ITSM Automation NG Express

Software release version: 2017.07

Troubleshooting Guide

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

Hewlett Packard Enterprise

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

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted rights legend

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

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation updates

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/.

1. Troubleshoot . . . 2
1.1 FAQs (CDF) . . . 2
1.2 FAQs (ITSMA) . . . 4
1.3 Troubleshoot HPE ITOM Container Deployment Foundation . . . 10
1.3.1 Manual verification commands . . . 11
1.3.2 Support toolset . . . 12
1.3.3 ITOM Container Deployment Foundation logs . . . 20
1.3.4 ITOM Container Deployment Foundation issues . . . 21
1.4 Troubleshoot the ITSMA suite . . . 25
1.4.1 ITSMA suite logs . . . 25
1.4.1.1 Check the ITSMA installation log . . . 26
1.4.2 ITSMA service dependency . . . 27
1.4.3 ITSMA suite issues . . . 33
1.4.4 Smart Analytics troubleshooting . . . 38
1.4.5 ITSMA suite configuration troubleshooting . . . 43

Troubleshoot

This section provides guidelines that can help you troubleshoot problems that you may encounter when installing and using ITOM Container Deployment Foundation (CDF) and the ITSMA suite.

- FAQs (CDF)
- FAQs (ITSMA)
- Troubleshoot HPE ITOM Container Deployment Foundation
- Troubleshoot the ITSMA suite

FAQs (CDF)

The following is a list of frequently asked questions about installing and using ITOM Container Deployment Foundation (CDF).

- Installation
- LDAP
- Login
- Miscellaneous

Installation

Do we have a way to upgrade the suite images in CDF?

Answer:
You will need to reinstall everything including ITOM Container Deployment Foundation and the suite if you want to upgrade the suite images.

Will the logs located in a pod be deleted when the pod is deleted?

All the suite logs are currently stored within the persistent volume so you will not lose them if the **Pod**s go down.

What is the recommended way to restart the CDF cluster (foundation components)?

Answer:

You can run the following command to delete pods (which will then be redistributed):
kubectl get pod --namespace=<namespace> | grep '<expression>' | cut -d " " -f1 - | xargs -n1 -P 10 kubectl delete pod --namespace=<namespace>

How do I deal with CDF installation failure related errors?

Answer:

- Check if Docker (including bootstrap Docker) daemon is running or not.
- Identify the relevant service.
- Check the Pods status on the back end.
- Run this command for more details: `kubectl describe pod podName -n namespaces`.
- Check the containers' logs of relevant pods.

What could I do when a worker node installation failed with flannel related error?

Answer:

- Double check whether the FQDN is resolved to correct IP address on master node or not.
- On master node, run `kube-restart.sh` under **$K8S_HOME/bin** directory.
- Reinstall the worker node from the ITOM Container Deployment Foundation management portal.

Why couldn't I install a suite on CDF with a 503 nginx error for the suite installer?

Answer:

Synchronize the time among the master and worker nodes.

When installing CDF on a node with multiple network adapters, the installation failed with a "fail to pull image" error

Answer:

Follow these steps below to reinstall the CDF on this node.

1. Uninstall CDF on this node with the **uninstall.sh** command.
2. Reboot the server.
3. Configure the FLANNEL_IFACE parameter in the install.properties file.
4. Follow the installation instructions to install CDF on that node.

LDAP

What is the difference between the LDAP in CDF and the LDAP in ITSMA?

Answer:

The LDAP feature in CDF was from IdM. CDF provides an LDAP interface to the suite team if they would like to share the IdM with CDF. The suite team can create the organization in the CDF UI, and use that for authentication; however, the suite team needs to do some integration with that. ITSMA uses its own IdM, and so LDAP in CDF is meaningless to ITSMA customers. In other words, for ITSMA, you need to configure LDAP in the ITSMA user interface (see Configure an external LDAP server).

Login

Why couldn't I see the CDF login page?

The Docker daemon could not be started with an error message: 'Gateway time out' when logging in to IDM.

Answer:

You need to start Kubernetes with the following commands:

```
cd $K8S_HOME/bin
./kube-start.sh
```

Why couldn't I log in to the CDF management portal?

Answer:

- Make sure you entered the correct URL and port: https://<master-hostname>:5443.
- Make sure you can access the host: `ping [master-hostname]`.
- Check your proxy settings in the browser.
- Check the installation logs in `/opt/kubernetes/install-<timestamp>.log`.
- Empty the NFS folder and then install the ITOM Container Deployment Foundation.

Why couldn't I log in to the CDF management portal? A login failure error: "The IDM service is not ready yet." and the **pod:autopass-lm-**, idm-, suite-installer- status are all CrasLoopBackOff ?

Answer:

1. Run the command: `kubectl delete -f autopass-lm.yaml; kubectl delete -f autopass-pg.yaml; kubectl delete -f idm.yaml; kubectl delete -f idm-pg.yaml; kubectl delete -f suite.yaml`
2. Delete the folder under autopass_db, idm_db, suite_db under the NFS folder $NFS_HOME/baseinfra-1.0.
3. Run the command: `kubectl create -f idm-pg.yaml; kubectl create -f idm.yaml; kubectl create -f autopass-pg.yaml; kubectl create -f autopass-lm.yaml; kubectl create -f suite.yaml`

Why couldn't I log in to Docker Hub?

Answer:

- Make sure the user name and password are correct.
- Make sure the Docker HTTP proxy is configured as below:
  /usr/lib/systemd/system/docker.service.d/http_proxy.conf
- Make sure the host HTTP proxy is configured as below:
  export http_proxy https_proxy

Miscellaneous

Are there exceptions required to be added to the firewall policy?

Answer:

Browsers access the HPE ITOM Container Deployment Foundation via HTTPS ports (TCP/5443). End users need to add it to the firewall exception policy.

What could I do when a pod is in pending/ImagePullBackOff or ErrImagePullstatus and when running command: kubectl describe pod podName -n namespaces, an error: image xxx not found"?

Answer:

Make sure the images are pushed into the private docker registry. Run the command: `docker pull [image name]` to double confirm.

What could I do when a pod is in ImagePullBackOff or ErrImagePull status and when running command: kubectl describe pod [pod name] -n [namespace], an error: Error while pulling image: Get http://localhost:5000/v1/repositories/xxx: dial tcp [::1]:5000: getsockopt: connection refused"?

Answer:

You can delete the Docker registry and registry proxy pods, and then restart them.

How can I configure the Docker proxy when Docker-Content-Trust is turned on?

Answer:

Follow the steps below to set the Docker proxy:

1. Run the **docker pull hello-world** command to check if you need to configure a proxy.
2. If you can pull the hello-world image, then you do not need to configure any proxy.
   If you cannot pull a hello-world image, you need to configure a proxy.
   a. To enable the service, run the **systemctl enable docker.service** command.
   b. Run the **mkdir -p /usr/lib/systemd/system/docker.service.d** command.
   c. Run the following commands:

      cat << EOF > /usr/lib/systemd/system/docker.service.d/http_proxy.conf
      [Service]
      Environment="HTTP_PROXY=<Your Proxy>" "HTTPS_PROXY=<Your Proxy>"
      EOF

      > You need to configure another proxy when you set the `Docker-Content-Trust on`. Run the following commands:
      >
      > **export http_proxy= <Your Proxy>**
      > **export https_proxy=<Your Proxy>**

   d. Run the **systemctl daemon-reload** command to reload the configuration.
   e. Run the **service docker restart** command to restart docker.
   f. Run the docker pull hello-world command to check if you can pull images from Docker Hub.

FAQs (ITSMA)

The following is a list of frequently asked questions about installing and using ITSMA.

- Installation
- Load balancing
- Database support
- Language support
- Master data onboarding
- LDAP/SSO
- Identity Manager (IDM)
- DNS
- Integrations
- Reverse proxy
- Miscellaneous

Installation

How do I find a service port that is exposed to external applications?

Answer:

The ports of some services of the suite are exposed to external applications. To get the port of a service:

1. Run the following command:

   ```
   kubectl describe svc {service-name} -n {namespace};
   ```

2. Find the nodePort value, which is the port.

You can access it by {EXTERNAL_ACCESS_HOST}:{service-nodePort}.

> {EXTERNAL_ACCESS_HOST} is defined in the install.properties file when you install ITOM CDF.

Do we support multiple master nodes for ITSMA?

Answer:

Yes. You can download the ITSMA 2017.07.001 patch from the HPE Software Support Online website. This patch includes an updated version of the CDF installation package to support using multiple master nodes in your ITSMA deployment. If you want to use multiple master nodes, do not download the CDF 2017.06 installation package from the HPE Software Entitlement Portal.

How do I know the image tag for pods?

Answer:

Run the following command: **kubectl describe pods <pod names> -n <namespace>**.

What shall I do if one of my pods fails?

Answer:

Run the following command: **kubectl delete pod <name of the pod that failed> -n <namespace>**.

Is there a way to remove and restart the ITSMA image installation?

Question details:

My ITSMA suite installation failed somewhere. I had issues with the image push and one image was not pushed.

Answer:

Re-import the images by running the image import script again. You can run the image import script multiple times. For more information, see Do wnload ITSMA images from Docker Hub to CDF.

Where are the docker proxy settings stored in the CDF environment?

Answer:

These settings are in both the /usr/lib/systemd/system/docker.service.d and /usr/lib/systemd/system/docker-bootstrap.service.d in the http_proxy.conf files.

[Service]

Environment= "http_proxy=http://<proxy>:<port>" "https_proxy=http://<proxy>:<port>" "no_proxy=localhost,127.0.0.1"

If the proxy is misconfigured, you may not be able to run downloadimages.sh successfully. After updating these proxy settings, run kube-restart.sh from /opt/kubernetes/bin, and then downloadimages.sh should work.

How do I know the namespace being used for ITSMA?

Answer:

ITOM Plaform supports only one instance of ITSMA. It automatically assigns a namespace for your ITSMA installatio**n: itsma1** for the first installation, **itsma2** for the second (you have to uninstall the first ITSMA instance first), and so on. To see the namespace being used for your current ITSMA installation, log in to CDF, and then navigate to **SUITE** > **Management**.

What shall I do if an suite update fails?

Question details:

When upgrading the suite images from the CDF user interface (**SUITE** > **Management** > More Actions icon (

) > **Update**), after I click NEXT on the storage configuration page, a message "the https://{{ingresshost}}/update might be temporarily down or it may have moved permanently to a new web address" is displayed.

Answer:

Follow the steps below to reload this page and finish the suite update.

1. Log on to the https://{{ingresshost}}/update page in a new tab with the same browser.
2. Reload the frame in the suite upgrade page where an error message is displayed.
3. Click **APPLY** on the page that opens to finish the suite update.

What shall I do if I cannot uninstall the suite from the CDF user interface?

Question details:

I want to uninstall ITSMA because the installation failed. However, the Uninstall menu option (**SUITE** > **Management** > More Action icon

> **Uninstall**) does not work. What shall I do?

Answer:

Recreate the suite-db pods with the following commands:

kubectl get pods -n core

kubectl delete pod <pod name of suite-db> -n core

Alternatively, you can restart the virtual machine where the suite -db installed.

Can I change the FQDN of the suite once installed?

Answer:

No. Changing the external access host name of the suite is not supported in this release. You need to uninstall the suite and then reinstall it.

Load balancing

What is the load balancing mechanism in Service Management?

ITSMA uses nginx as a bundled load balancer. It uses nginx located in a K8S Ingress controller for the load balancing of Service Management: in front of the Service Management web tier, there is a frontend nginx for load balancing web access to the web tier; in front of the Service Management Server, there is a backend nginx for load balancing http access from the web tier to the Server.

What if the customer plans to use a Citrix based load balancer outside of ITSMA? Will there be any impact on the ITSMA Load Balancer?

