

Ambari 2

Configuring Specific Ambari Views

Date of Publish: 2018-04-30

<http://docs.hortonworks.com>

Contents

Configuring Specific Views.....	3
Configuring Capacity Scheduler View.....	3
Create a Capacity Scheduler View instance.....	3
Set YARN Queue Manager View user permissions.....	6
Configure your cluster for Files View.....	10
Create and configure a Files View instance.....	11
Set up Kerberos for Files View.....	12
Configure local option for Files View.....	12
Configure custom option for Files View.....	13
Configuring Pig View.....	14
Configure your cluster for Pig View.....	14
Set up HDFS proxy user for Pig View.....	14
Set up WebHCat proxy user for Pig View.....	15
Set up WebHCat proxy user for the Ambari Server daemon account.....	15
Set up HDFS user directory.....	16
Create and configure a Pig View instance.....	16
Get correct configuration values for Pig View in a manually-deployed cluster.....	18
Set up user permissions for Pig View.....	20
Set up Kerberos for Pig View.....	21
Configuring SmartSense View.....	22
Configure your cluster for SmartSense View.....	22
Create a SmartSense View instance.....	22
Configure Workflow Manager View.....	24
Configure your cluster for Workflow Manager View.....	24
Set up HDFS proxy user.....	24
Set up HDFS user directory for Workflow Manager View.....	25
Set up Kerberos for Workflow Manager View.....	26
Set up proxy user for Oozie.....	26
Create and configure a Workflow Manager View instance.....	27

Configuring Specific Views

Ambari configures and deploys most views automatically, for each service added to a cluster.

Depending on your environment, each view (and associated service) may require some additional configuration or troubleshooting.

Configuring Capacity Scheduler View

If you have deployed your cluster manually, or if you need to re-configure the Ambari-created YARN Queue Manager View, you can use the following topics to create and configure a view instance.

Create a Capacity Scheduler View instance

When you deploy a cluster using Ambari, a Capacity Scheduler View instance is automatically created.

About this task

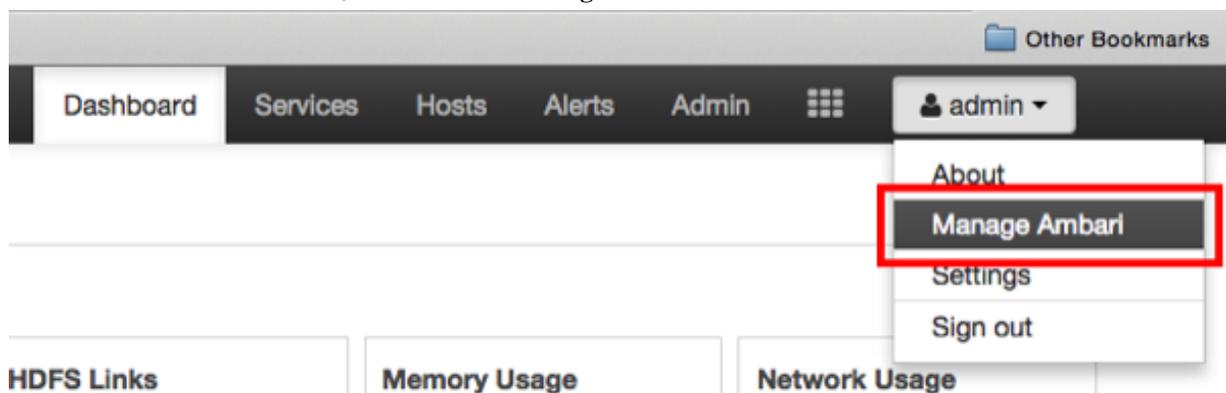
If you do not need to reconfigure the Ambari-created cluster, proceed to use the YARN Queue Manager View. If you have deployed your cluster manually, or if you need to re-configure the Ambari-created YARN Queue Manager View, you can use the following topics to create and configure a view instance. To set up a Capacity Scheduler / YARN Queue Manager view instance:

