Administer

Data Center Automation Premium 2017.05

Document Release Date: May 2017

Software Release Date: May 2017

Hewlett Packard Enterprise

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

AMD is a trademark of Advanced Micro Devices, Inc.

Intel® and Intel® Itanium® are trademarks of Intel Corporation in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/. 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 web site at: https://softwaresupport.hpe.com/. Most of the support areas require that you register as an HPE Passport user to sign in. Many also require a support contract. To register for an HPE Passport ID, click Register on HPE Support site or click Create an Account on the HPE Passport login page.

Table of Contents

Administer

As an IT administrator, you can perform the following tasks in ITOM CDF:

- Access ITOM CDF
- Password management
- Infrastructure management
- Nodes
- LDAP
- LWSSO
- User management
- License management
- Local registry
- Namespaces
- Workloads
- Services and discovery
- Persistent volume claims
- Configuration
- Logs
- Security
- Restart the ITOM CDF

Access ITOM CDF

To access the ITOM Container Deployment Foundation (CDF):

1. Launch the ITOM CDF from your browser:
   https://*<master node FQDN>*:5443

   You must use the master node host's FQDN instead of its IP address in this URL. That is, the name you specified for EXTERNAL_ACCESS_HOST in the **install.properties** file.

   Access the application using a supported web browser, from any computer with a network connection (intranet or Internet) to the servers. It is recommended to restore your browser settings to default. You will be asked to change the default password at first login.

2. Log in to ITOM CDF as the **admin** user.

   Use the out-of-box password if this is your first login or use the password that you specified at your initial login after installation. The default passwords are stored in the

DCA Premium 2017.05 Password Management document on the HPE Software Support website (https://softwaresupport.hpe.com/).

Log out

To log out, click the user name and select **Logout**. The application closes and the log in screen is displayed.

When you have completed your session, it is recommended that you log out to prevent unauthorized use.

Password management

The default passwords are stored in DCA Premium 2017.05 Password Management document on the HPE Software support website.

Change the default ITOM CDF password

1. On the ITOM CDF console, click on the user name and select **Change Password**.

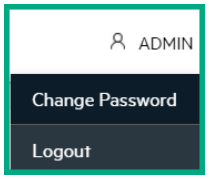

2. Enter the original password, the new password, and verify the new password.
   The password should have minimum 8 characters and must use the following categories:
   - Uppercase characters
   - Lowercase characters
   - Number
   - At least one special character
3. Click **UPDATE PASSWORD**.

Change the default UCMDB password

Important

- The password must have at least 8 characters.
- The password must have digits, uppercase letters, lowercase letters, and at least one special character from the following: , \ / . ,_ ? & % = + - [] () |

1. Get the namespace of your DCA deployment by running the `kubectl get ns` command. For example, the namespace can be dca1.
2. Get the name of the pod by running the `kubectl get po -n <namespace> | grep ucmdb-deploy` command.
3. Connect to the container by running the `kubectl exec -n <namespace> -it <podname> /bin/bash` command.
4. Run the `update_secret ADMIN_PASSWORD_VAULT_KEY <new_password>` command.
5. Update the credentials in the OO configuration:
   a. Connect to the OO portal (https://<hostname>:33445).
   b. Go to **Content Management** > **Configuration Items** > **System Accounts**.
   c. Update the `UCMDB_CREDENTIALS` and `UCMDBCredentials` entries with the new password.
6. Restart the containers so that the new password is used. For example, in the console of one of the nodes, run the `kube-restart.sh` command.

Change the SA password

You are required to provide the SA user name and password when installing DCA. If you change the password of the SA user, you must also update it in the OO configuration. To do this:

1. Connect to the OO portal (*https://hostname:33445*).
2. Go to **Content Management** > **Configuration Items** > **System Accounts**.
3. Update the `SA_CORE_CREDENTIALS` entry with the new password.

Change the Vertica password

You are required to provide the Vertica user name and password when installing DCA. If you change the password of the Vertica user, you must also update it in the DCA configuration so that the services continue working. To do this:

1. Get the namespace of your DCA deployment by running the `kubectl get ns` command. For example, the namespace can be dca1.
2. Get the name of the pod by running the `kubectl get po -n <namespace> | grep ucmdb-deploy` command.
3. Connect to the container by running the `kubectl exec -n <namespace> -it <podname> /bin/bash` command.
4. Run the `update_secret ANALYTICS_DB_PASSWORD_KEY <new_password>` command.
5. Restart the containers so that the new password is used. For example, in the console of one of the nodes, run the `kube-restart.sh` command.

Related topic

Manual verification commands

Infrastructure management

To view and monitor infrastructure resources, click **ADMINISTRATION** > **Admin.** This page displays information about the following:

UI element	Description
Namespaces	The list of the current default namespaces as well as the namespaces for the suites. Every suite on the same Kubernetes cluster is deployed in a different namespace.
Nodes	The composition of the Kubernetes cluster in terms of servers on which the cluster were installed (master and worker nodes, the physical servers or the VMs).
Persistent Volumes	The persistent volume configuration for one or more suites. These volumes contain the data that needs to live outside of the containers.

Related topics

User management

Namespaces

Nodes

Persistent volume claims

Nodes

The **Nodes** page provides the CPU and memory usage history of the selected namespace, a list of the predefined labels, and the list of nodes of the selected namespace.

When the CPU load is over 80%, it significantly impacts the efficiency of network transmission between the base infrastructure environment. It is recommended to control the CPU load so it is less than 80% by separating the suite instance into multiple worker nodes: adding more worker nodes and killing the pods on heavy-load nodes and deploying those pods on the newly added worker nodes.

View the existing nodes

To view existing nodes, click **ADMINISTRATION** > **Nodes**.

The page displays the CPU and memory usage of the selected namespace during the past 15 minutes, the list of the node labels, and the status, labels, readiness, and creation timestamp of the nodes corresponding to the selected namespace.

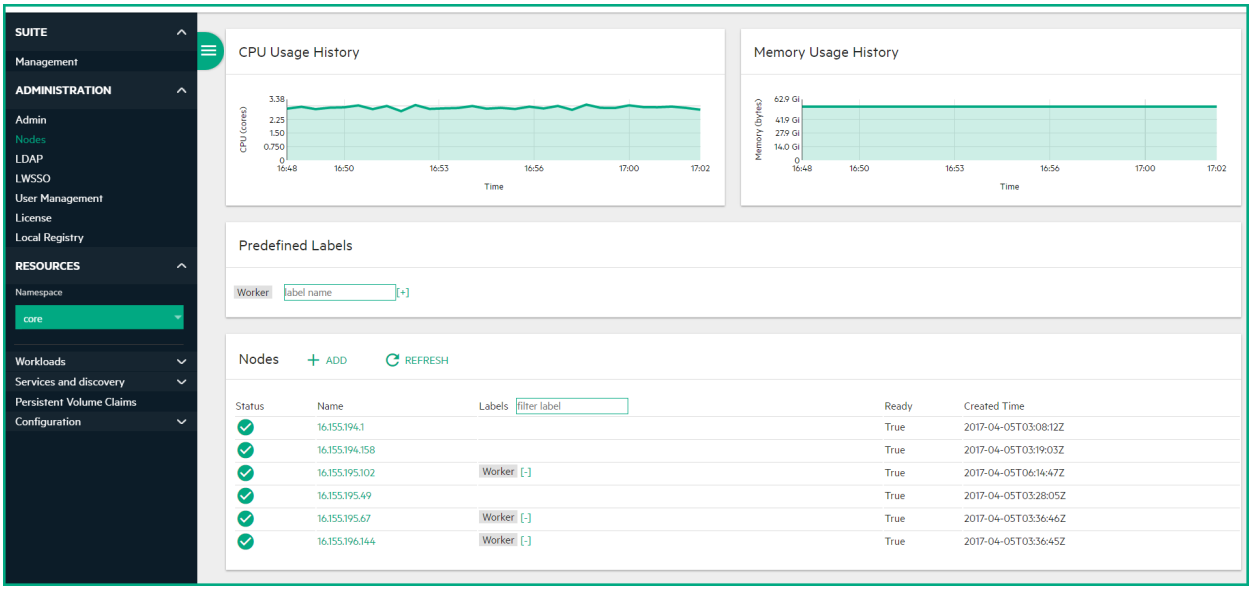

On this page, you can:

- Define a set of labels you want to use and then assign them to nodes by dragging them to the node. See Manage labels.
- Add a node. See Add a node.
- **REFRESH.** Click to refresh the display.
- Click the relevant node to see its details. See View node details.