Answer:

As ITSMA has its own load balancer already, it does not need an external Load Balancer.

Do we really need an external Load Balancer if ITSMA is providing a load balancer itself?

Answer:

As ITSMA has its own load balancer already, it does not need an external Load Balancer.

What happens if the load balancer fails within ITSMA?

Answer:

The 2017.07 release does not support high availability (HA) deployments. This release provides a self-recovery mechanism to restart the load balancer and therefore it is safer than classic ITSMA.

If an external Load balancer is also used, will the external load balancer connect with the ITSMA load balancer?

Answer:

As ITSMA has its own load balancer already, it does not need an external Load Balancer. In the event that you have already an external load balancer, it connects to the ITSMA load balancer.

Database support

How do I access the internal Postgres database for ITSMA CMDB?

Question details:

If CMDB uses the internal Postgres database, can I open the Postgres UI and browse through the contents of the database?

Answer:

The ITSMA CMDB's Postgres database has an exposed node port. You can use an external Postgres client to access the database server and browse through the contents of the database. For more information on the node port, see ITSMA node ports.

Does ITSMA support database case-sensitivity for Service Management?

Answer:

The current release does not provide case-sensitivity options.

Will all the applications that are part of ITSMA connect to their own specific Oracle database or will there be a single Oracle DB?

Answer:

ITSMA does not provide an Oracle database. Service Management and CMDB can share an external Oracle database with different schemas. Service Portal must use an external PostgreSQL database.

If Service Management and CMDB use the same external Oracle DB , will the Oracle table space requirements increase drastically?

Answer:

Since multiple applications can share the same Oracle instance with different schemas, they can share one table space or have separate table spaces.

What are the requirements for using Oracle databases for ITSMA?

Answer:

The configuration is the same as the classic configuration.

Will Service Portal (part of the ITSMA containers) use an external Oracle DB? Or will it use the internal PostgreSQL DB?

Answer:

Service Portal will use an external PostgreSQL database for production. The internal PostgreSQL is only for demonstration purposes.

Language support

What is the process for installing a language pack in the containerized Service Management in ITSMA?

Question details

ITSMA supports certain languages (for example, Brazilian Portuguese) in Service Portal, but not in Service Management (SM). This leads to strange behaviors in SM when the language is set to one of these languages. Take Brazilian Portuguese for example. Since SM 9.50 in non-containerized mode supports pt-br, it should be possible to install the pt-br language pack in the containerized SM?

Answer

You can use the following workaround to install a language pack in the containerized SM. The pt-br language pack is used here as an example.

1. Generate the 9.50 language unload file by installing the pt-br 9.50 language pack on a non-containerized Service Manager 9.52 server. This will generate a pt-Br.unl file under the Server installation folder. Its file size is about 22 MB.
2. Install and start the Service Manager 9.52 Windows client, and then connect it to the containerized Service Management by using the following connection settings:
   - Host: ***<EXTERNAL_ACCESS_HOST>***
   - Port: **31181**
   - Select the **Connect to External Load Balancer** checkbox on the **Advanced** tab
   - User account: **sysadmin**
3. Load the unload file into the containerized Service Management, from either the server side (recommended) or the client side.

   > If an unload file is large, we recommend you to load it from the server side. Loading a large unload file from the client side sometimes may cause the loading process to hang.

   Option 1: Load from the server side
   a. In the ITSMA containerized environment, copy the pt-Br.unl file to the following directory on the ITSMA NFS server: **<ITSMA global NFS share directory>/license/sm**. For example, **/var/vols/itom/itsma/itsma-itsma-global/license/sm**.
   b. In the Service Manager 9.52 Windows client, disable the **Client side load/unload** option.
   c. Restart the Windows client.
   d. Load the unload file by specifying the following unload file path: **/var/license/pt-Br.unl**.

   Option 2: Load from the client side
   a. Copy the pt-Br.unl file to a temporary directory on the Service Manager 9.52 Windows client host.
   b. In the Service Manager 9.52 Windows client, make sure that the **Client side load/unload** option is enabled.
   c. Load this unload file by browsing to the unload file on your local drive.

Master data onboarding

How do I push CIs from a standalone UCMDB to ITSMA CMDB?

Question details:

I have ITSMA installed and have the standalone UCMDB. I want to push the CIs from the standalone UCMDB to ITSMA CMDB. What are the URL and host details that I should give the standalone UCMDB?

Answer:

The current release does not support pushing CIs from a standalone UCMDB to ITSMA CMDB.

LDAP/SSO

In user information from multiple OUs brought into the containerized SM using CIT scenarios?

Answer:

In the ITSMA suite, we still use HPE Connect-IT (CIT) to synchronize data from multiple OUs to the Service Management database. This means there must be a CIT instance running outside of the container.

Can the existing SiteMinder agent work in the containerized ITSMA?

Answer:

Yes. This release of ITSMA supports LW-SSO and SAML 2.0 SSO. Therefore, if SiteMinder is configured to use SAML 2.0, it should be able to perform SSO with ITSMA.

Identity Manager (IDM)

What is the role of IDM in ITSMA?

Answer:

IDM is used for user authentication. IDM connects to the customer's LDAP server.

DNS

What is the role of DNS in ITSMA?

Answer:

Kubernetes DNS in ITSMA is only used for internal services.

Can we configure or use external DNS?

Answer:

No.

Integrations

Are Service Management integrations implemented by connecting to a specific server/SM Servlet?

Question details:

Due to load balancer issues, all integrations with Service Manager are implemented by connecting to a specific server/SM Servlet. How is this done in ITSMA Containerized version? Can we direct the connections to a specific node/SM Servlet in ITSMA or all the integrations have to be done through the ITSMA Load balancer?

Answer:

ITSMA NG Express also provides an nginx service for integration, which can handle all stateful and stateless requests. In ITSMA, all integrations

will be done through the ITSMA load balancer.

What happens if the ITSMA load balancer fails or crashes?

Answer:

ITSMA provides a self-recovery mechanism for the load balancer. It is safer than classic ITSMA. In the event that the load balancer is down, CIT/Webservice based integrations do not work anymore.

Can we configure external mail servers for notifications in ITSMA?

Answer:

Yes, ITSMA can work with an external mail server, which can be configured from the Suite Configuration user interface.

Reverse proxy

What is the role of Reverse Proxy?

Answer:

It is nginx based, the single gateway for incoming traffic. It works for routing, and authentication gateway.

Miscellaneous

Is there a way to reset the ITSMA admin password if I forgot it?

Answer:
The current release does not support resetting this password from the server. Be sure to save the password of the sysadmin user in a safe place.

What shall I do if a 502 error occurs when users attempt to log in to ITSMA?

Answer:

Restart the frontend ingress by running the following command: **kubectl delete pod**.

Is it possible to run Service Management background schedules on a separate pod?

Question details:

When we create any custom Service Management background schedulers, is it possible to run those schedules on a separate pod? How do I create a new pod to run custom schedulers with replication controllers features? One of the use cases is that, in a production environment, when there are huge backlogs for SLAs, we want to create custom schedulers and run them in a separate pod to balance the backlogs.

Answer:

No. Currently, ITSMA does not support creating a custom pod for schedulers.

Can I add more parameters to track Service Management performance?

Question details:

How do I add parameters such as **alertlimit**, and **alertlog** that are available in HPE Service Manager to capture long-running queries?

Answer:

Currently, ITSMA does not allow you to add such parameters. You can only specify debugging parameters that are exposed on the Suite Configuration user interface. For details, see Configure log level for debugging.

Troubleshoot HPE ITOM Container Deployment Foundation

This section provides information that can assist you in troubleshooting ITOM Container Deployment Foundation related issues.

- Manual verification commands
- ITOM Container Deployment Foundation logs
- Support toolset
- ITOM Container Deployment Foundation issues

Manual verification commands

You can use the following commands when troubleshooting ITOM Container Deployment Foundation. For example, you can use them to list the namespaces and services.

Command	Description
kubectl	The command to interact with K8S. To shorten this command to `kl` you can create a soft-link to the kubectl executable by running the following command: `ln -s /usr/bin/kubectl /usr/bin/kl`
kubectl cluster-info	Summarizes information about some of the services that are running on the cluster, including **Kubernetes master**, **KubeDNS** for service discovery, and the endpoints of the **KubeRegistry** (if you are running a registry).
kubectl get	To get various entities within K8S. You have to specify the resource type. **Example:** `kubectl get nodes`: to list all the nodes in the cluster. `kubectl describe nodes [node_IP]`: to get more specific information on the node such as labels, events, capacity, CPU, memory, the maximum of pods it can support, system information on the node, external IP address, the pods that are running, the list of namespaces, and resources, and more.
kubectl get pods	To list all pods in the default namespace (Used to separate the base ITOM Container Deployment Foundation services from the deployed suites).
kubectl get pods - - namespace=kube-system	To list all the pods that are running in the namespace where base IT OM Container Deployment Foundation and K8S pods are running.
kubectl get pods - - all-namespaces	To list all the pods that are currently running in the cluster.
kubectl describe pod <podName> --namespace=<namespace>	To view details about the specified pod in the specified namespace. A pod is a unit of scheduling inside K8S. A container is always run inside a pod. The details are the containers it is running, the image it is running, the port it is exposing, and the command (`/hyperkube`) that is running inside the container itself with their options, volumes, and more.
kubectl exec -t suite-installer -- namespace=kube-system -- sh	To view what is running in the suite installer pod on the specific namespace. This displays what is running inside the container. Use exit to get out of it.

kubectl get services -- all-namespaces	To view all the services running in the cluster. If a pod needs to expose what it is doing outside the pod, you need to create a service for it. Some services are internal and some are external. When you install a cluster there are a number of services that are automatically installed. **autopass-ln-svcis** an external service, **idmpostgresql-svc** used by the postgresql database that serves idm which is internal, **idm-svc** that is internal, **kube-dns** and **kube-registry** are also internal, **mng-portal** that is available externally, **postgresql-apln-svc** that serves the autopass also available internally, and the **suite-installer-svc** that is also available internally.
./kube-restart.sh	To restart the K8S cluster. In **/opt/kubernetes/bin**, run: `./kube-restart.sh` to restart the K8S cluster.
./kube-start.sh	To start the K8S cluster. In **/opt/kubernetes/bin**, run: `./kube-start.sh` to start the K8S cluster.
./kube-status.sh	To display the status of the K8S cluster. In **/opt/kubernetes/bin**, run: `./kube-status.sh` to view the status of the K8S cluster.
./kube-stop.sh	To stop the K8S cluster. In **/opt/kubernetes/bin**, run: `./kube-stop.sh` to stop the K8S cluster.

To `kubectl port-forward`, you must have the socat package installed prior.

Support toolset

ITOM Container Deployment Foundation (CDF) provides a support toolset that helps to collect information about Docker, Kubernetes, suites, commands, directories, and files as described below:

- Docker: containers, inspect, docker service systemd logs
- Kubernetes: nodes, pods, namespaces, images, containers, cluster-info, describe, and logs
- Suite: suite-db dump, suite data, modules, products deployments, and features
- Commands defined by users
- Directories and files defined by users

You can view the summary information on a console, and view detailed output information in an encrypted tar file.

How to use the support toolset

To use the support toolset, follow these steps:

1. Log in to the master node.
2. Run the following command:

cd $K8S_HOME/tools/support-tool

3. Run the following command:
   # ./support-dump [-c <dump-filename-with-path>] [-u <username> [-p <password>]] [-P <package_password>]
   Example usages:
   Create a dumpfile with a default file name in a default directory.
   # ./support-dump
   Create a dumpfile with a specified file name in a specified directory. For example, create a dump.des3 in /var/test:
   # ./support-dump -c /var/test/dump.des3
   Create a dumpfile with a specified user name and password. For example, create a dumpfile with a default file name in a default directory with password 'abcdef'. Connect the suite-installer with 'admin' as the user and '123456' as the password:
   # ./support-dump -u admin -p 123456 -P abcdef
4. Run the following command to unpack the dumpfile:
   # dd if=xxxx.des3 |openssl des3 -d -k <package_password>|tar zxf -

How to use configuration files

The support toolset provides a configuration file with some predefined `[commands], [files],` and `[dirs]` to specify information collection details. This default configuration file is: conf/supportdump.config.