Before you begin

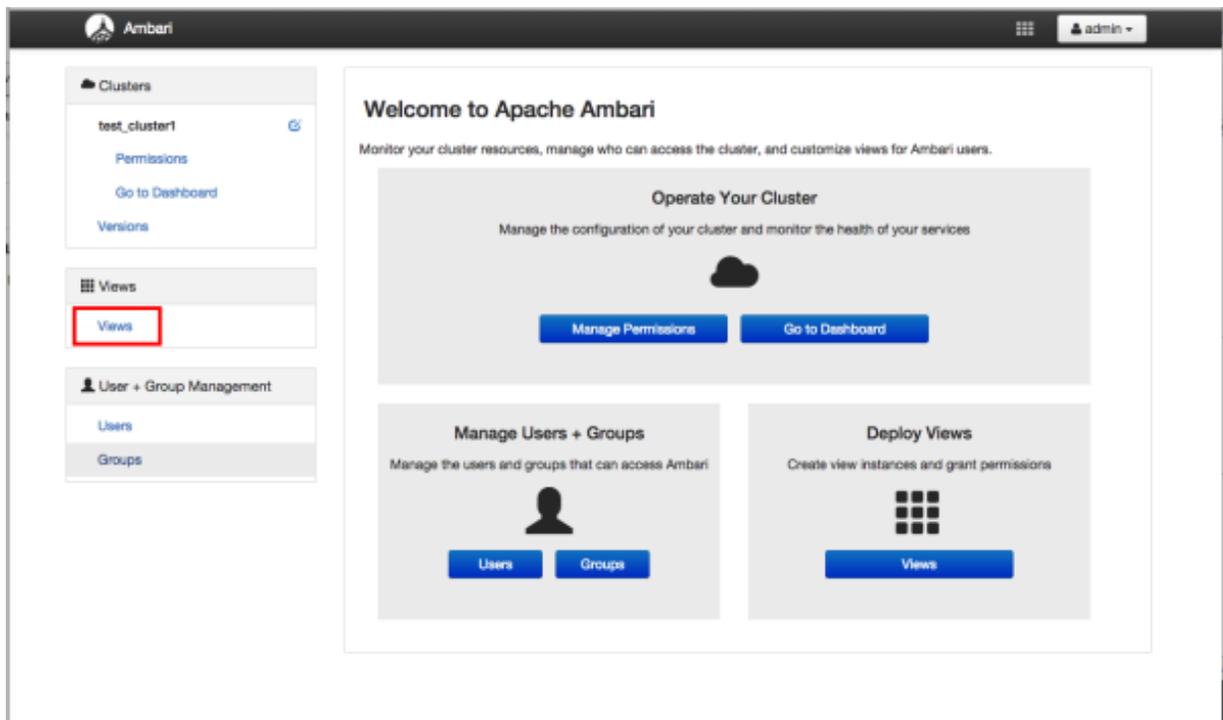
Capacity Scheduler View requires that the cluster is managed by Ambari. Capacity Scheduler View utilizes the Ambari Server API.

Procedure

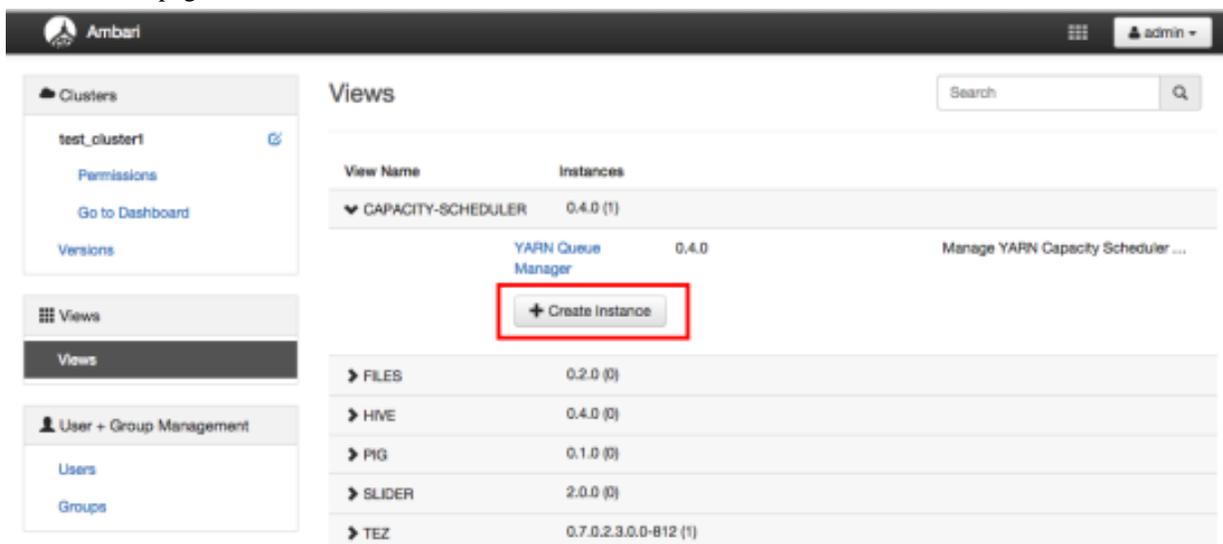
1. In the **Ambari Web** user menu, click **admin > Manage Ambari**.



2. On the **Ambari Admin** page, click **Views**.



3. On the **Views** page, click **CAPACITY-SCHEDULER**, then click **Create Instance**.



4. In the **Details** box on the **Create Instance** page, type an instance name, display name, and a description for the view.
The instance name cannot contain spaces or special characters.
5. In the **Cluster Configuration** box on the **Create Instance** page, configure the view to communicate with the HDP cluster.
 - For HDP clusters that are local (managed by the local Ambari Server), click the **Local Ambari Managed Cluster** option, then click the local cluster name.
 - To configure the view to work with HDP clusters that are remote (not part of this Ambari Server instance), click the **Custom** option, then enter the remote Ambari cluster API URL and the Ambari cluster user name and password.
6. Click Save at the bottom of the **Create Instance** page.

The screenshot displays the Ambari interface for creating a new instance of the CAPACITY-SCHEDULER view. The page is titled "Views / Create Instance". On the left sidebar, there are navigation options for Clusters, Views, and User + Group Management. The main content area is divided into three sections: "Details", "Cluster Configuration", and "Buttons".

