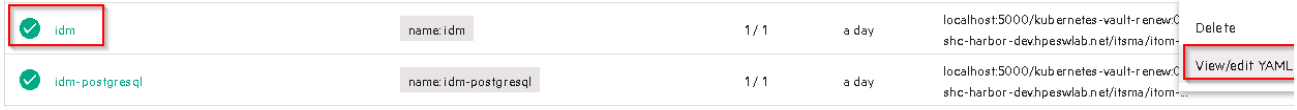


4. Search **jvm**, empty the value of JVM_MAX_SIZE and JVM_MIN_SIZE, type in target JVM heap size value, and click Update.

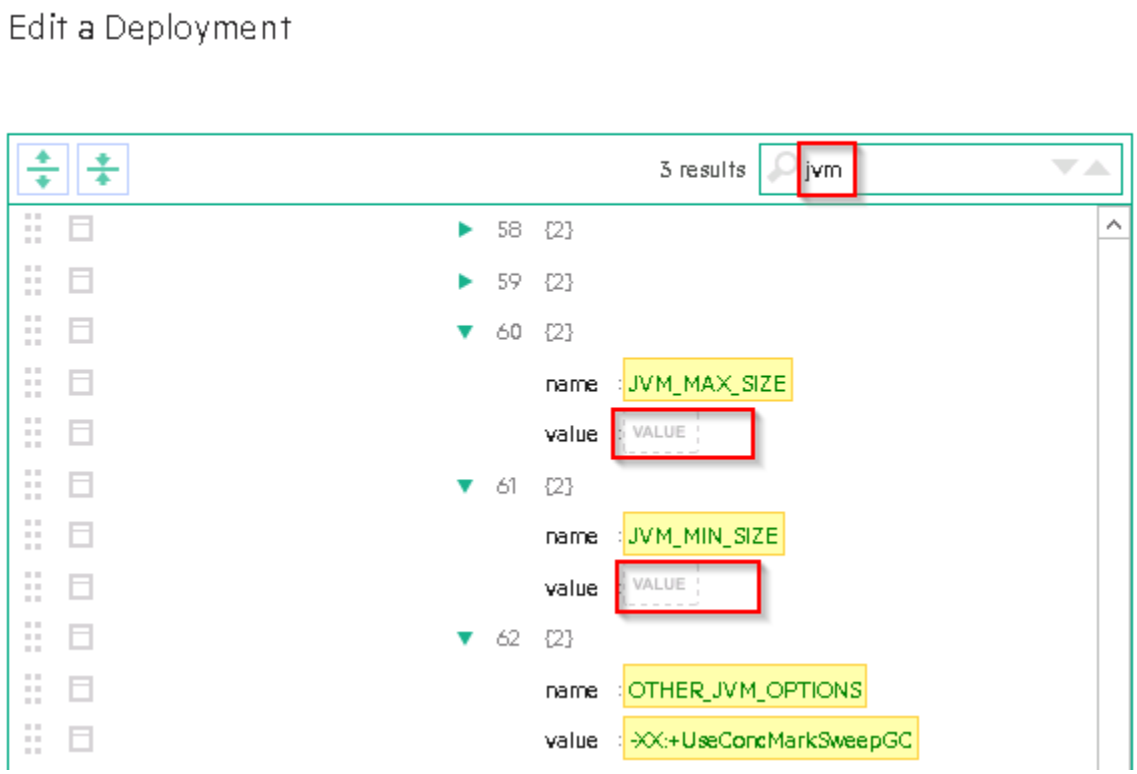


5. Wait for several minutes for the idm pod restarting.

### *Task 4: Configure additional SM_RTE_ARGS parameters*

Customers need to add some parameters manually after ITSMA installation. To do this, follow these steps:

1. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Go to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
3. Find pod with the name **sm-rte**, and click View/edit YAML.

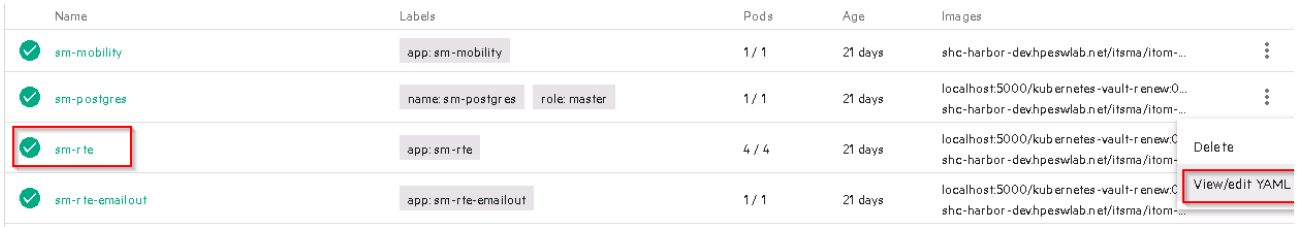


4. Add additional parameter "-ldapstats:0 -usedmemcompmode:2" in SM_RTE_ARGS, click Update.


5. Wait several minutes for the SM RTE pod to restart.
6. Find pod with the name **sm-rte-integration**, and click View/edit YAML.

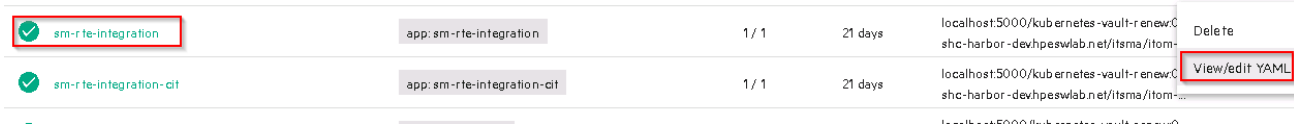


7. Add additional parameters "-ldapstats:0 -usedmemcompmode:2 -webservices_sessiontimeout:30" in SM_RTE_ARGS, click Update.


8. Wait several minutes for the SM RTE integration pod to restart.

### *Task 5: Remove internal PostgreSQL pods manually when utilizing external databases*

After customers choose to utilize external databases for the ITSMA installation, internal PostgreSQL pods may still exist there. Customers should remove them safely to release hardware resources. The following example show how to remove Service Manager PostgreSQL pod from ITSMA management console.

**ITSMA internal DB pod list**

```
Here is the internal DB pod name list:
1. idm-postgresql
2. itom-xservices-postgres
3. postgresql-ucmdb
4. propel-postgresql
5. sm-postgres
```

1. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Go to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
3. Find pod with the name **sm-postgres**, and click View/edit YAML.

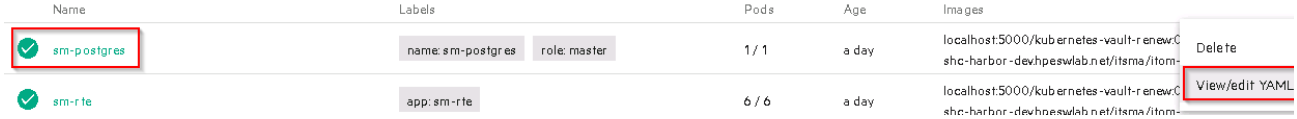


4. Set **replicas** to 0, and click update.

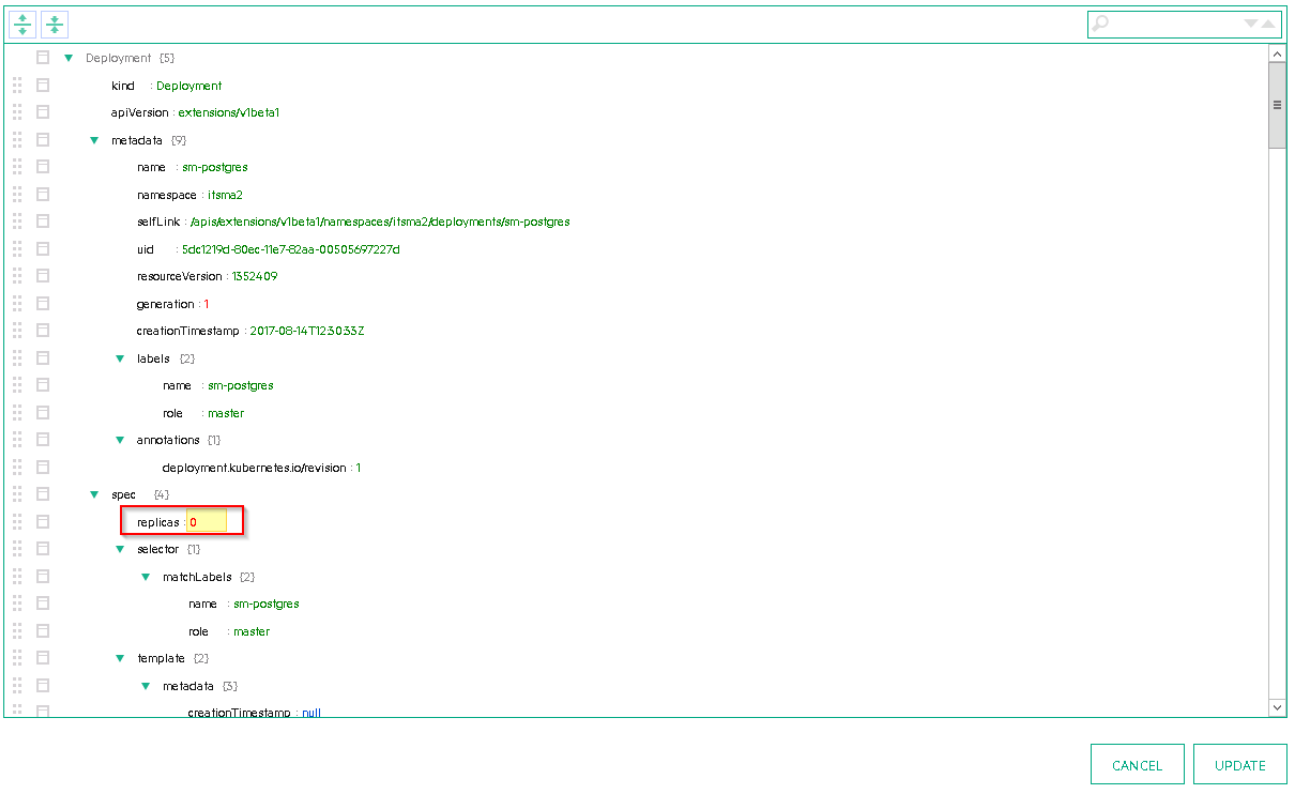


5. The Service Manager PostgreSQL pod will be removed.

## Scaling up and out



### Introduction

To accommodate the growth of users workload and concurrency, customers often need to extend their pods or hardware resource during daily operations. ITSMA provides very flexible and efficient ways to accomplish this purpose. Sometimes, customers may need to shrink the pods or work nodes resource because they un-checked some services provided by ITSMA.

The following scenarios are listed for reference and may need customers' scaling changes:

1. ITSMA installation process often requests more resources than actual usage after deployment, customers can choose to shrink hardware resource by removing worker nodes. This is especially true for extra small and small size deployment.
2. Customers may only need partial services provided by ITSMA instead of most or all of them after deployment, this way, customers can choose to remove relating pods / services from ITSMA, and remove worker nodes if applicable.
3. Customers may find some pods running slowly and always exhausting resources from Out-Of-Box limit configurations. This is especially true for common services / pods like IDM, Nginx, RabbitMQ, and so on. This way, customers can choose to scale up pods by assigning more CPU / Memory resource limit or increasing pods number.
4. With the growth of users load or workload, customers always find that the worker nodes' resource utilization is approaching 80% or more, this way, customers need to plan more capacity by adding more worker nodes.
5. Some pods / services are resource intensive, and need isolation. For example, customers choose internal database instead external for

their deployment. This way, they may need to hook these pods / services to dedicated worker nodes.

## Scaling master / worker nodes

### *Task 1: Remove worker nodes*

> **NOTES**
> 1. Remove worker nodes will result in some services downtime, implement the changes only in maintenance window in production.
> 2. Make sure there are enough vacant CPU / Memory resources to host pods / services migrated from reclaimed worker node.

The following detailed steps can be used to remove worker node server safely through management console:

1. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Go to ADMINISTRATION -> Nodes.
3. Choose target worker node, and click [-] to remove the labels on the node.

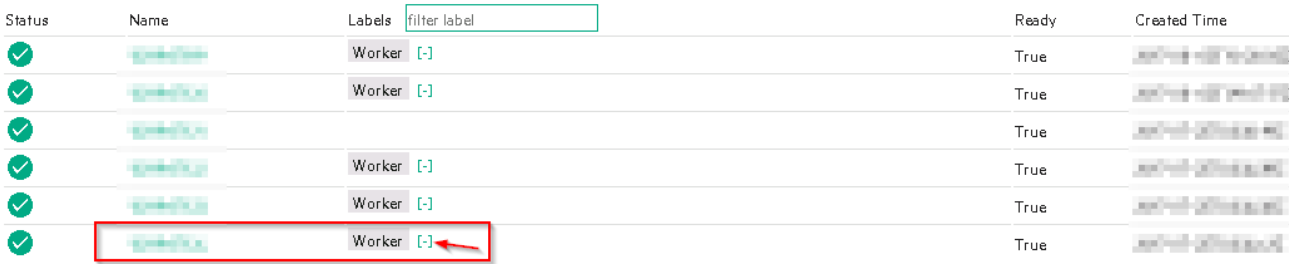


4. Wait several minutes because ITSMA needs to migrate pods / services to other worker nodes.
5. Reclaim the worker node server.

### *Task 2: Scale worker nodes out*

Once customers decide to extend capacity by adding more worker node server, they can use the following steps:

1. Prepare a new server which has the same configuration and size as existing worker node servers.
2. Add worker node through ITSMA management console. For details, please refer to how to add worker nodes.

### *Task 3: Scale master / worker nodes up*

Occasionally, customers may only need to add more CPU / RAM resource to the existing Master / worker nodes instead of scaling out by adding more servers. We do not suggest this kind of change especially when customers are utilizing physical machines.

For virtual servers, after enabling hot add RAM and CPU as mentioned in Task 5 - Enable hot add RAM and CPU when utilizing virtual servers, adding CPU / RAM become very easy without any downtime.

The following example shows how this is done with VMware.

1. Log on to data center through vSphere client.
2. Right-click the master or worker node server name, and click Edit Setting.
3. Add more memory.

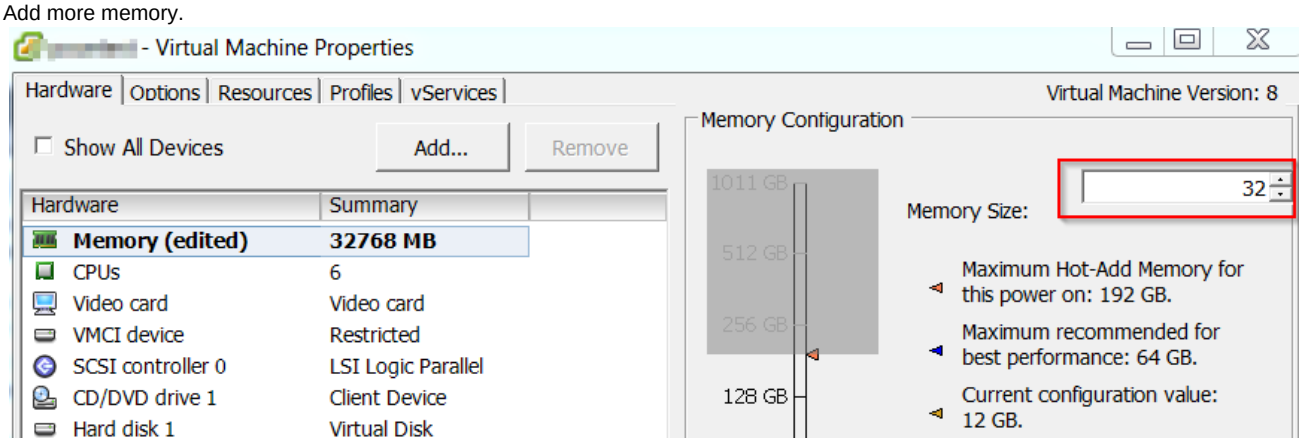


4. Add more CPU cores.

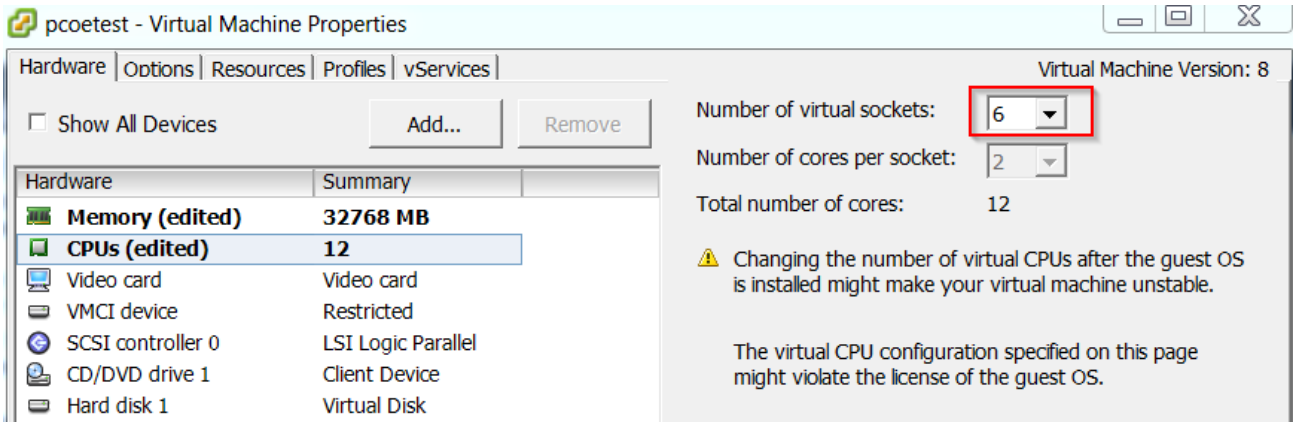


5. Save the changes.

## Scaling pods up and down

### *Task 1: Increase CPU / Memory resource limit*

CPU / Memory resource limit from Out-Of-Box configurations cannot cover every customer's scenarios. So during daily operations and maintenance, increasing CPU / Memory limit becomes very common for them. In ITSMA, implementing such kind of changes are very easy through ITSMA management console.

The following example shows how to change Propel RabbitMQ pod CPU / Memory limit:

1. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Go to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
3. Find pod with name **propel-rabbitmq**, click View /edit YAML.

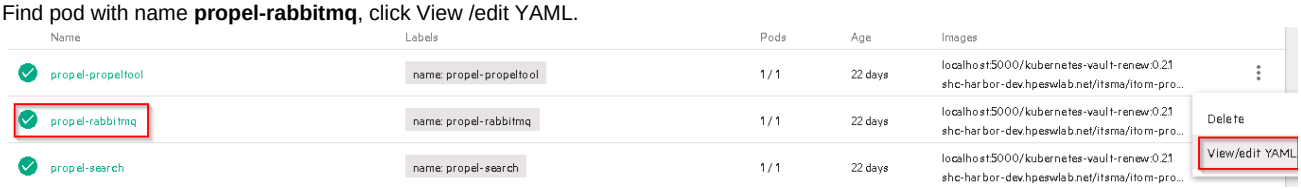


4. Increase the CPU / Memory resources limit.

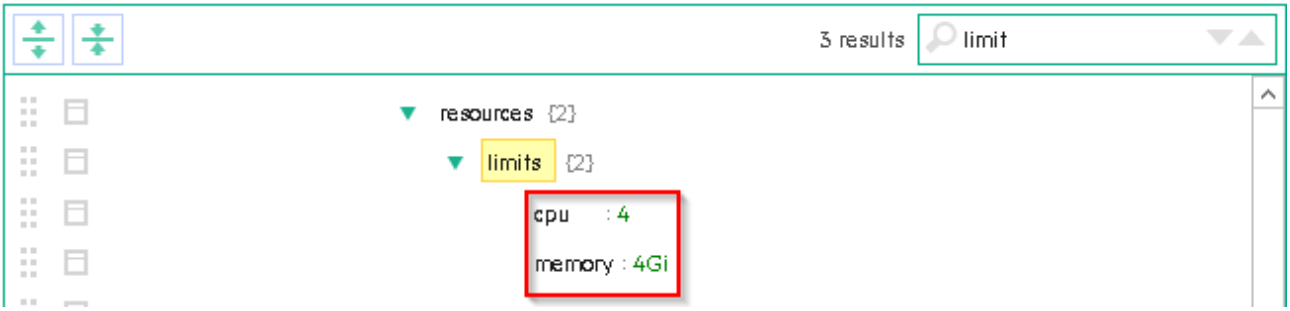


5. Save the changes, and wait several seconds for pod restarting.

### *Task 2: Change pods number*

Changes pods number is another way to extend or shrink services capability, it is highly depending on the actual business load requirements. Customers can achieve this purpose by updating replicas number of pods through ITSMA management console.

1. Login ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
2. Navigate to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
3. Find pod with name **sm-rte**, click View /edit YAML.

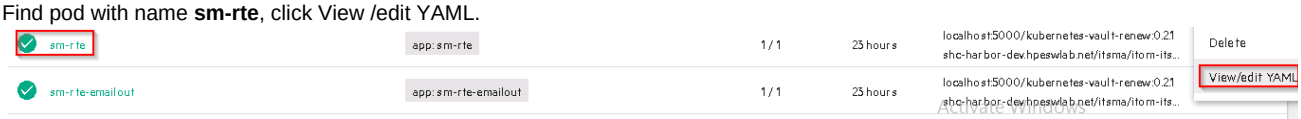


4. Increase the number of replicas.

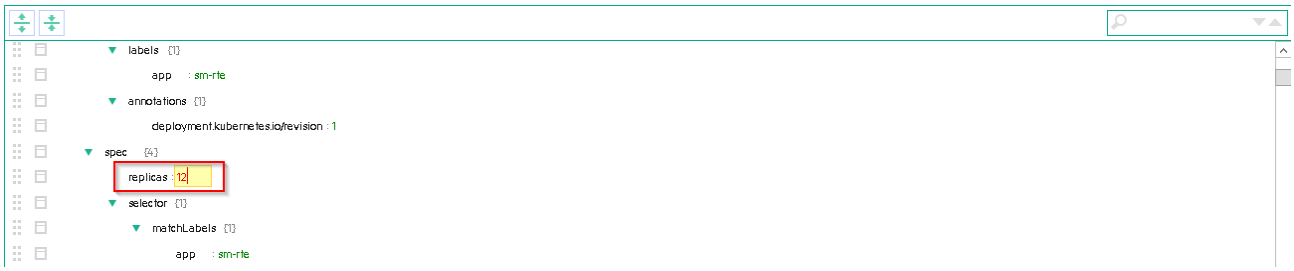


5. Save the changes, and wait for several minutes for the new pods up and running.

## Scaling rules

### *Rule 1: ITSMA worker node capacity*

Customers can extend system capacity by adding more worker node into ITSMA environment.

The following table lists some rough numbers for estimation, it is based on servers with 8 CPU cores / 32 GB RAM configuration.

| Users Type | Increased Capacity |
|---|---|
| ITSMA portal users | 400 ~ 600 concurrent users |
| ITSMA IT agents | 200 ~ 350 concurrent users |

### *Rule 2: ITSMA Service Manager RTE pod*

ITSMA Service Manager RTE pod is focusing on fulfilling IT agent users' requests. For each new 50 IT agent users, customers need to add one RTE pod accordingly.

| Users Type | Concurrent Users | CPU Usage | Memory Usage |
|---|---|---|---|
| ITSMA IT agents | 50 | 1 ~ 2 CPU cores | 4 GB |

### *Rule 3: ITSMA Service Manager RTE integration pod*

ITSMA Service Manager RTE integration pod is focusing on fulfilling integration users in which most users are from portal (other users like Chat, uCMDB services also utilize integration pods).

For each new 200 portal users, customers need to add one RTE integration pod accordingly.

| Users Type | Concurrent Users | CPU Usage | Memory Usage |
|---|---|---|---|
| ITSMA portal users | 200 | 1 ~ 2 CPU cores | 4 GB |

### *Rule 4: ITSMA Service Manager webtier pod*

| Users Type | Concurrent Users | CPU Usage | Memory Usage |
|---|---|---|---|
| ITSMA IT agents | 200 | 1 ~ 2 CPU cores | 2 GB |

### *Rule 5: ITSMA Smart Analytics pod*

| SMA component | Records number | CPU Usage | Memory Usage |
|---|---|---|---|
| SMA Content | 2 million | 2 ~ 4 CPU cores | 4 GB |

## Hook pods / services to dedicated worker nodes

Taking openldap pod as an example, the following steps shows how to put pods / services on dedicated worker nodes:

1. Add a new label to the dedicated worker node (in production, customers may need to prepare a new worker node first)
    a. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
    b. Go to ADMINISTRATION -> Nodes

c. In the 'Predefined Labels' section, type your new lab name and click '[+]' button, for example, an 'OpenLDAP' label to be added.

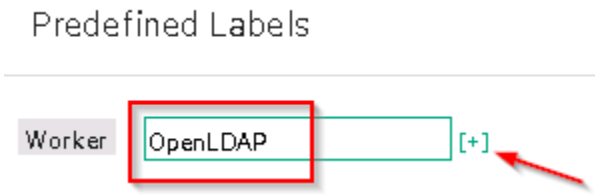


d. Choose the target worker node, click [-] to remove the Worker label of the dedicated worker node.

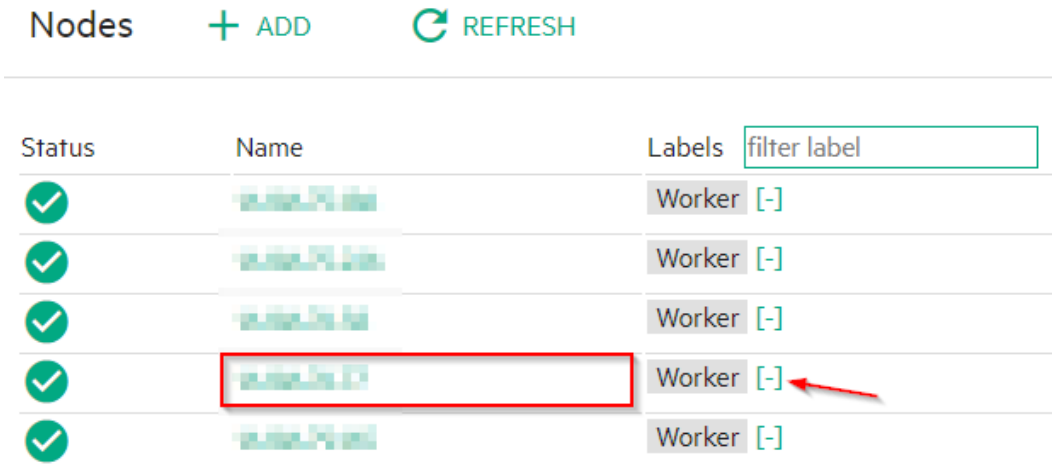


e. Drag the new created label in step 3 to the dedicated worker node.