You can define your own `[commands], [files],` and `[dirs]` in this configuration file. Additionally, you can create other configuration files in the same directory.

When using the configuration files, pay attention to the following:

- The outputs of the same command will be saved into one file. For example, all the outputs of the `cat` command will be saved to the `cat .out` file.
- All directories, files and outputs of commands are stored in the <local_ip>-<NodeType>/os directory.
- The wildcard characters can be used in a file name and directory name. For example, /etc/sysconfig/network-scripts/ifcfg-*
- Single environment variable is supported. For example, ${K8S_HOME}/log.
- A file or files (separated by spaces) followed a directory will be excluded from the support toolset collection.

Example usage:

${K8S_HOME}/cfg *_User.json

The support toolset collects all files and directories in the ${K8S_HOME}/cfg except the * _User.json file.

Below is the default configuration file (supportdump.config):

```
#(C) Copyright 2013-2017 Hewlett Packard Enterprise Development LP
[commands]
cat /proc/net/if_inet6
cat /proc/partitions
cat /sys/class/net/bond*/bonding/active_slave
crm_mon -rf1
df -a
df -h
dmidecode
drbd-overview
du -h --max-depth=5 / | sort -hr
ethtool eth0
fdisk -l
free -m
grep "model name" /proc/cpuinfo
head -n1 /etc/issue
hostname
ifconfig -a
iostat
ip -4 route show table all
ip -6 route show table all
ip addr show
ip route show
```

```
ipcs -l
ipcs -a
iptables -L
iptables -S
ip6tables -L
ip6tables -S
last reboot
lldpcli show neigh
ls -l /etc/pki/tls/certs/
ls -l /etc/pki/tls/private/
ls -l /etc/rc.d/rc3.d
lspci -vmm
lspci -vvvv
lvdisplay
mount | column -t
mpstat -P ALL
netstat -anp
netstat -g
ntpq -p
ps -e -ww -o pcpu,cpu,nice,state,cputime,args --sort pcpu | sed '/^ 0.0 /d'
ps -ef -ww
ps -ww -eo pcpu,pid,user,args | sort -k 1 -r | head -10
route -n
rpm -qa
rpm -qa hp-firmware\* | sort
top -n 1 -b -c
uname -a
vgdisplay
virsh list --all
virt-what
vmstat
vmware-toolbox-cmd -v
brctl show
sysctl -a
[files]
/etc/group
/etc/hosts
/etc/httpd/conf.d/ssl.conf
/etc/ntp.conf
/etc/partner_ntp.conf
/etc/passwd
/etc/system-release
/etc/resolv.conf
/etc/resolv.conf_alt
/etc/sysconfig/i18n
/etc/sysconfig/ip6tables
/etc/sysconfig/iptables
/etc/sysconfig/network
/etc/sysctl.conf
/mnt/usb/cic-manager-setup-config.json
/root/.virtinst/virt-install.log
/updatelogs/db_install_upgrade.log
/updatelogs/update.log
```

```
/usr/lib/rpm/rpm.log
/var/lib/rabbitmq/erl_crash.dump
/var/run/utmp
/var/tmp/usbAutoSetupTrace.log
/ci/logs/dynamic-configuration.log
/ci/etc/service-console
/ci/etc/pwreset/disable-root
/ci/etc/devlogin-enabled
/etc/pam.d/system-auth
${K8S_HOME}/install*.log
${K8S_HOME}/version.txt
[dirs]
/etc/sysconfig/network-scripts/ifcfg-*
/updatelogs/db_install_upgrade.log.*.gz
/updatelogs/failurelogs/failure*
/updatelogs/update.log.*
/var/cpq
/var/lib/pgsql/pg_logs
/var/log
/etc/sysctl.d/*
${K8S_HOME}/cfg *_User.json
${K8S_HOME}/log
${K8S_HOME}/manifests
${K8S_HOME}/objectdefs
```

```
${K8S_HOME}/runconf
${K8S_HOME}/scripts/*.*
${K8S_HOME}/ssl
```

Dump file

The default support dump file is: dmp/support_data_YYYYMMDD-hhmmss.des3. The dump file contains the support_data_YYYYMMDD-hhmmss .log of the running support toolset and the ITOM_Core_Platform directory for the dump files. The table below describes the dump files in the ITOM_Core_Platform directory.

Name	Description	Type
<local_ip>-<NodeType>	The directory of docker information and user defined information on the current node. • **docker** docker information dump files **bootstrap**: • **bootstrap_containers.out**: docker-bootstrap containers • **bootstrap_docker_inspect.out**: docker-bootstrap inspect • **journalctl_docker-bootstrap.out** : docker-bootstrap service log workload **workload**: • **containers.out**: docker containers • **docker_inspect.out**: docker inspect information • **journalctl_docker.out**: docker service log • **os**: user defined commands, directories and files • **commands**: directory of output files of commands defined in the `[commands]` section in .config files. The file name format: <command>.out. • **other directories**: directories and files defined in the `[files],` and `[dirs]`sections in .config files. The structure of directories will be reserved.	Directory

global	**deployment**: • **suite_data**: JSON files of all suites data • **suite-db_pgdump.tar**: PostgreSQL db dump file of suite-db pod • **suite_features.out**: suite deployment and features **kubernetes**: **cluster_info:**: • **cluster_info.out**: output of running command 'kubectl cluster_info dump' • **other dirs & files**: generated by command 'kubectl cluster_info dump' **kube_describe.out**: describe of all pods **kube_summary.out**: namespaces, pv, pvc, nodes, deployments, services, pods, ingress **platform**: • **containers_by_pod.out**: all containers, images on all pods	Directory

Example: Run the support toolset

Below is a sample of running the support toolset on a console.

```
[root@SGDLITVM0719 support-tool]# ./support-dump
Management Portal user:admin
Password:
Package password:
Retype package password:
################################################
ITOM Core Platform - Support Data Export
Date: 2017-03-08 13:30:11
Current node: 16.187.190.219
Node type: Worker
Docker: v1.12.6
Kubernetes: server-v1.4.3 client-v1.4.3
################################################
------------------------------------------------
Docker Containers
Export: containers.out
Comments: on Worker node 16.187.190.219
------------------------------------------------
CONTAINER ID IMAGE STATUS
9cbed51599a5 localhost:5000/nginx-ingress:0.8.3 Up 2 days
5efe4a972bb5 gcr.io/google_containers/pause-amd64:3.0 Up 2 days
...
------------------------------------------------
Docker bootstrap Containers
Export: bootstrap_containers.out
```

```
Comments: on Worker node 16.187.190.219
-----------------------------------------------
CONTAINER ID IMAGE STATUS
832bbd87fee1 gcr.io/google_containers/flannel-amd64:0.5.5 Up 2 days
-----------------------------------------------
Nodes
Export: kube_summary.out
-----------------------------------------------
NAME STATUS AGE
16.187.190.138 Ready 2d
16.187.190.219 Ready 2d
sgdlitvm0590.hpeswlab.net Ready 2d
sgdlitvm0598.hpeswlab.net Ready 2d
sgdlitvm0729.hpeswlab.net Ready 2d
-----------------------------------------------
Pods
Export: kube_summary.out
-----------------------------------------------
NAMESPACE NAME READY STATUS RESTARTS AGE IP NODE
core apiserver-16.187.190.138 1/1 Running 0 2d 16.187.190.138
16.187.190.138
core apiserver-sgdlitvm0590.hpeswlab.net 1/1 Running 0 2d 16.187.190.90
sgdlitvm0590.hpeswlab.net
...
-----------------------------------------------
POD Containers
Export: containers_by_pod.out
-----------------------------------------------
NAMESPACE POD NODE IMAGE CONTAINER CONTAINER_ID
core apiserver-16.187.190.138 16.187.190.138
gcr.io/google_containers/hyperkube:v1.4.3 apiserver 8e74804ee06e
...
core autopass-lm-116661985-hufdr 16.187.190.138
localhost:5000/autopass-lms:10.0 autopass-lm 5154fe13fb
...
-----------------------------------------------
Suite Deployment
Export: suite_features.out
-----------------------------------------------
SUITE VERSION NAMESPACE DEPLOYMENT_STATUS INSTALL_DATE NFS_SERVER
NFS_OUTPUT_PATH
itsma 2017.01 itsma1 DELETED 2017-03-06 14:25:40 SGDLITVM0668.hpeswlab.net
/var/vols/itom/itsma/itsma-itsma1
-----------------------------------------------
Suite Features
Export: suite_features.out
-----------------------------------------------
SUITE EDITION SELECTED FEATURE_SET FEATURE
itsma Express Edition true Smart Analytics Smart analysis
The Business Value of HPE ITSM Automation
Universal CMDB UCMDB
Service Manager <h2>Welcome to ITSMA</h2>
Faster Service. Better Experience. Lower Costs.
```

```
Docker inspect ..................... exported to docker_inspect.out &
bootstrap_docker_inspect.out
Docker cluster-info dump .......... exported to cluster_info.out
Docker describe ................... exported to kube_describe.out
Suite DB ........................... get dump file from
suite-db-4070668108-0c5pf failed!
Suite data JSON files ............. Not get suite-installer pod on current
node
Making OS commands & files list ... done
Running OS commands in list ....... done
Collecting OS files in list ....... done
Packing dump files ... package file is
/opt/kubernetes/tools/support-tool/dmp/support_data_20170322-150008.des3
```

```
Please use below command to uncompress the package file:
dd if=support_data_20170322-150008.des3 |openssl des3 -d -k
<your_password>|tar zxf -
```

ITOM Container Deployment Foundation logs

To view the ITOM Container Deployment Foundation (CDF) log files, follow these steps:

1. In the ITOM Container Deployment Foundation navigation pane, click **RESOURCES** > **Pods**.
2. Click the relevant pod.
3. Click **View logs** in the Pod area.

> All suite logs are currently stored within a persistent volume so that they are not lost even if the pods go down.

ITOM CDF log UI

UI element	Description
	Toggles the size of the font used in the log

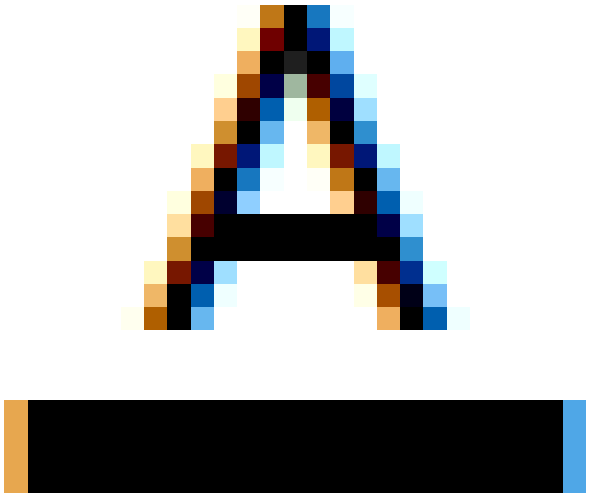	Toggles the colors of the log: white characters on black background or black characters on white background
Logs from 10/31/16 7:23 AM to 10/31/16 7:37 AM	Timestamp of the currently displayed log
	Buttons to navigate through logs

ITOM Container Deployment Foundation issues

The following is a list of ITOM Container Deployment Foundation related issues and their solutions.