Details Section:

- View:** CAPACITY-SCHEDULER
- Version:** 0.4.0
- Instance Name:** Capacity_Scheduler_1
- Display Name:** Capacity Scheduler 1
- Description:** Capacity Scheduler configuration 1
- Visible:**

Cluster Configuration Section:

- Local Ambari Managed Cluster:** (Selected)
- Cluster Name:** test_cluster1
- Custom:**
- Ambari Cluster URL*:** http://ambari.server:8080/api/v1/clusters/MyCluster
- Operator Username*:** djones
- Operator Password*:** [Redacted]

Buttons: Cancel, Save (highlighted with a red box)

Results

A Capacity Scheduler View instance is created, and the configuration page for the instance

The screenshot displays the Ambari interface for configuring a Capacity Scheduler View instance. The top navigation bar shows the Ambari logo and the user 'admin'. The main content area is titled 'Views / Capacity Scheduler 1' and includes a 'Delete Instance' button. The instance is identified as 'CAPACITY-SCHEDULER' with version '0.4.0'. The 'Details' section shows the instance name 'Capacity Scheduler_1', display name 'Capacity Scheduler 1', and description 'Capacity Scheduler configuration 1'. A 'Visible' checkbox is checked. The 'Permissions' section shows a table with 'Use' permissions and buttons for 'Add User' and 'Add Group'. The 'Cluster Configuration' section shows the 'Local Ambari Managed Cluster' selected with 'test_cluster1' as the cluster name, and 'Custom' configuration with fields for 'Ambari Cluster URL*', 'Operator Username*', and 'Operator Password*'.

appears.

What to do next

Set Capacity Scheduler (YARN Queue Manager) View user permissions.

Set YARN Queue Manager View user permissions

To add users and groups to a YARN Queue Manager view instance:

Procedure

1. On the Capacity Scheduler view instance configuration page, in **Permissions**, click **Add User**.

The screenshot displays the Ambari web interface for configuring the Capacity Scheduler View. The interface is organized into several sections:

- Navigation Sidebar:** Includes 'Clusters' (with 'test_cluster1' and links for 'Permissions', 'Go to Dashboard', and 'Versions'), 'Views' (with 'Views' selected), and 'User + Group Management' (with 'Users' and 'Groups' links).
- Header:** Shows 'Ambari' logo, a user profile 'admin', and a 'Delete Instance' button.
- View Information:** Displays 'View: CAPACITY-SCHEDULER' and 'Version: 0.4.0'.
- Details Section:** Contains fields for 'Instance Name' (Capacity_Scheduler_1), 'Display Name' (Capacity Scheduler 1), and 'Description' (Capacity Scheduler configuration 1). A 'Visible' checkbox is checked.
- Permissions Section:** Features a table with columns for 'Permission', 'Grant permission to these users', and 'Grant permission to these groups'. Under the 'Use' permission, there are 'Add User' and 'Add Group' buttons. The 'Add User' button is highlighted with a red box.
- Cluster Configuration Section:** Offers two options: 'Local Ambari Managed Cluster' (selected) with a 'Cluster Name' dropdown set to 'test_cluster1', and 'Custom' with an 'Ambari Cluster URL*' field containing 'http://ambari.server:8080/api/v1/clusters/MyCluster'.

2. In **Use**, enter user names, then click the blue check mark to add the users. You can use the same method to add groups in **Add Group**.

The screenshot displays the Ambari web interface for configuring the Capacity Scheduler View. The top navigation bar shows the Ambari logo and the user 'admin'. The left sidebar contains navigation links for Clusters, Views, and User + Group Management. The main content area is titled 'Views / Capacity Scheduler 1' and includes a 'Go to Instance' link and a 'Delete Instance' button. The configuration details are as follows:

- View:** CAPACITY-SCHEDULER
- Version:** 0.4.0
- Details:**
 - Instance Name:** Capacity_Scheduler_1
 - Display Name:** Capacity Scheduler 1
 - Description:** Capacity Scheduler configuration 1
 - Visible
- Permissions:**
 - Permission:** Use
 - Grant permission to these users:** bamith x djones x
 - Grant permission to these groups:** Add Group
- Cluster Configuration:**
 - Local Ambari Managed Cluster
 - Cluster Name:** test_cluster1
 - Custom
 - Ambari Cluster URL*:** http://ambari.server:8080/api/v1/clusters/MyCluster

3. After you have finished adding users and groups, click **Go to instance** at the top of the page to open the YARN Queue Manager view instance.

The screenshot shows the Ambari web interface for configuring the Capacity Scheduler View. The top navigation bar includes the Ambari logo, the user name 'admin', and a hamburger menu icon. The left sidebar contains navigation options for Clusters, Views, and User + Group Management. The main content area is titled 'Views / Capacity Scheduler 1' and includes a 'Go to Instance' button (highlighted with a red box) and a 'Delete Instance' button. Below the title, the 'View' is identified as 'CAPACITY-SCHEDULER' with a version of '0.4.0'. The 'Details' section shows the instance name 'Capacity_Scheduler_1', display name 'Capacity Scheduler 1', and description 'Capacity Scheduler configuration 1'. A 'Visible' checkbox is checked. The 'Permissions' section shows the 'Use' permission granted to users 'bsmith' and 'djones', and to the group 'product_management'. The 'Cluster Configuration' section is set to 'Local Ambari Managed Cluster' with the cluster name 'test_cluster1'. The 'Ambari Cluster URL*' is 'http://ambari.server:8080/api/v1/clusters/MyCluster' and the 'Operator Username*' is 'admin'.

The Capacity Scheduler view instance page

appears.

Configure your cluster for Files View

For Files View to access HDFS, the Ambari Server daemon hosting the view needs to act as the proxy user for HDFS.