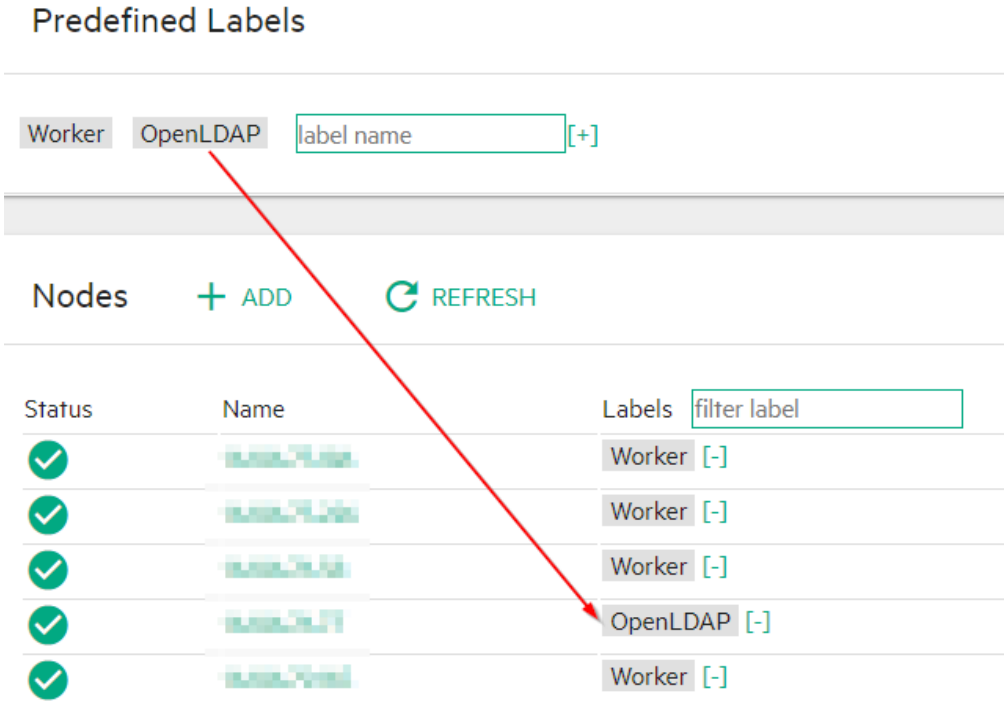


2. Update YAML through ITSMA managements console.
   a. Log on to ITSMA managements console with admin (https://<ITSMA_Master_Node>:5443).
   b. Go to Resources -> change Namespace to your ITSMA -> Expand Workloads -> click Deployments.
   c. Find the pod with the name **openldap-rc**, and click View/edit YAML.

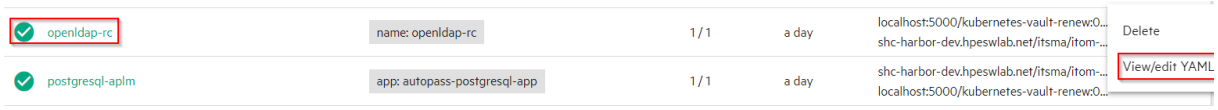


   d. Set **replicas** to 0.
   e. Search nodeSelector, Change the 'Worker' to the new label 'OpenLDAP', click update.

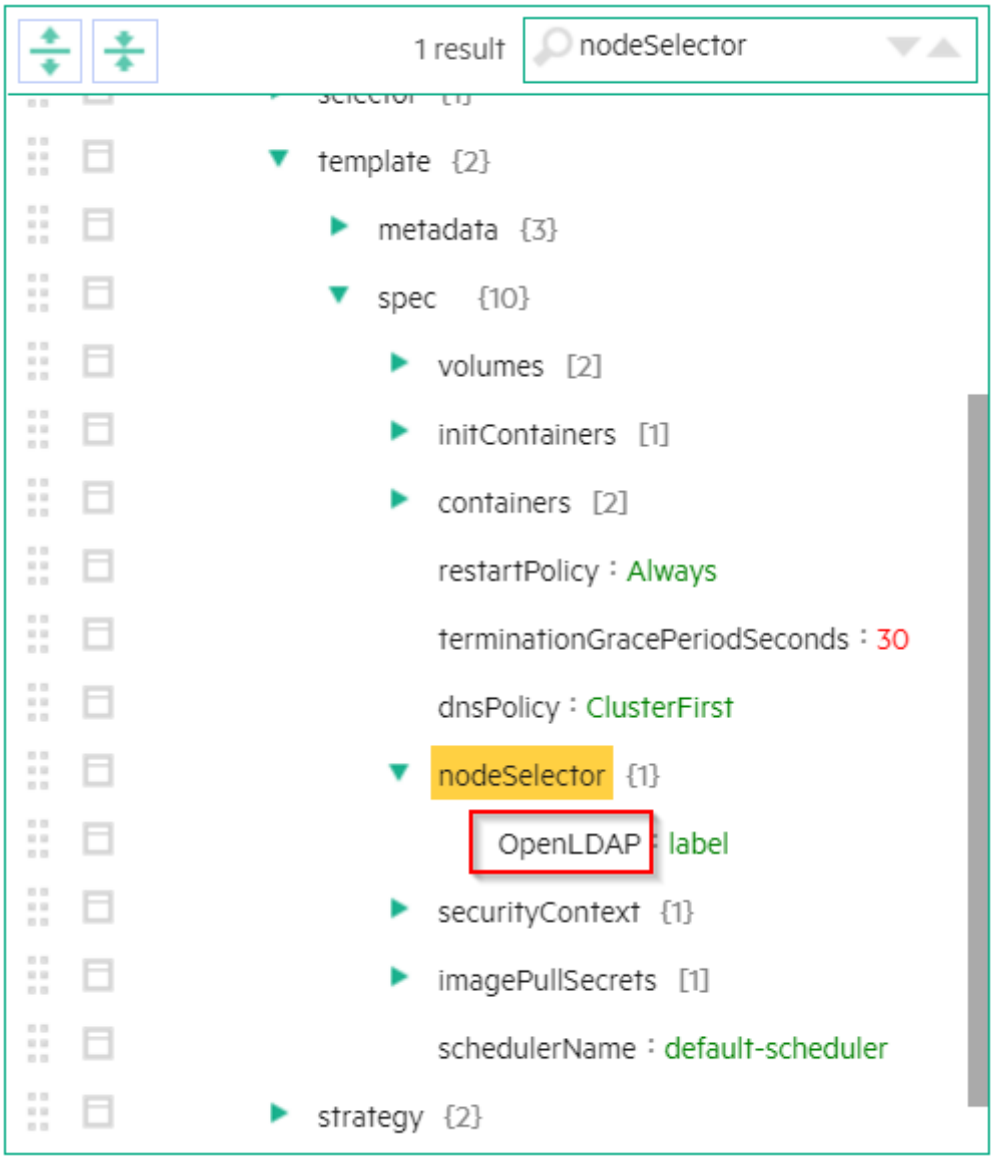


f. Find the pod with the name **openldap-rc**, and click View/edit YAML.
g. Set **replicas** to 1, and click Update.
h. Waiting several minutes for the openldap pod to start.

## 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 ITOM CDF configuration file (**install.properties**). Parameters in this file are described in the following table. Review the parameters and decide on which configurations are needed for your business needs.

> The database configuration parameters (DEFAULT_DB_TYPE, DEFAULT_DB_HOST, DEFAULT_DB_PORT, DEFAULT_DB_NAME, DEFAULT_DB_CONNECTION_URL) are used to specify a database for your Kubernetes cluster, not for the ITSMA suite. You will need to specify the suite databases during the suite installation (see Run the Suite Installer).

| Parameter | Description | Notes |
|---|---|---|

| | | |
|---|---|---|
| *MASTER_NODES* | Lists the cluster master nodes (IPv4 address or FQDN), separated by a blank and enclosed in double-quotes. Currently, CDF supports deployment with either one or three master nodes.<br><br>**Example:**<br><br>`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. If you want to use a master node as a worker node also, enter its IPv4 address or FQDN in the *WORKER_NODES* parameter.<br><br>Typically, a worker node runs the workload when you deploy a suite. By default, when you install a suite, you target a worker node.<br><br>**Example:**<br><br>`WORKER_NODES="10.10.10.20 10.10.10.21 10.10.10.22"` | Mandatory |
| *EXTERNAL_ACCESS_HOST* | Defines a fully-qualified hostname for external clients to access cluster services. The specified name must resolve to the IP address where the ingress is running.<br><br>Everything that runs on a cluster is actually on a private network and is not externally accessible. If you want to make any suite functionality available from outside the network (for example, a Help Desk operative on a client machine on another network), you must provide an ingress into the cluster to be able to access the functionality. This is done by configuring the *EXTERNAL_ACCESS_HOST* parameter.<br><br>If you are deploying multiple master nodes, you must set the *EXTERNAL_ACCESS_HOST* parameter to a fully-qualified domain name that resolves to the value of the *HA_VIRTUAL_IP* parameter.<br><br>**Example:**<br><br>`EXTERNAL_ACCESS_HOST=myd.XXXX.YYY.net` | Mandatory |
| *NFS_SERVER* | Specifies the IPv4 address or FQDN of the NFS server that serves the persistent volumes of the cluster services.<br><br>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 or databases, they must be located outside the container in a persistent volume provided by a Network File System (NFS). When you install CDF, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server.<br><br>**Example:**<br><br>`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.<br><br>When a master node is installed, it generates a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.<br><br>**Example:**<br><br>`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.<br><br>When a master node is installed, it generates a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.<br><br>**Example:**<br><br>`CLIENT_CERT_FILE=/tmp/client.crt` | Mandatory only for worker nodes |
| *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.<br><br>When a master node is installed, it generates a number of certificates and keys that are required when you install other master or worker nodes. You must copy these files and specify them in the install.properties file.<br><br>**Example:**<br><br>`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.<br><br>The VIP and each pool member must exist in the same sub-net. Since the VIP does not correspond to an actual physical network interface, users do not need to make any configuration. They only need to provide a virtual IP address and make sure that the IP address is not occupied before the installation. The requests to the API server should be sent to a VIP for higher availability.<br><br>**Example:**<br><br>`MASTER_NODES="10.10.10.10 10.10.10.11 10.10.10.12"`<br><br>`HA_VIRTUAL_IP=10.10.10.9` | Mandatory only if you are using multiple master nodes |

| | | |
|---|---|---|
| *PEER_CA_FILE* | Specifies the CA certificate for TLS authentication to the master nodes. The value of the parameter is the file name of the CA certificate, including the absolute path.<br><br>**Example:**<br><br>`PEER_CA_FILE=/tmp/ca/crt` | Mandatory only if you are using multiple master nodes |
| *PEER_CERT_FILE* | Specifies the certificate for TLS authentication to the master nodes. The value of the parameter is the file name of the certificate, including the absolute path.<br><br>**Example:**<br><br>`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.<br><br>**Example:**<br><br>`PEER_KEY_FILE=/tmp/server.key` | Mandatory only if you are using multiple master nodes |
| *NFS_FOLDER* | Specifies the root folder (fully-qualified directory) for the persistent volume that the NFS server exports.<br><br>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 CDF, you must install an NFS server that shares out the network volumes. The server can be a master node or an external server.<br><br>**Example:**<br><br>`NFS_FOLDER=/var/vols/itom/core` | 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.<br><br>When you install the CDF, 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.<br><br>**Example:**<br><br>`ROOTCA=/tmp/ca.crt` | Optional |

| | | |
|---|---|---|
| *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.<br><br>When you install the CDF, 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.<br><br>**Example:**<br><br>`ROOTCAKEY=/tmp/ca.key` | Optional |
| *NFS_STORAGE_SIZE* | Specifies the size of the NFS volume exported by the NFS server.<br><br>**Example:**<br><br>`NFS_STORAGE_SIZE=50 Gi` | Optional |
| *K8S_HOME* | Specifies the installation directory (fully-qualified directory) for the core platform binaries.<br><br>**Example:**<br><br>`K8S_HOME=/opt/kubernetes` | Optional |
| *MASTER_API_PORT* | Specifies the HTTP port for the Kubernetes (K8S) API server.<br><br>If you want to use K8S, you must dock to the K8S API server. The **kubectl** command line tool communicates with the K8S server.<br><br>**Example:**<br><br>`MASTER_API_PORT=8080` | Optional |
| *MASTER_API_SSL_PORT* | Specifies the HTTPS port for the K8S API server.<br><br>If you want to use K8S, you must dock to the K8S API server. The **kubectl** command line tool communicates with the K8S server.<br><br>**Example:**<br><br>`MASTER_API_SSL_PORT=8443` | Optional |

| | | |
|---|---|---|
| *THINPOOL_DEVICE* | Specifies the path to the Docker devicemapper storage driver. For more information, see the Docker documentation.<br><br>If this parameter is specified, the installation will use the devicemapper (direct-lvm) Docker storage driver. If it is not specified, the installation will use devicemapper(loop).<br><br>For production environments, HPE recommends using devicemapper(direct-lvm). For more information, see Meet the prerequisites (on-premises).<br><br>**Example:**<br><br>`THINPOOL_DEVICE= /dev/mapper/docker-thinpool` | Optional |
| *DEFAULT_DB_TYPE* | Specifies the type of database that CDF will use for the service connection.<br><br>• The `DEFAULT_DB_TYPE` value can be `EMBEDDED`, `EXTERNAL_PG`, and `EXTERNAL_ORA`. The values must be capitalized. By default, this parameter is set to EMBEDDED, which means CDF will use the embedded, containerized PostgreSQL database for installation.<br>• If `DEFAULT_DB_TYPE` is set to `EXTERNAL_PG` or `EXTERNAL_ORA`, the `DEFAULT_DB_HOST`, `DEFAULT_DB_PORT`, and `DEFAULT_DB_NAME` or the `DEFAULT_DB_CONNECTION_URL` must be specified during the CDF installation.<br>• If `EXTERNAL_ORA` is specified as the database type, you must enter the external default database `DEFAULT_DB_SCHEMA` during the CDF installation.<br><br>`EXTERNAL_ORA` is not supported in the AWS deployment mode.<br><br>**Example:**<br><br>`DEFAULT_DB_TYPE=EXTERNAL_PG` | Optional |
| *DEFAULT_DB_HOST* | Specifies the database host when the `DEFAULT_DB_TYPE` parameter is set to `EXTERNAL_PG` or `EXTERNAL_ORA`. You can enter the FQDN or the IP address of the database host.<br><br>**Example:**<br><br>`DEFAULT_DB_HOST=10.10.10.10` | Mandatory only if you choose `EXTERNAL_PG` or `EXTERNAL_ORA` for the `DEFAULT_DB_TYPE` parameter and the `DEFAULT_DB_CONNECTION_URL` parameter is not specified |
| *DEFAULT_DB_PORT* | Specifies the database port when choosing `EXTERNAL_PG` or `EXTERNAL_ORA` as the `DEFAULT_DB_TYPE`.<br><br>**Example:**<br><br>`DEFAULT_DB_PORT=5432` | Mandatory only if you choose `EXTERNAL_PG` or `EXTERNAL_ORA` for the `DEFAULT_DB_TYPE` parameter and the `DEFAULT_DB_CONNECTION_URL` parameter is not specified |

| | | |
|---|---|---|
| *DEFAULT_DB_NAME* | Specifies the database name when choosing EXTERNAL_PG or EXTERNAL_ORA as the DEFAULT_DB_TYPE.<br><br>**Example:**<br><br>DEFAULT_DB_NAME=CDFDB | Mandatory only if you choose EXTERNAL_PG or EXTERNAL_ORA for the DEFAULT_DB_TYPE parameter and the DEFAULT_DB_CONNECTION_URL parameter is not specified |
| *DEFAULT_DB_CONNECTION_URL* | Specifies the database connection URL when choosing the EXTERNAL_PG or EXTERNAL_ORA as the DEFAULT_DB_TYPE. This parameter is left empty when the DEFAULT_DB_HOST, DEFAULT_DB_PORT, and DEFAULT_DB_NAME are specified.<br><br>**Example:**<br><br>DEFAULT_DB_CONNECTION_URL=jdbc:oracle:thin:@IP:port:db_name | Mandatory only if you choose EXTERNAL_PG or EXTERNAL_ORA for the DEFAULT_DB_TYPE parameter and the DEFAULT_DB_HOST, DEFAULT_DB_PORT, and DEFAULT_DB_NAME are not specified. |
| *DOCKER_HTTP_PROXY*<br><br>*DOCKER_HTTPS_PROXY* | Specifies the proxy settings for Docker. Configure this parameter if access to the Docker hub or registry requires a proxy (by default, no proxy is used). The value of the parameter must be any valid HTTP proxy URL.<br><br>Specify these settings if you install suites and launch containers on Docker inside the Kubernetes cluster, or if you download images from the internet.<br><br>**Examples:**<br><br>DOCKER_HTTP_PROXY="<Your Proxy>"<br><br>DOCKER_HTTPS_PROXY="<Your Proxy>" | Optional |
| *DOCKER_NO_PROXY* | Specifies the IPv4 address or FQDN that does not need the proxy settings for Docker.<br><br>Example:<br><br>DOCKER_NO_PROXY=127.0.0.1 | Optional |
| *REGISTRY_ORGNAME* | Specifies the organization name (in string format) where the suite images are placed. The default name is hpeswitom.<br><br>**Example:**<br><br>REGISTRY_ORGNAME=hpeswitom | Optional |
| *FLANNEL_IFACE* | Specifies the interface for Docker inter-host communication as a single IPv4 address or interface name. This parameter is used when the nodes have more than one network adapter.<br><br>**Example:**<br><br>FLANNEL_IFACE=10.10.10.10 | Mandatory only if you install CDF on a node that has more than one network adapter. |
| *CLOUD_PROVIDER* | Specifies the cloud provider when installing CDF on a cloud server.<br><br>**Example:**<br><br>CLOUD_PROVIDER=AWS | Optional |

| | | |
|---|---|---|
| *AWS_REGION* | Specifies the AWS region to use when choosing AWS as the cloud provider. The default value of this parameter is an empty string.<br><br>**Example:**<br><br>`AWS_REGION=us-east-1` | Mandatory only if you choose AWS as the cloud provider when installing CDF on a Cloud server. |

# Prepare for installation (on-premises)

Before you can install ITOM Container Deployment Foundation (CDF) and ITSMA in an on-premises environment, perform the following tasks to prepare your environment.

- Enable your Docker Hub account
- Meet the prerequisites (on-premises)
- (Optional) Prepare databases for CDF and ITSMA (on-premises)
- (Optional) Set up Access Server for a DMZ network
- Download the CDF installation package (on-premises)

## Enable your Docker Hub account

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

To enable your Docker Hub account, follow these steps:

1. Create a Docker Hub account if you do not already have one:
    a. Go to https://hub.docker.com.
    b. Enter a Docker ID, your company email address, and then a password.
    c. Click **Sign Up**.
    d. You receive an email from Docker Hub, asking you to confirm your email address. Confirm your email address swiftly.
2. Log in to https://hub.docker.com with your Docker ID.
3. On the top right corner of the page, click **Settings** under your avatar and take a screenshot to include your Docker ID and the linked email address.
4. Send the following information together with the screenshot to the HPE software fulfillment and licensing team specific for your region to enable your Docker Hub account:
    - Your company name
    - Your HPE Customer SAID (must be valid and active)
    - HPE ITOM Suite edition (that is, ITSMA)

> Email addresses of the HPE software fulfillment and licensing teams for different regions:
>
> Americas region: dockersupport.ams@hpe.com
> APJ region: dockersupport.apj@hpe.com
> EMEA region: dockersupport.emea@hpe.com
>
> Once your Docker ID is enabled, you will receive a confirmation from HPE.

## Meet the prerequisites (on-premises)

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

- Meet the basic prerequisites
- Meet additional prerequisites (production only)
    - Task 1: Prepare logical volumes for the cluster nodes
    - Task 2: Add or extend a logical volume for a direct-lvm thin pool
        - Step 1: Add or extend a volume
        - Step 2: Prepare the volume
        - Step 3: Specify the THINPOOL_DEVICE parameter in install.properties

### Meet the basic prerequisites

Before the installation, be sure to meet the following basic prerequisites.

1. Use the **root** user or a user with sudo access to install ITOM CDF.
2. Make sure that the following ports, which are needed during the ITOM CDF installation, are not in use on the host VM:  2380, 4001, 5000, 5443, 8080, 8200, 8443, 10250, 10251, 10252, 10255.
3. Remove any NFS share folder if you installed ITOM CDF previously.
   For example, run the following command: **rm -rf /var/vols/itom/core**
4. Configure time synchronization on all of the hosts in the cluster with a tool. For example, Chrony or NTP.

   Chrony is installed by default on some versions of Red Hat/CentOS. However, if Chrony is not installed or not running on your system, do the following:
   a. Install Chrony with the following command:
      **# yum install chrony**
   b. Start Chronyd:
      **# systemctl start chronyd**
      **# systemctl enable chronyd**
   c. Verify that Chrony is operating correctly:
      **# chronyc tracking**
5. Make sure that all servers (the NFS servers, master nodes, and work nodes) are in the same subnet.
6. Make sure that the **/tmp** folder of the target system has enough free space (at least 2.5G), which is required for adding worker nodes from the CDF user interface (the "Management Portal"). See Install CDF on the worker nodes.
7. Install the following packages on the relevant hosts listed in the following table, by using the following command: **yum install [package name]**.

| Package | Host |
|---|---|
| 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 servers |
| systemd-libs (version >= 219) | master, workers |
| unzip | master, workers |

8. Make sure that each node has a static IP address.
9. On each of the master and worker nodes, add the following IP addresses to the `NO_PROXY` environment setting: the IP addresses or host names of the master nodes, which you will set in the `MASTER_NODES` parameter in the install.properties file.
10. Update the **/etc/hosts** file of each cluster node by adding a new line containing the IP address and fully qualified domain name of the node:
    ***<Node IP>  <Node FQDN>***
11. Disable SELinux on each cluster node.

## Meet additional prerequisites (production only)

Perform the following tasks to meet additional prerequisites, which are needed only for a production environment.

### Task 1: 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 CDF installation. 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 S upport 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 CDF by running the following command:

   **# mount *[logical volume path] [platform installation folder]***
   For example:
   **# mount /dev/core-platform/mylv /opt/kubernetes**

> The [platform installation folder] value must be the directory that you will configure in the **K8S_HOME** parameter in the install.properties file (for example: /opt/kubernetes). For more information, see Configure the install.properties file.

### Task 2: Add or extend a logical volume for a direct-lvm thin pool

In a production environment only, you need to prepare a data volume for direct-lvm thin pool 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 the following steps on all of the master and worker nodes in the cluster.

The CDF installation configuration file (install.properties) provides the **THINPOOL_DEVICE** parameter for this purpose (see CDF installation configuration).

#### *Step 1: Add or extend 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 new volume to the host server or extend an existing volume, and then restart the server. Usually, the thin pool for each standard ITSMA host server (master or worker) requires a disk space of 100 GB. For more information about hardware requirements for ITSMA, see Sizing.
   The following screenshot shows an example of extending an existing volume by using vSphere.