Problem	Solution

kube-proxy runs into a CrashLoopBackOff issue.	Increase the CPU limits of kube-proxy.
	1. In all master and worker nodes, modify /opt/kubernetes/runconf/kube-proxy.yaml as the following:
	…
	`resources:`
	`requests:`
	`cpu: 100m`
	`memory: 200Mi`
	`limits:`
	`cpu: 1`
	`memory: 200Mi`
	…
	2. Run the following command in all master and worker nodes to relaunch kube-proxy:
	sudo systemctl restart kubelet
	
	To make sure that the VM resource is not overloaded in the virtualization environment, set the CPU reservation value to at least 4GHz in each node.

Failed to start the IdM service of ITOM Container Deployment Foundation	To solve this problem, change the proxy settings in the autopass-lm.yaml file to an empty value: **env:** **- name: http_proxy** **value: ""** **- name: https_proxy** **value: ""**

Unable to connect to the server: invalid character "{" in host name?	Check if the firewall is disabled on the NFS server or not.
No Common Name set for a Pod certificate	To solve this problem: 1. Run the following command to update the vault token: **run ./update_kubevaulttoken** 2. Run the following command to recreate all pods: **run ./kube-redeploy.sh** 3. Run the following command to restart Docker: **docker ps|grep kube-registry-proxy|grep -v pause|awk '{print $1}'|xargs docker restart**
Token has expired	This problem may occur when an internal vault token has expired. To solve this problem: 1. On the master node, run the following command: **update_kubevaulttoken" under /opt/kubernetes/bin** 2. Run the following command to restart Docker: **systemctl restart docker** Another reason is that time is not synchronized between a worker and the master, or a POD and the VM. Fixing the time synchronization issue will solve this problem.
The autopass-lm-xxx-xxx pod in the core (namespace) crashes easily	The autopass-lm-xxx-xxx pod easily crashes and cannot return to normal after many restarts. This problem is caused by vault token expiration. To solve this problem: 1. Run the following command to update vault token: **./update_kubevaulttoken** 2. Run the following command to recreate all pods: **./kube-redeploy.sh**
When the vault token expires during the ITOM CDF installation process, the ITOM CDF installation process may fail.	To solve this problem: In the /opt/kubernetes/bin folder of the master node, run the **update_kubeVaultToken** command.
After you install ITSMA on ITOM CDF and log on to the ITSMA suite management portal, the **SUITE > Installation** page of the user interface does not work.	To solve this problem, manually clean the cookies and cache of the page.
When uninstalling ITOM CDF on a virtual machine, the virtual machine sometimes hangs.	To solve this problem, manually restart the virtual machine.
When using %, " , \ and space for the User Name in **ADMINISTRATION > User Management > ADD**, you cannot save the user name, and no error message is displayed.	To solve this problem, do not use the following characters for User Name: % " \ <blank>
An error may occur or a pod may enter the CrashLoopBackOff status when you migrate the pod from one node to another. This issue may occur because no vault token is generated.	If the pod is created by the deployment's replication controller (in this situation, there is a random string at the end of the pod name, such as idm-848511036-5ev30 or idm-t56ui), delete the pod and a new pod will be generated. Alternatively, run the kube-restart.sh file in the $K8S_HOME/bin folder of the node where the pod is located.

The following issues occur when you view ITOM CDF in Internet Explorer 11: • If you refresh the browser after you click **ADMINISTRATION > User Management > ADD** to add a new user, and then click **SAVE**, the newly added user is not in the user list. • If you refresh the browser after you click **ADMINISTRATION > User Management > > DELETE**, and then click **DELETE**, the newly deleted user is still in the user list. • If you refresh the browser after you click **ADMINISTRATION > Nodes > Predefined Labels**, enter a new label, and then click **[+]**, the newly added label is not in the label list. • If you refresh the browser after you click **ADMINISTRATION > Nodes > Predefined Labels,** and then click **[-]** to delete a label, the newly deleted label is still in the label list.	To solve this problem, open developer tools (press F12), and then refresh the page after you have added a record. Alternatively, use another browser.
When you access **SUITE > Management** > [more actions] icon > **Update** to upgrade the suite images and then click **NEXT** in the storage configuration page, the system displays an error message and shows that the content fails to load.	To solve this problem, log on to the https://{{ingresshost}}/update page in a new tab in the same browser. Then reload frame in the suite upgrade page where the error message appears.
When a suite failed to install on the ITOM Container Deployment Foundation, the **SUITE > Management** > [more actions] icon > **Uninstall** does not work.	To solve this problem, run the following commands to recreate the suite-db pods: **kubectl get pods -n core** **kubectl delete pod <pod name of suite-db> -n core** Alternatively, restart the virtual machine where the suite-db pod is installed.

When installing the ITOM Container Deployment Foundation, a "full crash report" error is found.	To solve this problem, update the kernal package, the system BIOS and network interface firmware of the Red Hat Enterprise to the latest version. Click here to see the detailed resolution on the Red Hat Enterprise solution page: https://access.redhat.com/solutions/43168. Resolution The NETDEV WATCHDOG message is the kernel's way of saying "This network device has not been transmitting data for a few seconds, even though I asked it to transmit data". The watchdog message does not tell us why the device stopped transmitting. It may be due to a hardware error or a software (kernel/driver/BIOS/firmware) bug. Red Hat Engineering have Private Bugs open for each individual driver where this issue is being seen. As the NETDEV WATCHDOG hang is a symptom of an issue, not an actual issue itself, the root cause of the NETDEV WATCHDOG hangs must be investigated on an individual basis. Please open a case with Red Hat Global Support Services, supplying a full sosreport, and as much of the following information as possible: Full dmesg (not just an excerpt with the NETDEV WATCHDOG message and the stack trace, for the reasons explained above). Information about the affected hardware (sosreport should be fine). Did the network interface recover automatically shortly afterwards? Or can connectivity be restored by doing ifdown followed by ifup? Or can connectivity be restored by rmmod followed by modprobe of the driver? Or is reboot the only way to make the device work again? How often does the issue occur? Does the occurrence of the issue seem to correlate with specific workloads? Is there a way to reproduce it, or at least to make it more likely to happen? Do any of these kernel boot parameters help?: Raw pcie_aspm=off (ASPM has been known to cause troubles in the past.) intremap=off (Interrupt remapping has been known to cause lost interrupts in conjuction with irqbalance on some platforms, e.g. bug 887006) Some additional troubleshooting steps which may help prevent the issue from re-occurring: Update the kernel package to the latest version, which will supply the latest available driver with fixes for all known issues. Update the system BIOS and network interface firmware. Product(s) Red Hat Enterprise Linux Component kernel Category Troubleshoot Tags hardware module network networking rhel_5 rhel_6 This solution is part of Red Hat’s fast-track publication program, providing a huge library of solutions that Red Hat engineers have created while supporting our customers. To give you the knowledge you need the instant it becomes available, these articles may be presented in a raw and unedited form.

When reinstalling the suite on the CDF management portal, the first page has been loading for a long time and still does not show up.	To solve this problem, clear the cached images and files, and clear cookies from **Clear the browsing data** of your browser before reinstalling the suite on the CDF management portal.
In Firefox 52, after clicking **Finish** to complete the suite installation on the CDF management portal, the system opens a page which may be refreshed and directed to a blank page.	To solve this problem, click **Management** to refresh the page. Alternatively, use a different browser.

Troubleshoot the ITSMA suite

This section provides information that can assist you in troubleshooting the ITSMA suite related issues.

- ITSMA suite logs
- ITSMA suite issues
- ITSMA service dependency

ITSMA suite logs

You can find ITSMA suite log information in the log files described in the following table.

> All suite logs are currently stored within a persistent volume so that they are not lost even if the pods go down. For more information about the persistent volumes of ITSMA, see Set up three NFS shares for ITSMA.

Component	Log and data directories	Notes
Service Management (SM)	• General: <*global NFS root folder*>/logs/sm-9.52 • Database: <*database NFS root folder*>/db/sm	The following is a list of the Service Management log sub-folders in <global NFS root folder>/logs/sm-9.52: • Scheduler: SM background scheduler responsible for running SM schedulers • Gossip: Docker container responsible for communications among RTE nodes • Irque: Docker container to write accumulated updates into the IR index files • Rte: SM RTE service • Integration: the integration service provided by the SM RTE docker container • Webtier: SM web tier service • Mobility: SM Mobility service • Emailout : SM backend email notification service • Integration-cit: dedicated service for Connect-IT (CIT) integration • Report-export: SM backend service for exporting reports • Sm-config: SM configuration pod • Sm-dataloading: SM OOB data loading pod
Service Portal	• <*global NFS root folder*>/logs/xservices	
UCMDB	• General: <*global NFS root folder*>/logs/ucmdb-<*10.3x*> • Database: <*database NFS root folder*>/db/ucmdb	
ITSMA suite landing page	• <*global NFS root folder*>/landing-page/logs	
Suite Configuration	• Yaml files: <*global NFS root folder*>/data/itsmaconfig/yamls • Logs: <*global NFS root folder*>/logs/itsmaconfig-1.0	

HPE Identity Manager (IdM)	• Logs: *<global NFS root folder>*/logs/idm • Data: *<global NFS root folder>*/data/idm/seeded_data • Database: *<database NFS root folder>*/db/idm	
Auth application	• *<global NFS root folder>*/logs/auth	
itsma-installer	• yaml files: *<CDF NFS share>*/suite-install/itsma/output	
Smart Analytics	• Logs: *<Smart Analytics NFS root folder>*/logs/idol • Config: *<Smart Analytics NFS root folder>*/config/idol • Data: *<Smart Analytics NFS root folder>*/data/idol	
Openldap	• Data: *<global NFS root folder>*/data/openldap • Logs: *<global NFS root folder>*/logs/openldap-2.4.41	
Ingress frontend and backend	• *<global NFS root folder>*/logs/ingress	
Chat	• *<global NFS root folder>*/logs/chatsvc • *<global NFS root folder>*/logs/openfire	

Check the ITSMA installation log

Follow these steps to check the ITSMA installation log:

1. Go the the suite management portal, and then log in as a system administrator.
2. Switch to "core" namespace in the side bar of management portal, then click **Services**.

