Dataplane Service Platform Administration

Date of Publish: 2018-05-18

http://docs.hortonworks.com

Contents

Administering DPS

Managing Clusters in DPS

From DPS Platform, you can manage the clusters you use with the services that plug into DPS. You can register and enable clusters, and view cluster details.

Register a cluster in DPS Platform

You must register Ambari-managed clusters with DPS Platform. You must register clusters with DPS before you can view or manage data on the clusters, or before you can enable the clusters for use with any DPS service.

About this task

Caution: After you register a cluster in DPS Platform, do not change the cluster name in Ambari. A cluster name change in Ambari does not currently propagate to DPS Platform. Therefore, jobs associated with that cluster name will fail, and datasets associated with that cluster name are not discoverable in the DPS services.

Before you begin

- You must be logged in using the DataPlane Admin (DPS Admin) role with valid LDAP credentials, to perform this task.
- If your clusters are configured for Knox Gateway and using a self-signed SSL certificate, you must have uploaded your certificate to DPS prior to starting this task (or you can disable certificate validation on registration). Refer to “Upload a certificate to DPS” for more information.
- Clusters must have been created using Apache Ambari.

  Clusters that are not managed by Ambari cannot be accessed by DPS Platform or associated services.
- All clusters must meet the requirements identified in DPS *Getting Started*.

Procedure

1. In the DPS Platform navigation pane, click the

   

   (Clusters) icon.
2. Click **Add Cluster**.
3. (Optional) Check **Ambari and Cluster Services Behind Knox Gateway** only if you are using Knox Gateway to proxy Ambari and DPS services.
4. (Optional) Check **Validate the SSL certificate and only allow trusted connections** only if your cluster is SSL-enabled.

   DPS validates the existing SSL certificate.

   If you plan to validate the SSL certificate but the certificate is self-signed, be sure to upload the certificate to DPS. Refer to “Upload a certificate to DPS” for more information.
5. Enter the URL of the Ambari host that manages the cluster.
   - If the cluster is behind a Knox Gateway, use the following format:

     ```
     https://knox_host_fqdn:knox_port/gateway_name/dp-proxy/ambari
     ```

     By default, the gateway_name is gateway.

- If the cluster is not using Knox Gateway, use the following format:

```
https://ambari_host_fqdn:ambari_port/
```

  You can use HTTP or HTTPS. HTTPS is used if Wire Encryption is enabled for Ambari.

You can also enter the IP address instead of the FQDN, but the FQDN is recommended.

The host name must either be in a DNS server or configured in the etc/hosts file.

Important:

DPS Platform host must be able to resolve and reach the Ambari URL, whether you are using the FQDN or the IP address. That is, you should be able to use curl or wget to access the Ambari URL from the DPS Platform host. If this requirement is not met, cluster registration fails.

If host names are resolved from /etc/hosts, you should explicitly register the cluster host names on the DPS container before the cluster is registered with DPS.

6. Click **Go**.

   The cluster name, IP address, and enabled HDP services display, similar to the following:

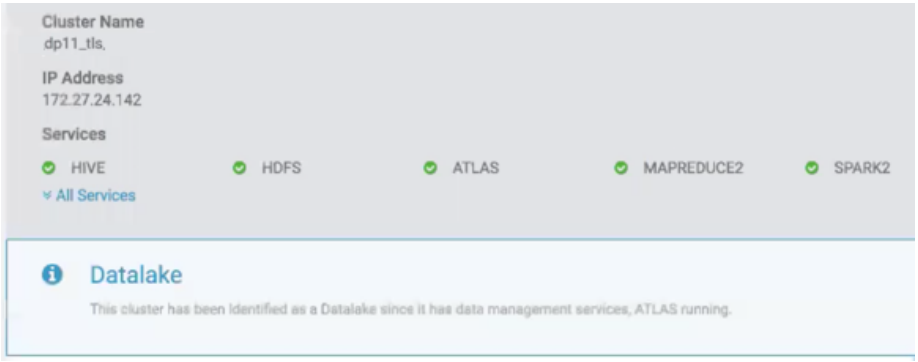

7. If the cluster is designated a data lake, add the following cluster details:

   Cluster Location
   Data Center
   Tags (optional)
   Description (optional)

What to do next

You can now enable the registered cluster for the DPS App you want to use it with.