About this task

This allows Ambari to submit requests to HDFS on behalf of Files View users.

If you are running views in an operational Ambari server (one that is operating the cluster) Ambari does this setup by default. You should verify that the setup described in the following subsections has been completed. If you are running views on a standalone server, you must setup proxy user settings manually.

To set up an HDFS proxy user for the Ambari Server daemon account, you need to configure the proxy user in the HDFS configuration. This configuration is determined by the account name the ambari-server daemon is running as. For example, if your ambari-server is running as root, you set up an HDFS proxy user for root with the following:

Procedure

1. In **Ambari Web**, browse to **Services > HDFS > Configs**.
2. On the **Advanced** tab, browse to the **Custom core-site** section.
3. Click **Add Property...**
4. Enter the following custom properties:

```
hadoop.proxyuser.root.groups="users"
hadoop.proxyuser.root.hosts=ambari-server.hostname
```

Notice the ambari-server daemon account name root is part of the property name. Be sure to modify this property name for the account name you are running the ambari-server as. For example, if you were running ambari-server daemon under the account name ambariusr, you would use the following properties instead:

```
hadoop.proxyuser.ambariusr.groups="users"
hadoop.proxyuser.ambariusr.hosts=ambari-
server.hostname
```

Similarly, if you have configured Ambari Server for Kerberos, be sure to modify this property name for the primary Kerberos principal user. For example, if ambari-server is setup for Kerberos using principal ambari-server@EXAMPLE.COM, you would use the following properties instead:

```
hadoop.proxyuser.ambari-server.groups="users"
hadoop.proxyuser.ambari-server.hosts=ambari-
server.hostname
```

5. Save the configuration change.

What to do next

Restart the required components as prompted by Ambari.

Create and configure a Files View instance

Ambari creates and configures a Files View instance automatically.

About this task

Only use this information if your Files View instance does not appear as expected.

Procedure

1. Browse to the **Ambari Admin** page.
2. Click **Views > Files View > Create Instance**.
3. In **Details**, enter the following values for View instance properties:

Table 1: Files View Properties

Property	Description	Value
Instance Name	This is the Files view instance name. This value should be unique for all Files view instances you create. This value cannot contain spaces and is required.	FILES_1
Display Name	This is the name of the view link displayed to the user in Ambari Web.	MyFiles
Description	This is the description of the view displayed to the user in Ambari Web.	Browse HDFS files and directories.
Visible	This checkbox determines whether the view is displayed to users in Ambari Web.	Visible or Not Visible

4. Complete the Files View configuration, using the following guidance:
Information that you provide in **Settings and Cluster Configuration** depends on your environment; specifically, whether:
 - your cluster is Kerberos-enabled or not
 - NameNode HA is enabled or not
 - your Files View instance resides in an operational or a standalone Ambari server

Table 2: Files View Configuration

Kerberos Enabled	NameNode HA Enabled	Operational Ambari Server	Standalone Ambari Server
No	No	Settings: defaults	Settings: defaults
No	Yes	Cluster Configuration: Local	Cluster Configuration: Custom
Yes	No	Settings: Kerberos Cluster Configuration: Custom	
Yes	Yes	Settings: Kerberos Cluster Configuration: Custom	

What to do next

Enable the Local Ambari Managed Cluster Configuration option in the Ambari Admin page, only if you are managing a cluster in an Operational Ambari Server.

Set up Kerberos for Files View

Requirements for operating Files View in a Kerberized cluster.

Before you begin

Set up Kerberos for Ambari by configuring the Ambari Server daemon with a Kerberos principal and keytab.

Procedure

- In **Files View > Settings**, enter the following properties:

Table 3: Kerberos Settings for Files View

Property	Description	Example Value
WebHDFS Username	This is the username the view will access HDFS as. Leave this default value intact to represent the authenticated view user.	\${username}
WebHDFS Authorization	This is the semicolon-separated authentication configuration for WebHDFS access.	auth=KERBEROS;proxyuser=ambari-server **This property is only needed if the view is Custom Configured or Ambari Server is Kerberized before 2.4.0.

With a Kerberos setup, the proxy user setting should be the primary value of the Kerberos principal for Ambari Server. For example, if you configured Ambari Server for Kerberos principal `ambari-server@EXAMPLE.COM`, this value would be `ambari-server`.

Configure local option for Files View

When you configure Files View using the Local option, Files View communicates with HDFS based on the `fs.defaultFS` property.

About this task

The **Local Ambari Managed Cluster Configuration** option is enabled on the **Ambari Admin** page if you are managing a cluster with Ambari. Ambari will automatically configure Files View based on how the cluster is configured.

Procedure

- When **Local Ambari Managed Cluster Configuration** is enabled, you can choose **Local Ambari Managed Cluster Configuration**.
When you configure Files View using the Local option, Files View communicates with HDFS based on the `fs.defaultFS` property.
- For example: `hdfs://namenode:8020`.
Files View also determines whether NameNode HA is configured and adjusts accordingly.

Configure custom option for Files View

There are certain properties that you must configure for Custom Files View.

Table 4: Properties Required for Custom Files View Option

Required Properties	Description	Example Value
WebHDFS FileSystem URI	The WebHDFS FileSystem URI in the format <code>webhdfs://[HOST]:[HTTP_PORT]</code>	<code>webhdfs://namenode:50070</code>

These properties are required if your cluster is configured for NameNode HA.