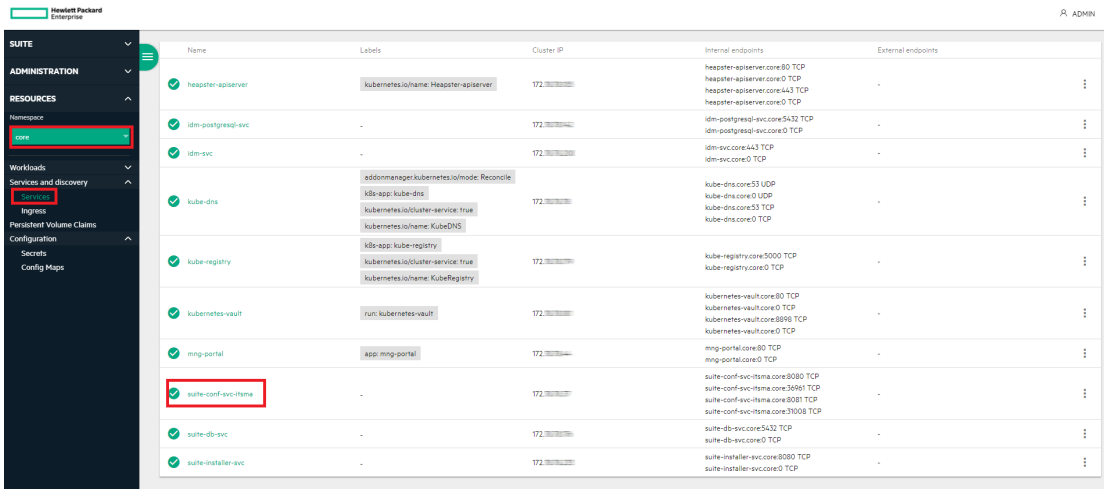

3. Click the "suite-conf-svc-itsma" service and get into the detail page.

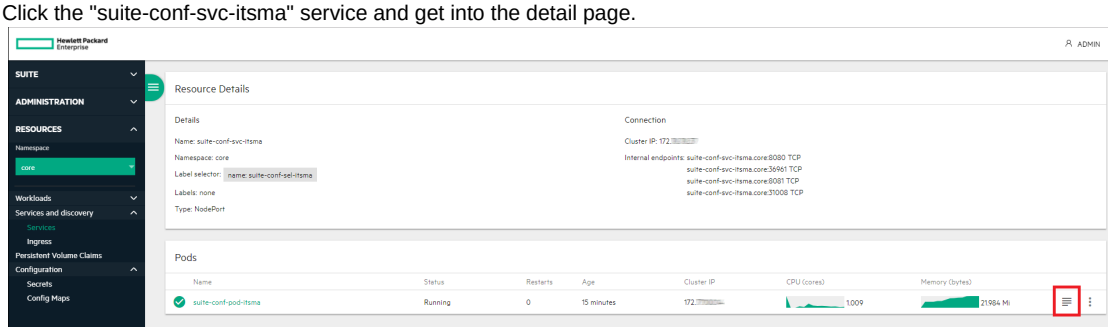

4. Click the **Log view** button to show the service log.

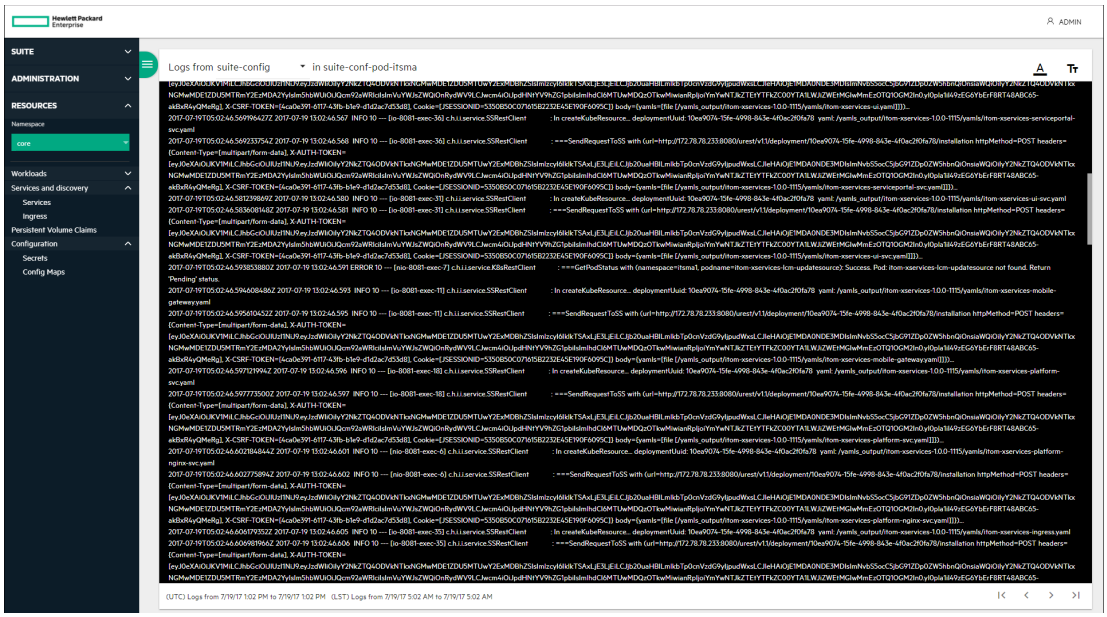

ITSMA service dependency

If one or multiple ITSMA services are not working correctly, you can use the following service dependency information to troubleshoot the possible root cause for service failures.

- ITSMA common services
- Service Management services
- Smart Analytics
- Service Portal
- Service Portal Admin
- CMDB services

ITSMA common services

Service	Dependency
auth	No hard dependency in general.
idm	It is recommended that you check the following services in this order:
idm-postgresql	idm-postgresql/openldap-rc > auth > idm
itom-nginx-ingress-deployment	Other services are independent of all the services.
itsma-config-deployment	
itsma-landing-page	
openldap-rc	

Service Management services

The start sequence of the Service Management services is as follows:

1. sm-postgres
2. sm-dataloading
3. sm-rte-gossip
4. sm-config
5. sm-rte
6. sm-rte-integration, sm-rte-irque, sm-rte-scheduler, sm-webtier, sm-mobility
7. sm-chat

Service	Dependency
sm-postgres	None
sm-dataloading	sm-postgres
sm-rte-gossip	sm-dataloading
sm-config	sm-rte-gossip
sm-rte	
sm-rte-integration	sm-rte
sm-rte-irque	
sm-rte-scheduler	
sm-webtier	
sm-mobility	
sm-chat	sm-rte-integration

Smart Analytics

All the Smart Analytics pods are started by the smarta-installer service.

Service	Dependency	Notes
smarta-admin-ui-backend	None	
smarta-idolfarm-config	None	
smarta-installer	None	
smarta-mix-l2proxy	smarta-ss-con-1a	Mixed mode only
	smarta-ss-con-1b	
	smarta-ss-con-2a	
	smarta-ss-con-2b	
smarta-mix-proxy	smarta-mix-l2proxy	
smarta-nginx	None	
smarta-search	None	
smarta-smsp-con-1a	None	
smarta-smsp-con-1b	None	
smarta-smsp-dah	smarta-smsp-con-1a	
	smarta-smsp-con-1b	
smarta-smsp-dih	smarta-smsp-con-1a	
	smarta-smsp-con-1b	
smarta-smsp-qms	smarta-ss-agent	
smarta-ss-agent	None	
smarta-ss-cfs	None	
smarta-ss-con-1a	None	
smarta-ss-con-1b	None	

smarta-ss-con-2a	None	
smarta-ss-con-2b	None	
smarta-ss-dah1	smarta-ss-con-1a	
	smarta-ss-con-1b	
	smarta-ss-con-2a	
	smarta-ss-con-2b	
smarta-ss-dih	smarta-ss-con-1a	
	smarta-ss-con-1b	
	smarta-ss-con-2a	
	smarta-ss-con-2b	
smarta-ss-imgsvr	None	
smarta-st-con-1	None	
smarta-st-con-2	None	
smarta-st-proxy	smarta-st-con-1	
	smarta-st-con-2	

Service Portal

The itom-xservices-infra services start first, and then itom-xservices start.

Service		**Dependency**
itom-xservices-infra	itom-xservices-postgres-svc	None
	itom-xservices-rabbitmq-svc	None
	itom-xservices-redis-svc	None
	itom-xservices-xmpp-svc	None
itom-xservices	itom-xservices-chat-ui-svc	sm-chatsvc-svc
	itom-xservices-ui-svc	itom-xservices-gateway-svc
	itom-xservices-gateway-svc	idm-svc itom-xservices-serviceportal-svc itom-xservices-mobile-gateway-svc itom-xservices-platform-svc
	itom-xservices-mobile-gateway-svc	itom-xservices-serviceportal-svc

	itom-xservices-serviceportal-svc	itom-xservices-platform-svc itom-xservices-redis-svc itom-xservices-rabbitmq-svc propel-search-svc propel-sx-svc propel-survey-svc propel-portal-svc propel-catalog-svc sm-rte-integration-svc smarta-search-svc
	itom-xservices-platform-svc	origin-itom-xservices-platform-svc
	origin-itom-xservices-platform-svc	idm-svc itom-xservices-xmpp-svc itom-xservices-postgres-svc itom-xservices-redis-svc itom-xservices-rabbitmq-svc

Service Portal Admin

The propel-postgresql service and propel-rabbitmq service start first, and then others can start at the same time.

Service		**Dependency**
itom-service-portal	propel-analytics-svc	propel-catalog-svc propel-postgresql-svc propel-rabbitmq-svc
	propel-autopass-svc	none
	propel-autopass-import-license-svc	none
	propel-bpm-svc	idm-svc propel-rabbitmq-svc propel-catalog-svc propel-sx-svc propel-postgresql-svc propel-ganglia-gmond-svc
	propel-bpm-ui-svc	propel-bpm-svc propel-launchpad-svc propel-autopass-svc propel-catalog-svc propel-sx-svc

propel-catalog-svc	propel-sx-svc propel-bpm-svc propel-postgresql-svc idm-svc propel-ganglia-gmond-svc propel-search-svc propel-rabbitmq-svc
propel-catalog-ui-svc	propel-launchpad-svc propel-autopass-svc propel-postgresql-svc propel-catalog-svc propel-bpm-svc propel-rabbitmq-svc propel-notification-svc
propel-diagnostics-svc	propel-search-svc propel-catalog-svc propel-sx-svc idm-svc propel-autopass-svc propel-launchpad-svc propel-portal-svc propel-sx-client-svc propel-subscription-ui-svc propel-catalog-ui-svc propel-sx-ui-svc propel-idm-admin-svc propel-autopass-ui-svc propel-bpm-svc propel-bpm-ui-svc
propel-diagnostics-ui-svc	propel-sx-svc propel-diagnostics-svc propel-autopass-svc propel-launchpad-svc propel-ganglia-ui-svc
propel-external-entities-ui-svc	propel-postgresql-svc propel-launchpad-svc propel-autopass-svc propel-catalog-svc propel-notification-svc propel-rabbitmq-svc

propel-ganglia-gmetad-svc	none
propel-ganglia-gmond-svc	none
propel-ganglia-ui-svc	none
propel-idm-admin-svc	propel-launchpad-svc propel-autopass-svc
propel-launchpad-svc	propel-autopass-svc propel-postgresql-svc propel-search-svc propel-catalog-svc propel-rabbitmq-svc propel-notification-svc propel-sx-svc propel-dashboard-svc propel-gateway-svc
propel-notification-svc	propel-rabbitmq-svc propel-postgresql-svc idm-svc
propel-portal-svc	propel-launchpad-svc propel-catalog-svc propel-autopass-svc propel-postgresql-svc propel-search-svc propel-rabbitmq-svc propel-notification-svc propel-sx-svc
propel-postgresql-svc	none
propel-propeltool-svc	idm-svc propel-sx-svc sm-rte-svc
propel-rabbitmq-svc	none
propel-search-svc	propel-idol-svc propel-catalog-svc
propel-subscription-ui-svc	propel-launchpad-svc propel-autopass-svc propel-postgresql-svc propel-catalog-svc propel-sx-svc propel-rabbitmq-svc propel-notification-svc

	propel-survey-svc	idm-svc sm-rte-svc
	propel-survey-ui-svc	propel-launchpad-svc propel-survey-svc propel-autopass-svc propel-rabbitmq-svc propel-notification-svc
	propel-sx-svc	propel-ganglia-gmond-svc propel-rabbitmq-svc propel-portal-svc propel-catalog-svc propel-bpm-svc
	propel-sx-ui-svc	propel-launchpad-svc propel-autopass-svc propel-diagnostics-svc propel-postgresql-svc propel-catalog-svc propel-sx-svc propel-ganglia-ui-svc propel-rabbitmq-svc propel-notification-svc

CMDB services

The postgresql-ucmdb service starts first, and then others can start at the same time.

Service	Dependency
postgresql-ucmdb	None
ucmdb	postgrsql-ucmdb
ucmdb-probe	ucmdb
Ucmdb-browser	
ucmdb-solr	None

ITSMA suite issues

The following tips help you to troubleshoot miscellaneous issues with the ITSMA suite.