Related Concepts

Configure Knox Gateway for DataPlane and HDP

Upload a certificate to DPS

Related Information

DPS Getting Started

LINK TO: Uploading a certificate to DPS, Enable a cluster for a service

Enable a cluster for a service

After registering clusters with DPS, you must enable each cluster with one or more DPS services. Each DPS service has specific configuration requirements that a cluster must meet before it can be used with the service.

About this task

When you enable a DPS service, a check is run to determine if the required service engine or agent is installed on any clusters. If the engine is installed but some configuration is still required, the cluster displays on the Services page with the button Enable. If the cluster meets all requirements for the service it is automatically enabled, and the enabled cluster can only be viewed on the Services page by selecting the Show All Clusters action for the service.

Before you begin

The DataPlane Admin role is required to perform this task.

Clusters must be managed by Apache Ambari and registered with DPS Platform.

The services you are associating with the clusters must already be enabled in DPS Platform.

Procedure

1. Click the
   
   (Services) icon in the DPS Platform navigation pane.
   The Services page displays. Services listed in the table have been enabled. Services identified by a tile icon are available to be enabled.
2. Click the row for a service.
   A list displays of any clusters that have the required service engine or agent installed but have not yet been configured for use with the service.

   Tip: If no clusters display for the service, verify that the clusters you expect to see have been registered with DPS Platform, and that the proper service engine or agent has been installed on the clusters.
3. Click **Enable** for the cluster you want to use with the service.

   A check runs to determine what configuration is required on the cluster for the service you selected. For example, a required service such as Apache Ranger might be installed on the clusters but not enabled in Apache Ambari.

   The Manual Install page displays, indicating what you need to configure on the cluster to make the cluster usable by the service.
4. Perform the actions stated on the Manual Install page.
   The required actions often involve enabling a service from Ambari. For example:

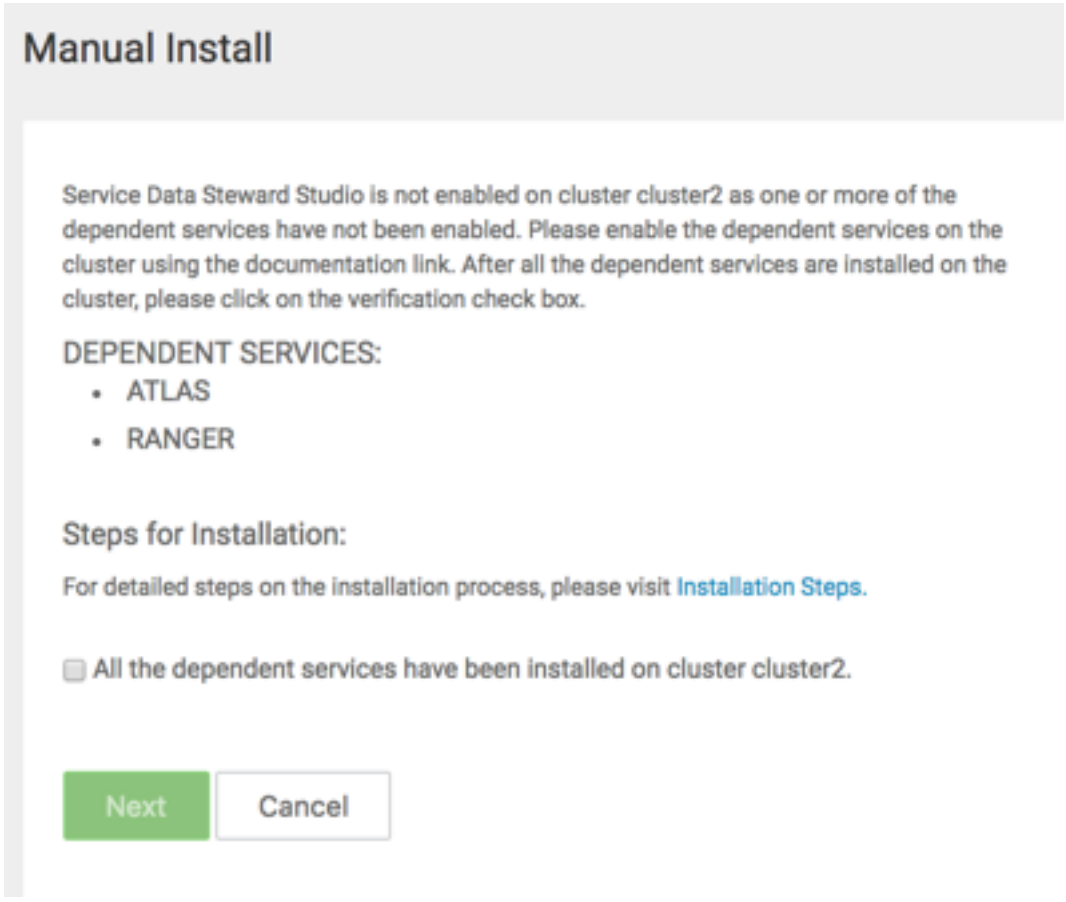

   You should complete the required actions before proceeding to the next step.

5. Click **All the dependent services have been installed...**, and then click **Next**.
   Another configuration check is run and if all requirements are met, a verification message displays. If the check fails, a message displays identifying the configuration tasks still to be completed.
6. Click the name of the service for which you enabled the cluster, and then enable **Show All Clusters**.
   The new cluster displays in the list on the Clusters page:

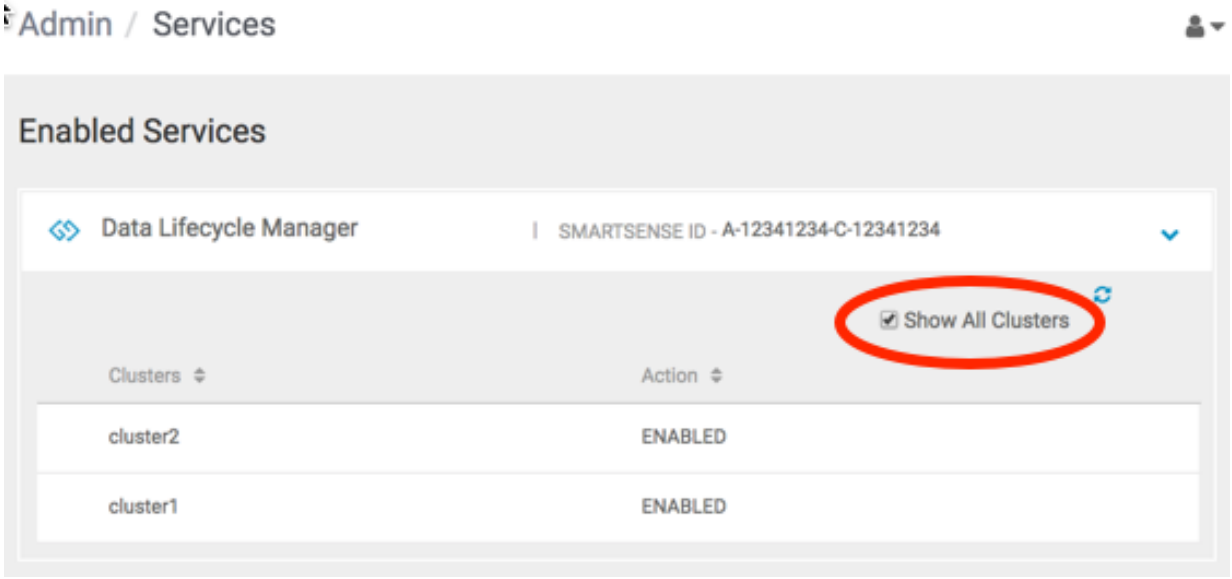

Edit a cluster in DPS Platform

You can modify the Cluster Location, Data Center name, Tags, or Description for any cluster registered in DPS Platform.

Before you begin

The DataPlane Admin role is required to perform this task.

About this task

Important: After you register a cluster in DPS Platform, do not change the cluster name in Ambari. A cluster name change in Ambari currently does not propagate to DPS Platform. Therefore, replication jobs in Data Lifecycle Manager associated with that cluster will fail, and information associated with that cluster will not be available for monitoring in Data Steward Studio.