Table 5: Properties Required for Custom Files View with NN HA

Property	Description	Example Value
Logical name of the NameNode cluster	Comma-separated list of nameservices.	<code>hdfs-site/dfs.nameservices</code> For example: <code>nameservice</code>
List of NameNodes	Comma-separated list of NameNodes for a given nameservice.	<code>hdfs-site/dfs.ha.namenodes</code> For example: <code>namenode1,namenode2</code>
First NameNode RPC Address	RPC address for first name node.	<code>hdfs-site/dfs.namenode.rpc-address.[nameservice].[namenode1]</code>
Second NameNode RPC Address	RPC address for second NameNode.	<code>hdfs-site/dfs.namenode.rpc-address.[nameservice].[namenode2]</code>
First NameNode HTTP (WebHDFS) Address	WebHDFS address for first NameNode.	<code>hdfs-site/dfs.namenode.http-address.[nameservice].[namenode1]</code>
Second NameNode HTTP (WebHDFS) Address	WebHDFS address for second NameNode.	<code>hdfs-site/dfs.namenode.http-address.[nameservice].[namenode2]</code>
Failover Proxy Provider	The Java class that HDFS clients use to contact the Active NameNode.	<code>hdfs-site/dfs.client.failover.proxy.provider.[nameservice]</code>

Configuring Pig View

To configure Pig View

Perform the following tasks:

Configure your cluster for Pig View

For Pig View to access HDFS, the Ambari Server daemon hosting the view must act as the proxy user for HDFS. The Ambari Pig View is deprecated in HDP 3.0 and later. Ambari does not enable Pig View. To enable Pig View in HDP 3.0 and later, you need to contact Hortonworks support for instructions that include how to install WebHCat using an Ambari management pack.

Ambari submits requests to HDFS on behalf of the users using the Pig View. This is critical since Pig View will store metadata about the user Pig scripts. This also means users that will access Pig View must have a user directory setup in HDFS. In addition, Pig View uses WebHCat to submit Pig scripts so the View needs a proxy user for WebHCat.

Set up HDFS proxy user for Pig View

To set up an HDFS proxy user for the Ambari Server daemon account, you must configure the proxy user in the HDFS configuration.

About this task

This configuration is determined by the account name the ambari-server daemon is running as. For example, if your ambari-server is running as root, you set up an HDFS proxy user for root .

Procedure

1. In **Ambari Web**, browse to **Services > HDFS > Configs**.
2. On the **Advanced** tab, navigate to the **Custom core-site** section.
3. Click **Add Property...** to add the following custom properties:

```
hadoop.proxyuser.root.groups="users"  
hadoop.proxyuser.root.hosts=ambari-server.hostname
```

Notice the ambari-server daemon account name root is part of the property name. Be sure to modify this property name for the account name you are running the ambari-server as. For example, if you were running ambari-server daemon under the account name ambariusr, you would use the following properties instead:

```
hadoop.proxyuser.ambariusr.groups="users"  
hadoop.proxyuser.ambariusr.hosts=ambari-server.hostname
```

Similarly, if you have configured Ambari Server for Kerberos, be sure to modify this property name for the primary Kerberos principal user. For example, if ambari-server is set up for Kerberos using principal ambari-server@EXAMPLE.COM, you would use the following properties instead:

```
hadoop.proxyuser.ambari-server.groups="users"  
hadoop.proxyuser.ambari-server.hosts=ambari-server.hostname
```

4. Save the configuration change.

What to do next

Restart the required components as prompted by Ambari.

Set up WebHCat proxy user for Pig View

You must set up an HDFS proxy user for WebHCat and a WebHCat proxy user for the Ambari Server daemon account.

About this task

To set up the HDFS proxy user for WebHCat :

Procedure

1. In **Ambari Web**, browse to **Services > HDFS > Configs**.
2. On the **Advanced** tab, navigate to the **Custom core-site** section.
3. Click Add Property... to add the following custom properties:

```
hadoop.proxyuser.hcat.groups=*  
hadoop.proxyuser.hcat.hosts=*
```

4. Save the configuration change and restart the required components as prompted by Ambari.

What to do next

set up a WebHCat proxy user for the Ambari Server daemon account.

Set up WebHCat proxy user for the Ambari Server daemon account

To setup a WebHCat proxy user for the Ambari Server daemon account, you need to configure the proxy user in the WebHCat configuration.

About this task

This configuration is determined by the account name the ambari -server daemon is running as. For example, if your ambari -server is running as root, you set up an WebHCat proxy user for root.

Procedure

1. In **Ambari Web**, browse to **Services > Hive > Configs**.
2. On the **Advanced** tab, navigate to the **Custom webhcat-site** section.
3. Click Add Property... to add the following custom properties:
4. Save the configuration change and restart the required components as indicated by Ambari.

```
webhcat.proxyuser.root.groups=*  
webhcat.proxyuser.root.hosts=*
```

Notice the ambari-server daemon account name root is part of the property name. Be sure to modify this property name for the account name you are running the ambari-server as. For example, if you were running ambari-server daemon under an account name of ambariusr, you would use the following properties instead:

```
webhcat.proxyuser.ambariusr.groups=*  
webhcat.proxyuser.ambariusr.hosts=*
```

Similarly, if you have configured Ambari Server for Kerberos, be sure to modify this property name for the primary Kerberos principal user. For example, if ambari-server is setup for Kerberos using principal ambari-server@EXAMPLE.COM, you would use the following properties instead:

```
webhcat.proxyuser.ambari-server.groups=*  
webhcat.proxyuser.ambari-server.hosts=*
```

Set up HDFS user directory

Is this task necessary for Pig View ??? seems Hive-specific

About this task

The Hive View stores user metadata in HDFS. By default, the location in HDFS for this metadata is `/user/${USER_NAME}` where `${USER_NAME}` is the user name of the currently logged in user that is accessing the Hive View.