- No Oracle JDBC driver found
- JVM crashes after updating to kernel-3.10.0-514.21.2.el7.x86_64
- The sysadmin user cannot log in with a new password after reinstallation
- Pods remain in a terminated state
- Unable to connect to the server: invalid character "{" in host name
- Suite is in the Pending state after a worker node reboot
- Some pods are in the ErrImagePull state during installation

- A 503 nginx error when you run the Suite Installer
- Suite pods have errors and cannot recover
- Failed to pull images from the localhost because of a "Status 503" error
- Log error: No License File, use InstantOn
- ITSMA is down
- Cannot reinstall ITSMA
- Service Management: an HTTP 500 error occurs
- Service Management: Smart Email does not work with the EWS protocol
- Performance: pods are not evenly distributed to workers
- Performance: when the concurrent users are more than 2000, the suite landing page takes more than 25 seconds to open
- Performance: the Linux buffer does not release
- Performance: an http 502 error might occur because of port exhaustion
- Cannot launch CMDB in Internet Explorer

No Oracle JDBC driver found

During the suite installation, if you see the "No Oracle JDBC Driver Found" error message when you test your external Oracle database, it means that you have not installed the required Oracle JDBC Driver.

In such case, follow these steps to install Oracle JDBC Driver (ojdbc6.jar) on the CDF NFS server host:

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
8. Click the **Cancel** button in the suite installation wizard.
9. Re-install the ITSMA suite.

JVM crashes after updating to kernel-3.10.0-514.21.2.el7.x86_64

This is causes by a Red Hat security fix in kernel 3.10.0-514.21.2. To solve this issue, upgrade to kernel-**3.10.0-514**.**26.2**.el7.x86_64 or above.

The sysadmin user cannot log in with a new password after reinstallation

This problem may occur when you uninstall and then reinstall ITSMA and specify a new password for the sysadmin user. For example, the password for the old installation is admin1234 and the new one is 1Qaz2wsx. You cannot log in as sysadmin with the new password.

The root cause is that after uninstalling ITSMA you did not remove the ITSMA NFS share directories, in which old installation data including the sysadmin user password is stored.

To resolve this issue, remove the old ITSMA NFS share directories and then reinstall ITSMA.

Pods remain in a terminated state

Run the following command to remove these pods:
kubectl delete pod [POD NAME] --grace-period=0

Unable to connect to the server: invalid character "{" in host name

To solve this problem, make sure that the firewall is disabled on the NFS server.

Suite is in the Pending state after a worker node reboot

It is a known issue that rebooting a worker node will cause the status of the ITSMA suite to change to Pending.

To avoid this issue, do not reboot a worker node.

Some pods are in the ErrImagePull state during installation

Use one of the following workarounds:

- Wait for a moment. The pod will retry an image pull.
- On the worker node where the pod is running, run the following command:

  ```
  docker pull XXX
  ```

A 503 nginx error when you run the Suite Installer

The following error occurs when you run the Suite Installer:

```
Warning FailedSync Error syncing pod, skipping: [failed to "InitContainer" for "install" with
RunInitContainerError: "init container \"install\" exited with 1", failed to "StartContainer" for "install"
with CrashLoopBackOff?
```

This issue occurs because time is not synchronized between the master node and worker node.

Solve this problem with one of the following methods.

Method 1:

Synchronize the time between the master and worker nodes, and then start over.

Method 2:

Make suit-config-pod-itsma run on the master node by adding the nodeSelector attribute to the suit-config.yaml file. To do this, follow these steps:

1. Open the yaml file:

   ```
   vi /var/vols/itom/core/suite-install/itsma/suite_config_template.yaml
   ```

2. Add the 'nodeSelector' attribute to the file, and configure it as follows:

   spec:
   ** nodeSelector:**
   **   master: "true"**
   containers:
   - name: suite-config
3. Run the following commands:

   cd /var/vols/itom/core/suite-install/itsma
   kubectl delete -f suite_config.yaml
   rm -rf suite_config.yaml
4. Reinstall the ITSMA suite by using the ITOM CDF UI.

Suite pods have errors and cannot recover

Reinstall the suite. To do this, follow these steps:

1. Run the following commands:
   kubectl delete namespace {yourNamespaces} && sleep 60
   rm /var/vols/itom/core/suite-install/itsma/suite_config.yaml
2. Reinstall the suite from the ITOM CDF UI.

Failed to pull images from the localhost because of a "Status 503" error

The images have been pushed to the local registry, but cannot be pulled from the localhost:5000 registry. After describing the pod, the following error is returned:

```
Error: Status 503 trying to pull……
```

This issue occurs because the kube-registry-v0-xxx pod is restarted, but the registry proxy is down on the server. Follow these steps to solve this problem:

1. Restart the pods kube-registry-proxy-xxx on all master nodes and worker nodes by running the following command:
   docker ps|grep kube-registry-proxy|grep -v pause|awk '{print $1}'|xargs docker restart
2. Run the following command to make sure that the local registry is connected:
   curl http://localhost:5000/v2/_catalog

Log error: No License File, use InstantOn

When installing the suite, many pods such as sm-rte-xxx are in the CrashLoopBackOff state. The following error occurs in the pod log:

```
No License File, use InstantOn.
```

To solve this problem, restart the server and increase the CPU and memory.

ITSMA is down

The suite is down and you cannot access ITSMA. The pods used to be running in parallel; however, all the pods are now running on a single node.

To solve this problem, check the node status by running the **kubectl describe node** command. Pay attention to the events and resource parts. To ensure a stable ITSMA deployment, we recommend using one master node and at least three worker nodes, depending on your suite size. For more information, see Sizing.

Cannot reinstall ITSMA

After you uninstall ITSMA and attempt to reinstall ITSMA, the installation fails because of an error that resembles the following:

Tolerations: <none>

Events:

FirstSeen LastSeen Count From SubobjectPath Type Reason Message

--------- -------- ----- ---- ------------- -------- ------ -------

11h <invalid> 134 {kubelet xx.xxx.xxx.xx} spec.initContainers{install} Normal Pulled Container ima already present on machine

11h <invalid> 133 {kubelet xx.xxx.xxx.xx} Warning FailedSync Error syncingnstall" with RunInitContainerError: "init container \"install\" exited with 1"

10h <invalid> 125 {kubelet xx.xxx.xxx.xx} spec.initContainers{install} Normal Created (events with common reason co

10h <invalid> 125 {kubelet xx.xxx.xxx.xx} spec.initContainers{install} Normal Started (events with common reason co

11h <invalid> 2937 {kubelet xx.xxx.xxx.xx} spec.initContainers{install} Warning BackOff Back-off restarting failed do

11h <invalid> 2913 {kubelet xx.xxx.xxx.xx} Warning FailedSync Error syncing pod, skipping: nInitContainerError: "init container \"install\" exited with 1", failed to "StartContainer" for "install" with CrashLoopBackOff: "Back-off 5m0s restarting failed container=install pod=propel-postgresql-0x"]

This problem occurs because time synchronization between the master and worker nodes is lost. To resolve this problem, configure time synchronization between the nodes, and then reinstall ITSMA.

Service Management: an HTTP 500 error occurs

You may encounter an HTTP 500 error when you access Service Management.

This happens if the SM RTE crashes in a containerized environment. However, the service will restart automatically. Try accessing Service Management a few seconds or minutes later.

Service Management: Smart Email does not work with the EWS protocol

To make Smart Email work with the EWS protocol (For example, https://<mail.server>/ews/exchange.asmx), do the following:

1. Browse to the Service Management yamls directory, and then run the **kubectl delete -f smsuite.yaml** command to stop the sm-rte pods.
   Here is a sample Service Management yamls directory: /var/vols/itom/core/suite-install/itsma/output/itom-sm-9.52.1009.RC005/yamls.
2. Save the customer's certificate (for example, smtp.cer) to the source folder.
   Here is a sample source folder: /var/vols/itom/itsma/itsma-itsma-global/certificate/source.
3. The system automatically re-generates itsma-truststore.jks and itsma-cer.pem in the ca-trust folder.
   Here is a sample ca-trust folder: /var/vols/itom/itsma/itsma-itsma-global/certificate/ca-trust.
4. Browse to the Service Management yamls directory, and then run the **kubectl create -f smsuite.yaml --validate=false** command to start the sm-rte pods.
   Here is a sample Service Management yamls directory: /var/vols/itom/core/suite-install/itsma/output/itom-sm-9.52.1009.RC005/yamls.
5. The system will restart the following sm-rte pods in several minutes:
   sm-rte
   sm-rte-emailout
   sm-rte-irque
   sm-rte-report-export
   sm-rte-scheduler
6. Configure Smart Email settings.

Performance: pods are not evenly distributed to workers

Occasionally, pods are not evenly distributed to workers even when some workers have only few-resource-consuming pods running.

In this case, you can set the same value for the *resource requests* parameter and the *resource limits* parameter so that Kubernetes guarantees the resource allocation.

Example in the .yaml file:

```
resources:

   limits:

      cpu: "4"

      memory: 4096M

   requests:

      cpu: "4"

      memory: 4096M
```

Performance: when the concurrent users are more than 2000, the suite landing page takes more than 25 seconds to open

In this case, you can add more pods to the itsma-landing-page service from the CDF Management Portal.

Performance: the Linux buffer does not release

The Linux buffer on master or worker nodes does not released during suite installation and operation even when the free memory is low.

To resolve the issue, use the following commands to release buffer manually:

```
echo 1 > /proc/sys/vm/drop_caches
echo 2 > /proc/sys/vm/drop_caches
echo 3 > /proc/sys/vm/drop_caches
sync
```

Performance: an http 502 error might occur because of port exhaustion

An http 502 error might occur because of port exhaustion when you have more than 20 concurrent users.

To avoid port exhaustion, set Linux kernel parameters on master and worker nodes to enable fast reuse and port recycling:

```
tcp_tw_recycle
tcp_tw_reuse
```

Sample script to enable the parameters:

```
echo 'net.ipv4.tcp_tw_recycle = 1' >> /etc/sysctl.conf
echo 'net.ipv4.tcp_tw_reuse = 1' >> /etc/sysctl.conf
/sbin/sysctl -p
```

Cannot launch CMDB in Internet Explorer

When launching CMDB, a blank page is displayed.

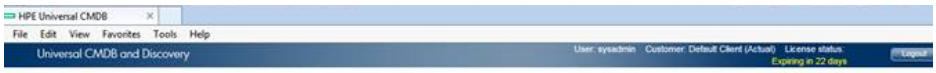

This problem occurs when the CMDB Jar files are not correctly downloaded to the client machine. These Jar files are stored in the **%TEMP%\UcmdbAppletJars** folder on your local machine. For example: **C:\Users\<user>\UcmdbAppletJars.**

To solve this problem, delete the **%TEMP%\UcmdbAppletJars** folder on your local machine and then launch CMDB again.

Smart Analytics troubleshooting

This topic provides guidelines that can help you troubleshoot problems that you may encounter when migrating to containerized Smart Analytics.

- Failed to execute the restore scripts due to lack of user permissions
- Failed to test connection on Smart Analytics
- Error: "Back up process is not started"
- Failed to back up Hot Topic Analytics and Smart Ticket Content in ITSMA suite before stopping the Smart Analytics component pods
- Parameters needed for the restore scripts “2-MergeConfiguration.sh” and “3-StartSMApodsRestoreContentData.sh”
- How to restart Smart Analytics component pods when the scale-out should be done before restore?
- The Smart Analytics pods fail to start when a node restarts abnormally during the IDOL indexing or training process
- Smart Analytics Content data restore has started, but failed in the operation process
- The script hangs when waiting for content data to be restored
- Failed to start Smart Analytics component pods
- Your classic Smart Analytics is installed on Windows operation system
- Error occurs when backing up the category file
- Failed to scale out the content server

Failed to execute the restore scripts due to lack of user permissions

You must log on to the ITSMA suite master node as a user who has the permissions to run the Kubernetes commands and has permissions to generate files. For example, the **root** or **itsma** user.