Manage labels

Add a label

To add a label:

1. Click **ADMIN** > **Nodes**.
2. To add a label in the **Predefined Labels** area, enter the **value** and click **[+]**. The label is added to the list.

Delete a label

To delete a label:

1. Click **ADMIN** > **Nodes**.
2. In the **Predefined Labels** area, click **[-]** for the relevant label.

Assign a label

To add a label to a node:

1. Click **ADMINISTRATION** > **Nodes**.
2. Drag the relevant label the **Predefined Labels** area to the relevant node in the **Nodes** area.

Unassign a label

To remove a label from a node:

1. Click **ADMINISTRATION** > **Nodes**.
2. In the **Nodes** area, click **[-]** for the relevant label and node.

Filter labels

To filter labels:

1. Click **ADMINISTRATION** > **Nodes**.
2. Enter the relevant string or keyword in the Labels box in the table header. The labels with names that include the relevant string are listed.

Add a node

To add a node:

1. Click **ADMINISTRATION** > **Nodes**.
2. In the Nodes area, click **+ ADD**.
3. Enter:
   a. the name of the node
   b. the name of a user that can remotely execute commands on the host - typically the root user
   c. the password of a user that can remotely execute commands on the host - typically the root user
4. Click **ADD** to remotely install the extra node.

   You can add multiple nodes simultaneously with **+ ADD**:

   - Enter the host names or IP addresses for each node.
   - Enter the user name and password for every node to access remotely on the host.

   ⚠ The host names or IP addresses should be separated by spaces.
   Those added nodes share the same user name and password.
   The installation of each node runs in parallel.

View node details

To view node details:

1. Click **ADMINISTRATION** > **Nodes**.
2. In the Nodes area, select a node name from nodes list.
    - The page displays the CPU and memory usage history of the selected node for the past 15 minutes.
    - The **Details** area displays details about the selected node as well as system information.
    - The **Allocated resources** area displays the minimum CPU requests, CPU limits, memory requests, and memory limits for the container as well as the percentage of <what is in use>/<what is available>. By default, pods run with unbounded CPU and memory limits. The format is: <what is in use>/<what is available>.
    - The **Conditions** area displays the type, status, last heartbeat and transaction time, reason, and message.
    - The **Pods** area displays the CPU and memory usage history of the pod for the past 15 minutes, the name of the pod, the status, number of restarts in the cycle, the amount of time passed since the pod has been created, the cluster IP, as well as the CPU and memory usage of the pod.
      You can:
        - Click a Pod name to open the Workloads - Pods page for the pod. See Pods.
        - Click to review the pod log.
        - Click and select **Delete** to delete the pod.
    - The **Events** area displays the message, source, sub-object, count, first seen and last seen information.

   You can:
    - Click a pod name to open the **Workloads - Pods** page for the pod. See Pods..
    - Click a pod to review the pod log.
    - Click a pod and select **Delete** to delete the pod.

Related topic

Workloads

LDAP

The LDAP page enables you to configure the integration of an external LDAP directory. The LDAP Configuration page enables you to add and edit users and groups, as well as add and edit their details.

You create and manage users on your LDAP server which is connected to ITOM CDF. The IT Administrator uses the **ADMIN** tab to define users and groups for the organization.

To configure LDAP, you first add an organization, then configure the LDAP parameters for the organization.

Add an organization

1. Logon to ITOM CDF using the LDAP user and password.
2. Click **ADMINISTRATION** > **LDAP**.
3. Click **ADD CONFIGURATION**.
4. Enter the name of the organization, the description (optional), and the type of organization (**CONSUMER**), and click **CREATE**. For example: test_org.
   The new organization is listed above Consumer and Provider.

Configure LDAP for the Organization

1. Click the organization name to enter the relevant LDAP information.
2. Click **ADD CONFIGURATION**.
3. Enter relevant information about the server:

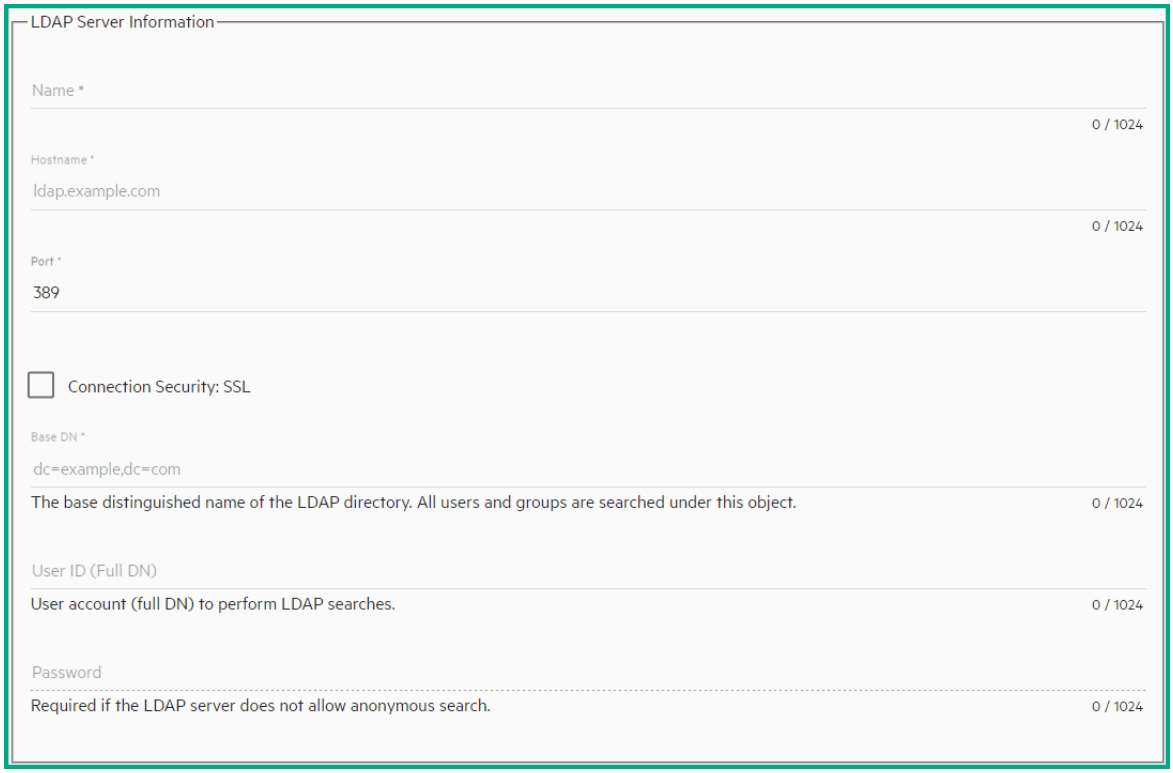

4. Enter relevant information about the authentication:

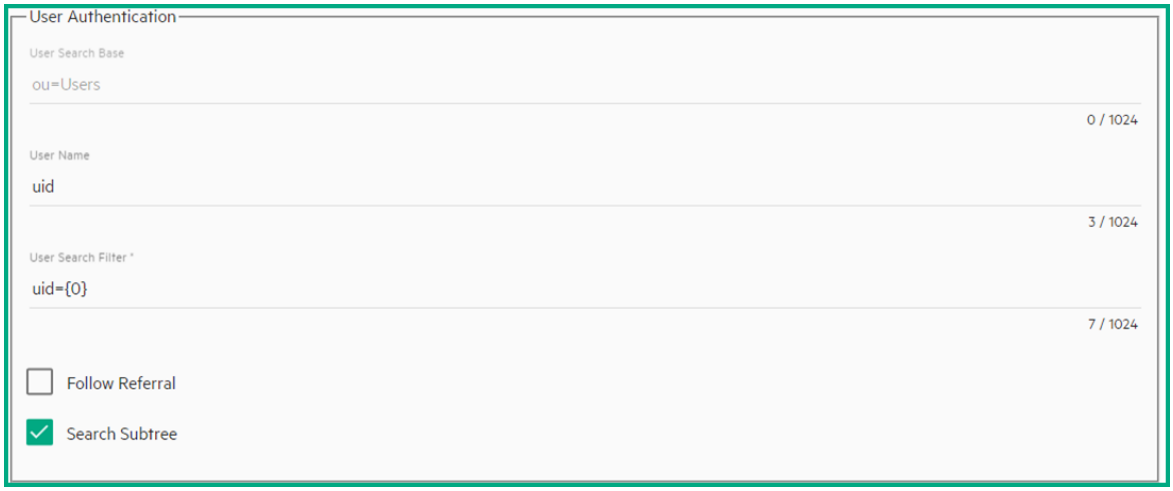

5. Enter relevant information about the attributes:

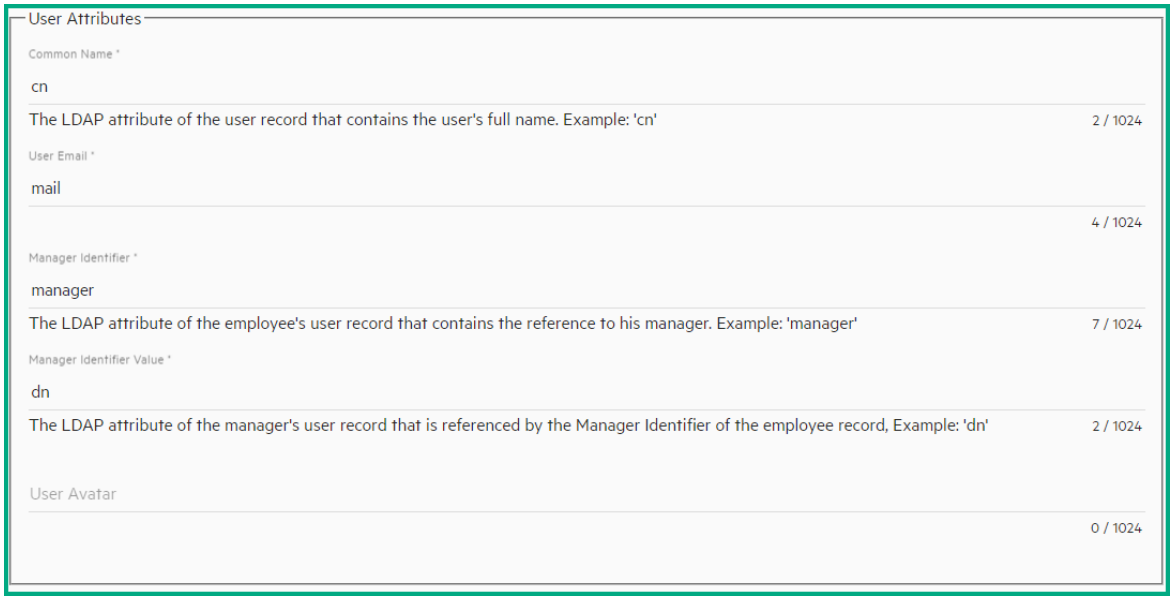

6. Enter relevant information about the user group:

7. Click **SAVE**.

Where to go from here

User management

LWSSO

Lightweight single sign-on (LWSSO) is a method of access control that enables a user to log on once and gain access to the resources of multiple software systems without being prompted to log on again.

The LWSSO page enables you to set up a single sign-on with other products.

> ⚠ The InitString and Domain of **LWSSO** have their own default values. You need to input the current user's password and then click **Show InitString** to see the InitString's default value. You can also change these default values according to your needs.

1. Click **ADMINISTRATION** > **LWSSO**.
2. Enter the InitString and Domain and click **UPDATE**.

> ✓ You can copy and paste the value of the InitString directly into other products for LWSSO integration.

User management

To access the User page, click **ADMINISTRATION** > **User Management**. The User page displays the user name, password, email, and user group. The ITOM CDF support two user groups: IT Administrators and Suite Administrators.

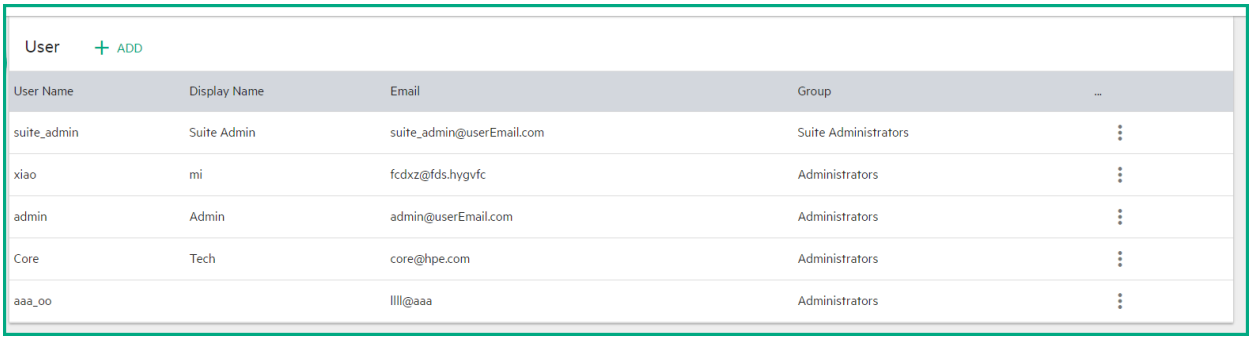

IT Administrator

Manages the shared services infrastructure and all suite products, as well as the grow/shrink functions, and adding and removing working nodes (machines). The IT Administrator is a super administrator. This user has ability to request or add resources and has wide access permissions.

The IT Administrator can create or delete users, and view or edit user information.

Suite Administrator

Manages a specific suite product. The Suite Administrator does not have access to the Admin menu and has the privileges with other operations only under a specific namespace. The Suite Administrator is responsible for the relevant suite deployment, configuration, health, images, and more.

Create a user

To create a user, click **ADD**. Enter the relevant information in the dialog box, and click **SAVE**.