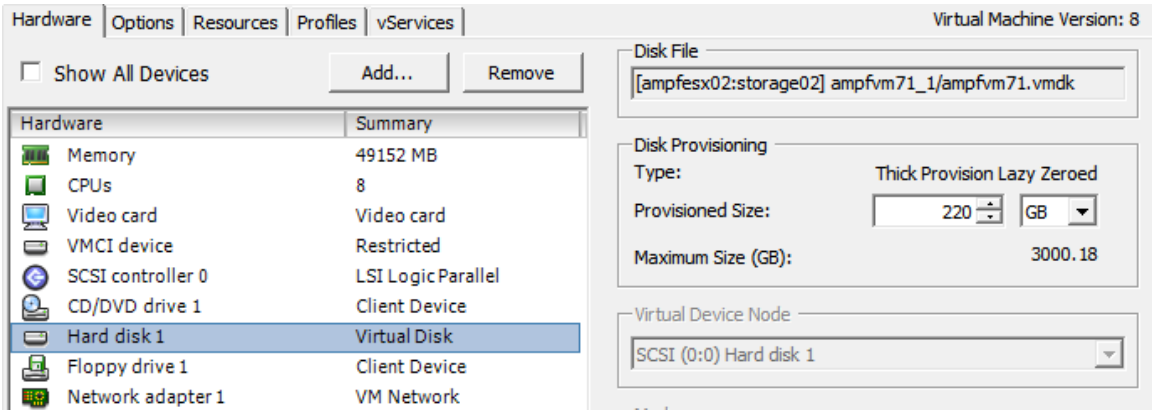


3. Type "fdisk -l" to check the partitions on your disk again. You should see more disk size is available.
4. Use "fdisk" to allocate the volume that you added or extended:
   a. Type "fdisk /dev/sda".
   b. Type "n" to create a new partition.
   c. Type "p" to use the primary.
   d. Choose the newly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

### *Step 2: Prepare the volume*

Prepare the new or extended 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
```

### *Step 3: Specify the THINPOOL_DEVICE parameter in install.properties*

In a production environment, before you set up the master nodes and worker nodes, specify the **THINPOOL_DEVICE** parameter in the install.properties file as follows (see Configure the install.properties file):

**THINPOOL_DEVICE=/dev/mapper/docker-thinpool**

For more information, see Tuning configuration.

## (Optional) Prepare databases for CDF and ITSMA (on-premises)

Database preparation is optional. Both ITOM Container Deployment Foundation (CDF) and ITSMA have an embedded PostgreSQL database. You can use their embedded PostgreSQL database for CDF and ITSMA; however, in a production environment, you may want to use your own databases for CDF and ITSMA.

For information about supported databases, see Support matrix (on-premises).

- Prepare a PostgreSQL or Oracle database for CDF
- Prepare databases for ITSMA
  - Prepare an Oracle database for ITSMA
  - Prepare one or more PostgreSQL database instances for ITSMA
    - Prepare PostgreSQL for Service Portal
    - Prepare PostgreSQL for IdM
    - Prepare PostgreSQL for Service Management

### Prepare a PostgreSQL or Oracle database for CDF

> If using an Oracle database, you will need to put a required Oracle JDBC Driver to the `HPESW_ITOM_Suite_Foundation_2017.06.00XXX/tools/drivers/jdbc` directory on the master node on which you will install CDF. See also Install CDF on the first master node.

You can prepare an external PostgreSQL or Oracle database for CDF installation.

1. Have the following information available, as you will need the information during installation:

| Item | Parameter in install.properties |
|---|---|
| Database server host and port | DEFAULT_DB_HOST<br>DEFAULT_DB_PORT |
| Database type | DEFAULT_DB_TYPE |
| Database name | DEFAULT_DB_NAME |

2. Have the database connection user name and password available. You will be prompted to enter this information during installation.

### Prepare databases for ITSMA

In ITSMA, the containerized Service Management, CMDB, and Service Portal capabilities, as well as the HPE Identity Provider (IdM) instance used by ITSMA require a database to work.

> - In a production environment, you are recommended to use external databases for ITSMA. The built-in PostgreSQL database in ITSMA is recommended for test environments. However, if you still want to use it in a production environment, see Meet the prerequisites (on-premises) for more information.
> - The containerized Service Management does not support case-insensitive PostgreSQL.

The following table lists the supported databases for each of these suite components and the configuration information required for them during the ITSMA installation.

| Component | PostgreSQL | Oracle | Configuration required during installation |
|---|---|---|---|
| Service Portal<br><br>IdM | | x | PostgreSQL for Linux:<br>• User Name<br>• Password<br>• PostgreSQL Server Host Name or IP Address<br>• PostgreSQL Server Port |

| | | | |
|---|---|---|---|
| Service Management | | | If PostgreSQL for Linux:<br>• User Name<br>• Password<br>• PostgreSQL Server Host Name or IP Address<br>• PostgreSQL Server Port<br>If Oracle:<br>• User Name<br>• Password<br>• Oracle server host name or IP<br>• Oracle server port<br>• Service Name |
| CMDB | x | | Oracle:<br>• User Name<br>• Password<br>• Oracle server host name or IP<br>• Oracle server port |

### Prepare an Oracle database for ITSMA

You can prepare an Oracle database to use for both Service Management and CMDB or to use for CMDB only. Later, you will be asked to enter the database connection settings when running the Suite Installer.

> Using an Oracle database requires that an Oracle JDBC Driver (ojdbc6.jar) be installed on the CDF NFS server; if the JDBC Driver is not found, the database connection test will fail during the suite installation. For details, see Run the Suite Installer.

To prepare your RDBMS, follow these steps:

1. Create an Oracle database.
   For the supported Oracle database version, see Support matrix (on-premises). For details about how to create an Oracle database, see the Oracle documentation.
2. If you will use this database for Service Management, 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. If you will use this database for CMDB, 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 one or more PostgreSQL database instances for ITSMA

- ITSMA supports PostgreSQL for Linux only.
- Make sure that the contrib package is installed on each PostgreSQL database server.
- The ITSMA suite does not support external PostgreSQL databases that have 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 ***<ITSMA database NFS share>*/propel** directory (for example, **/var/vols/itom/itsma/itsma-<namespace>/db/propel**) 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.

ITSMA allows you to use a separate database server for each component. You can create one or multiple PostgreSQL instances for ITSMA.

### *Prepare PostgreSQL for Service Portal*

Perform these steps:

1. Create a PostgreSQL database instance.
   For the supported PostgreSQL database version, see Support matrix (on-premises).
   For details about how to create a PostgreSQL database, see the PostgreSQL documentation.
2. In the postgresql.conf file, set the *max_connections* parameter to a value according to your suite size. For details, see Task 4: Tune database parameters.

   **Additional DB parameters**
   In a production environment, you need to tune additional parameters as described in Task 4: Tune database parameters.

3. Make sure that the PostgreSQL database uses the following encoding and locale settings:

   Encoding: UTF8
   LC_COLLATE: en_US.UTF-8
   LC_CTYPE: en_US.UTF-8
4. You are recommended to use the "postgres" user as the database user for Service Portal. Alternatively, you can create another user that has the same name as the database and has the create database/extension/user/schema privileges, and then set the database owner to this user.

   You will need to specify this user for Service Portal when configuring the database server during the suite installation. For details, see Run the Suite Installer.

5. Configure the pg_hba.conf file to allow the following users to connect to the following databases:

| Database | User |
|---|---|
| analytics | serviceportalbackend |
| bpmdb | serviceportalbackend |
| catalog | serviceportalbackend |
| dashboarddb | serviceportalbackend |
| jumpstart | serviceportalbackend |
| notificationdb | serviceportalbackend |
| sxdb | serviceportalbackend |
| xservices_ems | maas_admin |
| xservices_mng | maas_admin |
| xservices_rms | maas_admin |
| maas_admin | maas_admin |
| maas_template | maas_admin |

   For example:

   **host analytics serviceportalbackend 0.0.0.0/0 md5**

**host bpmdb serviceportalbackend 0.0.0.0/0 md5**
**host catalog serviceportalbackend 0.0.0.0/0 md5**
**host dashboarddb serviceportalbackend 0.0.0.0/0 md5**
**host jumpstart serviceportalbackend 0.0.0.0/0 md5**
**host sxdb serviceportalbackend 0.0.0.0/0 md5**
**host dashboarddb serviceportalbackend 0.0.0.0/0 md5**
**host xservices_ems maas_admin 0.0.0.0/0 md5**
**host xservices_mng maas_admin 0.0.0.0/0 md5**
**host xservices_rms maas_admin 0.0.0.0/0 md5**
**host maas_admin maas_admin 0.0.0.0/0 md5**
**host maas_template maas_admin 0.0.0.0/0 md5**

6. Restart the PostgreSQL database instance.

### *Prepare PostgreSQL for IdM*

Perform these steps:

1. Create a PostgreSQL database instance and configure the database to accept external connection requests over TCP/IP. For details, see the PostgreSQL documentation.
2. Make sure that the PostgreSQL database uses the following encoding and locale settings:

   ENCODING: UTF8
   LC_COLLATE: en_US.UTF-8
   LC_CTYPE: en_US.UTF-8
3. In the postgresql.conf file, set the *max_connections* parameter to a value according to your suite size. For details, see Task 4: Tune database parameters.

   > **Additional DB parameters**
   > In a production environment, you need to tune additional parameters as described in Task 4: Tune database parameters.

4. Configure the pg_hba.conf file to allow the "idm" user to connect to the "idm" database.
   For example, add the following line to this file:
   **host idm idm 0.0.0.0/0 md5**
5. Create a database user that has the following privileges:
   - Connect database
   - CRTEATE/DROP/UPDATE/QUERY Tables
   - CRATE/DROP/UPDATE PROCEDURE

   > You will need to specify this user for IdM when configuring the database server during the suite installation. For details, see Run the Suite Installer.

### *Prepare PostgreSQL for Service Management*

Perform these steps:

1. Create a PostgreSQL database.
2. Make sure that the PostgreSQL database uses the following encoding and locale settings:

   Encoding: UTF8
   LC_COLLATE: en_US.UTF-8
   LC_CTYPE: en_US.UTF-8
3. In the postgresql.conf file, set the *max_connections* parameter to a value according to your suite size. For details, see Task 4: Tune database parameters.

   > **Additional DB parameters**
   > In a production environment, you need to tune additional parameters as described in Task 4: Tune database parameters.

4. Configure the pg_hba.conf file to allow the "servicemanagement" user to connect to the "servicemanagement" database.
   For example, add the following line to this file:
   **host servicemanagement servicemanagement 0.0.0.0/0 md5**

   > When configuring the database server during your suite installation, you can specify the "postgres" user for Service Management. For details, see Run the Suite Installer.

## (Optional) Set up Access Server for a DMZ network

If your organization's network includes a demilitarized zone **(**DMZ), any access from the Internet must go through a DMZ device that acts as an reverse proxy. In this document, we assume that you use an Access Server as a DMZ device. The following diagram illustrates an ITSMA suite deployment in a corporate network with a DMZ.

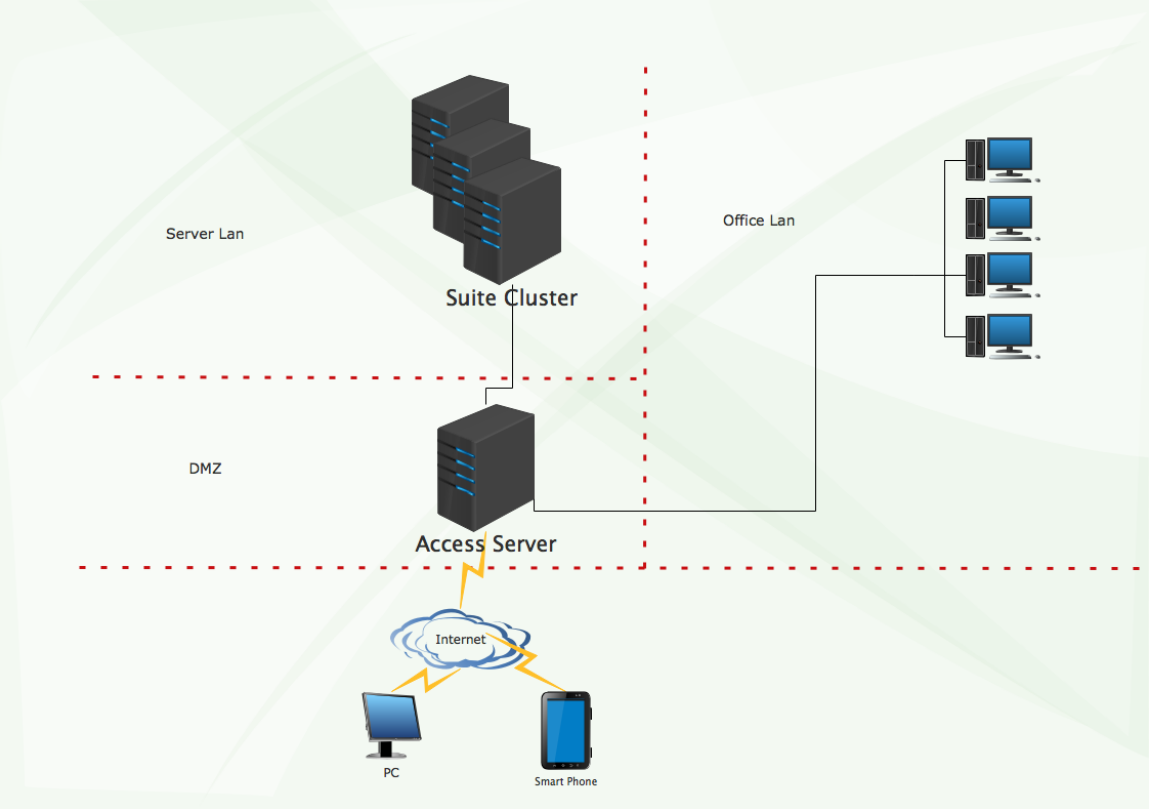


In this case, you need to complete the following tasks before the ITSMA suite installation:

1. Determine the Access Server public DNS record name. This record should be used by the ITSMA suite application exclusively. For example: itsma.apps.yourcompany.com.
2. Make sure the Access Server public DNS record is propagated.
3. Determine the master node address, which can be the FQDN or IP address. For more information, see CDF installation configuration.
4. Set up an L4 reverse proxy on the Access Server using any L4 reverse proxy solutions. Take Nginx with the stream module enabled as an example:

```
stream {
    server {
        listen 443;
        proxy_pass suite_443;
}
server {
        listen 5443;
        proxy_pass suite_5443;
    }



    upstream suite_443 {
        server <master node address or virtual IP address>:443;
}
upstream suite_5443 {
```