You must log on to the NFS server as a user who has the permissions to read and write files in the <Smart Analytics NFS root folder> folder.

Failed to test connection on Smart Analytics

You may encounter the following error message:

```
Send request to pod <Name> timeout. Please run this script again. Exit.
```

Follow these steps:

1. Access https://<EXTERNAL_ACCESS_HOST>/itsmaconfig/.
2. Click **Operation** > **Smart Analytics**, and then expand the **Smart Analytics Assistant** section.
3. Check the status of all components that are listed in the table and see if any component is in the Offline status.
   If there is any component in Offline status, log on to a master node of ITSMA suite and then run the **kubectl get pods --all-namespaces | grep smarta** command to display a list.
   For example:

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
itsma1	smarta-installer-3575148812-tk10x	1/1	Running	0	22h
itsma1	smarta-mix-l2proxy-3518347138-tlvkz	1/1	Running	0	21h
itsma1	smarta-mix-proxy-1162347695-p5511	1/1	Running	0	21h
itsma1	smarta-nginx-3439835978-2gwx8	1/1	Running	0	3h
itsma1	smarta-search-1816585793-skgb8	2/2	Running	0	21h
……	……	……	……	……	……

4. If the value in the STATUS column is not “Running”, or the value in the READY column is “0/1” or “1/2”, run the **kubectl logs <NAME> -n <NAMESPACE>** command to check the log.
5. If detailed log information is required, find the corresponding configuration file, locate the [Logging] section, set the value of the **logLevel** parameter to full, and then restart the component for troubleshooting purpose. You must change the value back to normal afterwards.

See the following table for the configuration file locations:

NAME	Configuration file location
smarta-mix-l2proxy-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/mixProxy/l2proxy.cfg
smarta-mix-proxy-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/st/mixProxy/proxy.cfg
smarta-smsp-con-1a-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/smsp/content1a/Content.cfg
smarta-smsp-con-1b-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/smsp/content1b/Content.cfg
smarta-smsp-dah-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/smsp/dah/dah.cfg
smarta-smsp-dih-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/smsp/dih/dih.cfg
smarta-smsp-qms-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/smsp/qms/qms.cfg
smarta-ss-agent-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/agentstore/agentstore.cfg
smarta-ss-cfs-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/CFS/CFS.cfg
smarta-ss-con-1a-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/content1a/Content.cfg
smarta-ss-con-1b-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/content1b/Content.cfg
smarta-ss-con-2a-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/content2a/Content.cfg
smarta-ss-con-2b-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/content2b/Content.cfg
smarta-ss-dah1-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/dah1/dah.cfg

smarta-ss-dih-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/dih/dih.cfg
smarta-ss-imgsvr-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/ss/imageserver/imageserver.cfg
smarta-st-con-1-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/st/content1/Content.cfg
smarta-st-con-2-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/st/content2/Content.cfg
smarta-st-proxy-<number>-<hash>	<SMART ANALYTICS FOLDER>/config/idol/st/proxy/agentstore.cfg

To restart a component, run the **kubectl delete pod <NAME> -n <NAMESPACE>** command. Wait for several minutes, and then run the **kubectl get pods --all-namespaces | grep smarta** command to check if the component is started.

Error: "Back up process is not started"

You may encounter the following error message when executing the 1-StopSMApods.sh restore script:

Back up process is not started. Exit.

The Hot Topic Analytics and Smart Ticket Content backup process is not started in the ITSMA suite before stopping the Smart Analytics component pods.

Follow these steps:

1. Make sure that all the Smart Analytics component pods are ready by checking the list in Smart Analytics Assistant on the ITSMA suite configuration page.
2. Run the following command:
   kubectl logs <HTA Content Pod Name> -n < NAMESPACE>
3. Check the logs to make sure that the Hot Topic Analytics and Smart Ticket Content process is running in the container.

Failed to back up Hot Topic Analytics and Smart Ticket Content in ITSMA suite before stopping the Smart Analytics component pods

Refer to the Smart Analytics help documents to fix this problem, and then run the 1-StopSMApods.sh restore script again to back up Hot Topic Analytics and Smart Ticket Content.

Follow these steps:

1. Find the cause of this problem.
   a. Access the help document in your external Smart Analytics installation folder, and then visit the local web page <Smart Analytics installation folder>\IDOL\help\IDOLServerReference\index.html. Search for the keyword "IndexerGetStatus", and then locate the "IndexerGetStatus Status Codes" section for details about the abnormal restore status.
   b. According to the information provided by the 1-StopSMApods.sh restore script, find the corresponding description message of the problem.
2. Solve the problem in the indexing process according to the explanation.
3. Run the 1-StopSMApods.sh restore script again to back up Hot Topic Analytics and Smart Ticket Content. If a backed up file is created, the script will continue to stop the Smart Analytics component pods.

Parameters needed for the restore scripts “2-MergeConfiguration.sh” and “3-StartSMApodsRestoreContentData.sh”

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. By default, there are two Hot Topic Analytics and Smart Ticket Content pods and every two Smart Search Content pods work together as a group. So the numbers must be even numbers.

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>***.

For the 2-MergeConfiguration.sh restore script, these two numbers are the actual numbers of Smart Analytics Content pods in ITSMA suite. After you run the 1-StopSMApods.sh restore script, values for these two numbers are shown in the result.

For the 3-StartSMApodsRestoreContentData.sh restore script, these two parameters are the numbers of Smart Analytics Content pods to be merged. After you run the 2-MergeConfiguration.sh restore script, values for these two numbers are shown in the result.

How to restart Smart Analytics component pods when the scale-out should be done before restore?

If your Smart Search Content pod group number in ITSMA suite is smaller than the number of your external Smart Search Contents, you can restart the Smart Analytics pods on the master node and then do Smart Search Content scale out on Suite Configure page. To start the Smart Analytics pods, run the following commands in the folder where 2-MergeConfiguration.sh is saved on the master node.

1. Run the following command:
   for i in `ls yamls`; do if [[$i = "ss"* || $i = "st"*]]; then kubectl create -f yamls/$i; fi done
2. Run the following command to check the pods' status and verify whether all pods are running:
   kubectl get pods --all-namespaces | grep smarta
3. Run the following commands:
   kubectl create -f yamls/l2proxy.yaml
   kubectl create -f yamls/proxy.yaml

The Smart Analytics pods fail to start when a node restarts abnormally during the IDOL indexing or training process

Information that resembles the following is displayed in the content server logs:

24/03/2017 08:13:02 [1] 00-Always: Reference index needs rebuild now
24/03/2017 08:13:02 [1] 00-Always: Generating new reference index
24/03/2017 08:13:14 [1] 00-Always: Creating new reference index: processed 5% of documents (elapsed: 11.72s)
24/03/2017 08:13:27 [1] 00-Always: Creating new reference index: processed 10% of documents (elapsed: 25.19s)
24/03/2017 08:13:41 [1] 00-Always: Creating new reference index: processed 15% of documents (elapsed: 39.17s)
24/03/2017 08:13:55 [1] 00-Always: Creating new reference index: processed 20% of documents (elapsed: 53.01s)

To resolve the issue, on each content server, perform the following steps

1. Locate and temporarily remove the following section in the .yaml file:
   livenessProbe:
   httpGet:
   path: /?action=encrypted&data=MzJ8mNJbn7G5d8dxZBSaBSMBM0dIdzC0Ovzv%2BTVHrfJi84M%3D
   port: 10010
   initialDelaySeconds: 600
2. Restart the pod.

Smart Analytics Content data restore has started, but failed in the operation process

Refer to the Smart Analytics help documents to fix this problem, and then run the 3-StartSMApodsRestoreContentData.sh restore script again to restore Smart Analytics Content data.

Follow these steps:

1. Find the cause of this problem.
   a. Access the help document in your external Smart Analytics installation folder, and then visit the local web page <Smart Analytics installation folder>\IDOL\help\IDOLServerReference\index.html. Search for the keyword "IndexerGetStatus", and then locate the "IndexerGetStatus Status Codes" section for details about the abnormal restore status.
   b. According to the information provided by the 3-StartSMApodsRestoreContentData.sh restore script, find the corresponding description message of the problem.
2. Solve the problem in the indexing process according to the explanation.
3. Run the 3-StartSMApodsRestoreContentData.sh restore script again to restore Smart Analytics Content data.

The script hangs when waiting for content data to be restored

When the execution of the 3-StartSMApodsRestoreContentData.sh restore script near its end, this script will wait for the restore process to be finished. When a Smart Analytics Content pod is performing the restore process, the response of checking restore status from the pod may be delayed and the system may display the following message:

```
..Hot Topic Analytics and Smart Ticket content data is restored.

..
```

Each dot stands for a pod in this message. This message indicates that the script hangs when the second Smart Search Content pod is being restored. You can run the following command to verify that weather the pod is performing the restore process or not.

kubectl logs -f <Name> -n < NAMESPACE>

If the logs show no error, you can wait a few minutes for the process to be completed.

Failed to start Smart Analytics component pods

When starting the Smart Analytics component pods during the execution of the 3-StartSMApodsRestoreContentData.sh restore script, this process may quit because the pod cannot start up normally. You can do the following to start the pod manually, and then run the 3-StartSMApodsRestoreContentData.sh restore script again.

Follow these steps:

1. Log on to the NFS server, and then change the directory to /var/vols/itom/core/suite-install/itsma/output/itom-smartanalytics-<version.tag>/yamls.
2. Copy and paste the following .yaml files to the master node of ITSMA suite:

   mix-mainproxy.yaml
   mix-level2proxy.yaml
   st-content-1.yaml
   st-content-2.yaml
   ss-content-1a.yaml
   ss-content-1b.yaml
   ss-content-2a.yaml
   ss-content-2b.yaml

   > The number of Smart Search Content pods' yaml file (ss-content-*.yaml) is determined by your ITSMA suite. The default .yaml files are listed in the example above. You must copy all of them to your production environment.

3. Run the **kubectl delete -f <component.yaml>** command to stop all Smart Analytics components.
4. After all pods are stopped, run the **kubectl get pods -n <namespace> | grep smarta** command to make sure that all pods you just deleted no longer exist.
5. Run the **kubectl create -f <component.yaml>** command to start Smart Analytics component pods. See the following list for the sequence to start the default pods:
   st-content-1.yaml
   st-content-2.yaml
   ss-content-1a.yaml
   sscontent-2a.yaml
   ss-content-2b.yaml
   mix-level2proxy.yaml
   mix-mainproxy.yaml

   > Before you start the level 2 proxy server pod, make sure that all content server pods show the "1/1 running" status. Run the **kubectl get pods -n <namespace> | grep 'smarta-s.-con** command to get all content server pod's status. Before you start the main proxy server pod, make sure that the level 2 proxy server pod has started and shows the "1/1 running" status. Run the **kubectl get pods -n <namespace> | grep 'smarta- mix-l2proxy** command to get level 2 proxy server pod's status.

Your classic Smart Analytics is installed on Windows operation system

If you want to migrate to containerized Smart Analytics from the Windows operation system, you must do full re-index for Hot Topic Analytics after the migration steps are finished.

Follow these steps:

1. From the System Navigator of your Service Manager Client, click **System Administration** > **Ongoing Maintenance** > **Smart Analytics** > **Hot Topic Analytics**.
2. Select a Hot Topic Analytics configuration record from the configuration list. For example, Incident.
3. Click the Start Index button to start indexing.
4. Repeat the previous step for each configuration record.

Error occurs when backing up the category file

The system displays the following error message when backing up the category file. As a result of the backup, the system generates two category files.

```
Failed to send request to 10.46.132.3:9020.
```

Follow these steps:

1. Go to the Smart Analytics installation folder, and then remove the SmartAnalyticsMigration folder.
2. Run the backup tool again.