Procedure

1. Click the
   
   (Clusters) icon in the DPS Platform navigation pane.
2. Optional: Enter a cluster name in the search field and press **Enter**.
   You can only search by cluster name. You can search by partial or full names.
3. In the cluster list, locate the row for the cluster you want to edit.
4. At the end of the row, click the

   (Actions) icon and then click **Edit**.
   The Edit Cluster page displays.
5. Modify the cluster details.
6. Click **Update**.

The Clusters page displays a list with the updated cluster.

View cluster details in DPS Platform

DPS Platform provides two levels of detail about each cluster. You can view general information about a cluster such as status, location, HDP version, number of nodes, and so forth, from the Clusters page. You can also view information such as the status of DataNodes and NodeManagers, heap size, disk space, and so forth, from the cluster Details page.

Before you begin

The DataPlane Admin role is required to perform this task.

Procedure

1. Click the
   
   (Clusters) icon in the DPS Platform navigation pane.
2. In the cluster list, locate the row for the cluster you want to edit and click the cluster name.
   The Details page displays more information about the selected cluster.
3. Refresh the values on the Details page by clicking the
   
   (Refresh) icon.
   Updated data is pulled from Ambari.
4. To view or edit cluster information in Ambari, click the

   (Actions) icon and then click **Go to Ambari**.
   A new browser tab opens to the login page for the Ambari host that manages the cluster.

Add host entries to the DPS environment

If you are using hosts that are not publicly addressable from a DNS server, you must add IP address and host name mapping to the /etc/hosts file of each associated *DPS container* in your DP instance. DPS provides a utility command to help with this to ensure that all hosts to be used with DPS are addressable.

About this task

You must use the specific method and format identified in this task to make the mapping information usable by DPS Platform.

Perform this task only after DPS Platform is installed and running.

> **Attention:**
>
> This task is required only for /etc/hosts handling. This setup is not required if Ambari, Knox, or hosts in your cluster are accessible via DNS.

Procedure

1. Log in as the root user to a terminal on the DPS host.
2. Type the following command to add the host to the /etc/hosts file of the container:

```
dpdeploy.sh utils add-host <ip> <host>
```

3. Repeat this procedure on each cluster registered with DPS Platform.

Upload a certificate to DPS

If you are registering clusters either through Ambari or a Knox Gateway proxy that are using a self-signed SSL certificate, you can upload the certificate to DPS. This enables DPS to validate the certificate during cluster registration. Alternatively, you can disable certificate validation during cluster registration.

Procedure

1. In the Navigation pane, click **Settings**.
2. Click **Upload** and complete the form.
   - Certificate Name: Create a name for the certificate.
   - Certificate File: Browse to the location of the certificate, such as a .pem file.

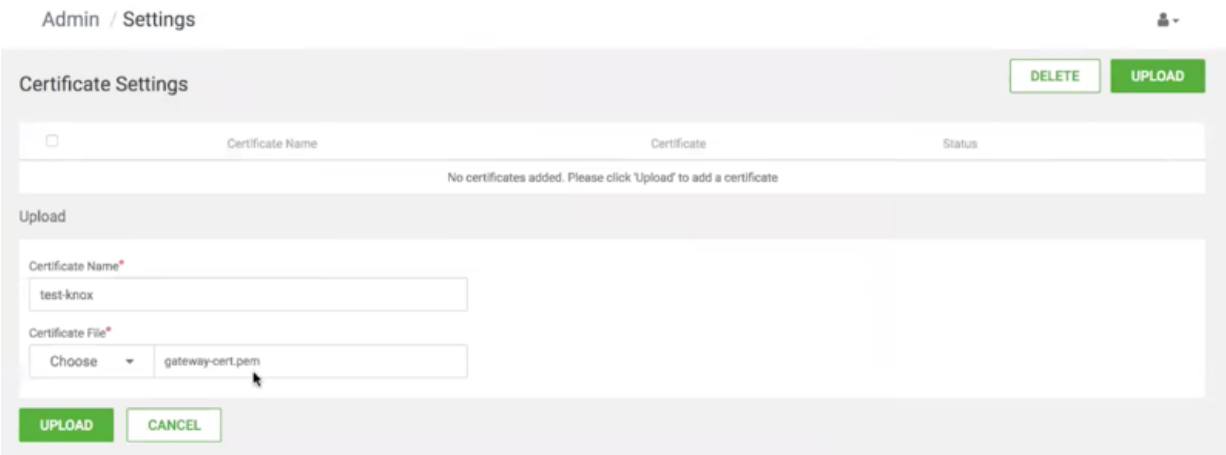

3. Click **Upload**.

Managing Users and Groups

From the DPS Platform UI, you can add and edit users and groups for all services associated with DPS. Users and groups must exist in the corporate LDAP directory before you can add them in DPS Platform.