Before you begin

Since many users leverage the default Ambari admin user for getting started with Ambari, the `/user/admin` folder must be created in HDFS. Therefore, be sure to create the admin user directory in HDFS using these instructions before using the view.

Procedure

1. Connect to a host in the cluster that includes the HDFS client.
2. Switch to the `hdfs` system account user.
`su - hdfs`
3. Using the HDFS client, make an HDFS directory for the user.
For example, if your username is `admin`, you would create the following directory:`hadoop fs -mkdir /user/admin`
4. Set the ownership on the newly created directory.
For example, if your username is `admin`, you would make that user the directory owner. `hadoop fs -chown admin:hadoop /user/admin`

Create and configure a Pig View instance

To create a Pig View Instance:

Procedure

1. Browse to the Ambari Admin interface.
2. Click **Views**, expand the **Pig View**, and click **Create Instance**.
3. On the **Create Instance** page, select **Version**.
If multiple Pig View jars are present, choose one.
4. Enter the Details and Settings.
The Instance Name appears in the URI, the Display Name appears in the Views drop-down list, and the Description helps multiple users identify the view:

View **PIG**

Version

Details

Instance Name

Display Name

Description

Visible

Settings

WebHDFS Username

WebHDFS Authentication

WebHCat Username

Scripts HDFS Directory*

Jobs HDFS Directory*

Meta HDFS Directory

5. Scroll down, and enter the Cluster Configuration information, which tells the Pig View how to access resources in the cluster.

For a cluster that is deployed and managed by Ambari, select **Local Ambari Managed Cluster**:

Cluster Configuration

Local Ambari Managed Cluster

Cluster Name:

Custom

WebHDFS FileSystem URI*

Logical name of the NameNode cluster:

List of NameNodes:

First NameNode RPC Address:

Second NameNode RPC Address:

First NameNode HTTP (WebHDFS) Address:

Second NameNode HTTP (WebHDFS) Address:

Fallover Proxy Provider:

WebHCat Hostname*

WebHCat Port*

6. Click **Save**.
7. Give permissions to the appropriate users and groups.

What to do next

Click **Go to instance** at the top of the page open instance of Pig View that you created.

Get correct configuration values for Pig View in a manually-deployed cluster

About this task

If you have manually deployed your cluster, you must enter cluster configuration values in the **Pig View > Create Instance** page. The following table explains where you can find cluster configuration settings in Ambari.

Scripts HDFS Directory*	<code>/user/\${username}/pig/scripts</code>
Jobs HDFS Directory*	<code>/user/\${username}/pig/jobs</code>
WebHDFS FileSystem URI*	Click HDFS > Configs > Advanced hdfs-site > dfs.namenode.http-address . When you enter the value in the view definition, pre-pend <code>webhdfs://</code> to the value

WebHCat Hostname*

you find in the advanced HDFS configuration settings.
For example, webhdfs://c6401.ambari.apache.org:50070

Click **Hive > Configs > Advanced > WebHCat Server > WebHCat Server host** to view the hostname.
For example, c6402.ambari.apache.org

WebHCat Port*

Click **Hive > Configs > Advanced > Advanced webhcat-site > templeton.port** to view the port number.
For example, 50111

For NameNode High Availability, the following values must be entered for primary and secondary NameNodes:

First NameNode RPC Address or Second NameNode RPC Address

Select the primary or secondary NameNode to view settings from that host in the cluster. When you enter the value in the view definition, pre-pend http:// to the value you find in the advanced **hdfs-site** settings. For example, http://c6401.ambari.apache.org:8020

First NameNode HTTP (WebHDFS) Address or Second NameNode HTTP (WebHDFS) Address

Click **HDFS > Configs > Advanced > Advanced hdfs-site > dfs.namenode.http-address**. When you enter the value in the view definition, pre-pend http:// to the value you find in the advanced **hdfs-site** settings. For example, http://c6401.ambari.apache.org:50070

To get First NameNode RPC Address values:

Procedure

1. In **Ambari Web**, browse to the **HDFS Summary** page. Click **NameNode** (primary) or **SNameNode** (secondary) to view the host page:
HDFS Service Page in Ambari:

The screenshot shows the Ambari HDFS Summary page. The left sidebar contains a navigation menu with items like HDFS, YARN, MapReduce2, Tez, Hive, HBase, Pig, Oozie, ZooKeeper, Storm, Ambari Infra, and Ambari Metrics. The main content area has tabs for Summary, Heatmaps, Configs, and Quick Links. The Summary tab is active, showing a table with columns for NameNode, SNameNode, and their status (Started) and alert status (No alerts). Below this table, there are several status indicators: DataNodes Status (3 live / 0 dead / 0 decommissioning), JournalNodes (0/0 JournalNodes Live), HDFS Gateways (0/0 Started), NameNode Uptime (6.12 days), NameNode Heap (102.9 MB / 1011.3 MB (10.2% used)), Disk Usage (DFS Used) (5.1 GB / 1.3 TB (0.40%)), and Disk Usage (Non DFS Used) (109.1 GB / 1.3 TB (8.50%)). On the right side, there are additional metrics: Disk Remaining (1.1 TB / 1.3 TB (81.10%)), Blocks (total) (738), Block Errors (0 corrupt replica / 0 missing / 0 under replicated), Total Files + Directories (947), Upgrade Status (No pending upgrade), and Safe Mode Status (Not in safe mode).