```
        server <master node address or virtual IP address>:5443;
    }


}
```

> The ngx_stream_core_module module is available only in version 1.9.0 or later.

## Download the CDF installation package (on-premises)

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

To download the CDF 2017.06 installation package, follow these steps:

1. Download ITSMA Containerized 2017.07.001, and then copy the CDF installation package to a temporary directory on the master node.

   > The 2017.07.001 CDF installation package adds support for high availability that is achieved by multi-master failover in on-premises and cloud-based ITSMA deployments. We recommend that you download this new version of the CDF installation package if this is a new installation. The previous CDF 2017.06 installation package (`CDF1706-00136-15000.zip`) released on the HPE Software Entitlement Portal does not support multi-master failover for ITSMA. If you have already the previous version of CDF deployed and want to redeploy a multi-master environment, you have to uninstall CDF first and then reinstall using the 2017.07.001 CDF installation package as described in the following steps. For instructions on how to uninstall CDF, see Uninstall CDF.

2. Unzip the HPESW_ITOM_Suite_Foundation_2017.06.00139_patched-keepalived.zip file, and go to the directory:

   `HPESW_ITOM_Suite_Foundation_2017.06.00139`

   > The installation package is signed.

   This directory includes installation files and directories. For details, see installation files and directories.

   Next, install CDF on the master node. For details, see Install CDF on the first master node.

Installation files and directories

The HPESW_ITOM_Suite_Foundation_2017.06.00139 folder includes the following files and directories.

| Name | Description | Type |
|---|---|---|
| **bin** | The bin directory includes:<br>• 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**).<br>• The scripts used to check CDF **(kube-restart.sh**, **kube-start.sh**, **kube-stop.sh**).<br>• The script to check that everything is running **(kube-status.sh**).<br>• 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 CDF and suite deployment. | Directory |

| | | |
|---|---|---|
| **cfg** | The initial user and role information that will be seeded into IDM to create user accounts (single sign on). | Directory |
| **cni** | This directory contains Container Network Interface (CNI) files | |
| **images** | All the core platform images and share services images | Directory |
| **install** | The binary that needs to be run to install C DF | 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. | Directory |
| **objectdefs** | | Directory |
| **rpm** | | Directory |
| **scripts** | | Directory |
| **tools** | | Directory |
| **uninstall.sh** | Used to uninstall CDF | File |
| **upgrade.sh** | Used to upgrade CDF | |
| **version.txt** | | File |
| **zip** | | Directory |

## Install CDF (on-premises)

The installation procedure of ITOM Container Deployment Foundation (CDF) varies depending on your deployment environment type: on-premises, or cloud-based.

> **AWS deployment**
> For information about the procedure for cloud deployments, see Install the suite on AWS.

Follow the procedure below to install CDF on premises (on physical or virtual hosts):

1. Prepare the systems to act as cluster nodes.
2. Download the CDF installation archive and uncompress it to the master node.
3. Install an NFS server, and create an NFS share directory for CDF.
4. Configure the installation for the master node using the install.properties file.
5. Copy the install.properties file to the master node.
6. Execute the CDF installation script of CDF on the master node.
7. Add worker nodes from the Management Portal.

Click the links below for the detailed steps.

- Set up an NFS share for CDF
- Configure the install.properties file
- Install CDF on the first master node
- (Optional) Install CDF on additional master nodes
- Install CDF on the worker nodes
- Verify the CDF installation
- Uninstall CDF

> After CDF is installed, you can still edit the current installation and modify your external CDF database configuration. For details, see Edit the current CDF installation and Modify the external database configuration.

## Set up an NFS share for CDF

If a container stops and is then restarted, all changes made inside the container are lost. If you want to save information such as configuration files or databases, the information must be located outside the container in a persistent volume provided by a Network File System (NFS).

The CDF installation configuration file (install.properties) contains two NFS parameters: NFS_SERVER (which specifies the NFS server host name or IP address) and NFS_FOLDER (which specifies the CDF NFS share root directory). For example:

- **NFS_SERVER=16.255.25.255**
- **NFS_FOLDER=/var/vols/itom/core**

You must set up the NFS server and then create a share directory as specified in the NFS_FOLDER parameter to store CDF data. For example, create the following directory: **/var/vols/itom/core**.

- The NFS server can be a master node or an external server. In a test environment, you can set up an NFS server on the master node; in a production environment, you need to set up a dedicated NFS server.
- You can install NFS v3 or v4. If you choose NFS v4, be sure to configure the /etc/exports file in NFS v3 style rather than in NFS v4 style.

To install an NFS server and set up an NFS share, 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 (on-premises)).
   If the package is not already installed, run the following command to install it:
   **yum install rpcbind**
3. Run the following command to install the NFS server:

   **yum install -y nfs-utils**
4. Run the following commands to enable the rpcbind and nfs-server services:

   **systemctl enable rpcbind**
   **systemctl start rpcbind**
   **systemctl enable nfs-server**
   **systemctl start nfs-server**
5. Run the following commands to create an NFS share directory for CDF and change the directory owner:

   **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 NFS root directory defined in the **NFS_FOLDER** parameter (see Configure the install.properties file).
6. In the **/etc/exports** file, expose the NFS directory that you created by adding the following line:

   **<CDF NFS share drectory> *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   For example:
   **/var/vols/itom/core *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
7. Run the following command to check what has been exported:

   **exportfs -ra**

Now, go to Configure the install.properties file.

## Configure the install.properties file

To correctly configure the Kubernetes cluster, you must configure parameters in the install.properties file. Once you have set up the properties file and installed a master node, you can reuse the file to install other nodes.

**Deprecated parameter**
The **INGRESS_HOST** parameter that was used in the previous release is deprecated and removed from the install.properties file.

- List the IP addresses or FQDNs for all the cluster nodes that you are going to install. In this case you can reuse this file for the installation of other nodes. Additionally, when setting FQDNs for the cluster nodes, make sure the FQDNs are resolved to the correct IP addresses and not to the loop back IP 127.0.0.1.
- About the value of the **EXTERNAL_ACCESS_HOST** parameter:
  - *FQDN*: this value must be a lowercase fully qualified domain name (FQDN).
  - *DMZ*: If your organization's network includes a DMZ, this value must be the FQDN of the Access Server . For more

information, see (Optional) Set up Access Server for a DMZ network.
    - *Mixed mode*: If you are going to install the ITSMA suite in mixed mode, this value must use the same domain as your external Service Manager and UCMDB systems. For information about mixed mode, see Deployment modes.
- Even though the worker nodes are specified in the **WORKER_NODES** property, no actual configuration of the worker nodes is performed during a master node installation. However, the installation executable will report an initialization failure if the property is left empty.

1. Make sure that you have already downloaded and unzipped the CDF installation package to a temporary directory on a master node. For details, see Download the CDF installation package (on-premises).
2. On the master node, go to the <ITOM CDF folder> directory, and then edit the `install.properties` file. The following table lists three sample configurations. For information about the number of nodes and hardware configuration required for your actual deployment, see Hardware sizing recommendations.

| Environment type | Sample parameter configuration | Notes |
|---|---|---|
| Demo | **MASTER_NODES="10.10.10.10"**<br>**WORKER_NODES="10.10.10.10 10.10.10.20"**<br>**EXTERNAL_ACCESS_HOST=myd.xxxx.yyy.net**<br>**NFS_SERVER=10.10.10.10**<br>**NFS_FOLDER=/var/vols/itom/core**<br>**DEFAULT_DB_TYPE=EMBEDDED** | This sample configuration uses the master node as a worker node and an NFS server. |
| Production (single-master) | **MASTER_NODES="10.10.10.10"**<br>**WORKER_NODES="10.10.10.20 10.10.10.21 10.10.10.22"**<br>**EXTERNAL_ACCESS_HOST=myd.xxxx.yyy.net**<br>**NFS_SERVER=10.10.10.31**<br>**NFS_FOLDER=/var/vols/itom/core**<br>**THINPOOL_DEVICE=/dev/mapper/docker-thinpool**<br>**DEFAULT_DB_TYPE=EMBEDDED** | This sample configuration uses a dedicated NFS server. |
| Production (multi-master) | **MASTER_NODES="10.10.10.10 10.10.10.11 10.10.10.12"**<br>**WORKER_NODES="10.10.10.20 10.10.10.21 10.10.10.22"**<br>**EXTERNAL_ACCESS_HOST=myd.xxxx.yyy.net**<br>**HA_VIRTUAL_IP=10.10.10.9**<br>**NFS_SERVER=10.10.10.31**<br>**NFS_FOLDER=/var/vols/itom/core**<br>**THINPOOL_DEVICE=/dev/mapper/docker-thinpool**<br>**DEFAULT_DB_TYPE=EMBEDDED** | This sample configuration uses three master nodes and a dedicated NFS server.<br><br>To support multiple master nodes, you must use the CDF installation package that is released with the ITSMA 2017.06.001 patch. See Download the CDF installation package (on-premises). |

- For a full description of the parameters in the install.properties file, see CDF installation configuration.
- Be sure to create and export the NFS folder (**NFS_FOLDER**) specified here when setting up an NFS share for CDF (see Set up an NFS share for CDF).
- This example uses an embedded PostgreSQL database for CDF. If you want to use an external PostgreSQL or Oracle database for CDF, you need to configure the **DEFAULT_DB_TYPE**, **DEFAULT_DB_HOST**, **DEFAULT_DB_PORT**, and **DEFAULT_DB_NAME** parameters. For details, see CDF installation configuration and (Optional) Prepare databases for CDF and ITSMA (on-premises). Additionally, the database configuration parameters (DEFAULT_DB_TYPE, DEFAULT_DB_HOST, DEFAULT_DB_PORT, DEFAULT_DB_NAME, DEFAULT_DB_CONNECTION_URL) are used to specify a database for the Kubernetes cluster, not for the ITSMA suite. Suite databases are specified during the suite installation (see Run the Suite Installer).

## Install CDF on the first master node

To install ITOM Container Foundation (CDF) on the first master node, follow these steps:

1. Log in to the first master node as root or a sudo user.
2. Run the following command:

   **export no_proxy="<*MASTER_NODE_IP*> <*HA_VIRTUAL_IP*>"**

   Replace **<*MASTER_NODE_IP*>** and **<*HA_VIRTUAL_IP*>** with the values that are specified in the *MASTER_NODES* and *HA_VIRTUAL_*

*IP* parameters in the install.properties file. If your deployment is using a single master node, include only the master node IP address.

3. Make sure that you have performed time synchronization on all master nodes and worker nodes. You can run the following command on each node to verify:
   **# chronyc tracking**
4. Update the **/etc/hosts** file of each cluster node by adding a line that contains the IP address and FQDN of the node:
   ***<Node IP> <Node FQDN>***
   For example:

   **10.10.10.10  pcoe001.mycompany.net**
5. Navigate to the CDF installation directory. To do this, run the following command:
   **cd HPESW_ITOM_Suite_Foundation_2017.06.nnnnn/**
6. Run one of the following commands:

   **./install** (if you are the root user)
   **sudo ./install** (if you are a non-root user)
7. If you have configured an external database for CDF by setting the *DEFAULT_DB_TYPE* parameter in the install.properties file to either *EXTERNAL_PG* or *EXTERNAL_ORA* (see Configure the install.properties file), follow the appropriate step:
   - If you set *EXTERNAL_PG*: enter the database user name and database password.
   - If you set *EXTERNAL_ORA*: enter the database user name, database schema, and database password.
8. The following messages indicate that the installation completed successfully:
   **Successfully added the node label.**
   **Successfully completed configuring the HPE ITOM Core Platform on this server!**

For information about what is installed, expand the "List of installed items" section.

List of installed items

The CDF installer installs the following items:

- The base installation files
- Docker
- Certificates
- etcd
- Flannel
- The internal network
- Vault
- The images
- The configuration for K8S
- The persistent volumes
- All the base CDF services, such as the postgresql for IDM and Management Portal
- More SSL certificates for Nginx used for proxying requests into CDF

To see what was installed, run the following commands:

**cd <*CDF installation directory*>**
**ls -l**

To see the installation log, run the following command:

**vi /tmp/install-<*timestamp*>.log**

The following table describes the files and directories that were installed.

| Name | Description | Type |
|---|---|---|

| | | |
|---|---|---|
| bin | The bin directory includes:<br>• 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, and docker-runc), the binary to access the distributed configuration database (etcdctl), and the runtime to interact with Kubernetes (kubectl).<br>• The scripts used to check CDF (kube-restart.sh, kube-start.sh, and kube-stop.sh).<br>• The script to check that everything is running (kube-status.sh).<br>• The script used during installation to create the configuration for Docker (mk-docker-opts.sh) and Vault. Vault is used for security purposes to store sensitive information and to generate and manage certificates for CDF and the suite deployment. | Directory |
| **cfg** | The cfg directory includes the Docker configuration. It includes docker, docker-bootstrap, and IdM. There are two Docker daemons running on each node. Only Docker physically runs on the host, and everything else is containerized. Services or programs that would typically run directly on the host are now run inside the docker-bootstrap container. It runs etcd and Flannel. There are two K8S running on two different sockets: docker and bootstrap-docker.<br>• To see what is running inside docker, run the following command: **docker ps**<br>• To see what is running inside the bootstrap docker, run the following command: **docker - H unix:///var/run/docker-bootstrap.sock ps**. The docker-bootstrap container runs Flannel, Vault, and etcd. Docker provides an abstraction layer from the host.<br>• To see what is running inside the bootstrap-docker container, which is a separate instance, pass the bootstrap-docker socket. To do this, run the following command: **ps -cf\|grep dockerd**. | Directory |

| | | |
|---|---|---|
| **data** | Contains runtime data that is generated by K8S.<br>To see what is in the data directory, run the following command:<br>**ls data/*** | Directory |
| **images** | Contains all the core platform images that have been imported locally.<br>To see what is in the images directory, run the following command:<br>**ls images** | Directory |
| **log** | Contains the logs of some of the currently running components.<br>• To see what is in the log directory, run the following command:<br>**ls log**<br>• To generate a recursive log, run the following command:<br>**ls -R log**<br>All the components include their running information in the logs. | Directory |
| **install-YYYY.MM.nnnnn.log** | The CDF installation log. | File |
| **jar** | Contains the Java files that are used to decrypt the IdM files. | Directory |
| **manifests** | Contains YAML files that describe how to deploy a container. These YAML files must run on every node (they are K8S components):<br>• **kube-apiserver.yaml**: For the K8S API server<br>• **kube-controller-manager.yaml**: Controls access to the K8S server<br>• **kube-proxy.yaml**: Contains proxy connections<br>• **kube-scheduler.yaml**: Schedules the node on which to execute a container<br>• **kube-registry-proxy.yaml**: Starts the kube registry proxy container | Directory |
| **objectdefs** | Contains YAML files for Autopass, IdM, persistent volumes, registry proxies, Vault, Management Portal, the Nginx controller, and the suite installer. | Directory |
| **rpm** | An installable package used to enable the installation of an NFS server. The NFS utility shares data via a networked volume. | Directory |

| | | |
|---|---|---|
| **runconf** | A transient directory used during the installation. | Directory |
| **ssl** | Contains all the certificates and the keys that have been generated by the running C DF. | Directory |
| **scripts** | Contains the scripts that are used in CDF in stallation and the suite image upload and download. | Directory |
| **uninstall.sh** | The uninstall script stops containers and removes containers, demons, and more. You must reboot the server after running the script. | File |
| **tools** | The support toolset helps you collect information about your Docker and Kubernetes environment to troubleshoot C DF. | Directory |
| **version.txt** | The version file shows the current version of the CDF package. | File |
| **zip** | Includes a subset of files used to install a new cluster node from Management Portal (using the **Add Node** functionality). | Directory |

- If you are installing a single master node deployment, follow the steps in Install CDF on the worker nodes.
- If you are installing a multiple master node deployment, follow the steps in (Optional) Install CDF on additional master nodes.

## (Optional) Install CDF on additional master nodes

> This step is optional and can be skipped if you want to deploy only one master node. If you need a multi-master deployment, install all master nodes before deploying the worker nodes.

To install CDF on multiple master nodes, follow these steps:

1. Make sure that you have installed the first master node by following the steps described in Install CDF on the first master node.
2. Download the installation package again to each additional master node (or, copy the installation package from the first master node to each additional master node), and then unzip it.
3. Make sure that you have performed time synchronization on each additional master node. To verify, run the following command on each additional master node:
   **# chronyc tracking**
4. Make sure that you have updated the **/etc/hosts** file of each cluster node by adding the following line:
   ***<Node IP> <Node FQDN>***
5. Follow these steps on each additional master node:
   a. Navigate to the installation directory. To do this, run the following command:

      **cd HPESW_ITOM_Suite_Foundation_2017.06.nnnn/**
   b. Initialize the environment variables. To do this, run the following command:

      **source /etc/profile**

      > You can also log out and log in again to initialize the environment variables.

   c. Copy the server certificate files (ca.crt, server.crt, and server.key in the $K8S_HOME/ssl directory) from the previously-installed master node to any local directory on the additional master node (for example: the /tmp directory).

      > The ca.crt, server.crt, server.key, client.crt, client.key are generated by the previously-installed master node under the $K8S_HOME/ssl folder (default is /opt/kubernetes/ssl).

   d. Copy the install.properties file from the previously-installed master node to each additional master node.
   e. Open the install.properties file and enter the following settings:
   - PEER_CA_FILE=/tmp/ca.crt
   - PEER_CERT_FILE=/tmp/server.crt
   - PEER_KEY_FILE=/tmp/server.key

   a. Run one of the following commands:

**./install** (if you are the root user)
**sudo ./install** (if you are a non-root user)

Next, follow the steps in Install CDF on the worker nodes.

## Install CDF on the worker nodes

Once you have installed ITOM Container Deployment Foundation (CDF) on the master nodes, you are ready to install CDF on the worker nodes. You do this by directly adding the worker nodes from the CDF user interface (the "Management Portal").

> When one worker node goes down, all services running on this worker node are re-created on the other worker nodes automatically.

### Log in to the Management Portal

At your initial login to the Management Portal, you must change the initial password for the **admin** user.

1. Launch the Management Portal from a supported web browser: https://<external_access_host>:5443.
   <external_access_host> is the EXTERNAL_ACCESS_HOST value from the `install.properties` file.
2. The Management Portal login page loads.
3. Log in with the following user account:
   User name **admin**
   Password: **cloud**.
4. Management Portal automatically navigates to the **Change Password** screen.
5. Provide the original password **cloud**.
6. Enter your new desired password.
7. Click **Update Password** to change the password. A message is displayed confirming that the password has been successfully changed.

### Add worker nodes from the Management Portal

To add worker nodes, follow these steps:

1. Click **ADMINISTRATION** > **Nodes**.
2. Enter the following information to add a worker node:
   - the node's host name
   - the name of a user that can remotely execute commands on the host
     For a non-root user, before clicking **ADD**, do the following:
     a. run the following command on the remote host that you are adding:
        **sudo visudo**
     b. Edit the file by adding the following line to the end:

        ```
        <username> ALL=(ALL) NOPASSWD: ALL
        ```

        For example, if the user name is **admin**, add the following line:

        ```
        admin ALL=(ALL) NOPASSWD: ALL
        ```

     c. Comment out the **Defaults requiretty** setting in the file if it is there:
        `# Defaults requiretty`
   - the password of that user
   - THINPOOL_DEVICE

     Only when you have already prepared a logical volume for a direct-lvm thin pool on this worker node (as described in Meet the prerequisites), set this parameter to **/dev/mapper/docker-thinpool**; otherwise leave this parameter empty. For example, leave this parameter empty for a demonstration environment, which requires no thin pool configuration.
   - FLANNEL_IFACE

     The FLANNEL_IFACE parameter specifies the interface for Docker inter-host communication as a single IPv4 address or interface name. This parameter is used only when the node has more than one network adapter so that Flannel can set up the correct routing table entries.
3. Click **ADD** to remotely install the worker node.

   > You can add multiple nodes simultaneously with **+ ADD**:
   >
   > - Enter multiple host names or IP addresses separated by a space.
   > - Enter the user name and password. These nodes share the same user name and password.
   >
   > The installation of each node runs in parallel.

4. 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. If the worker node is not displayed after quite a few minutes, check that you have updated the **/etc/hosts** file on all cluster nodes and then retry.

Next, proceed to the suite installation:

- If you are going to install ITSMA in containerized mode, go to Install ITSMA in fully containerized mode.
- If you are going to install ITSMA in mixed mode, go to Install ITSMA in mixed mode (scenario 1) or Install ITSMA in mixed mode (scenario 2).

> If you need to uninstall CDF, see Uninstall CDF.

### Verify the CDF installation

After you have deployed the master nodes and worker nodes, verify the CDF installation.

1. Log in to a master node as root or a sudo user.
2. Run the following command:
   **kubectl get pods --namespace core**
3. Verify that all pods are running and ready:
   Status: Running
   Ready: n/n (for example: **1/1**, **2/2**, or **3/3**...)

> If the Ready state is not n/n (for example, if the Ready state is **0/1** or **1/2**), the pod is not ready.

### Uninstall 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. Run the following command:

   **./local_backup.sh localhost:5000**

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

#### Uninstall ITOM CDF

> You must uninstall the worker nodes first before you uninstall the master nodes.

1. Run the following command on the worker nodes:
   **uninstall.sh**
   The uninstallation process stops containers, and removes containers, daemons, and so on.
2. Once the worker nodes are all uninstalled, run the following command on the master nodes:
   **uninstall.sh**
   The uninstallation process stops containers, and removes containers, daemons, and so on.
3. Clear the CDF NFS folder.
4. Reboot the servers.

## Install the ITSMA suite (on-premises)

Once ITOM Container Deployment Foundation (CDF) is installed, you are ready to install the ITSMA suite. Your ITSMA installation may use one of the following modes:

- Install ITSMA in fully containerized mode
- Install ITSMA in mixed mode (scenario 1)
- Install ITSMA in mixed mode (scenario 2)

To learn more about the two deployment modes, see Deployment modes.

The uninstallation process is the same for both modes. For details, see Uninstall the ITSMA suite.

## Install ITSMA in fully containerized mode

> **ITSMA installation modes**
> You can deploy ITSMA in one of the following modes: fully containerized, or mixed mode (including scenarios 1 and 2). For instructions on how to install ITSMA in mixed mode scenario 1 and scenario 2, see Install ITSMA in mixed mode (scenario 1) and Install ITSMA in mixed mode (scenario 2).

In this mode, all ITSMA components are containerized. Additionally, once ITSMA is installed, all suite components (Service Management, Service Portal, CMDB, CMDB Browser, Smart Analytics, and Chat) are pre-integrated and work with each other seamlessly.

### Prerequisite

Before you proceed, make sure that you have already ITOM Container Deployment Foundation (CDF) installed in your environment. For details, see Install CDF (on-premises).

### Installation procedure

The installation procedure of this mode comprises the following steps:

1. Download ITSMA images from Docker Hub to CDF
2. Set up three NFS shares for ITSMA
3. Run the Suite Installer

## Download ITSMA images from Docker Hub to CDF

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 Enable your Docker Hub account.
> - 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
  - If the master node can reach Docker Hub
  - If the master node cannot reach Docker Hub
- Task 2: Upload the suite Images to CDF
- Task 3: Verify the images in CDF

### Task 1: Download the images from Docker Hub

Log on to a master node as root or a sudo user and perform the following steps according to whether your master node has access to Docker Hub.

#### *If the master node can reach Docker Hub*

1. On the master node, run the following command to make sure that you can pull images from Docker Hub:
   **docker pull hello-world**
2. If you cannot pull the hello-world image, you need to configure a proxy 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, run the following commands:
**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. Run the following command to make sure that you can pull images from Docker Hub:
**docker pull hello-world**

3. Execute the downloadimages.sh script:
   a. Ensure that the master node is connected to the Internet.
   b. Go to the **$K8S_HOME/scripts** directory.
   c. Run the following command:**./downloadimages.sh -s itsma -u <username> -p <password>**

> Use your Docker Hub account (username and password) that has been enabled by HPE (see Enable your Docker Hub account). 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 **2017.07** or **latest**.

The script starts the downloading process. When the script has finished execution, the following message is displayed: **Succes sfully downloaded the ITSMA suite version: 2017.07 ...**

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**).

> The downloading of the suite images may take quite long. You can download part of the suite images at one time, and upload them into your local registry; after that, you can continue to download the suite images from the break-point. You can also download the suite images from the CDF user interface: **SUITE** > **Management**.

### *If the master node cannot reach 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 uploadimages.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 a computer ( free disk space is more than 100G) that can reach the Docker Hub.
   a. Run the **uname -r** command to get your current kernel version.
   b. Check if your OS is 64-bit, Linux kernel version is 3.10 or higher.
2. Install docker on the machine that can reach Docker Hub. Refer to https://docs.docker.com/engine/installation/ for more details.
   a. Configure yum proxy to download:
      i. Run the **vi /etc/yum.conf** command.
      ii. Add the following line:
      proxy= <Your Proxy>
   b. Run the **yum list** command to list the package version in the system.
   c. Run the following commands to add the yum repo:

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

   d. Run the **yum update --skip-broken -y** command to update the source information.
   e. Run the **yum install -y docker-engine** command to install Docker.
   f. Run the **systemctl enable docker.service** command to enable the service.
   g. Run the following commands to configure the proxy so you can download official images:
   **cat << EOF > /usr/lib/systemd/system/docker.service.d/http_proxy.conf**

**[Service]**
**Environment="HTTP_PROXY=<Your Proxy>" "HTTPS_PROXY=<Your Proxy>"**
**EOF**

> You also need to run the following commands to configure another proxy, if Docker-Content-Trust is turned on:
> **export http_proxy=<Your Proxy>**
> **export https_proxy=<Your Proxy>**

h. Run the **systemctl daemon-reload** command to reload configuration.
i. Run the **service docker restart** command to restart docker.

3. Download the suite images:
   a. Copy the `downloadimages.sh` script, the jq file, image-list.json file, and deployments.json (located in: `$K8S_HOME/scripts`) from your master machine to the machine that can pull images from Docker Hub.
   b. Move the jq file to **/usr/local/bin/** using the following commands:
   **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 -s itsma -u <username> -p <password>**
   Where: <username> and <password> are your Docker Hub account that has been enabled by HPE.
   d. Select the ITSMA suite version by entering **2017.07** 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: 2017.07 ...**
   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 all the files that you downloaded to the following directory on the master node: **/var/opt/kubernetes/offline/suite_images**.

### Task 2: Upload the suite Images to CDF

1. Log on to the master node as the root user or a sudo user.
2. Go to the following directory where the suite image tars are located:
   **/var/opt/kubernetes/offline/suite_images**
3. Get the IP address or host name of the master node where you will run the `uploadimages.sh` script . For a single master node environment, skip this step.
   **source /etc/profile**
   **\$K8S_HOME/scripts/uploadimages.sh --check**
4. Execute the `uploadimages.sh` script in the '**$K8S_HOME/scripts** directory:
   **./uploadimages.sh**

> By default, the suite image files are located in the **/var/opt/kubernetes/offline/suite_images** directory.

> You can run the **./uploadimages.sh --help** command to get the usages. Also, you can run the script with out any parameter.

### Task 3: Verify the images in CDF

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

Now, set up three NFS shares, which are required for you to successfully run the Suite Installer. See Set up three NFS shares for ITSMA.

## Set up three NFS shares for ITSMA

If a container stops and is then restarted, all changes made inside the container are lost. If you want to save data such as configuration files or databases, the data must be stored outside the container in a persistent volume provided by a Network File System (NFS). This release of ITSMA uses three separate NFS persistent volumes ("NFS shares") :

- Global volume: this volume stores global data of the suite (configuration files, logs, SSL certificate files, and so on).
- Smart Analytics (SMA) volume: this volume stores Smart Analytics data.
- Database volume: this volume stores database data.

> If you need to reinstall ITSMA, be sure to remove the NFS share directories that you have created for the old installation before you start the new installation. If you fail to do this, unexpected problems may occur.

**One NFS server only**
In this release, ITSMA and CDF must share the same NFS server. You must export three share directories for ITSMA by using the CDF NFS server. For information about how to set up an NFS server for CDF, see Set up an NFS server.

On the CDF NFS server, set up three volumes to store global data, Smart Analytics data, and database data of ITSMA, respectively. The following are example mount paths of the NFS volumes:

- /var/vols/itom/itsma/itsma-itsma1-smartanalytics
- /var/vols/itom/itsma/itsma-itsma1-db
- /var/vols/itom/itsma/itsma-itsma1-global

To set up the NFS share directories, follow these steps:

1. Log on to the CDF NFS server host.
2. Create three directories under the CDF NFS root folder:

   **mkdir -p *<ITSMA global NFS volume>***
   **mkdir -p *<ITSMA Smart Analytics NFS volume>***
   **mkdir -p *<ITSMA database NFS volume>***

   For example, if the CDF NFS root folder is **/var/vols/itom/**:

   **mkdir -p /var/vols/itom/itsma/itsma-itsma1-global**
   **mkdir -p /var/vols/itom/itsma/itsma-itsma1-smartanalytics**
   **mkdir -p /var/vols/itom/itsma/itsma-itsma1-db**

   You are recommended to use this structure for each directory: **/var/vols/itom/itsma/itsma-itsma<n>-xxx** (n=**1**, **2**..., and xxx=**global**, **smartanalytics**, or **db**) so that you can easily identify them.

3. Grant the user (UID=1999) the right permissions:

   If the group (GID=1999) and user (UID=1999) already exist, you do not need to perform the **groupadd** and **useradd** commands.

   **sudo exportfs -ra**
   **groupadd -g 1999 itsma**
   **useradd -g 1999 -u 1999 itsma**
   **chown -R 1999:1999 *<ITSMA global NFS volume>***
   **chown -R 1999:1999 *<ITSMA Smart Analytics NFS volume>***
   **chown -R 1999:1999 *<ITSMA database NFS volume>***

   For example:

   **sudo exportfs -ra**
   **groupadd -g 1999 itsma**
   **useradd -g 1999 -u 1999 itsma**
   **chown -R 1999:1999 /var/vols/itom/itsma/itsma-itsma1-global**
   **chown -R 1999:1999 /var/vols/itom/itsma/itsma-itsma1-smartanalytics**
   **chown -R 1999:1999 /var/vols/itom/itsma/itsma-itsma1-db**
4. Configure the NFS shares in the /etc/exports file by adding the following lines:

   ***<ITSMA global NFS volume>* *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   ***<ITSMA Smart Analytics NFS volume>* *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   ***<ITSMA database NFS volume>* *(rw,sync,anonuid=1999,anongid=1999,all_squash)**

   For example, add the following lines:

   **/var/vols/itom/itsma/itsma-itsma1-global *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   **/var/vols/itom/itsma/itsma-itsma1-smartanalytics *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   **/var/vols/itom/itsma/itsma-itsma1-db *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
5. Run the following command to check what has been exported:

   **exportfs -ra**

## Run the Suite Installer

Once the ITOM Container Deployment Foundation (CDF) user interface (the "Management Portal") is available, you are ready to install the ITSMA suite. From the Management Portal, you can launch an installation wizard, which will guide you through a process that includes the following major steps:

- Step 1: Select services to install and specify a suite size
- Step 2: Specify NFS shares created previously

- Step 3: Configure external database servers
- Step 4: Specify an initial password for the seeded user **sysadmin**
- (Optional) Step 5: Configure an external LDAP server
- Step 5: Install ITSMA

- **Deployment mode**: during installation, all containerized ITSMA services are selected by default (that is, the default deployment mode is fully containerized mode). This document describes the installation process for fully containerized mode. For information about the installation process for mixed mode, see Install ITSMA in mixed mode (scenario 1) and Install ITSMA in mixed mode (scenario 2).
- **NFS shares**: you must have already configured NFS shares for ITSMA before running the Suite Installer. During installation, you will be asked to specify the NFS servers and NFS share directories that you have configured.
- **Databases**: the ITSMA suite has an embedded PostgreSQL database, and you have the option to configure an external database for each relevant suite component. If using an external Oracle database, make sure that the required Oracle JDBC driver is placed under the required directories, otherwise the database connection test will fail. For details, see the **Oracle JDBC Driver** warning in this topic.
- **sysadmin user**: this is a seeded user with super admin rights for ITSMA. Be aware that there is no way to reset its password if you forget the initial password.
- **LDAP configuration**: the ITSMA suite is bundled with an OpenLDAP server, which you can use in a test environment. In a production environment, you need to configure an external LDAP server. You can skip this step during installation if you want and perform LDAP configuration after installation. For more information, see Configure an external LDAP server.
- **Ports**: before you proceed, make sure that the ports to be used by the suite are not in use on the master and worker nodes. See ITSMA node ports.
- **Namespace:** 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 installation. If you uninstall the first instance and then reinstall ITSMA, the namespace becomes "itsma2", and so on.
- **Page refresh**: 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.

**Oracle JDBC Driver**
If you are using an external Oracle database for ITSMA, you must make sure that Oracle JDBC Driver ojdbc6.jar is placed under the CDF NFS share and the ITSMA global NFS share directories. Perform the following steps before you run the Suite Installer:

1. Download the Oracle JDBC Driver ojdbc6.jar from here.
2. Log in to the CDF NFS server host, and create the following directory if it does not already exist: ***<CDF NFS share>*/suite-install/itsma/output**.
   For example: **/var/vols/itom/core/suite-install/itsma/output**.
3. Run the following command to change the **output** folder owner to the **itsma** user:
   **chown itsma:itsma output**
4. Place the ojdbc6.jar file to the ***<CDF NFS share>*/suite-install/itsma/output** directory, and run the following command to change the owner of jar file to **itsm**a:
   **chown itsma:itsma ojdbc6.jar**
5. Log on to the ITSMA global NFS share server, and create the ***<ITSMA global NFS share>*/jdbc** directory if it does not already exist. For example:
   **/var/vols/itom/itsma/itsma-itsma-global/jdbc**
6. Run the following command to change the **jdbc** folder owner to the **itsma** user:
   **chown itsma:itsma jdbc**
7. Place the ojdbc6.jar file to the ***<ITSMA global NFS share>*/jdbc** directory, and run the following command to change the owner of jar file to **itsm**a:
   **chown itsma:itsma ojdbc6.jar**

If you skip this step, when you test database connection settings during the suite installation, your test will fail with a "No Oracle JDBC Driver Found" error.

**ITSMA installation log**
if the ITSMA installation fails, you can check the installation log for troubleshooting. For details, see Check the ITSMA installation log.

**Remove old NFS shares**
If you installed ITSMA previously, be sure to remove the NFS share directories that you have created for the old installation before you start the new installation. If you fail to do this, unexpected problems may occur. For details, see Uninstall the ITSMA suite.

To install the ITSMA suite, follow these steps:

1. Log in to the ITOM CDF UI ("Management Portal") as **admin**:
   https://*<EXTERNAL_ACCESS_HOST>*:5443
   The initial password for **admin** is **cloud** if this is the first login.

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

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

3. Review and accept the end user license agreement and HPE privacy policy, and then click **Next**.
4. The suite version selection page displays version **2017.07**, which is detected and automatically selected. Click **Next**.
5. On the **Configure the Suite Storage** page, specify the server name (FQDN or IP address) and mount path for each of the following NFS shares that you have created for ITSMA (see Set up three NFS shares for ITSMA):
    - Global data volume (global-volume)
    - Database data volume (db-volume)
    - Smart Analytics data volume (smartanalytics-volume)

    a. click the Edit icon on the far right and then enter the following information in the Configure Volume popup window:

    - **Server**: enter the FQDN or IP address of the NFS server that provides the volume. Once the server is specified, all NFS shares that are configured on this server are automatically displayed in the **Mount Path** drop-down list.
    - **Mount Path**: select the appropriate NFS share directory from the drop-down list.

    b. Click **Apply**.

    The following figure shows an example configuration.

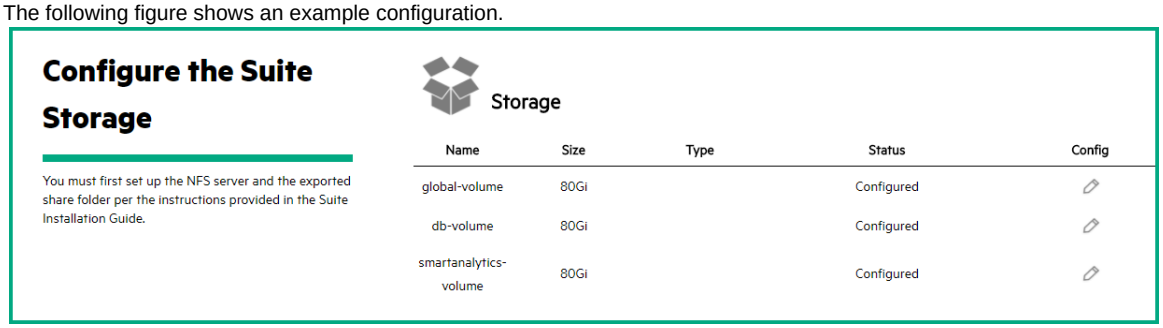


> The volume size of 80G is hardcoded and provided only as the minimum recommended volume size. A total volume size of at least 240G is recommended for the three volumes.

6. On the Suite Installation Configuration page, do the following:
    a. Make sure that all services are selected. This is the default setting.
    b. Select an appropriate suite size: **Demo**, **Extra Small**, **Small**, **Medium**, or **Large**.
    c. Click **Next**.
7. The Suite Installer starts loading the suite images, which you have imported to the local registry (see Download ITSMA images from Docker Hub to CDF). This process may take a while. Wait for the image loading to complete.
8. On the database configuration page, switch on the **External** button to configure your Oracle and PostgreSQL database connection settings.

> - In a test environment, you can skip this step to use a built-in PostgreSQL database by not switching on the **External** b utton. The built-in database is recommended for demontration environments only.
> - The containerized Service Management does not support case-insensitive PostgreSQL.

    a. For Service Management, select **Oracle** or **PostgreSQL**, and configure the database connection information accordingly. Oracle:

| Setting | Description |
|---|---|
| **Oracle Server Host Name or IP Address** | Enter the fully-qualified domain name or IP address of the external Oracle database server. |
| **Oracle Server Port** | Enter the communications port of the external Oracle database server. |

| | |
|---|---|
| **Service Name** | Enter the service name of the Oracle database.<br><br>A service name is 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. |
| **User Name** and **Password** | Enter the user name and password of the user that Service Management uses to connect to the external Oracle database.<br>This user must have the privileges required by Service Management, as described in (Optional) Prepare databases for CDF and ITSMA (on-premises). |

PostgreSQL:

| Setting | Description |
|---|---|
| **PostgreSQLServer Host Name or IP Address** | Enter the fully-qualified domain name or IP address of the external PostgreSQL database server. |
| **PostgreSQL Server Port** | Enter the communications port of the external PostgreSQL d atabase server. |
| **Username** and **Password** | Enter the user name and password of the PostgreSQL user for Service Management. For more information, see (Option al) Prepare databases for CDF and ITSMA (on-premises).<br><br>**Service Management PostgreSQL DBA account**<br>This user is a PostgreSQL DBA account, which wil l be used to create a database and a database user for Service Management, both of which are named **servicemanagement**. |

b. Configure Oracle database connection information for CMDB:

| Setting | Description |
|---|---|
| **Oracle Server Host Name or IP Address** | Enter the fully-qualified domain name or IP address of the external Oracle database server. |
| **Oracle Server Port** | Enter the communications port of the external Oracle database server. |
| **User Name** and **Password** | Enter the user name and password that CMDB uses to connect to the external Oracle database.<br><br>This user must have the privileges required by CMDB, as described in (Optional) Prepare databases for CDF and ITSMA (on-premises). |

- Configure PostgreSQL database connection information for Service Portal:

| Setting | Description |
|---|---|
| **PostgreSQLServer Host Name or IP Address** | Enter the fully-qualified domain name or IP address of the external PostgreSQL database server. |
| **PostgreSQL Server Port** | Enter the communications port of the external PostgreSQL database server. |

| | |
|---|---|
| **Username** and **Password** | Enter the user name and password of the user that you have already created for Service Portal (see (Optional) Prepare databases for CDF and ITSMA (on-premises)).<br><br>**Service Portal PostgreSQL DBA account**<br>This user is a PostgreSQL DBA account, which will be used to create the a number of databases and database users for the Service Portal frondend and backend:<br>- Frondend databases: xservices_ems, xservices_mng, xservices_rms,maas_admin, maas_template<br>- Frondend database user: maas_admin<br>- Backend databases: analytics, bpmdb, catalog, dashboarddb, jumpstart, notificationdb, postgres, sxdb<br>- Backend database user: serviceportalbackend |

a. Configure PostgreSQL database connection information for IdM (the HPE Identity Manager (IdM) instance used by ITSMA):

| Setting | Description |
|---|---|
| **PostgreSQLServer Host Name or IP Address** | Enter the fully-qualified domain name or IP address of the external PostgreSQL database server. |
| **PostgreSQL Server Port** | Enter the communications port of the external PostgreSQL database server. |
| **Username** and **Password** | Enter the user name and password of the user that you have already created for IdM. For more information, see (Optional ) Prepare databases for CDF and ITSMA (on-premises).<br><br>**IdM PostgreSQL DBA account**<br>This is a PostgreSQL DBA account, which will be used to create a database and a database user, both of which are named **idm**. |

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

> If the connection test returns a "No Oracle JDBC Driver Found" error message, follow the instructions in "No Oracle JDBC Driver Found" here.

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** use r 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 ITSMA suite administrator password).
10. On the LDAP configuration page, do the following:
    a. Switch on the **External** button and then configure your external LDAP server as described in the following tables.

> In a demo environment, you can skip LDAP configuration to use an internal OpenLDAP server, and you can add more LDAP users if you want (see Configure users in the internal LDAP server); in a production environment, you must configure an external LDAP server, because the internal one is for demonstration purposes only. In a production environment, you can also skip LDAP configuration at this point and instead configure LDAP after the installation. You may want to do this when you are uncertain about the correct LDAP settings and do not want the installation process to be blocked by incorrect LDAP configuration. Once ITSMA is installed, you can configure your external LDAP server from the Suite Configuration user interface. For details, see Configure an external LDAP server.

**LDAP Server Settings**

| Field | Description | OpenLDAP default value |
|---|---|---|
| Host | The fully-qualified domain name (server.domain.com) or IP address of the LDAP server. | |
| Port | The port used to connect to the LDAP server (by default, 389). | 389 |
| Base DN | Base distinguished name. The Base DN is the top level of the LDAP directory that is used as the basis of a search. | dc=itsma,dc=com |
| User ID (Full DN) | The fully distinguished name of any user with authentication rights to the LDAP server. If the LDAP server does not require a User ID or password for authentication, this value can be omitted. | cn=admin,dc=itsma,dc=com |
| Password | Password of the User ID. If the LDAP server does not require a User ID or password for authentication, this value can be omitted. | |
| Enable SSL | If your LDAP server is configured to require ldaps (LDAP over SSL), select the **Enable SSL** checkbox. | |
| Search Subtree | When a user logs in, the LDAP directory is queried to find the user's account. The **Search Subtree** setting controls the depth of the search under User Searchbase.<br><br>If you want to search for a matching user in the User Searchbase and all su btrees under the User Searchbase, make sure the **Search Subtree** checkb ox is selected.<br><br>If you want to restrict the search for a matching user to only the User Searchbase, excluding any subtrees, unselect the **Search Subtree** checkbox . | |

**LDAP User Attributes**

| Field | Description | OpenLDAP example value |
|---|---|---|
| User Base DN | Base distinguished name for the User object. The User Base DN is the top level of the LDAP directory that is used as the basis of a search for the User object. | ou=people,dc=itsma,dc=com |
| User Class | Value of objectClass that is used to identify the user. | inetOrgPerson |

| | | |
|---|---|---|
| User Filter | Specifies the general form of the LDAP query used to identify users during login. It must include the pattern {0}, which represents the user name entered by the user when logging in.<br><br>The filter must use the following format: (&(objectclass=*)(cn=falcon)) | (objectclass=inetOrgPerson) |
| First Name | **Optional field.**<br><br>First name of the user. | givenName |
| Last Name | **Optional field.**<br><br>Last name of the user. | sn |
| User Display Name | The display name of the user. | cn |
| User Name Attributes | The name of the attribute of a user object that contains the username that will be used to log in. The value for this field can be determined by looking at one or more user objects in the LDAP directory to determine which attribute consistently contains a unique user name. Often, you will want a User Name Attribute whose value in a user object is an email address. | uid |
| User Email | The email address of the user. | mail |
| Phone Number | **Optional field.**<br><br>Business phone number of the user. | telephoneNumber |
| User Avatar | **Optional field.**<br><br>The LDAP attribute whose value is the URL to a user avatar image that is displayed for the logged-in user. If no avatar is specified, a default avatar image is used. | jpegPhoto |
| Manager Identifier | **Optional field.**<br><br>The name of the attribute of a user object that identifies the manager of the user. | manager |

| | | |
|---|---|---|
| Manager Identifier Value | **Optional field.**<br><br>The name of the attribute of a user object that describes the value of the Manager Identifier's attribute. For example, if the value of the Manager Identifier attribute is a distinguished name (such as cn=John Smith, ou=People, o=xyz.com) then the value of this field could be dn (distinguished name). Or, if the Manager Identifier is an email address (such as **admin@xyz.com**) then the value of this field could be email. | dn |
| Last Modified | **Optional field.**<br><br>The LDAP attribute that stores the timestamp when an object was last updated. | modifyTimestamp (for OpenLDAP)<br><br>whenChanged (for Active Directory) |

**LDAP Group Attributes**

| Field | Description | OpenLDAP example value |
|---|---|---|
| Group Base DN | Base distinguished name for the Group object. The Group Base DN is the top level of the LDAP directory that is used as the basis of a search for the Group object. | ou=groups,dc=itsma,dc=com |
| Group Class | Value of objectClass that is used to identify the Group object. | groupOfUniqueNames |
| Group Filter | Specifies the general form of the LDAP query used to identify user groups during login. It must use a standard search filter syntax for your LDAP server. | (objectclass=groupOfUniqueNames) |
| Group Display Name | The display name of the user group. | cn |
| Group Membership | The name of the attribute(s) of a group object that identifies a user as belonging to the group. If multiple attributes convey group membership, the attribute names should be separated by a comma. | uniqueMember |

**Group Mappings**

These fields are available only when you configure LDAP during installation.

| Field | Description |
|---|---|

| | |
|---|---|
| Self Service User Group | This is a user group that will be mapped to the Service Portal **Consumer** group. |
| Administrator Group | This is a user group that will be mapped to the Service Portal **Administrators** group. |

b. If you selected **Enable SSL** option, before you apply the LDAP configuration, copy the LDAP CA certificate file to the **{itsma_global_volume}/certificate/ca-trust/** folder, where **{itsma_global_volume}** is the ITSMA global NFS volume (see Set up three NFS shares for ITSMA). For example: /var/vols/itom/itsma/itsma-itsma-global/certificate/ca-trust/. Note that UCMDB only supports importing a certificate file whose extension is ".crt", which means you need to convert or rename the certificate file to that format.
c. **Click** Test to test your LDAP configuration, and then click **Apply**.
d. Click **Next**.

11. A warning is displayed, asking you to confirm if you want to start the installation process. 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.

12. Wait until a Congratulations page is displayed. The ITSMA suite has been successfully installed.

> **ITSMA installation log**
> if the ITSMA installation fails, you can check the installation log for troubleshooting. For details, see Check the ITSMA installation log.

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

       **kubectl get ns**

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

    d. Log in to ITSMA as the **sysadmin** user:

       https://*<EXTERNAL_ACCESS_HOST>*/main (for the suite administrator "sysadmin")

       > 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.

**Post-install tasks**
If your selected suite size is Small or above, set the webservices_sessiontimeout parameter to 30s from the Suite Configuration user interface. For details, see Configure log level for debugging.

Next, you can log in to ITSMA as **sysadmin** to activate a suite license or perform additional configurations. The following are examples:

- Install an ITSMA suite license
- Configure an external LDAP server
- Replace the certificate for ITSMA
- Configure Email
- Configure SAML SSO
- Configure the Service Portal mobile app
- Import master data

If you need to uninstall or reinstall the suite, see Uninstall the ITSMA suite.

## Install ITSMA in mixed mode (scenario 1)

> **ITSMA installation modes**
> You can deploy ITSMA in one of the following modes: fully containerized, or mixed mode (including scenarios 1 and 2). For instructions on how to install ITSMA in fully containerized mode or in mixed mode scenario 2, see Install ITSMA in fully containerized mode and Install ITSMA in mixed mode (scenario 2).

During the transition phase from classic deployment to containerized deployment, you may want to keep your existing implementation for some capabilities that were previously deployed in the classic manner. The support of mixed mode enables you to use an external classic CMDB system, Service Manager system, or both, together with containerized components (Chat, Service Portal, and Smart Analytics) in ITSMA NG Express.

Mixed mode scenario 1 is intended for customers who want to adopt ITSMA Service Portal, Smart Analytics and Chat on top of existing Service Manager and CMS without re-implementation. In this scenario, the external Service Management (SM) and CMDB systems are installed, configured, and integrated in the classic way. When running the Suite Installer, you need to select to not install the containerized Service Management and CMDB.

> **Licensing**
>
> For licensing information about mixed mode, see Install an ITSMA suite license.

> **Load balancer**
> In this scenario, you can optionally use an F5 hardware load balancer between ITSMA and your external Service Manager system. The Service Manager software load balancer cannot load balance requests sent from ITSMA.



### Prerequisite

Before you proceed, make sure you have already ITOM Container Deployment Foundation (CDF) installed. For details, see Install CDF (on-premises).

> **One domain**
> To install the ITSMA suite in mixed mode scenario 1, the EXTERNAL_ACCESS_HOST parameter value in the CDF installation configuration file (install.properties) must use the same domain as your external Service Manager and UCMDB systems; otherwise the ITSMA suite will not work correctly.

### Installation procedure

To install and configure ITSMA to support external SM and CMDB, complete the following tasks in this topic.

#### Task 1: Download ITSMA images from Docker Hub to CDF

For detailed instructions, see Download ITSMA images from Docker Hub to CDF.

#### Task 2: Configure NFS shares for ITSMA

For detailed instructions, see Set up three NFS shares for ITSMA.

### Task 3: Install ITSMA without SM and CMDB

Run the suite installer without selecting **Service Management** and **CDMB**.

> On the Suite Installation Configuration page, you must clear the **Service Management** and **CDMB** check boxes to exclude SM and CMDB from suite installation. By doing this, SM and CMDB will not be installed in the containerized environment and you will need to configure the suite to integrate with external SM and CMDB later after the suite installation.
>
> Service Portal, Chat, and Smart Analytics are included in suite installation by default and you must install them.

For detailed instructions, see Run the Suite Installer.

### Task 4: Configure integration with external Service Manager

> Make sure that the system time including the timezone is identical on your external Service Manager system and the suite nodes.

This task includes the following steps.

#### Step 1: Apply unload files to Service Manager

> If you are using Service Manager 9.41, we recommend that you upgrade to SM9.41.p6HF2 to improve performance. Please find the SM9.41.p6HF2 installation files from CRyPT.

Before you proceed with setting up the integration with your external Service Manager server, you must apply CompatibleForNG_SM941to952.unl and CompatibleForNG_Plus_SM941.unl by using Unload Manager. These unload files are used to enable the APIs at the SM side for suite integration. Follow these steps:

1. Download the following unload files from HPE Marketplace and save them to your computer:
   - Service Manager 9.41: CompatibleForNG_SM941to952.unl and CompatibleForNG_Plus_SM941.unl.
   - Service Manager 9.5x: CompatibleForNG_SM941to952.unl
2. Log on to Service Manager as a system administrator.
3. Go to **System Administration** > **Ongoing Maintenance** > **Unload Manager**.
4. Double-click **Apply Unload**. A wizard opens.
5. Select the following unload files, also specify a backup file, and then click **Next**. Details of the unload file appear.
   - For Service Manager 9.41, load CompatibleForNG_Plus_SM941.unl first, and then load CompatibleForNG_SM941to952.unl.
   - For Service Manager 9.5x, load CompatibleForNG_SM941to952.unl only.
6. Double-click a conflicting object in the table to open the merge tool:
   a. Merge the object, and then select the **Reconciled** check box.
   b. Click Save to go back to the wizard.
7. Click **Next** after all the conflicting objects are reconciled.
8. Click **Yes** on the confirmation window to apply the unload.
9. Click **Finish**. The unload has been applied and at the same time your old data backed up.

#### Step 2: Create an integration account in Service Manager

After you apply the unload files to Service Manager, the system automatically creates the smIntgAdmin operator account for SM integration with the ITSMA suite. The containerized components use this account to call SM web services APIs. However, this account is locked out by default. You must unlock this account and its password must not expire.

> Using the smIntgAdmin operator account is recommended. If you want to use a different integration account, make sure that this integration account has the same privileges and configurations as smIntgAdmin. In addition, this account must be a dedicated account only for the use of integration.

Follow these steps to activate the smIntgAdmin operator account and set its password:

1. Click **System Administration** > **Ongoing Maintenance** > **Operators**.
2. Search for the **smIntgAdmin** operator.
3. Click the **Security** tab, and then clear the **Administrative Lockout** check box.
4. Click **More** on the toolbar and then select **User Lockout Reset**.
5. Set a password for the smIntgAdmin operator by using either of the following methods:
   - Method 1 (Recommended): Use the **Reset Operators Password** menu option (this method requires that the password reset policy be set to “Prompt for Value” in the System Information Record):
     a. Click **More** > **Reset Operators Password**.
     b. Click **Yes** to confirm your operation, type a new password, and click **OK**.

c. Retype the new password to confirm it.
- Method 2: Type a new password in the Password field.

6. Click **Save**.

### Step 3: Create a new dedicated SM RTE integration instance or clustering for Service Portal

Create a new dedicated SM RTE integration instance or clustering for Service Portal. It includes the following specific startup parameters:

1. Add *webservices_sessiontimeout:30* to the startup command line.
2. Add *connectionTimeout:60000* to the startup command line if you are working with Service Manager 9.51 and 9.52.

### Step 4: Configure SM integration in the suite

Follow these steps to configure integration with an external Service Manager server:

1. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user. A Get Started wizard is displayed if you did not install Service Management during the suite installation.
2. On the **External Service Management Connection Settings** page, enter the connection information for the external SM server.

> If SSL in enabled in the external Service Manager system, perform *Task 1: Replace the out-of-box ITSMA certificate* and *Task 2: Configure one-way SSL for integration with external Service Manager* in Configure SSL for ITSMA in mixed mode instead.

| Setting | Description |
|---|---|
| URL | The URL (including the port number) of the external Service Manager server.<br><br>If you use an F5 external hardware load balancer between ITSMA and the Service Manager Server, enter the URL of the F5 virtual server, which directs requests to those dedicated SM RTE debug nodes. |
| User Name | The integration account that you created for accessing and integrating the external Service Manager server. HPE recommends you to use the smIntgAdmin operator account. |
| Password | The password for the integration account. For example, the s mIntgAdmin operator account. |

When the configuration is complete, click **Test** to make sure you can successfully connect to the external SM server and then click **N ext**. The system opens the LW-SSO configuration page.

3. Configure LW-SSO. For details, see Configure LW-SSO.
4. On the mixed mode configuration dashboard, continue to configure and enable Service Portal, Smart Analytics, and Chat in the mixed mode, as described after this task.

## Task 5: Update to the containerized Service Portal URL

The ITSMA suite installation includes containerized Service Portal by default. To use the containerized Service Portal with an external SM system, you must update to the containerized Service Portal URL in Service Manager.

Follow these steps to update to the new Service Portal URL in Service Manager:

1. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user.
2. Click **Configuration** > **Mixed Mode**.
3. Click **Configure** on the Mixed Mode Configuration Dashboard page.
4. On the **Service Portal Migration** page, click **Finish**.
   The suite will automatically update to the new Service Portal URL in Service Manager. The external Service Manager system uses this URL to send Service Desk emails, Survey emails, and knowledge article links in Collaboration.
5. If you do not want to use the containerized Smart Analytics for Service Portal search, install and configure the Solr plugin to enable km search in Service Portal without Smart Analytics. For details, see Install Solr plugin for Service Portal search. If you plan to use the the containerized Smart Analytics, skip this step and follow the instructions in Task 6: Connect to containerized Smart Analytics.

## Task 6: Connect to containerized Smart Analytics

The ITSMA suite installation includes containerized Smart Analytics by default. No matter whether you are using the Solr search engine or Smart Analytics with an external SM system, you can migrate your existing investment to the containerized Smart Analytics.

This task is required if you want to use the containerized Smart Analytics in ITSMA. If you still want to use the Solr search engine, skip this task and instead install and configure a Solr plugin to automatically push Solr indexed data to IDOL in the containerized environment, which is used by Service Portal search. See Install Solr plugin for Service Portal search for detailed instructions.

If you want to enable SSL for Smart Analytics so that your external Service Manager connects to the containerized Smart Analytics over SSL, perform *Task 3: Configure SSL for Smart Analytics (SMA)* in Configure SSL for ITSMA in mixed mode first before proceeding.

### *Option 1: Purchase a Smart Analytics license and enable containerized Smart Analytics*

Suppose you are not using Smart Analytics with an external SM system. Once you have purchased and applied a Smart Analytics module license, you can enable Smart Analytics, configure the connections to containerized Smart Analytics, and then index data to containerized Smart Analytics.

Follow these steps to purchase and apply a Smart Analytics module license:

1. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user.
2. Click **Configuration** > **Mixed Mode**. The system opens the Mixed Mode Configuration Dashboard page.
3. In the Feature Enablement section, click the **Buy** button for Smart Analytics. The system will direct you to the https://saas.hpe.com/en-us/contact website where you can purchase a Smart Analytics module license.
4. Browse to the <*Service Manager server installation path*>/RUN/ directory, and then add your Smart Analytics module license to the LicFile.txt file.
5. Click the **Configure** button for External SM Connection. The system opens the External Service Management Connection Settings page.
6. Click **Test** and **Apply**.

Now the **Buy** button for Smart Analytics changes to **Enable** on the Mixed Mode Configuration Dashboard page. You can proceed with the next steps.

Follow these steps to connect to containerized Smart Analytics:

#### Step 1: Enable Smart Analytics in Service Manager

Do the following:

1. Log on to the external Service Manager server as an administrator, and then click **System Administration** > **Ongoing Maintenance** > **Smart Analytics** > **Configuration**.
2. Click **Enable Smart Analytics**.
   After you click this button, a message is displayed to state that once you migrate to IDOL, you cannot use the Solr search engine any longer and you have to log out and re-log in to Service Manager before Smart Analytics is applied.
3. Click **Yes** for confirmation. Your account logs out automatically and you need to re-log in to Service Manager.

After you complete this step, access the **Enable Smart Analytics** page in suite configuration, and then click **Finish** under **Manual Migration Step 1**.

To access the **Enable Smart Analytics** page, access the suite configuration interface, click **Configuration** > **Mixed Mode**, and then click **Enable** for Smart Analytics.

#### Step 2: Set up Smart Analytics connections in Service Manager

Do the following:

1. Log on to the external Service Manager server as an administrator, and then click **System Administration** > **Ongoing Maintenance** > **Smart Analytics** > **Configuration**.
2. Enter the address and port (http(s)://<EXTERNAL_ACCESS_HOST>:31370) for the Smart Analytics server, and then click **Test Connection**.
3. Enter the address and port (http(s)://<EXTERNAL_ACCESS_HOST>:31360) for the default CFS server, and then click **Test Connection**. This default CFS server is used for Service Manager attachment index.
4. Enter the address and port (http(s)://<EXTERNAL_ACCESS_HOST>:31395) for the Image Server, and then click **Test Connection**.
5. Click **Save**.

After you complete this step, access the **Enable Smart Analytics** page in suite configuration, and then click **Finish** under **Manual Migration Step 2**.

#### Step 3: Configure Smart Analytics features

To configure the Smart Ticket, Hot Topic Analytics, and Smart Search features, see the *Service Manager Help Center* for your specific version at https://docs.software.hpe.com/.

After you complete this step, access the **Enable Smart Analytics** page in suite configuration, and then click **Finish** under **Manual Migration Step 3**.

### *Option 2: Migrate from external Smart Analytics to containerized Smart Analytics*

If you are already using Smart Analytics with an external SM system, you can migrate the existing investment to containerized Smart Analytics.

These configuration and customization includes the existing definition, tuning, training, data cleansing, stop phrase, stop words, IDOL configuration file modification, and so on.

- Prepare enough disk space before you start step 3 and step 5.
- If your Smart Ticket and HTA content is installed on Windows, you must manually re-index HTA after migrate to containerized Smart Analytics.
- This release does not support the migration for Smart Ticket and HTA content scale out.

Follow these steps to migrate from external Smart Analytics to containerized Smart Analytics:

### Step 1: Stop the scheduled processes and the SMIS tasks

Do the following:

1. To stop scheduled processes:
   a. Click **System Status**.
   b. Locate the KMAttachment, KMUpdate and KMReindex processes in the list.

      The KMReindex process exists only in Service Manager 9.51 and later versions.

   c. Type **k** in the Command column of each process, and then click **Execute Commands**.
2. To stop the SMIDOLxxx integration instance:
   a. Type **smis** in Service Manager command line, and then press Enter.
   b. Select an instance whose name starts with "SMIDOL". For example, SMIDOL0.
   c. Click the **Disable** button on the left.
   d. Repeat the previous two steps to disable all integration instances whose name starts with "SMIDOL".

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Stop the scheduled processes and the SMIS tasks**, and then click **Next**.

To access the Smart Analytics Migration page, log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user. Click **Configuration** > **Mixed Mode**, and then click **Migrate** for Smart Analytics.

### Step 2: Scale out the content groups

Do the following only if you scaled out Smart Search content in your external Service Manager system.

1. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user.
2. Click **Operation** > **Smart Analytics**.
3. Click the **Add New Content Group** button, and then follow the wizard to add new content groups.

Do the following to calculate the number of new content groups to be added:

1. Log on to Service Manager as a system administrator.
2. Type **saa** in the Service Manager command line, and then press Enter.
3. Locate the Smart Search Content section, and record the number of Smart Search Content. For example, 4.

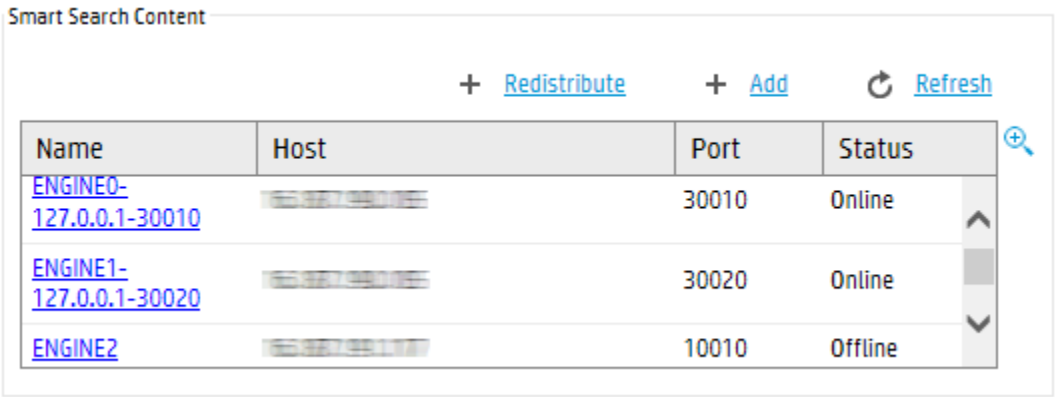


You can click the magnifier button to view all Smart Search Content.

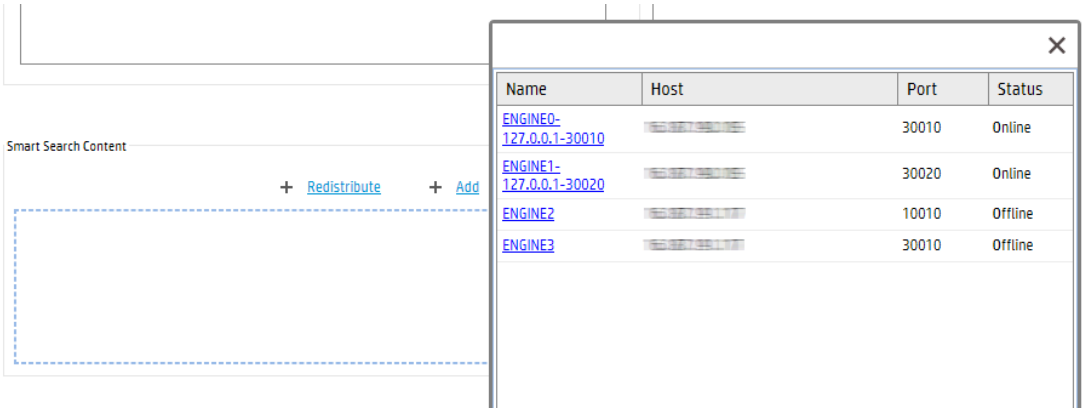


4. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user.
5. Click **Operation** > **Smart Analytics**.
6. Expand the **Content Group Scaling** section. There is an image which displays the visualized capacity under **Service Management Content Group**. For example:

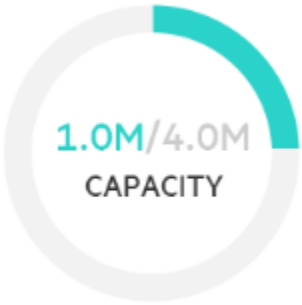


7. Divide the total capacity (the number after the slash) by 2. In this example, the quotient is 2.
8. The number of new content groups to be added = the number of Smart Search Content in SM - the total capacity/2. In this example, 4-4/2=2. This indicates that you need to add at least two content groups.

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Scale out the content group**, and then click **Next**.

### Step 3: Execute the script to backup data

Do the following:

1. Copy the backup scripts from the <Smart Analytics NFS root folder>/data/idol/SmartAnalyticsMigration folder in container to your Service Manager Server host. The backup scripts are as follows:
   The backup.bat backup script for Service Manager on Windows.
   The backup.sh backup script for Service Manager on Linux.
   The mixed-idol-migration-1.0.0.jar file.
2. Run the backup script on your Service Manager Server host. The backup result will be displayed after script execution.

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Execute the script to backup data**, and then click **Next**.

### Step 4: Upload the backup files to the container

Do the following:

1. Manually upload the configuration files that you backed up to the <Smart Analytics NFS root folder>/data/idol/SmartAnalyticsMigration folder in container. The configuration files that you backed up are located under the SmartAnalyticsMigration folder in the same path as the backup.bat script in Service Manager Server host.
2. Manually upload other data files (and folders) that you backed up to the <Smart Analytics NFS root folder>/data/idol/SmartAnalyticsMigration folder in container. The detailed locations of these files are displayed on the summary of the backup script execution.

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Upload the backup files to the container**, and then click **Next**.

### Step 5: Execute the restore scripts

Before you execute the restore scripts, ensure that all the Smart Analytics component pods are ready by checking Smart Analytics Assistant from the ITSMA suite configuration page.

Do the following:

1. Copy the 1-StopSMApods.sh and 3-StartSMApodsRestoreContentData.sh scripts from the <Smart Analytics NFS root folder>/data/idol/SmartAnalyticsMigration folder to the master node of ITSMA.
2. Run 1-StopSMApods.sh on the master node of ITSMA to stop the Smart Analytics component pods.
3. Run 2-MergeConfiguration.sh under <Smart Analytics NFS root folder>/data/idol/SmartAnalyticsMigration to merge the configurations.
4. Run 3-StartSMApodsRestoreContentData.sh on the master node to start the Smart Analytics component pods and restore the data.

> 1. When you run the 2-MergeConfiguration.sh script and the 3-StartSMApodsRestoreContentData.sh script, you must append the number of Hot Topic Analytics and Smart Ticket Content pods, and the number of Smart Search Content pods in suite to the parameter. For example ***./2-MergeConfiguration.sh <Number of Hot Topic Analytics and Smart Ticket Content pods in Suite> <Number of Smart Search Content pods in Suite>***.
> 2. If errors occur when you are executing these migration scripts, refer to Smart Analytics troubleshooting for more information. You can also roll back the data and then perform the migration again. see Roll back from containerized Smart Analytics for detailed instructions.

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Execute the restore scripts**, and then click **Next**.

### Step 6: Configure Service Manager

Do the following:

1. Log on to Service Manager as a system administrator.
2. From the System Navigator, click **System Administration** > **Ongoing Maintenance** > **Smart Analytics** > **Configuration** to open the Smart Analytics Configuration page.
3. Configure the addresses list as follows:
   a. Set http(s)://<EXTERNAL_ACCESS_HOST>:31370/ for the Smart Analytics Server.
   b. Set http(s)://<EXTERNAL_ACCESS_HOST>:31360/ for the Default CFS Server.
   c. Set http(s)://<EXTERNAL_ACCESS_HOST>:31395/ for the Image Server.
4. Click **Save**.

After you complete this step, access the Smart Analytics Migration page. Click **Finish** under **Configure Service Manager**, and then click **Next**.

### Step 7: Start the scheduled processes and the SMIS tasks

Do the following:

1. To start scheduled processes:
   a. Click **System Status**.
   b. Click **Start Scheduler**. HPE Service Manager displays a list of processes that you can start.
   c. Find the processes of KMAttachment, KMUpdate and KMReindex in the list. Double-click these processes to start them.

   > The KMReindex process exists only in Service Manager 9.51 and later versions.

2. To start the SMIDOLxxx integration instance:
   a. Type **smis** in Service Manager command line, and then press Enter.
   b. Select an integration instance whose name starts with "SMIDOL". For example, SMIDOL0.
   c. Click the **Enable** button on the left.
   d. Repeat the previous two steps to disable all integration instances whose name starts with "SMIDOL".

After you complete this step, access the Smart Analytics Migration page, and then click **Finish** under **Start the scheduled processes and SMIS tasks**.

### Step 8: (Optional) Migrate configuration for OMNI Group Server

If you are using the SharePoint connector and the OMNI Group server to manage its user accessibility, you must edit the configuration file of mixed main proxy and restart it manually after you execute the backup and restore scripts.

Do the following:

1. Log on to the NFS server.
2. Browse to the <Smart_Analytics_NFS>/config/idol/st/mixProxy directory, and then open proxy.cfg with a text editor.
3. Locate the following two lines:
   //GroupServerHost=localhost
   //GroupServerPort=5057
   Update them as follows:
   GroupServerHost=<OMNI_GROUP_SERVER_HOST>
   GroupServerPort=<OMNI_GROUP_SERVER_PORT>
4. Save your changes and close this file.
5. Log on to the master node of ITSMA suite, and then run the following command:
   **kubectl get pods -n <NAMESPACE> | grep mix-proxy | awk '{print $1}' | xargs kubectl delete pod -n <NAMESPACE>**
   Wait for several minutes and the mixed main proxy will be restarted.

### Step 9: (Optional) Migrate configuration for CFS

If you are using the external connector to retrieve data from Wiki, SharePoint or a file system, and the connector is not installed on the same server with Smart Analytics main server, you must edit the CFS configuration after migrating to containerized Smart Analytics. After that, the system can index the data into the IDOL server in ITSMA suite.

> Skip this step if the connector and the Smart Analytics main server are installed on the same computer.

Do the following:

1. Log on to the server where the external connector is installed.
2. Browse to the directory where CFS is installed, and open CFS.cfg with a text editor.
3. Locate the [IdolServer] section, and then update the host and port as follows:
   [IdolServer]
   Host=<EXTERNAL_ACCESS_HOST>
   Port=31370
4. Save your changes and close this file.
5. Restart the CFS service.

## Task 7: Connect to containerized Collaboration (Chat)

The ITSMA suite installation includes the containerized Collaboration by default. The following diagram illustrates how the Chat functionality works in mixed mode (scenario 1).

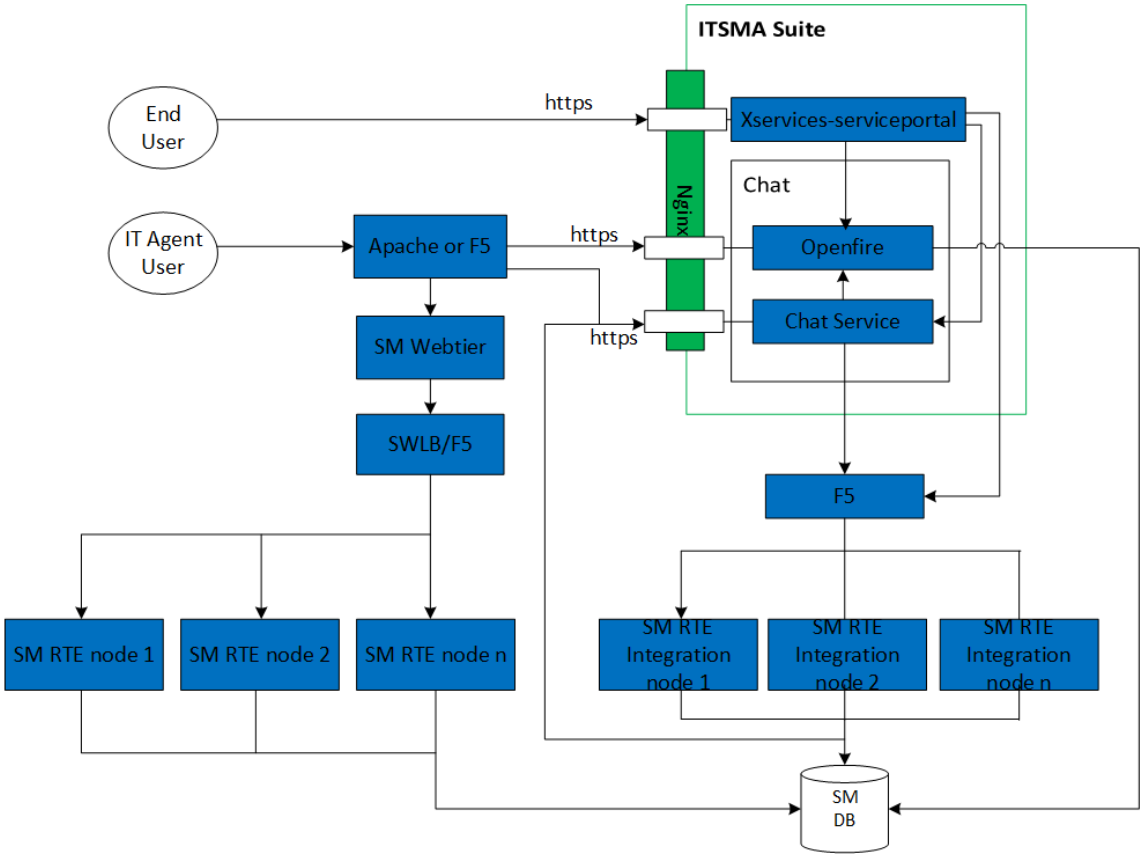


In this diagram:

- **Nginx**: frontend Nginx for the ITSMA suite
- **Apache or F5**: an Apache server or F5 load balancer for the SM web tier
- **SWLB/F5**: the Service Manager software load balancer or an F5 load balancer
- **SM RTE node (1, 2...n)**: Service Manager Server nodes for the web tier
- **SM DB**: Service Manager database
- **XserviceServicePortal**: Service Portal frontend (based on HPE Service Anywhere)
- **Openfire**: Openfire server for Chat
- **Chat Service**: Chat service for Chat
- **F5**: a dedicated F5 hardware load balancer to handle requests from the ITSMA suite. It is not needed if one dedicated Service Manager

Server node can handle all requests sent from the suite.
- **SM RTE Integration node** (1, 2,...n): Service Manager Server nodes for integration with the ITSMA suite

> - In SM 9.41, the HPE Service Manager Collaboration solution provides the IT Collaboration function only. If you are working with SM 9.41, you can continue using the IT Collaboration function with an external SM 9.41 system.
> - To enable End User Chat for Service Portal users, your must upgrade your external Service Manager system to version 9.5x (9.50 or later).
> - This release does not support the Lync integration when migrating to containerized Collaboration.

In order for the external Service Manager system to work with the containerized Collaboration, perform the following steps.

### Step 1: Import the CA certificate

If the ITSMA frontend NGINX uses a certificate that is not signed by a public CA, you need to import the NGINX's CA certificate to the trust store that is used by the Service Manager Server integration nodes (see the Chat functional diagram in this topic). By default, the frontend Nginx CA certificate is /opt/kubernetes/ssl/ca.crt.

- If SSL is not enabled in the Service Manager Server integration nodes (that is, their sm.ini file does not include any ssl parameters), you need to import the ITSMA frontend Nginx's CA certificate to <Service Manager Server installation directory>/RUN/jre/lib/security/cacerts.
- If SSL is enabled in these Service Manager Server integration nodes, you need to import the ITSMA frontend Nginx's CA certificate to their trust store file specified in the **truststoreFile** parameter in their sm.ini file.

Follow these steps to import the ITSMA frontend Nginx's CA certificate:

1. Copy the /opt/kubernetes/ssl/ca.crt file to the <Service Manager installation directory>\RUN\jre\lib\security directory.
2. Run the **keytool -keystore cacerts -importcert -alias mysuite -file ca.crt** command.
3. Configure the ssl parameters in the sm.ini file. For exmaple:

```
keystoreFile:server.keystore
keystorePass:{SERVER_KEYSTORE_PASSWD}
ssl_trustedClientsJKS:trustedclients.keystore
ssl_trustedClientsPwd:{TRUSTEDCLIENTS_KEYSTORE_PASSWD}
ssl:1
ssl_reqClientAuth:0
sslConnector:1
trustedsignon:1
truststoreFile:cacerts
truststorePass:{CACERT_PASSWD}
```

### Step 2: Configure database connection for containerized Collaboration

Follow these steps to configure database connection for containerized Collaboration:

1. Log in to the ITSMA Suite Configuration user interface (https://<EXTERNAL_ACCESS_HOST>/itsmaconfig) as the **sysadmin** user.
2. Click **Configuration** > **Mixed Mode**.
3. On the **Collaboration Configuration** page, enter the database connection information of your external Service Manager system.

| Setting | Description |
|---|---|
| Database Host / IP | Specify the host name or the IP address of the database. |
| Database Port | Specify the port to connect to the database. |
| Database Name | Specify the name of the database. |
| Database Account | Specify the user name to log on to the Service Manager database. This account must have permission to create tables if Collaboration is not configured before. |
| Database Password | Specify the password to log on to the Service Manager database. HPE suggests that you use a strong password. |

> If your external Service Manager system is using an Oracle database, you must download the Oracle JDBC Driver (ojdbc6.jar), save this file to the <ITSMA global NFS volume>/jdbc directory before testing connection to the Oracle database, and then run the **chown itsma:itsma ojdbc6.jar** command to change the jar file owner to be itsma. For example, save ojdbc6.jar to the /var/vols/itom/itsma/itsma-itsma-global/jdbc directory. After you prepare the ojdbc6.jar file, you need to restart the itom-itsma-config-deployment-*xxxxxx* pod.

When the configuration is complete, click **Test** to make sure you can successfully connect to the database.

4. Click **Apply** to apply the database connection settings. This will also apply the Chat service URL and other Collaboration settings to

your external Service Management system.

5. Verify Service Manager Collaboration is enabled.
    a. Log on to the external Service Manager system as a system administrator.
    b. Click **System Administration** > **Ongoing Maintenance** > **Collaboration** > **Configuration** to open the Collaboration Settings form.
    c. Make sure the **Enable Collaboration** check box is selected.
    d. Specify field values as described in the following table.

| Field | Value | Description |
|---|---|---|
| Maximum Participants Per Conversation | 200 | The maximum number of participants in a conversation. The default value is 200. |
| Notification Delay Time (Seconds) | 30 | The maximum time that an online participant has to wait to receive the live conversation notifications. The default value is 30. Notification delay is disabled if this value is set to 0 or minus. |
| Chat Service URL | https://*<FQDN of the Apache HTTP server>*/chatservice | The chat service URL. |

### Step 3: Deploy the Apache HTTP server

If you have not configured a web server for Collaboration before, you need to deploy and configure the web server for Service Manager Colla boration. See the following Apache deployment steps as an example:

1. Download Apache with OpenSSL (for example, httpd-2.4.xx-x64.zip for Apache 2.4) from here.
   You can also download a pre-configured Apache 2.4 from ITOM Marketplace.
   Extract the zip file to C:\. This unzip process creates a new C:\Apache24 directory or a new C:\Apache22 directory.
2. Navigate to the C:\Apache24\conf folder.
3. Make a copy of the httpd.conf file and save it as httpd_OOB.conf.
4. Open the httpd.conf file with a text editor.
5. Locate `httpd-vhosts.conf`, and then uncomment `Include conf/extra/httpd-vhosts.conf`.
6. Save and close the httpd.conf file.
7. Navigate to the C:\Apache24\conf\extra directory.
8. Make a copy of the httpd-vhosts.conf file and save it as httpd-vhosts_OOB.conf.
9. Navigate to the C:\Apache24\bin folder.
10. Double-click httpd.exe to start the Apache server.
    The httpd.exe window opens. Click the minimize button to minimize this window.
11. In your web browser, type http://localhost and press Enter. The following page is displayed, indicating Apache has started successfully.

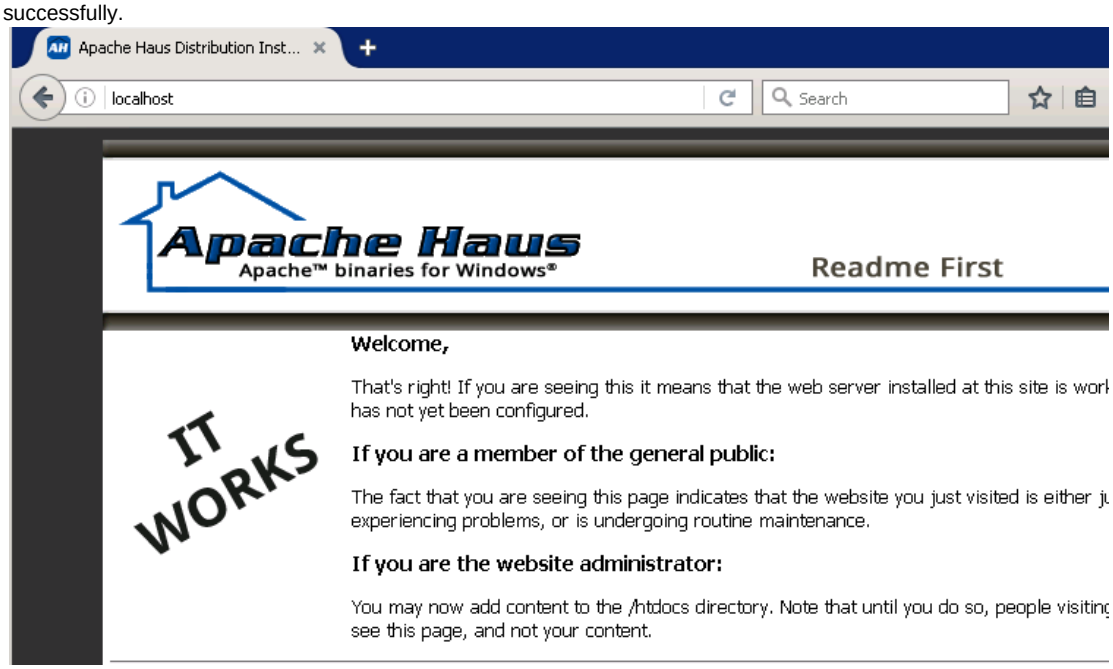


12. Close the browser.
13. Close the Apache httpd.exe window. The steps below will install Apache as a Windows service.
14. Navigate to the C:\Apache24\bin folder. Open a DOS command prompt and change the directory to C:\Apache24\bin.

**cd C:\Apache24\bin**

15. Run the **httpd –k install** command to install the Windows service.