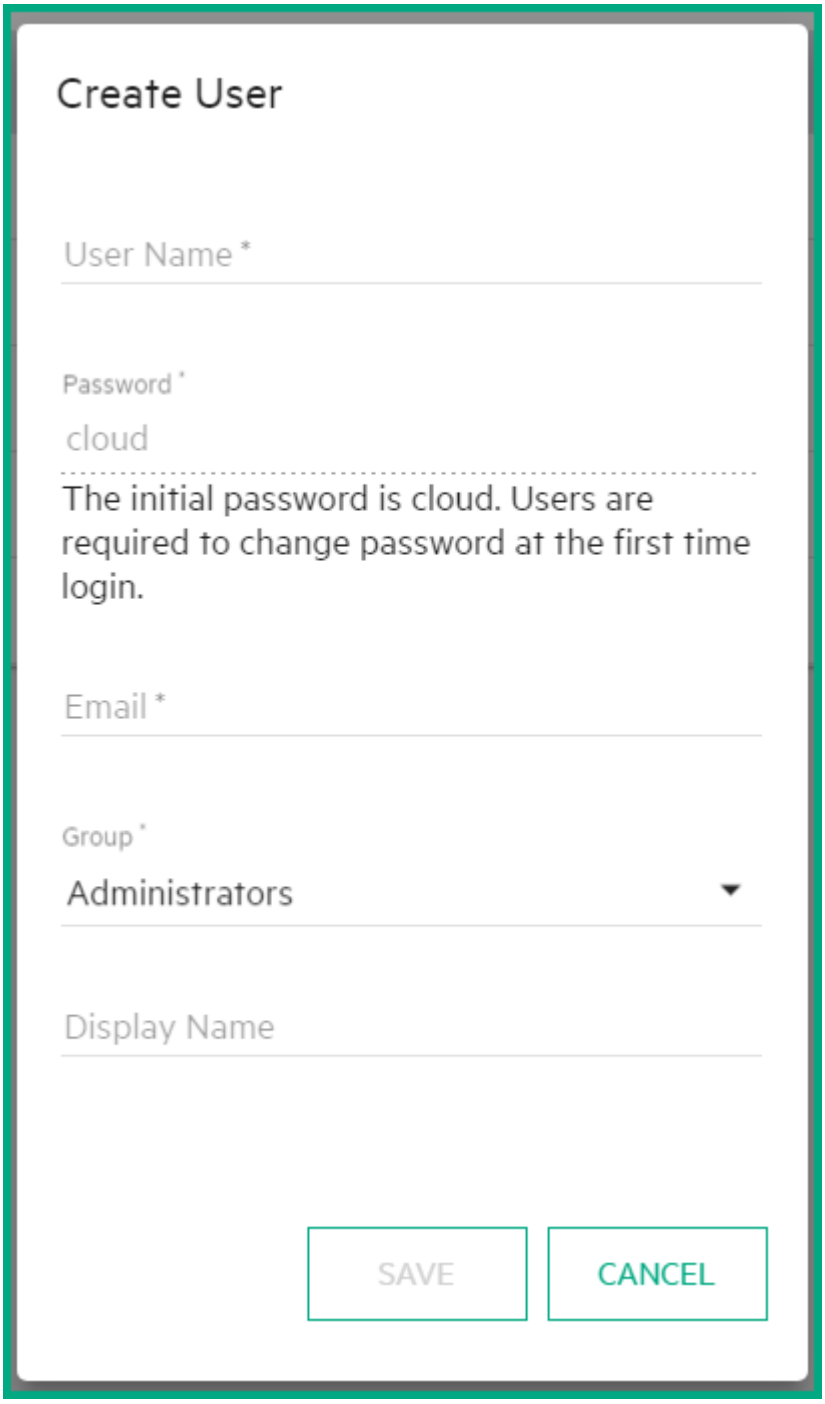

Delete a user

To delete a user, click the right-side ACTION icon for the user, select **Delete**, and then click **Delete** again to confirm the deletion.

Edit or view user information

To edit or view a user information, click the ACTION icon for the user, and then select **View/Edit**.

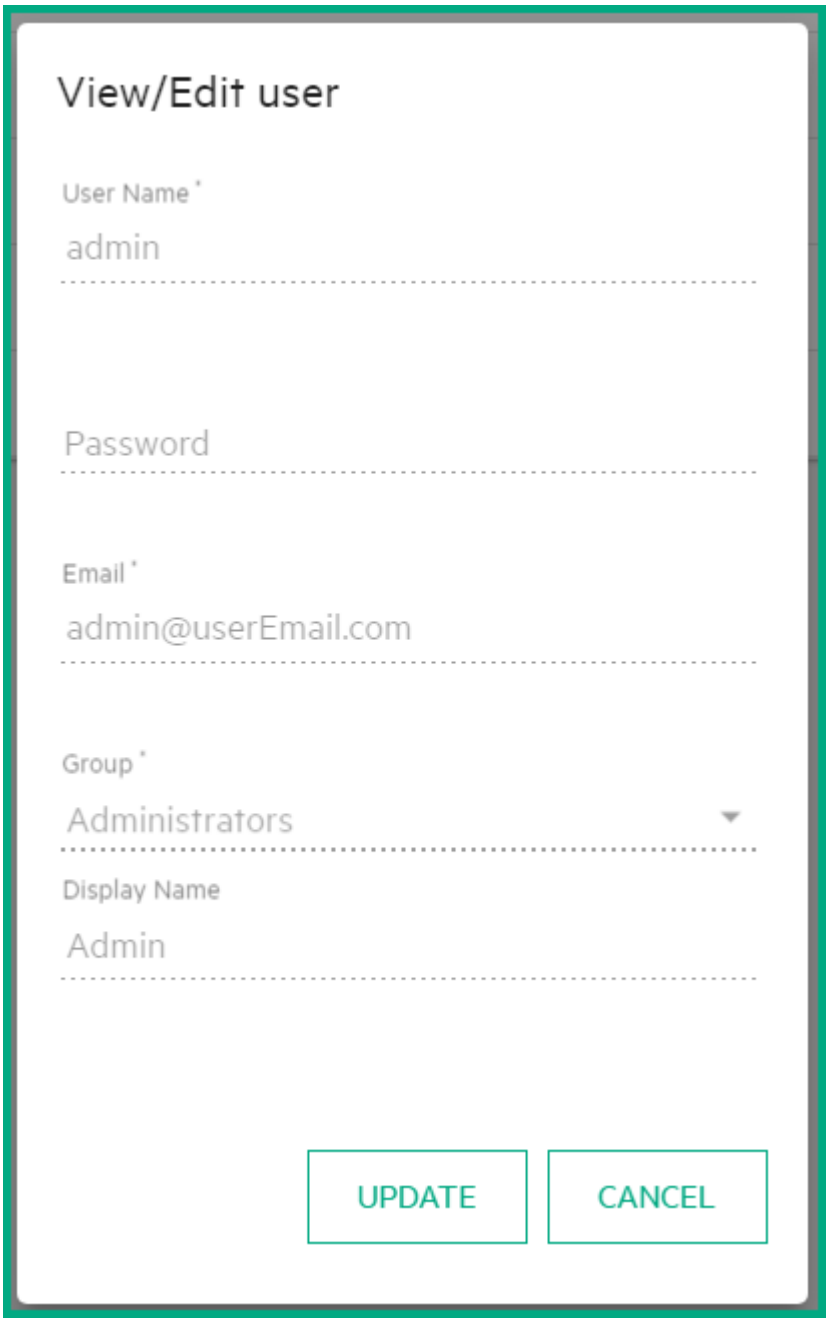

Related topic

Infrastructure management

License management

Evaluation license

An evaluation license of 90 days is installed during DCA installation. During the evaluation period all features and components can be used without any restriction. An alert is displayed on the top of the DCA Console to inform you of the evaluation license status.

On completion of the evaluation period, you will not be able to perform any task from the DCA console. Contact your administrator or HPE Software support with the **Lock Code** (lock ID) for procuring a DCA Premium 2017.05 license.

On expiry of the evaluation license, the following screen is displayed when you log in to the DCA console:

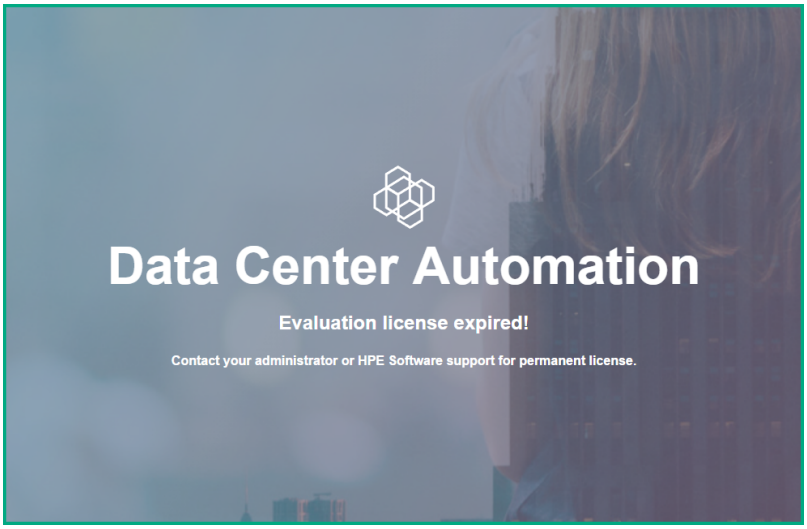

Permanent license

Once you have a permanent DCA Premium 2017.05 license from your administrator or HPE Support, you can install it from the License page on the ITOM CDF management console.