2. On the host page, click **Configs > Advanced**.
3. Enter **rpc** in the Filter search well at the top right corner of the page or navigate to the **Advanced hdfs-site** settings to find the **dfs.namenode.rpc-address** value that you can enter into the Pig View definition.
Here is an example of using Filter to search settings in **Advanced hdfs-site**.

The screenshot shows the Ambari configuration page for Pig View. At the top, there are tabs for 'Summary', 'Heatmaps', 'Configs', 'Quick Links', and 'Service Actions'. Below these is a 'Group' dropdown set to 'Default (3)' and a 'Manage Config Groups' button. A search box contains 'rpc'. A notification banner shows a user 'admin' logged in 7 days ago on HDP-2.5. Below the notification is a 'Settings' section with an 'Advanced' tab. Under 'Advanced hdfs-site', the 'dfs.namenode.rpc-address' is set to 'c6401.ambari.apache.org:8020'. A tooltip explains that this is the 'RPC address that handles all clients requests'.

Set up user permissions for Pig View

Before you begin

Save the Pig View instance definition

Procedure

1. Grant view permissions for all Pig View users.
2. On the Pig View instance configuration page, in **Permissions**, click **Add User**.

Views / My Pig View [Go to instance](#) Delete Instance

View **PIG**

Version

Details Edit

Instance Name

Display Name

Description

Visible

Permissions

Permission	Grant permission to these users	Grant permission to these groups
Use	<input type="text" value="Add User"/>	<input type="text" value="Add Group"/>

Settings Edit

3. In **Use**, enter user names, then click the blue check mark to add the users. You can use the same method to add groups in **Add Group**.
4. After you have finished adding users and groups, click **Go to instance** at the top of the page to open the Pig View instance.

Set up Kerberos for Pig View

Before you begin

Set up basic Kerberos for the Ambari views server.

Procedure

- Manually set the following Pig View property:

WebHDFS Authentication

`auth=KERBEROS;proxyuser=[AMBARI_PRINCIPAL_NAME]`

The screenshot shows the 'Properties' configuration page for SmartSense View. The 'WebHDFS Authentication' field is highlighted with a red box and contains the value 'auth=KERBEROS;proxyuser=ambariuser'. Other fields include 'WebHDFS FileSystem URI*', 'WebHDFS Username', 'WebHCat URL*', 'WebHCat Username', 'Dataworker Username', 'Scripts HDFS Directory*', 'Jobs HDFS Directory*', and 'Meta HDFS Directory'.

This property is only needed if the view is Custom Configured or Ambari Server is Kerberized before 2.4.0.

Configuring SmartSense View

When you deploy a cluster with Ambari, a SmartSense View instance is automatically created as long as an Ambari Agent is deployed on the host running the Ambari Server.

If necessary, review and complete the following procedures:

Configure your cluster for SmartSense View

If an Ambari Agent is not installed on the Ambari Server host, the view will not be automatically created, and you will have to add a SmartSense instance manually.

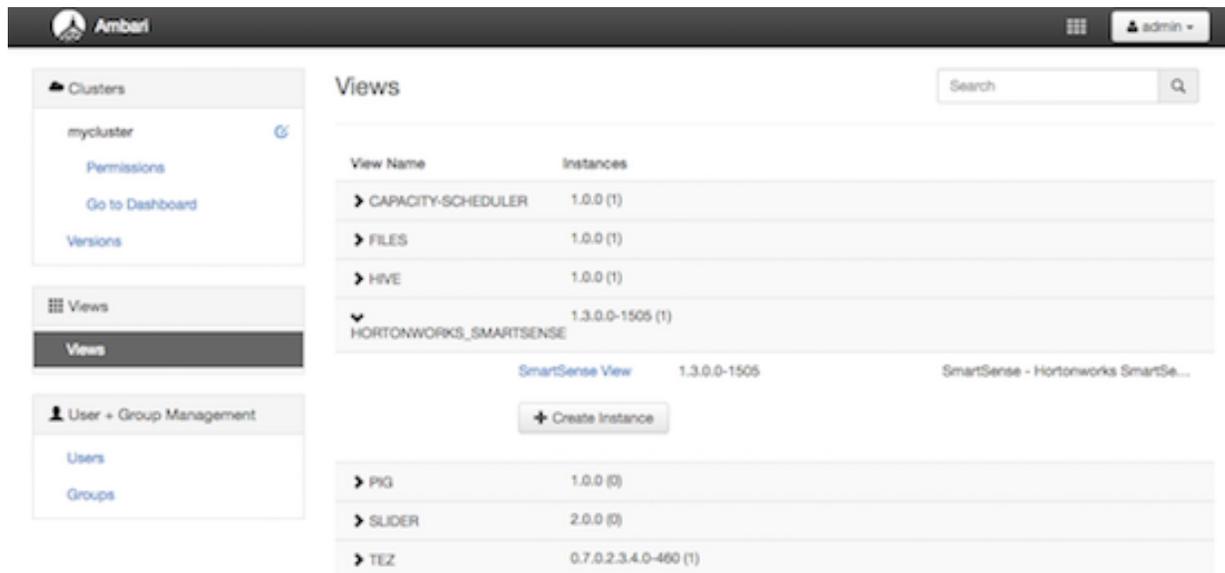
Before accessing SmartSense View, you should enter your SmartSense user ID, account name (both are available in the Hortonworks support portal in the Tools tab), and email address in the SmartSense service configuration properties.