```
C:\Users\Administrator>cd C:\Apache24\bin

C:\Apache24\bin>httpd -k install
Installing the 'Apache2.4' service
The 'Apache2.4' service is successfully installed.
Testing httpd.conf....
Errors reported here must be corrected before the service can be started.

C:\Apache24\bin>
```

If you see an error here, navigate to the logs directory and check the error.log file. Depending on the error, you may need to repeat the steps above. To verify whether the error still exists, type **httpd –k start** to start Apache from the command line.

16. Go to Windows Services, and start the newly installed Apache2.4 service.

### Step 4: Connect Apache to Tomcat

If you have not configured a web server for Collaboration before, you need to set up Apache to connect to Tomcat through the AJP port. Consequently, Secure Sockets Layer (SSL) is open by default. You can perform this step rather than enable full SSL on the Service Manager environment.

Follow these steps:

1. Navigate to the C:\Program Files\Apache Software Foundation\Tomcat 8.0_SMWeb\conf directory.
2. Open the server.xml file with a text editor.
3. Make sure that the AJP 1.3 Connector port is set to `8009`.

```
<!-- Define an AJP 1.3 Connector on port 8009 -->
<Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />
```

4. Save and close the server.xml file.

Steps for Apache 2.4:

> If you are using the pre-configured Apache downloaded from ITOM Marketplace, skip step 1 to 15 and start with step 16.

1. Navigate to the C:\Apache24\conf directory.
2. Open the httpd.conf file with a text editor.
   The next few steps describe how to uncomment a number of `LoadModule` codes in the httpd.conf file.
3. Locate `lbmethod`.

```
#LoadModule lbmethod_bybusyness_module modules/mod_lbmethod_bybusyness.so
#LoadModule lbmethod_byrequests_module modules/mod_lbmethod_byrequests.so
#LoadModule lbmethod_bytraffic_module modules/mod_lbmethod_bytraffic.so
#LoadModule lbmethod_heartbeat_module modules/mod_lbmethod_heartbeat.so
#LoadModule ldap_module modules/mod_ldap.so
```

4. Uncomment two lines as shown below:

```
#LoadModule lbmethod_bybusyness_module modules/mod_lbmethod_bybusyness.so
LoadModule lbmethod_byrequests_module modules/mod_lbmethod_byrequests.so
LoadModule lbmethod_bytraffic_module modules/mod_lbmethod_bytraffic.so
#LoadModule lbmethod_heartbeat_module modules/mod_lbmethod_heartbeat.so
#LoadModule ldap_module modules/mod_ldap.so
```

5. Locate the following section by searching for `proxy_module`.

```
#LoadModule proxy_module modules/mod_proxy.so
#LoadModule proxy_ajp_module modules/mod_proxy_ajp.so
#LoadModule proxy_balancer_module modules/mod_proxy_balancer.so
#LoadModule proxy_connect_module modules/mod_proxy_connect.so
#LoadModule proxy_express_module modules/mod_proxy_express.so
#LoadModule proxy_fcgi_module modules/mod_proxy_fcgi.so
#LoadModule proxy_ftp_module modules/mod_proxy_ftp.so
#LoadModule proxy_html_module modules/mod_proxy_html.so
#LoadModule proxy_http_module modules/mod_proxy_http.so
#LoadModule proxy_scgi_module modules/mod_proxy_scgi.so
```

6. Uncomment 8 lines as shown in the following:

```
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_ajp_module modules/mod_proxy_ajp.so
LoadModule proxy_balancer_module modules/mod_proxy_balancer.so
LoadModule proxy_connect_module modules/mod_proxy_connect.so
LoadModule proxy_express_module modules/mod_proxy_express.so
#LoadModule proxy_fcgi_module modules/mod_proxy_fcgi.so
LoadModule proxy_ftp_module modules/mod_proxy_ftp.so
LoadModule proxy_html_module modules/mod_proxy_html.so
LoadModule proxy_http_module modules/mod_proxy_http.so
#LoadModule proxy_scgi_module modules/mod_proxy_scgi.so
```

7. Locate the following section by searching for `slotmem_shm`.

```
LoadModule setenvif_module modules/mod_setenvif.so
#LoadModule slotmem_plain_module modules/mod_slotmem_plain.so
#LoadModule slotmem_shm_module modules/mod_slotmem_shm.so
```

8. Uncomment the following line:

```
LoadModule slotmem_shm_module modules/mod_slotmem_shm.so
```

9. Locate the following section by searching for `xml2enc_module`.

```
#LoadModule version_module modules/mod_version.so
#LoadModule vhost_alias_module modules/mod_vhost_alias.so
#LoadModule watchdog_module modules/mod_watchdog.so
#LoadModule xml2enc_module modules/mod_xml2enc.so
<IfModule unixd_module>
```

10. Uncomment the following line:

```
LoadModule xml2enc_module modules/mod_xml2enc.so
```

11. Locate the following section. You may need to search for `mod_proxy_html` twice.

```
# Configure mod_proxy_html to understand HTML4/XHTML1
<IfModule proxy_html_module>
Include conf/extra/httpd-proxy-html.conf
</IfModule>
```

12. If the `Include` line does not contain `Include conf/extra/httpd-proxy-html.conf`, change the `Include` line to `Include conf/extra/httpd-proxy-html.conf`.
13. Browse to the end of the file, and then add the line in bold:

```
<IfModule http2_module>
 ProtocolsHonorOrder On
 Protocols h2 h2c http/1.1
</IfModule>
Include conf/httpd-proxy_ajp_loadbalanced.conf
```

14. Comment out the lines in bold by inserting # in front of each line:

```
#<IfModule http2_module>
#ProtocolsHonorOrder On
#Protocols h2 h2c http/1.1
#</IfModule>
Include conf/httpd-proxy_ajp_loadbalanced.conf
```

15. Save and close the httpd.conf file.
16. Navigate to the C:\Apache24\conf directory, and then create a new file called httpd-proxy_ajp_loadbalanced.conf.

```
<Proxy balancer://smcluster>
BalancerMember ajp://localhost:8009 route=161652175430301
Require all granted
</Proxy>
<Location /webtier-9.52>
Options FollowSymLinks
Require all granted
ProxyPass balancer://smcluster/webtier-9.52 stickysession=JSESSIONID|jsessionid nofailover=On
</Location>
<Location /chatui>
Options FollowSymLinks
Require all granted
ProxyPass balancer://smcluster/chatui stickysession=JSESSIONID|jsessionid nofailover=On
</Location>
```

- You must paste `ProxyPass balancer://smcluster/webtier-9.52stickysession=JSESSIONID|jsessionid nofailover=On` in one line.
- You must paste `ProxyPass balancer://smcluster/chatui stickysession=JSESSIONID|jsessionid nofailover=On` in one line.

17. The script in the previous step assumes that the web tier directory is webtier-9.52 and the chat UI directory is chatui(see the line below). If your web tier or chat UI uses another name, update the httpd-proxy_ajp_loadbalanced.conf file with the actual name of your web tier.
18. In *step 3* in this task, you configured the `AJP 1.3 Connector` port in the server.xml file. If this port is 8009, continue with the next step; if the connector listens on another port, update the following line in the httpd-proxy_ajp_loadbalanced.conf file with that port number.
    If Apache is deployed on the same computer in the all-in-one example described in this document, use `ajp://localhost:8009`. Otherwise, you need to update this value to the correct IP of Tomcat.

    ```
    BalancerMember ajp://localhost:8009 route=161652175430301
    ```

19. Access Apache's link with Apache's FQDN, and then log on to Service Manager as a system administrator. The system displays the administrator's To Do Queue.
    If you are directed to a Logout Successful page, there may be some issues with the LW-SSO setup. Check all your files from the previous tasks and then try again.

    From now on, you must use HTTPS and the fully qualified domain name (FQDN) in the web tier URL when logging on to the Service Manager web client.

20. Log out from Service Manager.

### Step 5: Configure the reverse proxy on an external Web server

You need to configure the reverse proxy on an external web server which is located between the Service Manager web application server and the browser, such as Apache, F5, or IIS. See the following Apache configuration steps as an example:

1. Navigate to the C:\Apache24\conf\extra directory, and then open the httpd-ahssl.conf file with a text editor.
2. Locate the section of "Virtual host".
3. Update reverse proxy and point to the containerized Collaboration.