- Install permanent license
- View existing licenses
- Archive a license
- Restore an archived license
- Delete a license from the License Manager
- View the Licenses Report
- Alerts

Install permanent license

1. Click **ADMINISTRATION** > **License**.

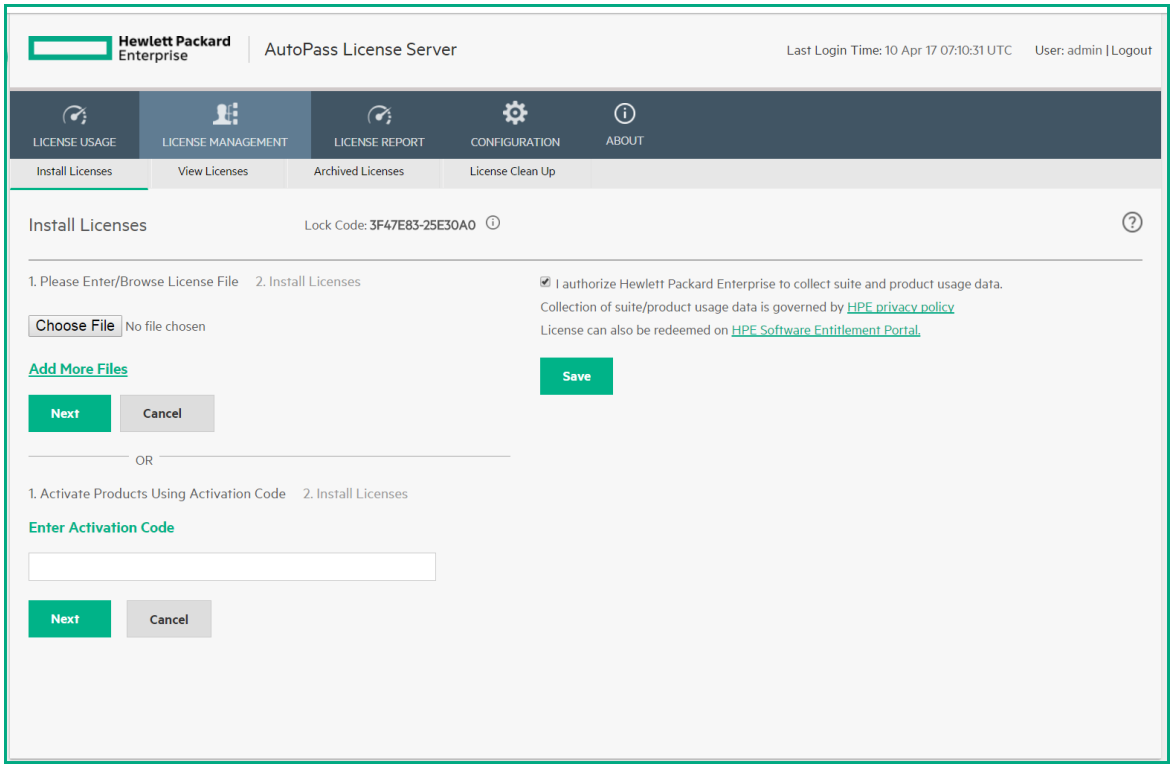

   The **Lock Code** (lock ID) is displayed on the top.
2. Click **Choose file** to select the license file in your local system.
3. Click **Add More Files** to select another license file in your local system.
4. Click the boxes to agree to the HPE End User License Agreement and authorize the suite and product usage data collecting.
5. Click **Next**.
   You can select the license keys and click **Install Licenses** to install the licenses (you can also select to go back to the previous Install Licenses page by clicking the **Back** button).
   The page displays the licenses in the selected file. You must select the licenses you want to install out of the displayed licenses.
6. After selecting, click **Install Licenses.**

> ⚠ The license consumption is tracked and displayed on the DCA UI on reaching the limit or exceeding it. Contact your administrator or HPE Software support for additional license units.

View existing licenses

> **Note**
>
> By default, there is no license installed. You can view the existing licenses only after you have installed them.

1. Click **ADMINISTRATION** > **License**.
2. Select the relevant product in **Select Product**.
   The page displays the feature ID: version, product number,capacity, start date, expiry Date, the date when it was installed, and who installed it, as well as the Lock Code.

Archive a license

1. In the View Licenses tab, select the unused licenses you want to archive.
2. Click **Archive**.

The licenses are removed from the list of installed licenses in the License Management table and become unavailable for customers to fetch and activate the products.

Restore an archived license

1. In the Archived License tab, select the product whose archived licenses you want to restore.
2. Select the relevant licenses that you want to restore.
3. Click **Restore**.

The licenses are again displayed in the License Management pane.

> **Note**
>
> You cannot restore an ID locked license that was auto archived, unless all the licenses that are locked to a lock value and belong to the same feature are deleted or archived.

Delete a license from the License Manager

1. In the Archived Licenses tab, select the product whose license(s) you want to delete.
2. Select the license to delete.
3. Click **Delete** and confirm the deletion.

View the Licenses Report

1. Click **ADMIN** > **LICENSE REPORT**.
   The license report page tracks and displays the licenses currently installed and used on the License Manager.
   It also displays specific check out information about a feature license including the product name and version, the requester ID, and the timestamp of when it was accessed last.
   You can export the license report details to Excel.
   You can also search a license with the product name, product version or requester IP address.

Alerts

The license consumption is tracked and appropriate messages are displayed in the console.

The **Learn More** link provided in the alerts lead you to the current documentation page.

The following messages are displayed in the Alerts section:

- **Evaluation license**:
  Trial license expires in XX days (where, XX is calculated from the day of installation).
- Permanent license:
  Your license consumption has reached the limit. Contact HPE Software support.
  Your license consumption has exceeded the limit by <number of> units. Contact HPE Software support.

Related topic

Overview of DCA

Local registry

This section provides information about the images that are in the local registry.

View the existing images

To view the existing images, click **ADMINISTRATION > Local Registry**. The Local Images page is displayed.

Related topic

Install DCA

Namespaces

Kubernetes supports multiple virtual clusters backed by the same physical cluster called Namespaces.

Select the namespace

Select a namespace to filter the information in the pages of the user interface and display only the items related to the namespace.

1. Click **RESOURCES** > **Namespace** and select the relevant namespace. The resources will be displayed filtered by the specific namespace. The page shows the CPU and memory usage history for the selected namespace, for the past 15 minutes, the name of the namespace, its labels, pods, the timestamp of the creation of the namespace and its images.
2. Click the relevant namespace to display additional details. See View the namespace details.

View namespace details

To view the namespace details, click **RESOURCES** > **Namespace** and select the relevant namespace. You can also click **Workloads** > **Namespaces**, and click the relevant namespace. The page shows the CPU and memory usage history for the selected namespace for the past 15 minutes, the name of the namespace, its labels, pods, the timestamp of the creation of the namespace and its images.

Related topics

Resources

Workloads

Workloads

Click **RESOURCES** > **Workloads** to display information about resources, including deployments, replica sets, jobs, and pods in one page, filtered by the selected namespace.

The Workloads page displays:

- The CPU and memory usage of the selected namespace during the past 15 minutes.
- The list of replication controllers linked to the selected namespace. See Replication controllers.
- The list of pods linked to the selected namespace. See Pods.

Namespaces

See Namespaces.

Deployments

You can create and manage sets of replicated containers (actually, replicated pods) using Deployments. A deployment provides declarative updates for pods and replica sets (the next-generation Replication controllers). A deployment ensures that a specified number of pod “replicas” are running at any time. If there are too many, it will kill some; if there are too few, it will start more. You can deploy a containerized app, or select another namespace.

View deployments

To view deployments:

1. Click **RESOURCES** > **Workloads** > **Deployments**. The page displays the CPU and memory usage history of the selected namespace during the past 15 minutes, the names of the available deployments, their labels, the number of pods, the creation timestamp of the deployment, and its images.
2. You can:
   - Click a deployment to display its details.
   - Click and **Delete**, to delete the deployment.
   - Click and **View/edit YAML**, to view or edit a deployment.

View deployment details

To view deployment details, click **RESOURCES** > **Workloads > Deployments**, and then click the relevant deployment. The page displays the CPU and memory usage history of the selected deployment during the past 15 minutes. It also displays details about the new replica set, the old replica sets, and the events that have taken place.

Replica sets

Replica set is the next-generation replication controller. The only difference between a replica set and a replication controller right now is the selector support. Replica sets support the new set-based selector requirements as described in the labels user guide whereas a replication controller only supports equality-based selector requirements.

The replica set ensures that a specified number of pod “replicas” are running at any given time. However, a deployment is a higher-level concept that manages replica sets and provides declarative updates to pods along with a lot of other useful features. Use deployments unless you require custom update orchestration or do not require updates at all.

For more information, see http://kubernetes.io/docs/user-guide/replicasets/.

This section displays information about replica sets of the selected namespace.

View replica sets

To view replica sets:

1. Click **RESOURCES** > **Workloads** > **Replica Sets**. The page shows the name of the available replica sets for the selected namespace, its labels, pods, images and creation timestamp.
2. Click a replica set to display its details. The details page shows details about the selected replica set, the services (see Services), pods (see Pods), and events related to the replica set.
3. Click and select **Delete** to delete the replica set.