Add a user or group

When you install Hortonworks DPS, the DataPlane Admin role is created. This role has access to DPS Platform and all DPS services. You should add additional users with permissions limited to the service or services you want the user to access.

About this task

You can also add groups, which enable you to more easily manage users. You might configure groups for all users who can perform specific tasks, such as creating replications, creating datasets, and so forth.

> **Tip:** User-level assignments override group-level assignments.

Before you begin

User and group accounts must already exist within your corporate LDAP directory.

The DataPlane Admin role is required to perform this task.

Procedure

1. Click the
   
   (Users) icon in the DPS Platform navigation pane.
2. Click **Add User**.
3. In **Usernames**, enter the name of a user from your corporate LDAP directory, and then click the name when it pops up.

   Tip:

   You must click the name of the user when it displays and ensure it appears in the Username field on a dark background.

   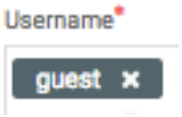

   If the name appears on a white background, it means the name is not recognized and the action fails.
4. Select one or more roles to assign to the user.
   - DataPlane Admin (or DPS Admin)

     Can perform all actions in DPS Platform, and can access and perform all actions in the UI of enabled services.
   - Infra Admin

     Can perform all actions in the Data Lifecycle Manage (DLM) service UI, and can manage DLM-enabled clusters in DPS Platform.
   - Data Steward

     Can perform all actions in the Data Steward Studio (DSS) service UI, and can manage DSS-enabled clusters in DPS Platform.
5. Click **Save**.
   The new user displays in the list on the Users page.
6. Click the **Groups** tab, and then click **Add Group**.
7. In **Group Names**, enter the name of a group from your corporate LDAP directory, and then click the name when it pops up.

   Tip:

   You must click the name of the group when it displays and ensure it appears in the Group Name field on a dark background.

   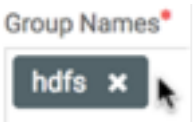

   If the name appears on a white background, it means the name is not recognized and the action fails.
8. Select the roles to assign to the group.
9. Click **Save**.
   The new group displays in the list on the Groups page.

Edit a user or group

You can edit any user or group that has already been added to DPS Platform.

Before you begin

The DataPlane Admin role is required to perform this task.

About this task

User-level assignments override group-level assignments.

Procedure

1. Click the
   
   (Users) icon in the DPS Platform navigation pane.
   The Users page displays a list of existing users, their roles, and status.
2. Locate the user you want to edit by browsing the user list or entering a user name in the search field.
   You can only search by user name. You can search by partial or full names.
3. Select the user to edit by doing one of the following:
   - Click a user name in the list of users.
   - Click

     (Actions icon) and Edit.

   The user's information displays in a slide-out panel.
4. Change the user status, or add or delete roles.
5. Click **Save**.
   The modifications are shown in the Users list.
6. Click the **Groups** tab.
7. Locate the group you want to edit by browsing the group list or entering a group name in the search field.
   You can only search by group name. You can search by partial or full names.
8. Select the group to edit by doing one of the following:
   - Click a group name in the list of groups.
   - Click

     (Actions icon) and Edit.

   The group's information displays in a slide-out panel.
9. Add or delete roles assigned to the group.
10. Click **Save**.
    The modifications are shown in the Groups list.

Edit LDAP settings

If the corporate LDAP administrator Bind DN and password are changed, such as for security purposes or policy requirements, or the LDAP server URL changes, such as due to hardware issues, the DataPlane Admin can modify those settings in DPS so that users can continue to log in to DPS.

Before you begin

You must be logged in as the DataPlane Admin to complete this task.

Procedure

1. In the DPS Platform UI, click the
   
   (Users) icon.
2. To the right of the LDAP URL, click the

   (Actions) icon, and then click **Edit**.