```
<VirtualHost _default_:443>
...
SSLProxyEngine On

ProxyPass /of-http-bind https://<EXTERNAL_ACCESS_HOST>/of-http-bind
ProxyPassReverse /of-http-bind https://<EXTERNAL_ACCESS_HOST>/of-http-bind
ProxyPass /of-plugins https://<EXTERNAL_ACCESS_HOST>/of-plugins
ProxyPassReverse /of-plugins https://<EXTERNAL_ACCESS_HOST>/of-plugins
ProxyPass /chatservice https://<EXTERNAL_ACCESS_HOST>/chatservice
ProxyPassReverse /chatservice https://<EXTERNAL_ACCESS_HOST>/chatservice
...
</VirtualHost>
```

### Step 6: Make sure LW-SSO is configured

Make sure that LW-SSO is configured in ITSMA and the external Service Manager (both the Server and web tier). For details, see Configure LW-SSO.

## Install Solr plugin for Service Portal search

If you do not use Smart Analytics, you can install and configure the Solr plugin to automatically push Solr indexed data to IDOL in container, which is used by Service Portal search.

> This task is only required if your external Service Manager uses the Solr search engine instead of Smart Analytics and you do not want to migrate to Smart Analytics in the mixed mode.

To install Solr plugin for Service Portal search, follow these steps:

1. Make sure that you have installed KM Solr search engine in Service Manager and completed KM indexing.

   > For detailed instructions on installing KM Solr search engine and indexing KM knowledgebases, see the help center for your external Service Manager systems on https://docs.software.hpe.com.

2. Download the the `propel-solr-plugin.zip` file from HPE Marketplace.
3. Copy the `propel-solr-plugin.zip` file to the Solr search server for Service Manager and then unzip the file. The contents are:
   jackson-mapper-asl-1.9.13.jar
   jackson-core-asl-1.9.13.jar
   jasypt-1.9.2.jar
   KMExtAccess.unl
   propel-solr-plugin-1.1.0.jar
4. Copy the `.jar` files to your primary search server. That is, copy `propel-solr-plugin-1.1.0.jar`, `jackson-mapper-asl-1.9.13.jar`, `jackson-core-asl-1.9.13.jar`, and `jasypt-1.9.2.jar` to <*`Primary_Search_Server`*>`\Search_Engine\tomcat\webapps\KMCores\WEB-INF\lib\`.
5. Edit the <*`Primary_Search_Server_Home`*>`\Service Manager9.xx\Search_Engine\kmsearchengine\KMCores\kmcore\conf\solrconfig.xml` file to add an `updateRequestProcessorChain`:

   ```
   <updateRequestProcessorChain name="propelSearch" default="true">
   <processor class="com.hp.propel.solr.plugin.PropelPushUpdateFactory">
   <str name="baseUrl">https://{hostname}:31600/api/search/v1/article</str>
   <str name="username">searchTransportUser</str>
   <str name="password">{Password}</str>
   <str name="tenant">Provider</str>
   </processor>
   <processor class="solr.RunUpdateProcessorFactory"/>
   </updateRequestProcessorChain>
   ```

   Where:
   - *`Hostname`* is the FQDN of the external access host for the suite (normally the master node FQDN if only one master node is used). The port 31600 is the dedicated port for the Service Portal search service (propel-search) for the suite.
   - *`Password`* is the password for `searchTransportUser`. (The default password is `searchTransportUser`.)

6. Update the same `solrconfig.xml` and modify the `requestHandler`.

   ```
   <requestHandler name="/update" class="solr.XmlUpdateRequestHandler">
   <lst name="defaults">
   <str name="update.processor">propelSearch</str>
   </lst>
   </requestHandler>
   ```

7. In Service Manager, apply the following unload files (available on HPE Marketplace):
   - For SM 9.41: First apply **CompatibleForNG_Plus_SM941.unl**, and then apply **CompatibleForNG_SM941to952.unl**.
   - For SM 9.5x: Apply **CompatibleForNG_SM941to952.unl**.
8. Restart KM Solr search engine.
9. Restart Service Manager.
10. In Service Manager, reindex KM:
    a. Select **Knowledge Management** > **Administration** > **Environment**.
    b. Check **SRC**.
    c. Select the **Search Server Name**.
    d. Click **Full Reindex.**
11. In Service Manager, reindex KM Libraries:
    a. Select **Knowledge Management** > **Knowledgebases**.
    b. Click **Knowledge Library**, and then click **Full Reindex**.

## Install ITSMA in mixed mode (scenario 2)

**ITSMA installation modes**
You can deploy ITSMA in one of the following modes: fully containerized, or mixed mode (including scenarios 1 and 2). For instructions on how to install ITSMA in fully containerized mode or in mixed mode scenario 1, see Install ITSMA in fully containerized mode and Install ITSMA in mixed mode (scenario 1).

During the transition phase from classic deployment to containerized deployment, you may want to keep your existing implementation for some capabilities that were previously deployed in the classic manner. The support of mixed mode enables you to use an external classic CMDB system, Service Manager system, or both, together with containerized components (Chat, Service Portal, and Smart Analytics) in ITSMA NG Express.

Mixed mode scenario 2 is intended for customers who want to adopt ITSMA NG Express in phases. In this scenario, Service Management is containerized and only CMDB is still classic. When running the Suite Installer, you need to select to not install the containerized CMDB.



## Prerequisite

Before you proceed, make sure you have already ITOM Container Deployment Foundation (CDF) installed. For details, see Install CDF (on-premises).

**One domain**
To install the ITSMA suite in mixed mode scenario 2, the EXTERNAL_ACCESS_HOST parameter value in the CDF installation configuration file (install.properties) must use the same domain as your external UCMDB system; otherwise the ITSMA suite will not work correctly.

## Installation procedure

To install and configure ITSMA to support external CMDB, complete the following tasks.

### Task 1: Download ITSMA suite images from Docker Hub to CDF

For detailed instructions, see Download ITSMA images from Docker Hub to CDF.

### Task 2: Prepare your databases

For detailed instructions, see (Optional) Prepare databases for CDF and ITSMA (on-premises).

### Task 3: Set up three NFS shares for ITSMA

For detailed instructions, see Set up three NFS shares for ITSMA.

### Task 4: Install ITSMA without CMDB

Run the suite installer without selecting CDMB.

On the Suite Installation Configuration page, you must clear the **CDMB** check box to exclude CMDB from suite installation. By doing this, CMDB will not be installed in the containerized environment and you will need to configure the suite to integrate with external CMDB later after suite installation.

You cannot clear the **CDMB** check box without selecting **Service Management**.

For detailed instructions, see Run the Suite Installer.

### Task 5: Install an ITSMA suite license

> **Licensing**
>
> In this scenario, an ITSMA suite license is needed to activate the containerized capabilities: Service Management, Smart Analytics, and Service Portal.

For detailed instructions, see Install an ITSMA suite license.

### Task 6: Configure integration between containerized SM and external CMDB

> **Configure SSL**
> If you want to configure SSL for this scenario, configure the integration by following the instructions in the *Scenario 2* section in Configure SSL for ITSMA in mixed mode.

To set up the integration between containerized SM and external CMDB, follow these steps:

1. Download the **ServiceManagerEnhancedAdapter9-41.zip** from HPE Marketplace.

   > Downloading and deploying this adapter only when it is not already installed in your UCMDB environment.

2. Deploy the adapter in your external CMDB:
   a. Click the  button to open the Deploy Packages to Server dialog box.
   b. Click the  button to open the Deploy Packages to Server (from local disk) dialog box.
   c. Select the **ServiceManagerEnhancedAdapter9-41.zip** file, and then click **Open**. All the resources are selected by default.
   d. click **Deploy.** A status report appears indicating whether the deployment was successful for each resource selected.
3. In the containerized Service Management, add the UCMDB and UCMDB Browser Connection Information. For instructions, see the "add the UCMDB and UCMDB Browser Connection Information" section in the SM-UCMDB Integration Guide.
4. In the external CMDB, create an integration point with the following settings:

   Hostname/IP: <EXTERNAL_ACCESS_HOST>
   Port: 31190
   User name: intgAdmin
   Password: Admin_1234
   Adapter: ServiceManagerEnhancedAdapter9-41 (this is the adapter that you downloaded and installed in steps 1 and 2)

   > - The **intgAdmin** user account (default password: Admin_1234) is an out-of-box integration user in the containerized Service Management. If you want to use another user account, you can create a new one based on **intgAdmin**. For information about how to create a user account in Service Management, see the How to Create an Integration Account .
   > - For more information about how to set up an integration point in your external CMDB, see How to Create an Integration Point in UCMDB.

5. Click **Test Connection** to make sure that a successful connection has been established.
6. Click **Save** to save the integration point.

## 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://<*EXTERNAL_ACCESS_HOST*>:5443
2. Click **Suite** > **Management**.
3. Click the action icon for the ITSMA suite instance, and then select **Uninstall**.

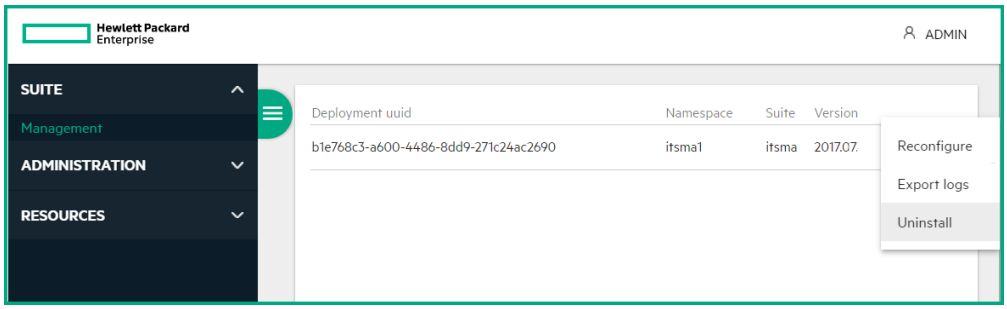


### Steps to reinstall ITSMA

Follow these steps to reinstall ITSMA:

1. Uninstall ITSMA as described previously.
2. Log on to the ITSMA NFS server, and then delete the old NFS volumes. For example:

   **rm -rf /var/vols/itom/itsma/itsma-itsma-smartanalytics**
   **rm -rf /var/vols/itom/itsma/itsma-itsma-db**
   **rm -rf /var/vols/itom/itsma/itsma-itsma-global**

   You can also delete the old log and yamls if necessary:
   **rm -rf /var/vols/itom/core/suite-install/itsma/output/**

   > Before a new installation, be sure to remove any existing NFS shares to remove old installation data. If you fail to do this, unexpected problems may occur.

3. Recreate the NFS shares:

   **mkdir -p /var/vols/itom/itsma/itsma-itsma-smartanalytics**
   **mkdir -p /var/vols/itom/itsma/itsma-itsma-db**
   **mkdir -p /var/vols/itom/itsma/itsma-itsma-global**
   **exportfs -ra**
4. Make sure that the itsma user is the owner of the NFS folder:
   **chown -R itsma:itsma /var/vols/itom/itsma/**
5. Log on to the NFS server as root and run the following command to check the NFS volumes:
   **# showmount -e**
6. Run the Suite Installer. For details, see Run the Suite Installer.

# Install the suite on AWS

ITSMA NG Express leverages container technology from Docker and Kubernetes. Docker provides a means 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.

ITSMA can be deployed on on-premises servers (that is, physical or virtual servers) or cloud servers (using Amazon Web Services). This section describes the steps for a cloud-based deployment. For information about the procedure to install CDF and ITSMA in an on-premises (physical or virtual) environment, see Install the suite on premises.

HPE provides an AWS deployment toolkit (.zip), which contains HashiCorp Packer and Terraform scripts that automate the installation of CDF on a cluster of nodes. Packer scripts automate the creation of machine images, and are used here to create the necessary Amazon Machine Image (AMI) instance. Terraform scripts create and modify datacenter infrastructure. The Terraform scripts in the deployment toolkit deploy CDF of the AMI instance created by Packer. For more information about Packer and Terraform, refer to the HashiCorp website.

The suite installation process remains the same as for on-premises deployments; however, the deployment toolkit can automatically download ITSMA suite images from Docker Hub and import the images to the local registry (you can also download the images manually). Once the images are ready, you can install ITSMA from the CDF user interface, guided by an installation wizard.

- Cloud-based deployment supports only fully containerized mode.
- Because the AWS deployment toolkit automates the deployment process, you do not have to configure the **install.properties** fi le to install CDF.
- Both CDF and ITSMA must use an external database (PostgreSQL or Oracle).
- Amazon Elastic File System (Amazon EFS) storage is automatically configured for CDF, and therefore you do not need to

manually configure an AWS NFS share for CDF. However, you must do this for ITSMA.
- To run the deployment scripts on a Windows-based machine, you must install Git Bash, Cygwin, or MingW.

## Installation procedure

The cloud-based ITSMA NG Express installation procedure comprises the three stages illustrated in the following workflow.

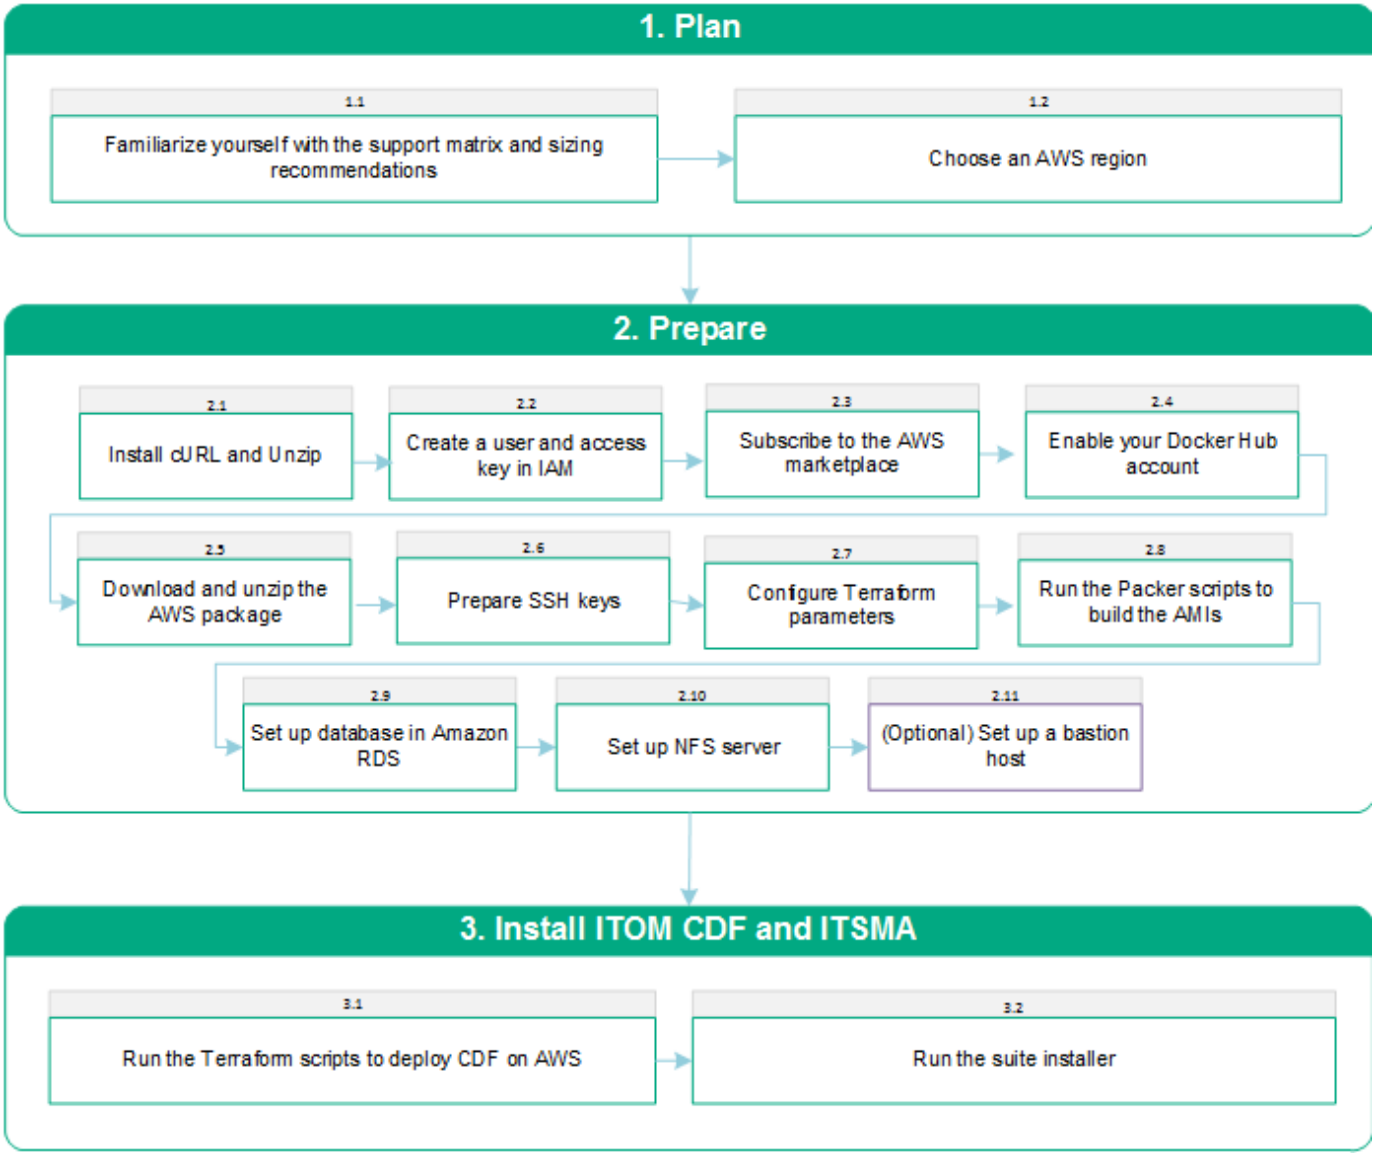


The following table lists detailed information about the steps illustrated in the workflow diagram. You can use the checklist in the following table to keep track of completed installation steps. You can also click the link for each step to view detailed instructions.

| Stage | Steps |
|---|---|
| Plan | 1.1 Read the support matrix and sizing recommendations<br>1.2 Choose an AWS region |

| | |
|---|---|
| Prepare | 2.1 Install cURL and Unzip<br>2.2 Subscribe to the AWS marketplace<br>2.3 Create an AWS access key ID and secret access key in AWS IAM<br>2.4 Enable your Docker Hub account<br>2.5 Download and unzip the AWS package<br>2.6 Prepare SSH keys<br>2.7 Configure the Terraform parameters<br>2.8 Run the Packer scripts to build the AMIs<br>2.9 Set up databases in Amazon RDS<br>2.10 Set up an NFS server<br>2.11 (Optional) Set up a bastion host to access the CDF and ITSMA instances on AWS EC2 |
| Install ITOM CDF and ITSMA | 3.1 Run the Terraform scripts to deploy CDF on AWS<br>3.2 Run the suite installer |

## Plan a cloud-based suite deployment

Before you install ITSMA, read the support matrix and sizing recommendations, and choose an AWS region.

### Support matrix (cloud-based)

- Supported environments
- Operating systems
- Required network identification (FQDN)
- Databases
- Supported deployment mode
- Browsers
- Language support
- Other requirements

#### Supported environments

- Cloud environment (Amazon Web Services)

#### Operating systems

- 64-bit CentOS 7.2, 7.3 (the Linux kernel version must be 3.10.0-514.26.2.el7.x86_64 or above)

> The master node, worker nodes, and the NFS server hosts must use the same operating system.

#### Required network identification (FQDN)

- IPv4

#### Databases

Cloud-based ITSMA production environments require an external database instance created in Amazon Relational Database Service (Amazon RDS). The database instance must be either Oracle or PostgreSQL.

#### Supported deployment mode

Mixed mode is not supported. Only containerized mode is supported if you deploy ITSMA in the cloud. For more information about mixed mode, see Deployment modes.

## Browsers

Use the the following browsers to access the CDF Management Portal and ITSMA:

| Interface | Browser |
|---|---|
| Management Portal | • Microsoft Internet Explorer 11<br>• Google Chrome 48 or later<br>• Mozilla Firefox 45 ESRWindows version<br>• Apple Safari 10.1 |
| • Service Management<br>• CMDB<br>• CMDB Browser<br>• Suite Configuration | • Internet Explorer (IE) 11<br>• Latest Firefox<br>• Latest Chrome |
| Service Portal | • Internet Explorer (IE) 11<br>• Latest Firefox<br>• Latest Chrome<br>• Edge |

- 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 Internet Explorer 11 (see Known issues, limitations, and workarounds).
- To access the CMDB capability UI from the latest Chrome or Firefox, enable JNLP by following instructions here. This is a one-time operation.

## Language support

The following tables provides language support information about CDF (Management Portal) and ITSMA.

| Capabilities/Modules | Supported languages | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | English | Arabic | Brazilian Portuguese | Danish | Dutch | French | German | Greek | Italian | Japanese | Korean | Norwegian | Polish | Turkish | Russian | Simplified Chinese | Spanish |
| Service Management | | | | | | | | | | | | | | | | | |
| CMDB | | | | | | | | | | | | | | | | | |
| • Management Portal<br>• Suite Configuration | | | | | | | | | | | | | | | | | |
| Service Portal | | | | | | | | | | | | | | | | | |

- Except for Service Portal, the displayed language for ITSMA capabilities can be switched by changing the language settings in your browser. For Service Portal, the displayed language is set according to the setting in the user profile in Service Portal as s pecified by individual users.
- In Service Portal, when you use a language that is only supported by Service Portal, you might see some mixed languages or languages remained in English in some forms as labels, action labels, drop-down options, and so on. This is because such data comes from other ITSMA capabilities (such as Service Management) that do not support the current language specified in Service Portal user profile.
- For suite configuration, when you access the Service Portal theme setting and feature setting pages, you might see a different language than the supported four suite configuration languages. This is because the Service Portal theme setting and feature setting pages are part of the Service Portal application in which the displayed language is specified in the user profile in Service Portal instead of in your browser language setting. To avoid this, set the language in Service Portal user profile to the same language that you use in suite configuration.
- Be aware that, if using Firefox or Windows 7 and Internet Explorer 11, you must select <*Language*> instead of "<*Language*> (<*Country*>)" when you set your browser language.

### 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 | Android 7.x, 6.x, 5.x with Android browser | For accessing the Mobility capability of the ITSMA suite |

## Sizing (cloud-based)

When you install ITSMA, you will need to select a suite size: Demo, Extra Small, Small, Medium, or Large. Different suite sizes require different hardware configurations.

For sizing information including hardware requirements, refer to Sizing.

**Sizing differences**
The sizing recommendations apply for both on-premises and cloud-based deployments, except that in the Tuning configuration section only the following tasks are needed for cloud-based deployments:

- *Tuning before installation*

  Task 4: Tune database parameters

- *Tuning after installation*

  Task 1: Browser SSL certificate
  Task 3: Increase ITSMA IDM JVM heap size
  Task 4: Configure additional SM_RTE_ARGS parameters
  Task 5: Remove internal PostgreSQL pods manually when utilizing external databases

## Choose an AWS region

CDF requires Amazon Elastic File System (EFS), but this feature is not available in all AWS regions. When you select an AWS region, you must choose one that supports EFS (if you do not, the Terraform scripts will fail to install CDF).

At the time of release, the following regions support EFS.

| Global area | Regions that support EFS |
|---|---|
| Americas | 1. Northern Virginia<br>2. Ohio<br>3. Oregon |
| EMEA | 1. Ireland<br>2. Frankfurt |
| Asia Pacific | 1. Sydney |

For the latest information and list of AWS regions that support EFS, refer to the following Amazon website:

AWS region table

# Prepare for installation (cloud-based)

Before you can install ITSMA on AWS, you must perform the following preparatory tasks in the order they are listed:

- Download CDF and AWS packages
- Bind your AWS elastic IP address to your public FQDN

- Subscribe to the AWS marketplace
- Create an AWS access key ID and secret access key in AWS IAM
- Install cURL and Unzip
- Enable your Docker Hub account (cloud-based)
- Prepare SSH keys
- Run the Packer scripts to build the AMIs
- Configure the Terraform parameters
- Set up databases in Amazon RDS
- Set up an NFS server
- (Optional) Set up a bastion host to access the CDF and ITSMA instances on AWS EC2

## Download CDF and AWS packages

Before starting using this solution, you need to set up a provisioning environment. The environment can be Windows (with bash support), Linux or Mac. The scripts provided in this solution are tested on Centos, Ubuntu and Windows Git Bash (MINGW emulator).

Once your provisioning environment is ready, download the CDF installation package and the AWS package to this environment.

### Download the CDF installation package

To deploy ITSMA on AWS, you need to download the CDF installation package to the provisioning environment that you have prepared.

The CDF installation package has two versions:

- The CDF installation package that you can download from the HPE Software Entitlement website. This CDF package supports only one-master deployments.
- The CDF installation package that is shipped with the ITSMA Containerized 2017.07.001 patch. This CDF package supports both one-master and multi-master deployments. If you have not installed CDF yet, we recommend that you download this version.

For details, see Download the CDF installation package (on-premises).

The CDF installation package includes a HPESW_ITOM_Suite_Foundation_2017.06.00xxx.zip file ("CDF master zip") and a HPESW_ITOM_Suite_Foundation_Worker.zip file ("CDF worker zip", which is under the **zip** subfolder). You need to place the two zip files in a temporary directory in your provisioning environment. For example:

- /tmp/HPESW_ITOM_Suite_Foundation_2017.06.00xxx.zip
- /tmp/HPESW_ITOM_Suite_Foundation_Worker.zip

> Later, you will need to specify CDF_MASTER_ZIP_PATH and CDF_WORKER_ZIP_PATH when running the Packer scripts to to build the Amazon Machine Images (AMIs). See Run the Packer scripts to build the AMIs.
>
> For example:
>
> CDF_MASTER_ZIP_PATH="/tmp/HPESW_ITOM_Suite_Foundation_2017.06.00xxx.zip"
>
> CDF_WORKER_ZIP_PATH="/tmp/HPESW_ITOM_Suite_Foundation_Worker.zip"

### Download the AWS package

Download the AWS package to the provisioning environment that you have prepared.

> You can deploy one or three master nodes on AWS. To deploy three master nodes, you must download the latest AWS package (itsma-aws-provisioning-2017.07.001_20170921.100124-17.zip).

1. Click the following link to download the AWS package from the HPE Marketplace:
   https://marketplace.saas.hpe.com/itom/content/itsma-containerized-suite-2017-07-express-edition

   The AWS package has two versions. Choose one of them depending on whether you want to deploy multiple master nodes:
   - If you want to deploy only one master node, download the CDF itsma-aws-provisioning-xxxx.zip package. This package supports only one-master deployments.
   - If you want to deploy multiple master nodes, download the itsma-aws-provisioning-2017.07.001_20170921.100124-17.zip package. This package supports multi-master deployments, and must be used with the CDF installation package shipped with the ITSMA Containerized 2017.07.001 patch.
2. Run one of the following commands to unzip the package:
   **unzip itsma-aws-provisioning-xxxx.zip**

   **unzip itsma-aws-provisioning-2017.07.001_20170921.100124-17.zip**

> "xxxx" is a placeholder for the version number.

3. The unzipped package contains the following directories:
   - Packer
   - Terraform

## Bind your AWS elastic IP address to your public FQDN

To set up ITSMA on AWS, you must bind your AWS elastic IP address to a public FQDN that you want to use for your ITSMA deployment.

1. Obtain an elastic IP address from AWS.

   > You will need to specify this IP address in the *aws_eip_public_ip* parameter in the variables.tpl file (see Configure the Terraform parameters).

2. Buy a public domain name from AWS. For example: itsma.ng.net.
3. Create a DNS record and bind it to your AWS elastic IP address.
   a. Log in to your domains page at AWS: https://console.aws.amazon.com/route53/home.
   b. Select **Hosted zones**.
   c. On the Hosted zones page, select the name of the public domain that you want to use.
   d. Click **Create RecordSet**.

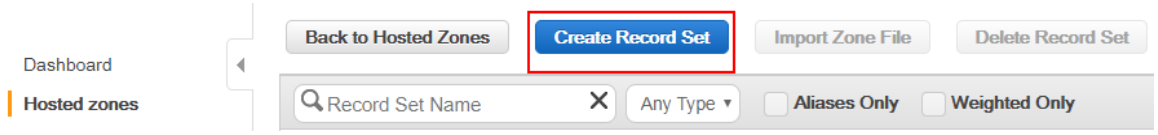


   e. Enter a DNS name in the Name field, and enter your AWS elastic IP address in the Value field.

   > You will need to specify this FQDN in the external-access-fqdn parameter in the variables.tpl file (see Configure the Terraform parameters). For example: *external-access-fqdn=myitsma.itsma-ng.net*

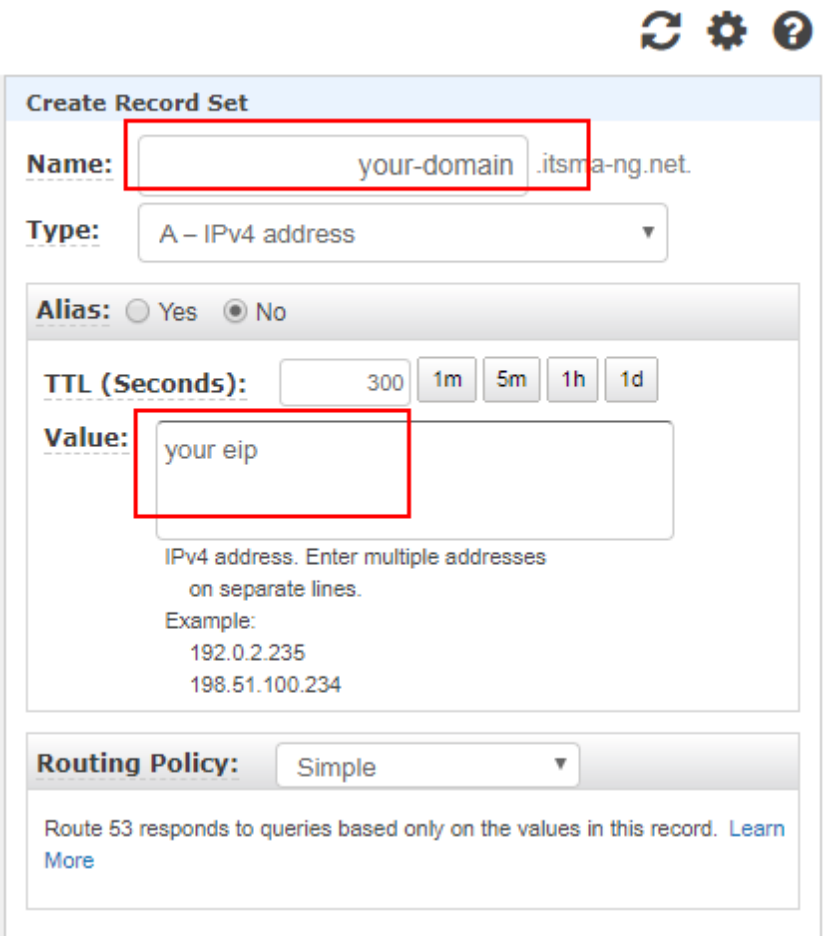


## Subscribe to the AWS marketplace

To subscribe to the AWS marketplace, you must launch and then terminate an AMI. To do this, follow these steps:

1. Visit the Amazon EC2 console. If you do not already have an AWS account, create one.
2. On the EC2 dashboard, click **Launch Instance**.

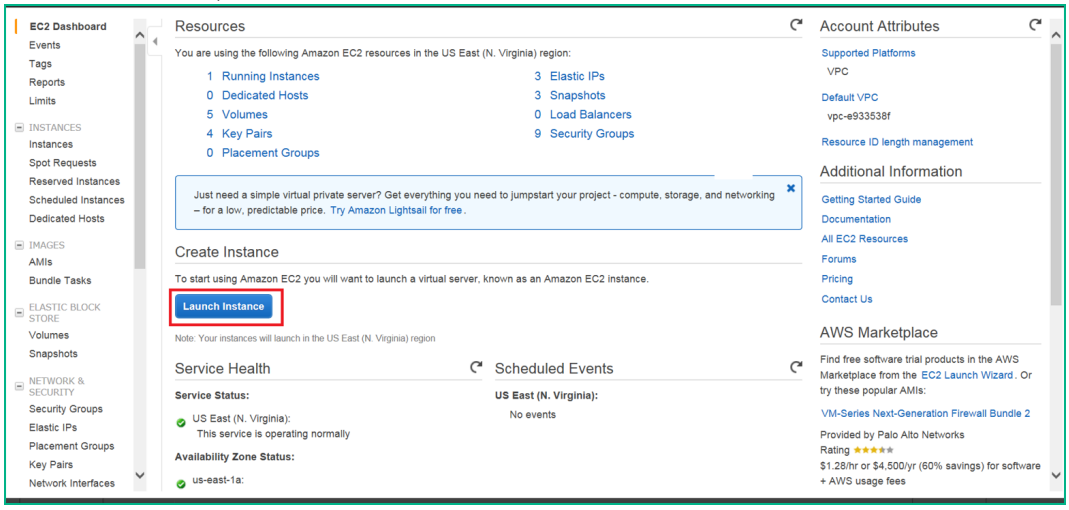


3. Select AWS Marketplace, and then search for and select the **CentOS 7 (x86_64) - with Updates HVM** image.

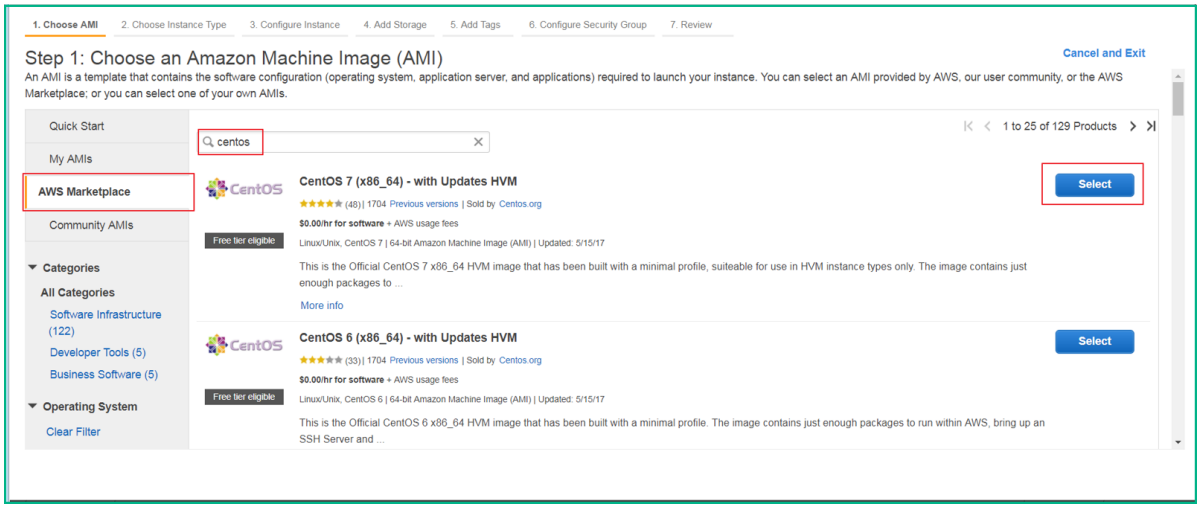


4. Click **Review and Launch**, and then click **Launch**.
5. Select or create a key pair, and then click **Launch Instances**.
6. Click **View Instance**. Make sure that the instance is displayed before you proceed.
7. In the navigation pane, select **Instances**.
8. Select the instance, select **Actions**, click **Instance State**, and then click **Terminate**.
9. When prompted, select **Yes, Terminate**.

> For more information about the CentOS 7 (x86_64) - with Updates HVM image, see the following Amazon website:
>
> CentOS 7 (x86_64) - with Updates HVM

## Create an AWS access key ID and secret access key in AWS IAM

> You may reuse an existing access key ID and secret access key if you have already created them.

1. Click the following link to log on to the AWS Identity & Access Management (IAM) console: https://console.aws.amazon.com/iam/
2. In the navigation pane, select **Users**, and then enter a user name.

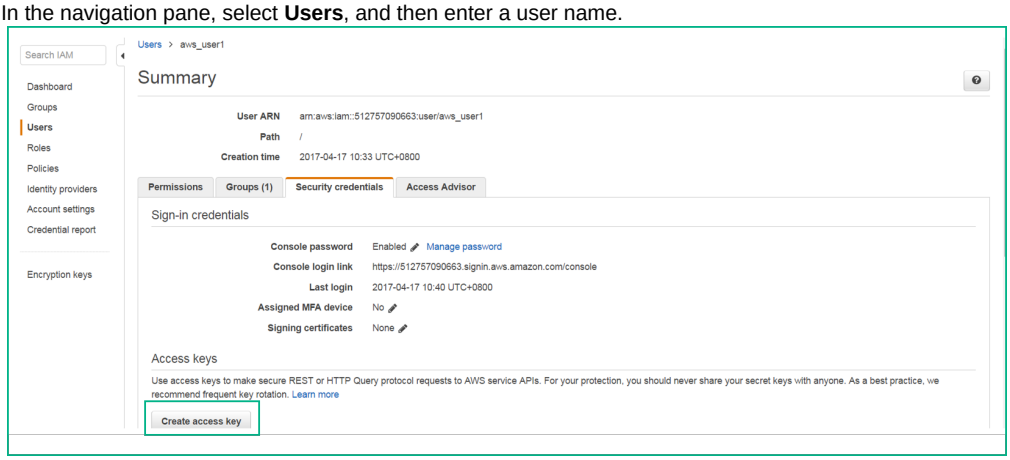


3. On the **Security credentials** tab, click **Create access key**.
4. Click **Download .csv file** to save the access key ID and secret access to your local directory.

> Store this file in a secure location. You will not have access to the secret access key after this dialog box closes.

5. After you have downloaded the .csv file, click **Close**.

## Install cURL and Unzip

To install cURL and Unzip, run the following command:

**$ sudo yum install -y curl unzip**

To verify that cURL and Unzip are installed correctly, run the following commands:

**$ curl --version**

**$ unzip -hh**

## Enable your Docker Hub account (cloud-based)

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

To enable your Docker Hub account, follow these steps:

1. Create a Docker Hub account if you do not already have one:
    a. Go to https://hub.docker.com.
    b. Enter a Docker ID, your company email address, and then a password.
    c. Click **Sign Up**.
    d. You receive an email from Docker Hub, asking you to confirm your email address. Confirm your email address swiftly.
2. Log in to https://hub.docker.com with your Docker ID.
3. On the top right corner of the page, click **Settings** under your avatar and take a screenshot to include your Docker ID and the linked email address.
4. Send the following information together with the screenshot to the HPE software fulfillment and licensing team specific for your region to enable your Docker Hub account:
    - Your company name
    - Your HPE Customer SAID (must be valid and active)
    - HPE ITOM Suite edition (that is, ITSMA)

> Email addresses of the HPE software fulfillment and licensing teams for different regions:
>
> Americas region: dockersupport.ams@hpe.com
> APJ region: dockersupport.apj@hpe.com
> EMEA region: dockersupport.emea@hpe.com
>
> Once your Docker ID is enabled, you will receive a confirmation from HPE.

## Prepare SSH keys

SSH keys are required in order to connect to AWS EC2 instances. You must generate new SSH keys if you do not have any, or modify existing keys if you wish to reuse them. Expand the relevant section to view detailed steps.

### Create new SSH keys

Follow the steps below to generate the SSH keys:

1. Run the following command to navigate to the Terraform directory: **$ cd terraform**
2. Generate the SSH keys with the following command: **$ ssh-keygen -f itom-key**
3. The SSH-keygen tool will prompt you to enter a passphrase. Enter the passphrase. If there is no passphrase, press Enter twice.