Failed to scale out the content server

After you click **Add New Content Group** on the ITSMA configuration page, **i**f the "Adding a new content group" status hangs for more than 10 minutes, you can log on to the ITSMA suite master node, and then run the **kubectl get pods --all-namespaces | grep smarta-ss-con** command to check if the content group is already added and started.

Follow these steps if there is no content group added:

1. Run the **kubectl delete -f /opt/kubernetes/objectdefs/suite.yaml** command.
2. Run the **kubectl create -f /opt/kubernetes/objectdefs/suite.yaml** command.
3. Run the **kubectl get pods -n core** command, and then wait until all pods' status are running.
4. Log on to the NFS server, and then open the <Smart_Analytics_NFS>/config/idol/count file. Search for *"content":{"ss":<NUMBER>* and update it to *"content":{"ss":0*.
5. Log on to ITSMA suite master node, run the **kubectl get pods -n <NAMESPACE> | grep smarta-installer | awk '{print $1}' | xargs kubectl delete pod -n <NAMESPACE>** command.
6. Run the **kubectl get pods -n <NAMESPACE> | grep smarta-installer** command, and then wait until the status is running.
7. Perform scale-out again.

Follow these steps if the content group is already added and started:

1. Run the **kubectl get pods -n core** command to see if there is a pod whose name is in "kube-proxy-<IP>" format and its status is "CrashLoopBackOff".
2. Log on to the failed node, and the run the **/opt/kubernetes/bin/kube-restart.sh** command.
3. Wait for several minutes till the process is finished.

ITSMA suite configuration troubleshooting

When you are configuring ITSMA, you may encounter the following issues:

- After you change any configurations and click **Apply**, the pop-up window displays a list of "abnormal objects" in red. This means these pods, deployments or configmaps' status is not "Running" or "Completed".
- Instances are displayed as unhealthy. See the following screen shot as an example:

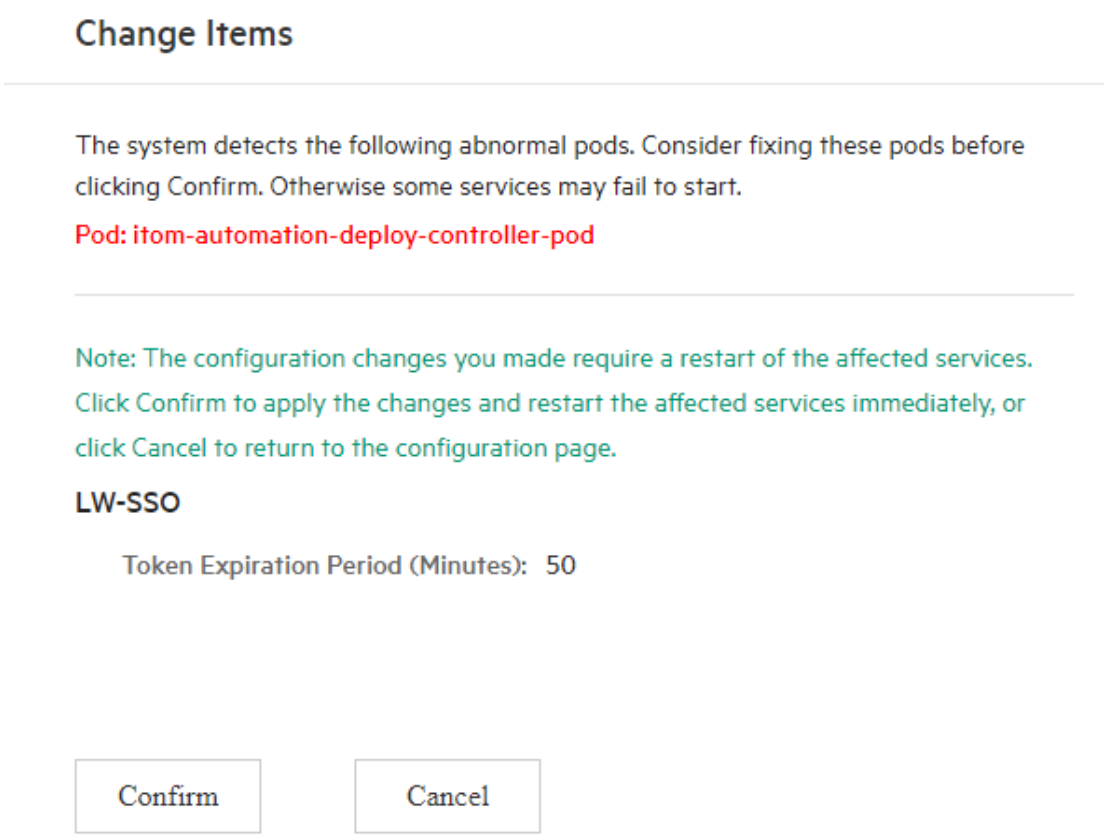

If you failed to apply your configurations, you can access the suite configuration log file at /var/vols/itom/itsma/itsma-<namespace>/logs/itsmaconfig-1.0, and then refer to the following table for solutions:

Error code	Solution
errorCode = csrf-400	This error code means the system failed to get the authentication token and the CSRF token from the suite platform. You may find that the Config UI cannot load any data. Please check or restart the pod suite-installer in Core namespace.
errorCode = etcd-400	This error code means the system failed to get configuration from ETCD. You may find that the Config UI cannot load any data. Please check or restart the pod suite-installer in Core namespace.
errorCode = get-500	This error code means an internal error occurred during getting configuration in Config backend. Please check the log file for error messages.
errorCode = prop-204	This error code means no valid property changes are sent to Config backend, or the property changes cannot be found in any configmap. Hence, no action will be triggered by Config backend. Please check the payload from Config UI.
errorCode = save-400	This error code means the system failed to save configurations to ETCD. You may find that Config UI may load old data. Please check or restart the pod suite-installer in Core namespace.
errorCode = apply-204	This error code means an **Apply** process is still running and no more Apply action can be triggered now. Please wait for the Apply process completed.

errorCode = pp-404	The following screen shot shows a correct system response if you change the Smart Analytics log level only. However, in other cases, this response is an error which means that the Apply action is blocked because no related point products services or deployments are found. 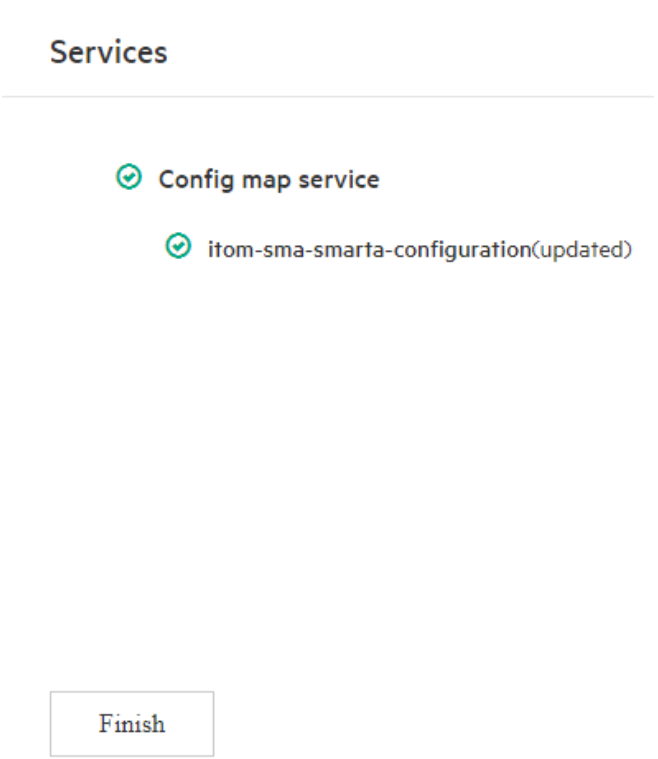
errorCode = apply-400	This error code means an internal error occurred when applying configuration in Config backend. Please check the log file for error messages.
errorCode = getConfigMap-400	This error code means an internal error occurred when getting a Configamp. Please check the request body for api /configuration/getConfigMapData. The format of request body should be displayed as follows: `{ "configMapName": "xxxx", "data": { "key1":"value1", "key2":"value2", "key3":"value3", ... } }`

errorCode = propel_tool_status_404	This error code means the system failed to get configmap propel-configmap or failed to find key propel_tool_status in the configmap data. Please check the configmap propel-configmap.
errorCode = propel_tool_status_exception	This error code means an internal error occurred when getting key propel_tool_status from configmap propel-configmap. Please check the configmap propel-configmap.
Error message showed in config UI for Applying process	As shown in the following screen shot, the LW-SSO configmap is missing. Hence, the Apply process is blocked at the first step and the other four steps are also skipped. 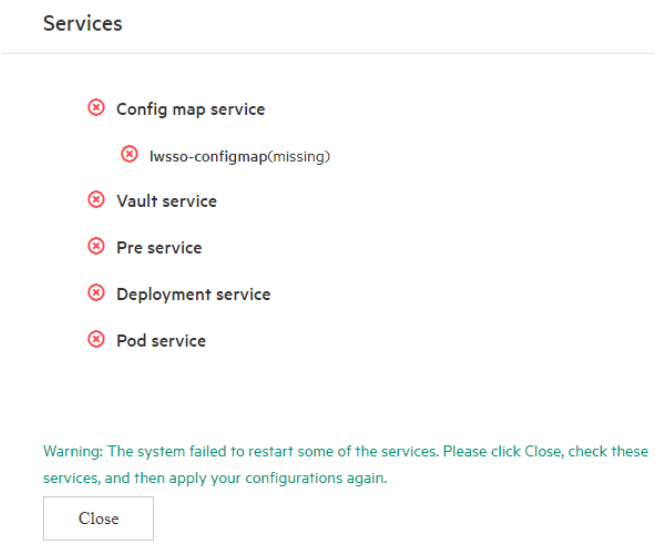 As shown in the following screen shot, the SM RTE deployment is missing. Hence, the Apply process is blocked at the deployment step and the rest deployment are also skipped. The status for each deployment show its status. 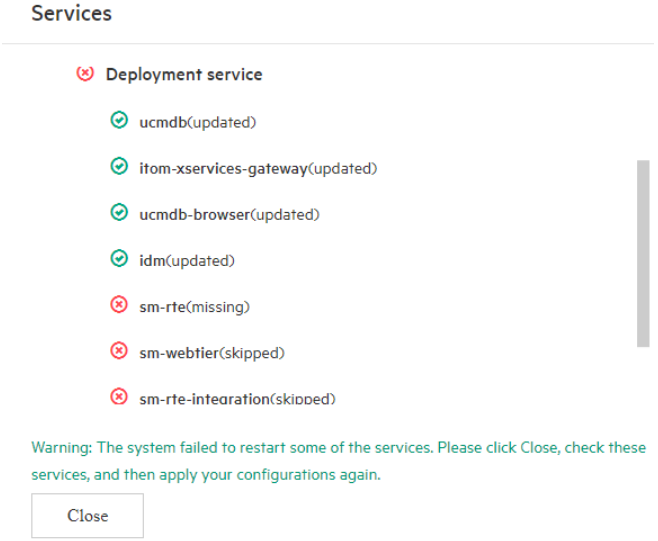 As shown in the following screen shot, the sm-rte-integration-2389502142-tk5kk pod failed to restart. Hence, the Apply process is blocked at the last step. The Config UI will hang for about 30 minutes for system restart.

Services

sm-webtier(updated)

sm-rte-integration(updated)

Pod service

sm-webtier-4143951259-r25jg

idm-1525767860-0147q

ucmdb-browser-3536406496-jpw1s

sm-rte-integration-2389502142-tk5kk

ucmdb-2449047849-lg4g3

Note: The system is applying your changes.