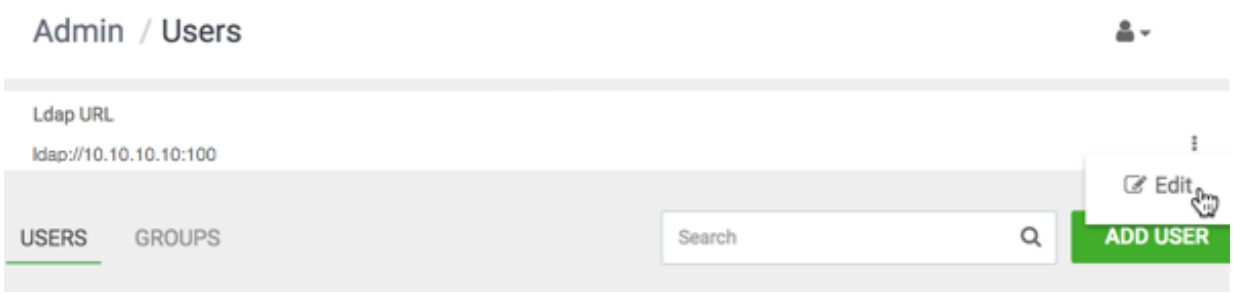

3. In the Edit LDAP page, you can modify the following settings:
   - URL
   - Administrator Bind DN
   - Administrator Password
   - Upload a new certificate
4. Click **Save**.

Managing DPS Service Apps

DPS supports multiple services. You can install and enable any combination of supported services.

Enable services

You must enable, through DPS Platform, any DPS service you want to use. Before enabling a service, you must have properly installed and configured the service app on the DPS host, as well as the management pack for the service engine or agent on each cluster.

Before you begin

The DataPlane Admin (DPS Admin) role is required to perform this task.

Before enabling a service, you must have properly installed and configured the service and the associated service engine.

You must have a SmartSense ID available. You can retrieve the SmartSense ID from the Hortonworks Support Portal under the Tools tab.

Procedure

1. Click the
   
   (Services) icon in the DPS Platform navigation pane.
   The Services page displays. Services listed in the table have been enabled. Services identified by a tile icon have not yet been enabled.
2. Move the cursor over the tile for the service you want to enable.
   - If an **Install** button displays, you must install the service before you can enable it.

     Clicking the button opens the installation documentation for the service. Install the service, then return to DPS Platform to enable the service.
   - If an **Enable** button displays, click it.

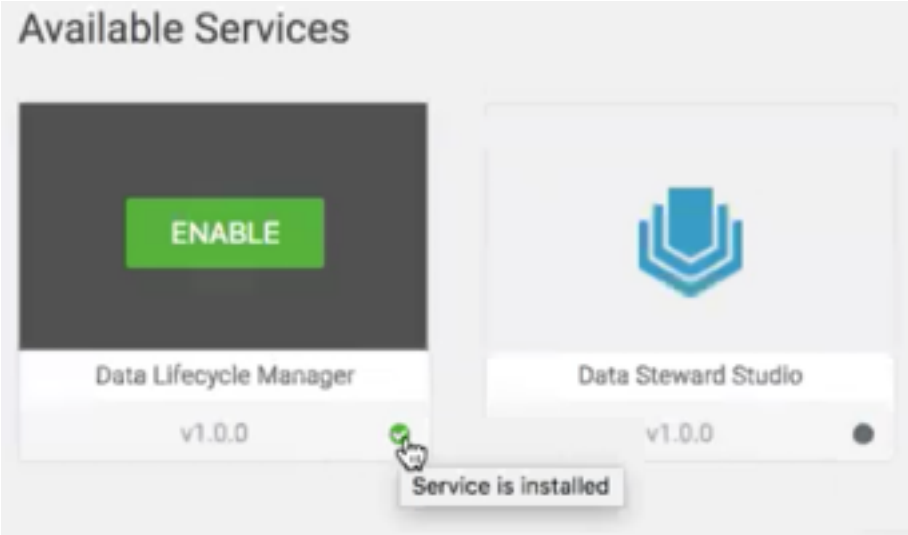

3. Enter the SmartSense ID and click **Verify**.

   The ID is case-sensitive.
4. Click **Next**.