```
# [root@mcimaster1 terraform]# ssh-keygen -f itom-key

Generating public/private rsa key pair.:

Enter passphrase (empty for no passphrase):
```

4. The SSH keys are stored in the ./terraform/itom-key file.

### Reuse existing SSH keys

To reuse existing SSH keys, you must copy your private key and corresponding public key to the ./terraform directory, and change the file names to *itom-key*. When you are finished, the directory layout will resemble the following:

```
terraform
 itom-key

 itom-key.pub

 deploy_cdf_on_aws.sh

 destroy_cdf_on_aws.sh

 variables.tpl

 tpl/
```

## Run the Packer scripts to build the AMIs

To build the Amazon Machine Images (AMIs) by running the Packer scripts, follow these steps:

1. Run the following command to navigate to the Packer directory:
   **$ cd Packer**
2. Open the ami.properties file, and then set values for the following parameters:
   - *AWS_ACCESS_KEY_ID*
   - *AWS_SECRET_ACCESS_KEY*
   - *AWS_DEFAULT_REGION*
   - *AWS_DEFAULT_ZONE*
   - *CDF_VERSION*
   - *CDF_SOURCE_AMI_ID*
   - *CDF_MASTER_ZIP_PATH*
   - *CDF_WORKER_ZIP_PATH*

   For example, you configure the parameters as follows:

   ```
   AWS_ACCESS_KEY_ID="AKIAICCGXXXXXXXXXX"
   AWS_SECRET_ACCESS_KEY="ijEJD761vgxBhWo2z1aXXXXXXXXXX"
   AWS_DEFAULT_REGION="us-east-1"
   AWS_DEFAULT_ZONE="us-east-1b"
   CDF_VERSION="2017.06.0068"
   CDF_SOURCE_AMI_ID="ami-6d1c2007"
   CDF_MASTER_ZIP_PATH="/tmp/HPESW_ITOM_Suite_Foundation_2017.06.00xxx.z
   ip"
   CDF_WORKER_ZIP_PATH="/tmp/HPESW_ITOM_Suite_Foundation_Worker.zip"
   ```
3. Build the master and worker AMIs. To do this, run the following commands:
   - **$ sudo chmod +x ./create_amis.sh**
   - **$ sudo ./create_amis.sh**
4. When the execution is complete, make a note of the worker and master AMI IDs.

## Configure the Terraform parameters

To correctly deploy CDF and ITSMA on AWS, you must first configure the parameters in the variables.tpl file (located in the /itsma-aws-provisioning-xxxx/terraform directory) and in the variables.main.tpl file (in the /itsma-aws-provisioning-xxxx/terraform/tpl directory).

**HA support**
You can select to deploy one or three master nodes on AWS. To deploy three master nodes, make sure that you have downloaded the appropriate CDF and AWS packages. See Download CDF and AWS packages.

### Parameters in the variables.tpl file

| Parameter | Description |
|---|---|
| aws_access_key_id | The *aws_access_key_id* parameter defines the AWS access key ID that is used to deploy CDF.<br>Example:<br>`aws_access_key_id="AKIAICCGXXXXXXX"` |
| aws_secret_access_key | The *aws_secret_access_key* parameter defines the AWS secret access key that is used to deploy CDF.<br>Example:<br>`aws_secret_access_key="ijEJD761vgxBhWoXXXXXXX"` |
| aws_resource_prefix | The *aws_resource_prefix* parameter defines the prefix for all AWS resources, such as the VPC, subnet, and security groups.<br>Example:<br>`itom-itsma1` |
| aws_eip_public_ip | The *aws_eip_public_ip* parameter defines the AWS Elastic IP (EIP) that is dedicated to the deployment. |
| region | The *region* parameter specifies the geographical AWS region.<br>Example:<br>`region ="us-east-1"` |
| zone | The *zone* parameter specifies the geographical AWS zone.<br>Example:<br>`zone ="us-east-1b"` |

| | |
|---|---|
| sizing_profile | The *sizing_profile* parameter defines the size profile of the deployment. You can enter one of the following values (case-insensitive): **Demo**, **Extra Small**, **Small**, **Medium**, or **Large**.<br><br>Example:<br><br>`demo` |
| enable_multiple_master | Enables HA support. If set to true, three master nodes will be deployed; if set to false, one master will be deployed.<br><br>Example:<br><br>`true` |
| external-access-fqdn | The *external-access-fqdn* parameter defines the fully-qualified hostname that external clients use to access cluster services. The specified name must resolve to the IP address on which the ingress is running.<br><br>Example:<br><br>`external-access-fqdn=myd.XXXX.YYY.net` |
| enable-suite | The *enable-suite* parameter enables the suite installer. The possible values for this parameter are "true" and "false". By default, the value is "false".<br><br>If the *enable-suite* parameter is set to "true", you must also configure the *suite-name*, *suite-version*, *dockerhub-username*, and *dockerhub-password* parameters.<br><br>Example:<br><br>`enable-suite = false` |
| dockerhub-username | The *dockerhub-username* parameter specifies the user name of the Docker hub account that is used to access the suite images.<br><br>Example:<br><br>`dockerhub-username = "mydockerhubname"` |

| | |
|---|---|
| dockerhub-password | The *dockerhub-password* parameter specifies the password of the Docker hub account that is used to access the suite images.<br><br>Example:<br><br>`dockerhub-password = "mydockerhubpassword"` |

### Parameters in the variables.main.tpl file

| Parameter | Description |
|---|---|
| amis-master | The *amis-master* parameter defines the AMI ID that is used to deploy the CDF master node. This is a "map" type parameter.<br><br>Example:<br><br>`us-east-1 = "ami-c40757d2"` |
| amis-worker | The *amis-worker* parameter defines the AMI ID that is used to deploy the CDF worker nodes. This is a "map" type parameter.<br><br>Example:<br><br>`us-east-1 = "ami-c50e5ed3"` |

## Set up databases in Amazon RDS

Both ITOM Container Deployment Foundation (CDF) and ITSMA have an embedded PostgreSQL database. However, in a production environment, we recommend that you set up your own database instances in Amazon Relational Database Service (Amazon RDS).

For information about supported databases, see Support matrix (cloud-based).

### Create a VPC for the database

Before you can create a database instance in Amazon RDS, you must create a publically-accessible Amazon Virtual Private Cloud (Amazon VPC) in which to put it.

- You must create two subnets in the VPC, each with a different availability zone.
- You must create a database subnet group for the VPC.
- You must create a new security group for the VPC, or add a custom TCP rule to an existing VPC security group. The new group or rule should allow source 0.0.0.0/0 to access port 1521 (for an Oracle database) or port 5432 (for a PostgreSQL database).

For more information about how to do this, refer to the following Amazon website:

Working with an Amazon RDS DB Instance in a VPC

### Create a database instance in RDS

To create a database instance in Amazon RDS, follow these steps:

1. Sign in to the AWS Management Console, and then open the Amazon RDS console.
2. Click **Launch a DB Instance**, and then select a database that is supported by ITSMA.
3. Follow the instructions for your selected database.

#### Create an Oracle database instance

1. Click **Select** beside Enterprise Edition or Standard Edition, depending on your choice of database version.

For more information about pricing and the differences between versions, see Amazon RDS for Oracle Database.


2. On the **Production?** page, select **Production** or **Dev/Test**, depending on whether you will use the database instance in a production or test environment, and then click **Next Step**.
3. Use the information in the following table to populate the fields in the **Specify DB Details** page, and then click **Next Step**.

| Parameter | Value |
|---|---|
| DB Engine Version | Oracle 12.1.0.2.v8 |
| DB Engine Class | db.m4.4xlarge (or larger) |
| Multi-AZ Deployment | Yes/No, depending on whether HA is required for the database |
| Storage Type | General Purpose (SSD) |
| Allocated Storage | Value depends on your deployment plan; at least 100 GB |
| Settings | Value depends on your deployment plan |

4. Use the information in the following table to populate the fields in the **Configure Advanced Settings** page, and then click **Launch DB Instance**.

| Setting | Value |
|---|---|
| VPC | The VPC that you created in Create a VPC for |
| Subnet Group | The subnet group that you created in Create a database |
| Public Accessible | Yes |
| Availability Zone | Value depends on your deployment plan |
| VPC Security Group(s) | The security group that you created in Create a database |
| DB Parameter Group | Refer to the right-hand sidebar in AWS for help |
| Option Group | Refer to the right-hand sidebar in AWS for help |

5. When the instance status is "available", make a note of the endpoint. You will need this information to configure the database connection when you run the suite installer.

   Here is an example endpoint URL: **itsma-201707-xxx-rds-oracle.ctz1ruxhq5vf.us-west-2.rds.amazonaws.com:1521**.

### Create a PostgreSQL database instance

> A database user named "**postgres**" with the "**rds_superuser**" role must already exist prior to ITSMA installation, and you must use this "**postgres**" user when configuring the PostgreSQL database during your ITSMA installation. Otherwise, your ITSMA installation may fail when using an RDS postgres service.
>
> If you are creating an RDS postgres service from scratch, set the **master username** to "postgres".

1. Click **Select** beside PostgreSQL.
2. On the **Production?** page, select **Production** or **Dev/Test**, depending on whether you will use the database instance in a production or test environment, and then click **Next Step**.
3. Use the information in the following table to populate the fields in the **Specify DB Details** page, and then click **Next Step**.

| Parameter | Value |
|---|---|
| DB Engine Version | PostgreSQL 9.6.2-R1 |
| DB Engine Class | db.m4.4xlarge (or larger) |
| Multi-AZ Deployment | Yes/No, depending on whether HA is required for the database |
| Storage Type | General Purpose (SSD) |
| Allocated Storage | Value depends on your deployment plan; at least 100 GB |
| Settings | Value depends on your deployment plan |

4. Use the information in the following table to populate the fields in the **Configure Advanced Settings** page, and then click **Launch DB Instance**.

| Setting | Value |
|---|---|
| VPC | The VPC that you created in Create a VPC for |
| Subnet Group | The subnet group that you created in Create a database |
| Public Accessible | Yes |
| Availability Zone | Value depends on your deployment plan |
| VPC Security Group(s) | The security group that you created in Create a database |
| DB Parameter Group | Refer to the right-hand sidebar in AWS for help |

5. When the instance status is "available", make a note of the endpoint. You will need this information to configure the database connection when you run the suite installer.

   Here is an example endpoint URL: **201707-xxx-rds-postgres.ctz1ruxhq5vf.us-west-2.rds.amazonaws.com:5432**.

For more detailed information about using Amazon RDS, refer to the following Amazon websites:

Creating an Oracle DB Instance and Connecting to a Database on an Oracle DB Instance

Creating a DB Instance Running the Oracle Database Engine

Creating a PostgreSQL DB Instance and Connecting to a Database on a PostgreSQL DB Instance

## Set up an NFS server

If a container stops and is then restarted, all changes made inside the container are lost. If you want to save data such as configuration files or databases, the data must be stored outside the container in a persistent volume provided by a Network File System (NFS). This release of ITSMA uses three separate NFS persistent volumes ("NFS shares") :

- Global volume: this volume stores global suite data (configuration files, logs, SSL certificate files, and so on).
- Smart Analytics (SMA) volume: this volume stores Smart Analytics data.
- Database volume: this volume stores the database.

These persistent volumes can be provided by one or more NFS servers. You must set up an NFS server and a shared directory for each of the NFS volumes. To install a cloud-based NFS server and to set up a shared directory, follow these steps.

> We strongly recommend using dedicated NFS servers in a production environment. Use the master node as the NFS server only in a test environment.

### Prepare an EC2 instance

> - The following steps apply only if you do not have any existing private AMI images.
> - VPC peering connection is currently not available. The NFS server must be created on the same VPC that runs the master and worker nodes.

1. Log on to the AWS management console.
2. Click **Launch Instance** and select the **CentOS 7 (x86_64) - with Updates HVM** image from AWS Marketplace.
3. Select the desired type of Instance (we recommend **m4.xlarge** or larger).
4. Configure the instance according to your needs.
    - You must select the **Disable for Auto-assign Public IP** option.
    - The volume type of the Root Volume must be **IO1**.
    - The volume size must be at least 200 GB with IOPS above 1280.
    - Add Tags if necessary.
5. Create new security groups.
6. Create a new key pair, download the private key (.pem) file, and then click **Launch**.
7. On your Instances panel, select the newly-created instance, and then click **Connect**. Follow the instructions and make sure that you can connect to your instance.
8. Click **Actions** to make any additional customized configurations.

### Install an NFS server

If you use the same NFS server for both ITOM CDF and ITSMA, you should have already an NFS server set up and should therefore skip this step. If you use another NFS server for all of the three volumes, this task is a one-time effort. If you use different NFS servers for the volumes, repeat this task for each of the servers.

To install an NFS server, follow these steps:

1. Log on to the NFS server.
2. Make sure that the **rpcbind** package is installed on the NFS server (see Meet the prerequisites (on-premises)). 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**

### Set up a shared directory

Once the NFS server is ready, set up the NFS volume. Repeat this task three times to set up three volumes to store global data, Smart Analytics data, and database data of ITSMA.

The following are example mount paths of the NFS volumes:

- /var/vols/itom/itsma/itsma-itsma1-smartanalytics
- /var/vols/itom/itsma/itsma-itsma1-db
- /var/vols/itom/itsma/itsma-itsma1-global

Repeat the following steps to set up each of the ITSMA NFS volumes:

1. Log on to the NFS server.

2. Create a directory to store relevant suite information:

   **mkdir -p <ITSMA NFS volume>**

   The following is an example for the global NFS volume:

   **mkdir -p /var/vols/itom/itsma/itsma-itsma1-global**

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

3. Configure the NFS share in the /etc/exports file by adding the following line:
   **<ITSMA NFS volume> *(rw,sync,anonuid=1999,anongid=1999,all_squash)**
   For example, add the following line:
   **/var/vols/itom/itsma/itsma-itsma1-global *(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 volume>**

   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-global**

### (Optional) Set up a bastion host to access the CDF and ITSMA instances on AWS EC2

You can access the Amazon Elastic Compute Cloud (EC2) CDF instance in either of the following ways:

- An access URL provided by Terraform or from the Amazon S3 console after CDF is deployed
- A bastion host that is connected to the EC2 CDF instance

To set up a bastion host, follow these steps:

1. Run the following commands to copy the private key file (itom-key) to your bastion host:

   **$ chmod 400 itom-key**

   **$scp -i itom-key./itom-key ec2-user@<*bastion_public_ip*>:~**
2. Run the following command to connect the private key file to your bastion host:

   **$ssh -i itom-key ec2-user@<*bastion_public_ip*>**
3. Run the following command to connect the bastion host to your EC2 instances:

   **$ sudo ssh -i -itom-key centos@<*ip of master1*>**

For example, you run the following commands to copy the private key file, connect the private key file to your bastion host, and connect the bastion host to your EC2 instances:

**$scp -i itom-key./itom-key ec2-user@34.200.237.87:~**

**$ssh -i itom-key ec2-user@34.200.237.87**

**[ec2-user@ip-10-0-0-100 ~] $ sudo ssh -i -itom-key centos@10.0.0.10**

## Install CDF and ITSMA (cloud-based)

The process to install CDF and ITSMA on AWS comprises the following tasks:

- Run the Terraform scripts to deploy CDF on AWS
- Run the suite installer (cloud-based)

Before you start this process, make sure that you have planned your deployment and prepared the deployment environment.

### Run the Terraform scripts to deploy CDF on AWS

To deploy the CDF master node and worker nodes on AWS by using the Terraform scripts, follow these steps:

1. Run the following command to navigate to the Terraform directory:
   **$ cd terraform**
2. Run the following command to deploy the master node and worker nodes on AWS:
   **$ ./deploy_cdf_on_aws.sh**

> The number of worker nodes is dependent on the sizing profile. The sizing profile is specified by the value of the *sizing-profile* paramete

r. For more information, see Configure the Terraform parameters and Sizing (cloud-based).

When the CDF deployment is completed, text that resembles the following is displayed:

```
$ ./deploy_cdf_on_aws.sh

aws_efs_mount_target.cdf: Still creating... (2m30s elapsed)

aws_efs_mount_target.cdf: Still creating... (2m40s elapsed)

aws_efs_mount_target.cdf: Creation complete (ID: fsmt-4e089807)

Apply complete! Resource: 24 added, 0 changed, 0 destroyed.

The state of your infrastructure has been saved to the path below. This
state is required to modify and destroy your

infrastructure, so keep it safe. To inspect the complete state use the
'terraform show' command.

State path:

Outputs:

bastion_public_ip = 34.200.237.87

elastic_ip = 34.230.170.78

region = us-east-1

[3/3] terraform apply finished.

The terraform template is saved in [dist-20170704-134813]

Wait a moment...

Access URL: https://<external-access-fqdn>:5443
```

The access URL uses the <external-access-fqdn> value that you defined in the variables.tpl file. This FQDN must be a public FQDN (not the AWS default FQDN).

The script will first generate a local folder whose name is formatted "dist-YYYYMMDD-HHMMSS" (for esxample, ist-20170607-131028). This folder contains all the runtime data needed by Terraform. This folder is also required when you destroy the environment by using Terraform. Note that the Terraform solution is still working in the background when this folder is created, and that you must wait for some time before it completes.

The installation log files are written to a newly-created s3 bucket. To check the progress of the installation, follow these steps:

1. Visit the following website and sign in to the Amazon S3 console: https://console.aws.amazon.com/s3/home.
2. Browse to the newly-created S3 bucket.
3. Click the Refresh icon in the top right corner, and then check the status of the current files.
   The following is an example for a one-master deployment:

```
status.instance.master.1.ok
status.instance.worker.10.0.0.107.ok
status.instance.worker.10.0.0.108.ok
```

The following is an example for a multi-master deployment:

```
status.instance.master.1.ok
status.instance.master.2.ok
status.instance.master.3.ok
status.instance.worker.10.0.0.107.ok
status.instance.worker.10.0.0.108.ok
```

## Run the suite installer (cloud-based)

Once the ITOM Container Deployment Foundation (CDF) is already deployed on AWS, you are ready to install the ITSMA suite. The suite installation steps in a cloud-based environment are basically the same as in an on-premises environment (see Run the Suite Installer); however, there are some slight differences between an on-premises suite deployment and a cloud-based one, which are described in the following.

**Deployment mode**:
You can only select fully containerized mode when deploying ITSMA on AWS. That is, all suite components must be containerized. Mixed mode is not supported.

**Database configuration**:

During the suite installation, keep the following in mind when configuring databases for the suite:

- Oracle and PostgreSQL database server host and port settings: you must specify the these settings based on the endpoint URLs that are generated when you set up the database instances in Amazon RDS. For more information, see Set up databases in Amazon RDS. The following are two example endpoint URLs:
  - PostgreSQL: **itsma-201707-xsp-rds-oracle.ctz1ruxhq5vf.us-west-2.rds.amazonaws.com:5432**
  - Oracle: **itsma-201707-xsp-rds-oracle.ctz1ruxhq5vf.us-west-2.rds.amazonaws.com:1521**
- PostgreSQL user name and password settings: when configuring PostgreSQL databases, you must specify a "**postgres**" user. This database user must already exist and be assigned the "**rds_superuser**" role. For more information, see Set up databases in Amazon RDS.