View replica set details

To view replica set details:

1. Click **RESOURCES** > **Workloads > Replica Sets**.
2. Click the relevant replica set.
   The page shows details about the selected replica set, the services (see Services), pods (see Pods), and events related to the replica set.

Replication controllers

A replication controller ensures that a specified number of pod replicas are running at any given time. It both allows for easy scaling of replicated systems and handles re-creation of a pod when the machine it is on reboots or otherwise fails.

See http://kubernetes.io/docs/user-guide/replication-controller/ for more information.

View replication controller

To view replication controllers:

1. Click **RESOURCES** > **Workloads** > **Replication Controllers** to display the available replication controllers. The page displays the following information:
   - **CPU usage history:** The CPU usage of the selected namespace during the past 15 minutes.
   - **Memory usage history:** The memory usage of the selected namespace during the past 15 minutes.
   - **<Replication controllers>:** The name, labels, pods, age, and images of the replication controllers associated with the selected namespace.
2. Click and select:
   - **View details** to view the details of a replication controller. You can also click the relevant replication controller.
   - **Scale.** See Scale the number of pods linked to the replication controller.
   - **View/edit YAML** to edit a replication controller.
   - **Delete to delete a** replication controller.

Scale the number of pods linked to the replication controller

To scale the number of pods:

1. Click **RESOURCES** > **Workloads** > **Replication Controllers**.
2. Click and then select **Scale**.
3. Enter the relevant number of pods and click **OK**.

View the replication controller details

To view replication controller details:

1. Click **RESOURCES** > **Workloads** > **Replication Controllers**.
2. Click and select **View details**, or click the relevant replication controllers. It displays the CPU and memory usage history of the selected replication controller for the past 15 minutes, the details of the selected replication controller, and the services provided by the selected replication controller.

Daemon sets

A Daemon set ensures that all (or some) nodes run a copy of a pod. As nodes are added to the cluster, pods are added to them. As nodes are removed from the cluster, those pods are garbage collected. Deleting a demon set will clean up the pods it created.

Use a daemon set if you are running clustered Kubernetes and are using static pods to run a pod on every node. Static pods are managed directly by kubelet daemon on a specific node. Static pods are always bound to one kubelet daemon and always run on the same node with it. Kubelet automatically creates a mirror pod on Kubernetes API server for each static pod.

See http://kubernetes.io/docs/admin/daemons/ for more information.

The **Daemon Sets** page provides information about the daemon sets for the selected namespace.

View daemon sets

1. Click **RESOURCES** > **Workloads** > **Daemon Sets** to display the current daemon sets.
2. Click the relevant daemon set to view its details.

Pet sets

A pet set is a controller that provides a unique identity to its pods. It provides information about the ordering of deployment and scaling.

View pet sets

1. Click **RESOURCES** > **Workloads** > **Pet Sets** to display the available pet sets.
2. Select a pet set to view its details.

Jobs

A job creates one or more pods and ensures that a specified number of them successfully terminate. As pods successfully complete, the job tracks the successful completions. When a specified number of successful completions is reached, the job itself is complete. Deleting a job will clean up the pods it created. Create a job object in order to run the pods to completion. The job object will start a new pod if the first pod fails or is deleted (for example due to a node hardware failure or a node reboot). A job can also be used to run multiple pods in parallel.

View jobs

1. Click **RESOURCES** > **Workloads** > **Jobs** to display the current jobs.
2. Select a job to view its details.

Pods

A pod is a co-located group of containers and volumes. They are scheduled onto the same host. Pods serve as units of scheduling, deployment, and horizontal scaling/replication. Pods share fate and share some resources, such as storage volumes and IP addresses.

The **Pods** page provides information about the pods that are currently running or that have been running for the past 15 minutes. You can also access details about a specific pod as well as its log. By default, pods run with unbounded CPU and memory limits. This means that any pod in the system will be able to consume as much CPU and memory on the node that executes the pod. You can restrict the resources a single pod may consume.

See http://kubernetes.io/docs/user-guide/pods/ for more information.

View pods

To view pods:

1. Click **RESOURCES > Workloads** > **Pods**. The page displays the CPU and memory usage history of the namespace the pod belongs to, namespace the pod belongs to, the name, status, number of restarts during the lifecycle of the pod, the amount of time passed since the creation of the pod, the IP address of the pod, the CPU and memory usage of the pod itself in the last 15 minutes.
2. Click  to display the log of a pod.
3. Click a pod to display its details.

View logs

1. Click **RESOURCES** > **Workloads** > **Pods**.
2. Click the relevant pod.
3. Click  in the Pod page, **View logs** in the Pod Details page, or click **View logs** in the Container area. The page displays the stdout/stderr information for the pod.

Tool	Description
	Toggles to change the size of the font used in the log.
	Toggles to change the colors of the log: white characters on a black background or black characters on a white background.
	The timestamp of the currently displayed log.
	Use the relevant buttons to navigate between logs.

View pod details

Click **RESOURCES** > **Workloads** > **Pods**, and then click the relevant pod. This displays the CPU and memory usage history of the pod in the last 15 minutes, the pod details, and the network details.

To display the log of the pod, see View log. The page also displays information about the pod containers such as the name, image, environment variables, commands, arguments, and more. To display the log of the container, see View log.

Services and discovery

Click **Services and discovery** to display information about the following:

- Services
- Ingress

Services

A service defines a set of pods and a means by which to access them, such as single stable IP address and corresponding DNS name (such as a web service or API server) that directs and loads-balances traffic to the set of pods that it covers.

See http://kubernetes.io/docs/user-guide/services/.

View services

To view services, click **RESOURCES** > **Services and Discovery** > **Services**.

The Service page displays the names of the services attached to the selected namespace, the labels assigned to the service, the IP of the related cluster, and the internal and external endpoints.

You can:

- Click and select **Delete** to delete the service.

- Click and select **View/edit YAML** to edit the service.
- Click the relevant service to display its details.

View service details

To view service details, click **RESOURCES** > **Services and Discovery** > **Services**, and then click the relevant service. The **Resource Details** page displays details about the service and the connection, as well as information about the related pods.

Ingress

An Ingress is a collection of rules that allow inbound connections to reach the cluster services.

It can be configured to give services externally-reachable URLs, load balance traffic, terminate SSL, offers name-based virtual hosting, etc. Users request ingress by POSTing the Ingress resource to the API server. An Ingress controller is responsible for fulfilling the Ingress, usually with a load balancer, though it may also configure your edge router or additional frontends to help handle the traffic in a high availability manner.

See http://kubernetes.io/docs/user-guide/ingress/.

View ingress

To view ingress, click **RESOURCES** > **Services and Discovery > Ingress**.

The **Ingress** page displays the names of the ingresses attached to the selected namespace, the labels assigned to the ingress, the IP of the related cluster, and the internal and external endpoints.

You can:

- Click and select **Delete** to delete the ingress.
- Click the relevant ingress to display its details.

View ingress details

To view ingress details, click **RESOURCES** > **Services and Discovery > Ingress**, and then click the relevant Ingress. The page displays details of the selected ingress and its related pods.

Persistent volume claims

The **Persistent Volume Claims** page displays information about the currently running persistent volumes.

A Persistent Volume (PV) is a piece of networked storage in the cluster that has been provisioned by an administrator. It is a resource in the cluster just like a node is a cluster resource. PVs are volume plugins like volumes but have a lifecycle independent of any individual pod that uses the PV. This API object captures the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system.

A Persistent Volume Claim (PVC) is a request for storage by a user. It is similar to a pod. Pods consume node resources, whereas PVCs consume PV resources. Pods can request specific levels of resources (CPU and

memory). Additionally, claims can request specific size and access modes too (for example, can be mounted once read/write or many times read-only).

See http://kubernetes.io/docs/user-guide/persistent-volumes/.

View persistent volume claims

To view persistent volume claims, click **RESOURCES** > **Persistent Volume Claims**.

The **Persistent Volume Claims** page displays the name of the persistent volume, the volume it belongs to, the labels, and the timestamp of the creation of the persistent volume.

Each suite will have at least one persistent volume but may have more depending on the suite. You can click the relevant volume to display its details.

View persistent volume claim details

To view persistent volume claim details, click **RESOURCES** > **Persistent Volume Claims**, and then click the relevant Persistent Volume Claims. This page displays detailed information about the persistent volume claim.