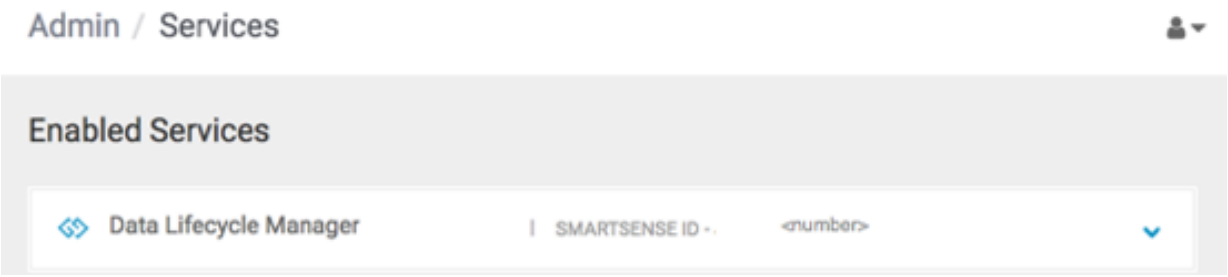

The enabled service displays in the Enabled Services list.

Navigating between services

You can access any service for which you have been assigned the proper role. The DPS Admin has access to all DPS services.

Before you begin

The DPS Admin must have assigned you the required role for any service you want to access.

Procedure

1. Click the

   

   (Service Navigation) icon in the upper left corner of any page in DPS Platform.
2. Click the tile for the service you want.

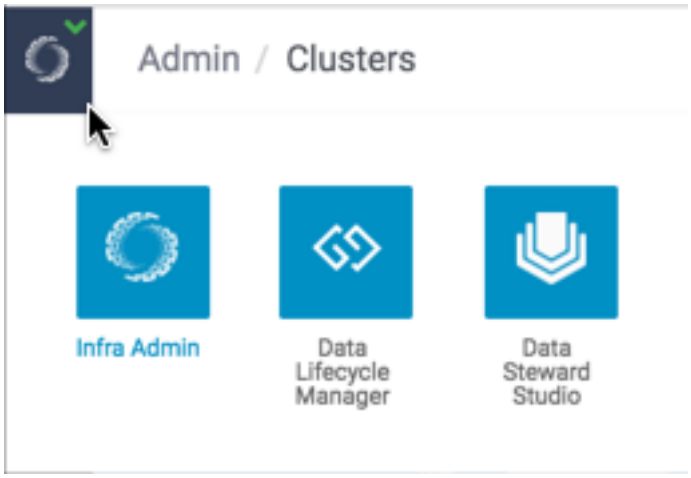

If the service you want to access is not displayed, either the service is not enabled or you have not been assigned the role required to access the service. The DPS Admin can enable services and assign roles.

Disable and enable data telemetry

As part of the installation process, data collection using cookies and other telemetry mechanisms is turned on by default. This topic describes how to disable tracking and telemetry, and enable or check the status of telemetry.

About this task

Use the dpdeploy help command to learn more about dpdeploy command options.

Procedure

1. To disable tracking, in a terminal enter the following command:

```
./dpdeploy.sh utils disable-config dps.ga.tracking.enabled
```

You see the following output:

```
UPDATE 1
Config value for key dps.ga.tracking.enabled was updated successfully with
 value false
```

2. To enable tracking, in a terminal enter the following command:

```
./dpdeploy.sh utils enable-config dps.ga.tracking.enabled
```

You see the following output:

```
UPDATE 1
Config value for key dps.ga.tracking.enabled was updated successfully with
 value true
```

3. To enable or disable getting tracking status, in a terminal enter the following command:

```
./dpdeploy.sh utils get-config dps.ga.tracking.enabled
```

You see the following output if enabling tracking status:

```
dps.ga.tracking.enabled: Enabled
```

You see the following output if disabling tracking status:

```
dps.ga.tracking.enabled: Disabled
```

DPS Reference Information

Roles required to work with DPS Services

Access to DPS Services and functionality within those services requires a different role or set of roles for each service.