To see the contents of itom-vol, go to the master node (the NFS server) and run the command **cd /var/vols/itom/**. It contains the **baseinfra-<version-number>** and the **suite-install** sub-directories.

- Run the command **ls -R baseinfra-<version-number>** to display the **PrivateRegistry**.
- Run the command **ls -R suite-install** to display information about the containers that includes the configuration information to deploy the supported suites.

Configuration

The configuration files are standard pod definitions in .json or .yaml formats in a specific directory.

The following configurations are available:

- Secrets
- Config maps

Secrets

A secret stores sensitive data, such as authentication tokens, which can be made available to containers upon request.

See http://kubernetes.io/docs/user-guide/secrets/.

View secrets

To view secrets, click **RESOURCES** > **Configuration > Secrets**. This page displays the list of secrets and their age. You can click the relevant secret for displaying its details.

View secret details

To view details of a secret, click **RESOURCES** > **Configuration > Secrets** and select the relevant secret. This page displays the details of the selected secret and its data.

Config maps

The Config Map API resource holds key-value pairs of configuration data that can be consumed in pods or used to store configuration data for system components such as controllers. Config maps are similar to secrets, but designed to more conveniently support working with strings that do not contain sensitive information.

See http://kubernetes.io/docs/user-guide/configmap/.

View config maps

To view config maps, click **RESOURCES** > **Configuration** > **Config Maps**. This page displays the names of the configuration map and its labels, and the amount of time passed since the configuration map was created.

You can:

- Click on a config map and select **Delete** to delete the config map.
- Click and select **View/edit YAML** to edit the config map.
- Click the relevant config map to display its details.

View config map details

To view details of a config map, click **RESOURCES** > **Configuration > Config Maps** and select the relevant config map. This page displays the selected config map details and its related data.

Logs

This section describes the logs.

The maximum filesize of a log file is set to 10 MB, by default. If a log file increases to 10 MB, a backup of the file is created and .1 is appended to the file name, for example, asdwq-json.log.1. A maximum of 5 back up files with total file size of 50 MB are created. If the number of files reaches 5, the oldest log file is deleted and a new log file is created.

To view the logs of a particular pod:

1. Click **RESOURCES** and select the namespace from the list of namespaces.
2. Click **Workloads** > **Pods**.
3. Click the relevant pod.
4. Click **View logs** in the Pod area. The stderr/stdout of the pod you selected is displayed as shown in the following example:

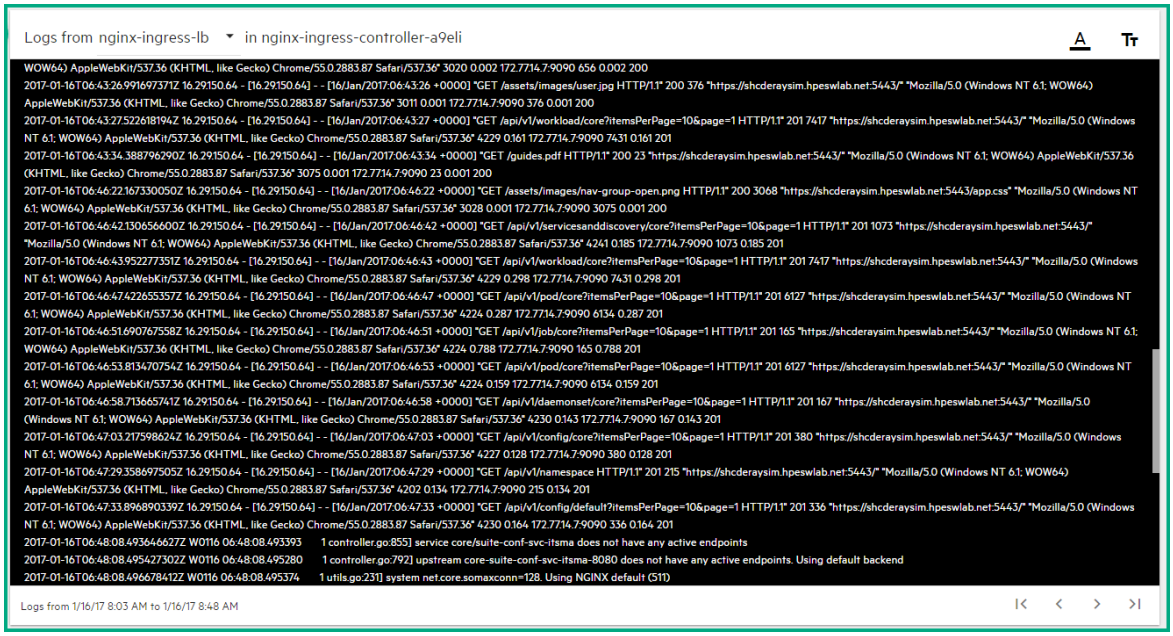

You can use the following tools on this page:

Tool	Description
Tт	Toggles to change the size of the font used in the log.
A	Toggles to change the colors of the log: white characters on black background or black characters on white background.
I< < > >I	Use the relevant buttons to navigate between logs

Logs from 10/31/16 7:23 AM to 10/31/16 7:37 AM is the timestamp of the currently displayed log.

You can also view the logs using the command prompt. Using the command prompt you can view:

- All the DCA logs
- Logs of a particular container
- Location of the log files

To view all logs of all the DCA containers:

1. Log on to the DCA host as the root user.
2. Run the following command:

```
find /opt/kubernetes/data/docker/containers -name *.log -exec grep -l
\"io.kubernetes.container.name\":\"dca- {} \; | uniq
```

Returns a list of log files of all the containers in DCA.

To view logs of a particular container:

1. Log on to the DCA host as the root user.
2. Run the following command:

```
kubectl logs <pod_name> [container_name] --namespace=<namespace>
```

Example

kubectl logs dca-analytics-deploy-2300857514-ijrd3 dca-analytics --namespace=dca1

Returns the log file of the DCA Analytics container.

Note

The ITOC logs, except debug logs, are available in the itoc-core container located at **/opt/wildfly/standalone/log/**.

To view the log file location of a particular container:

1. Log on to the DCA host as the root user.
2. Run the following command:

```
find /opt/kubernetes/data/docker/containers -name *.log -exec grep -l
\"io.kubernetes.container.name\":\"<container>\" {} \; | uniq
```

Example

```
 find /opt/kubernetes/data/docker/containers -name *.log -exec grep
-l \"io.kubernetes.container.name\":\"dca-api\" {} \; | uniq
```

Returns the location of the DCA API container. If more than one log file exists, the list of all log files is displayed.

Note

The log files of uCMDB and Operations Orchestration are located at:

- (uCMDB) /ucmdb/runtime/log/*.log

- (Operations Orchestration) /usr/local/hpe/oo/central/var/logs/*.log

If you want to view logs of a particular pod of uCMDB, run the following command:

```
kubectl --namespace=<namespace> get pods | grep ucmdb-deploy
kubectl --namespace=<namespace> exec -it <pod-name> /bin/sh
```

If you want to view logs of a particular pod of Operations Orchestration, run the following command:

```
kubectl --namespace=<namespace> get pods | grep dca-suites-oo-central
kubectl --namespace=<namespace> exec -it <pod-name> /bin/sh
```

Related topic

Manual verification commands

Export suite export logs

You can export suite logs by gathering them from the suite persistent volume(s) and zipping them up.

To export suite logs, click **SUITE > Management** > **Export Logs**. The log package is downloaded to your local disk.

⚠ You can export logs after you have installed a suite.

Security

This section is intended for ITOM CDF implementers and system administrators who need to implement the ITOM CDF environment securely.

This section includes the following information:

- Secure implementation and deployment
- ITOM CDF security parameters
- Installation security
- Network and communication
- Authorization
- Data integrity
- Encryption
- Log and trace

- Network and communication security

Secure implementation and deployment

This section provides information on implementing and deploying the ITOM CDF securely.

Technical system landscape

The ITOM CDF is a container that integrates with other suites. The ITOM Platform platform is written in Java, JavaScript, and Go.

For more information about typical deployment schemes and options, see Deployment scenarios.

Security in ITOM CDF configurations

The ITOM CDF configurations may be deployed in the following three implementations.

- Single mode.
- Distributed mode 1 (one master node and multiple worker nodes).