To perform actions in DPS Platform and associated services, you must be logged in as DPS administrator (DPS Admin), infrastructure administrator (Infra Admin), or data steward administrator (Data Steward). In addition, to perform actions in Apache Ambari that impact DPS (such as creating clusters, changing configuration settings for services, and so forth), you must be an Ambari administrator or a cluster administrator.

Other roles might be required during installation, depending on your configuration. See the installation instructions for roles required during installation.

DataPlane Admin role

The DataPlane Admin (DPS Admin) has access to DPS Platform and administrative permissions to perform all actions in DPS Platform. A DPS Admin role is created during installation, so you can initially log in to DPS Platform.

The DPS Admin has the following capabilities and restrictions:

- Can access DPS Platform and perform all actions in DPS Platform related to clusters, users, and enabling services.
- Can access all services enabled with DPS Platform, and perform the same actions as each administrator role assigned to the enabled services, such as Infra Admin, Data Steward, and so forth.

Roles required for installation and troubleshooting

You need the Ambari Admin or Cluster Admin roles to install Hortonworks DPS, add clusters to Ambari, troubleshoot cluster issues, and so forth.

See Apache Ambari Administration for further details about these roles.

Ambari Admin role

The Ambari Admin has full control over all aspects of Ambari. It includes all capabilities of the Ambari Cluster Admin role, plus additional capabilities. This role can be used to troubleshoot issues with your clusters and to restart or reconfigure the Beacon engine that is associated with DLM.

The Ambari Admin role has the following capabilities and restrictions:

- Can install HDP by using the Ambari installation wizard
- Can install the DLM Engine (Beacon) management pack and configure the DLM Engine
- Can start and stop the DLM Engine and troubleshoot cluster problems
- Cannot access DPS Platform or any enabled service in DPS Platform

Cluster Admin role

The Cluster Admin is an Ambari role that has control over a cluster, its hosts, and services. This role can be used to troubleshoot issues with your clusters and to restart or reconfigure the Beacon engine that is associated with DLM.

The Cluster Admin role has the following capabilities and restrictions:

- Can create clusters to be registered with DPS Platform
- Can start and stop the DLM Engine and troubleshoot cluster problems in Ambari
- Cannot access DPS Platform or any enabled service, such as DLM, in DPS Platform

Required roles by task

The following tables indicate the roles required to perform various tasks in DPS Platform and the associated services.

DPS Platform tasks and required roles

The following table shows tasks you can perform that are related to DPS Platform and the roles required to perform the tasks. Except for installation of DPS, and unless otherwise noted, all tasks that require the DataPlane Admin role are performed within the DPS Platform UI.

Task	DataPlane Admin	Infra Admin	Data Steward	Ambari Admin	Cluster Admin
Install HDP using Ambari				X	
Install DPS Docker image (CLI)	X				
Install service (DLM, DSS) Docker image	X				
Configure LDAP	X				X
Enable DPS services (DLM, DSS)	X				
Manage DPS users	X				
Manage DPS clusters (register, delete, etc.)	X				
Monitor clusters in DPS	X	X	X		
Create a cluster				X	X

Troubleshooting DPS

Endpoint not accessible

DLM does not support updating any cluster endpoints (HDFS, Hive, Ranger, or DLM Engine).

If an endpoint must be modified, contact Hortonworks Support for assistance.

Logging in using the DataPlane local admin role

The local admin role allows you to perform administrative activities and troubleshoot problems when access through LDAP and Knox is not available. The local admin is also the role you use to log in to DataPlane the first time, before LDAP is configured in DataPlane for SSO.

About this task

When you log in as the local DataPlane Admin, you bypass Knox.

For login, the default username is “admin”. The password you use to log in is set during the installation process.

Procedure

Log in by appending /sign-in to the DataPlane login URL, for example:

```
http://dataplane-host-name/sign-in
```

Ranger UI does not display deny policy items

When a policy with deny conditions is created on Ranger-Admin for a replication relationship, the Policy Details page in Ranger should show the deny policy items also. However, the deny policy items do not display on the Ranger admin Policy Details page.

Procedure

1. Enable deny conditions for policies from **Ambari**>**Ranger**>**Configs**>**Advanced**>**Custom ranger-admin-site.**
2. Add ranger.servicedef.enableDenyAndExceptionsInPolicies=true.
3. Restart target ranger-admin.