All of these implementations share the same basic out-of-the-box security configuration options.

1. In an out-of-the-box default installation, the Transport Layer Security/Secure Socket Layer (TLS/SSL) security is enabled between the browser and the ITOM CDF server by default.
2. In an out-of-the-box default installation, the ITOM CDF requires users to enter username and password credentials to gain access to the application.

External authentication

With additional configuration, it is possible to supplement or replace the default authentication & authorization provider for the ITOM CDF by using a variety of industry-standard protocols and tools such as LDAP and Single Sign-On.

Common security considerations

The ITOM CDF can only be deployed on supported operating systems.

It is recommended to follow vendor-provided best practices and security hardening guides for each of the third-party components used in support of your ITOM CDF deployment, which includes Docker, Kubernetes, Vault and Nginx, NFS. Below are some resources that can serve as a starting point for researching these recommended security considerations:

- Docker security tips: https://www.docker.com/docker-security
- Kubernetes security tips: http://kubernetes.io/docs/troubleshooting/
- Vault Security Tips: https://www.hashicorp.com/security.html
- Nginx Security Tips: http://nginx.org/en/security_advisories.html
- NFS security tips: http://www.cert.org/historical/advisories/

ITOM CDF security parameters

This section contains reference to some of the ITOM CDF parameters that are relevant to security.

Secure file storage

The ITOM CDF allows users to upload files (suite installation binary) to the ITOM Container Deployment Foundation Server. All files uploaded to the server must be validated, since they can contain viruses, malicious code, or Trojans.

As a result, it is strongly recommended to implement proper antivirus protection for the file storage.

Installation security

The following aspects of installation security are available:

- Supported operating systems
- Database security recommendations
- Application server security recommendations

Supported operating systems

See Support matrix.

1. Harden SSH on OS.
2. On each node, the SSH server is configured with weak cipher and weak KexAlgorithms by default.
3. Set the values of **KexAlgorithms**, **Ciphers** and **MACs** in file: /etc/ssh/sshd_config as follows:

- - **KexAlgorithms** ecdh-sha2-nistp521,ecdh-sha2-nistp384,ecdh-sha2-nistp256
  - **Ciphers** aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr
  - **MACs** hmac-sha2-256

Database security recommendations

See http://www.openscg.com/postgresql-security-guidelines/ for information about PostgreSQL database security solutions.

Application server security recommendations

- Always change default passwords.
- Always use the minimal possible permissions when installing and running ITOM Platform.

Action	Prerequisites
Installing or running the ITOM CDF	You must install and run root permissions using the **sudo** command.

Network and communication

This section provides information on network and communication security.

Secure topology

The ITOM CDF is designed to be part of a secure architecture and can meet the challenge of dealing with security threats to which it could potentially be exposed.

To securely deploy ITOM CDF, it is recommended that you use the TLS/SSL communication protocol.

Replace the certificate

Users can replace the certificate and private key of **Ingress Service** with a customized certificate and private key by performing the following steps:

1. Generate a certificate and private key for the **host name**, where the Ingress service is running. Save this on the master node.
2. On master node, delete a secret by running the following command:

   ```
   kubectl delete secret nginx-default-secret -n core
   ```

3. On the master node, recreate the secret with a new certificate and private key. Run the following commands:

   ```
   echo "
   apiVersion: v1
   kind: Secret
   metadata:
   name: nginx-default-secret
   namespace: core
   data:
   tls.crt: `base64 <certificate file name with absolute path> |tr -d \"\n\"`
   tls.key: `base64 <private key file name with absolute path> |tr -d \"\n\"`
   " | kubectl create -f -
   ```

4. On the master node, delete and recreate the ingress service by running the following commands:
    - `kubectl delete -f ${K8S_HOME}/objectdefs/nginx-ingress.yaml`
    - `kubectl create -f ${K8S_HOME}/objectdefs/nginx-ingress.yaml`

Authorization

This section provides information related to user authorization in ITOM CDF.

Authorization model

Access to the ITOM CDF resources is authorized based on the user's following settings:

- User name
- Session and Inactivity timer timeouts

FAQ

Question

Can the ITOM CDF inherit users' information and authorization profiles from an external repository, such as LDAP?

Answer

No.

Data integrity

The database server is used as a simple data store and is responsible for all persistent storage. While the database contains definitions describing business logic, no processing is actually performed in this tier, other than create, read, update, and delete (CRUD) operations in response to requests from ITOM CDF. Referential integrity is enforced by the application, thereby protecting transactions. In addition, the database captures a complete audit log of all changes to data.

The data backup procedure is also an integral part of data integrity and while the ITOM Platform does not provide native backup capabilities, the following guidelines should be considered:

- Database backup is especially important before critical actions such as upgrades.
- Backup files should be stored properly according to the industry best practices to avoid unauthorized access.
- Since database backup can be a resource intensive process, it is strongly recommended to avoid running backups during peak demand times.

Encryption

This section provides information on data encryption in ITOM CDF.

TLS/SSL data transmission

The ITOM CDF is configured to use TLS/SSL to transmit data between the server and browsers.

You can change the default value of SSL CIPHER by performing the following steps:

1. On the master node, change the ssl-ciphers value in file $K8S_HOME/objectdefs/nginx-ingress.yaml file.
2. Recreate the ingress container by running the following commands:
   ```
   kubectl delete -f $K8S_HOME/objectdefs/nginx-ingress.yaml
   kubectl create -f $K8S_HOME/objectdefs/nginx-ingress.yaml
   ```

Encryption of stored database fields

The ITOM CDF uses proprietary algorithms when encrypting data stored in the database and uses Identity Manager (IDM) to manage user passwords.

Log and trace

This section provides information related to logs.

Log and trace model

Recommendations:

- Pay attention to the log level and do not leave tracing or debug parameters enabled unnecessarily.
- Pay attention to log rotation/switching.

FAQ

Question

Are exceptions required to be added to the firewall policy?

Answer

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

Network and communication security

Add the iptables rule as listed below.

⚠ Apart from the listed ports, all other ports must be blocked at the localhost level.

Required ports	Service	Add rules on Server	Direction	Short description
111	NFS	NFS server	Nodes ->NFS Server	NFS server port access by all nodes
2049	NFS	NFS server	Nodes ->NFS Server	NFS server port access by all nodes
2380	Etcd	Master Node	Master<-> Master	Etcd service port for etcd cluster communication
4001	Etcd	Master Node	Nodes -> Master	Etcd service port for connection from client
4194	Kubernetes	All Nodes in Cluster	Localhost only	Cadvisor for local kubelet
5000	Private Registry	All Nodes in Cluster	Localhost only	Registry port for localhost
5443	MngPortal	Ingress Node	All -> Ingress Node	The port exposed on ingress node. all clients could access this port
8200	Vault	Master Node	Nodes->Master	Vault port for client connection
8443	kubernetes	Master Node	Nodes->Master	API server port for client connection
10250	Kubernetes	All Nodes in Cluster	Nodes->Nodes	Kubernetes port for internal communication
10251	Kubernetes	—	Nodes->Nodes	Kubernetes port for internal communication
10252	Kubernetes	—	Nodes->Nodes	Kubernetes port for internal communication
10255	Kubernetes	—	Nodes->Nodes	Kubernetes port for internal communication
20048	NFS	NFS server	Nodes ->NFS Server	NFS server port access by all nodes

To add iptable rules to port 8443 on the master node, perform the following:

The master node is installed on 10.10.10.10.

```
iptables -I INPUT 1 -p tcp -m tcp -s 0.0.0.0/0 --dport 8443 -j DROP
iptables -I INPUT 1 -p tcp -s 127.0.0.1 --dport 8443 -j ACCEPT
iptables -I INPUT 1 -p tcp -s 10.10.10.10 --dport 8443 -j ACCEPT
```

Related topic

Manual verification commands

Restart the ITOM CDF

Perform the following steps to stop the ITOM CDF:

- On each master node, run the following command:

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

- On each worker node, run the following command:

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

Perform the following steps to start the ITOM CDF:

- On each master node:

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

- On each worker node:

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

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email.

Add the following information in the subject line: Feedback on Data Center Automation 2017.05 - Premium

Just add your feedback to the email and send your feedback to docs.feedback@hpe.com.

We appreciate your feedback.