Create a SmartSense View instance

To create a SmartSense view instance manually:

Procedure

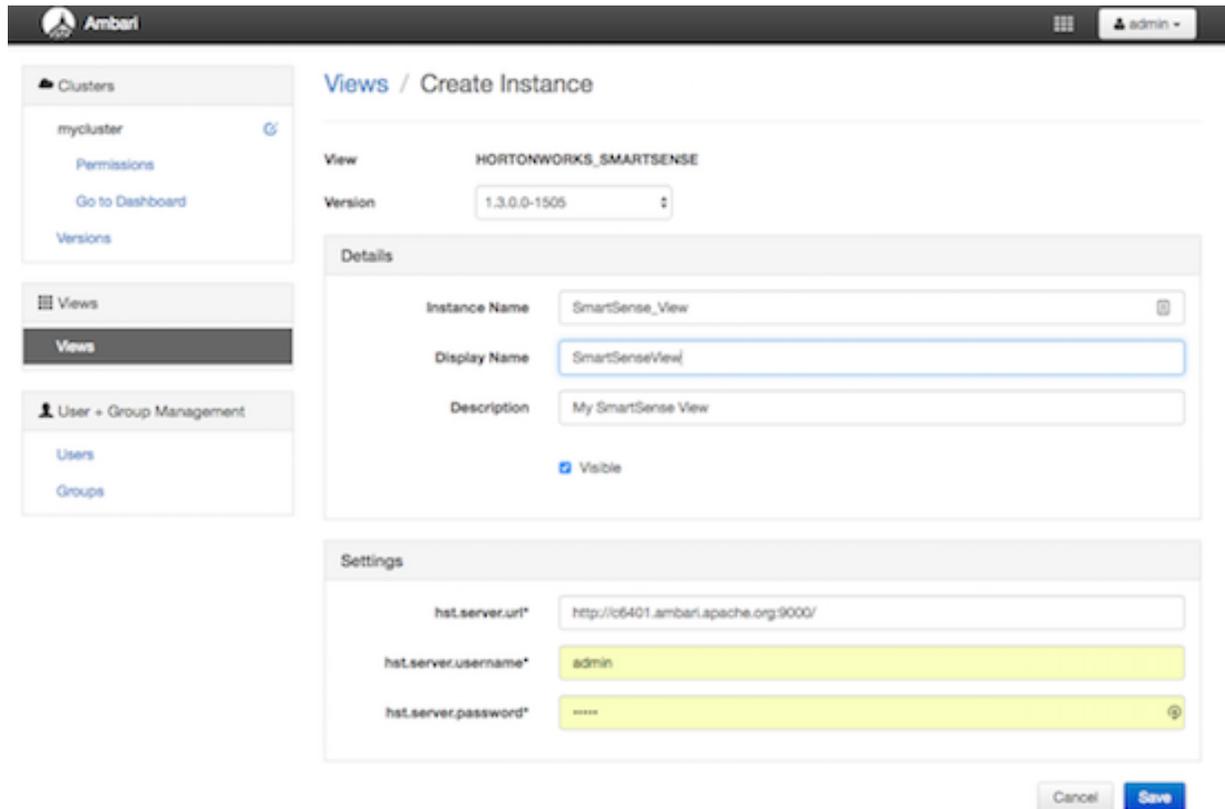
1. Browse to the Ambari Admin user interface.
2. Click **Views > HORTONWORKS_SMARTSENSE**, expand the menu, and click **Create Instance**.



The screenshot shows the Ambari interface with the 'Views' section selected. A table lists various views and their instances:

View Name	Instances
▶ CAPACITY-SCHEDULER	1.0.0 (1)
▶ FILES	1.0.0 (1)
▶ HIVE	1.0.0 (1)
▼ HORTONWORKS_SMARTSENSE	1.3.0.0-1505 (1)
SmartSense View	1.3.0.0-1505
SmartSense - Hortonworks SmartSe...	
+ Create Instance	
▶ PIG	1.0.0 (0)
▶ SLIDER	2.0.0 (0)
▶ TEZ	0.7.0.2.3.4.0-460 (1)

- On the **Create Instance** page, select the **Version**.



The screenshot shows the 'Views / Create Instance' page in Ambari. The 'View' is set to 'HORTONWORKS_SMARTSENSE' and the 'Version' is '1.3.0.0-1505'. The 'Details' section contains the following fields:

- Instance Name: SmartSense_View
- Display Name: SmartSenseView
- Description: My SmartSense View
- Visible:

The 'Settings' section contains the following fields:

- hst.server.url*: http://c6401.ambari.apache.org:9000/
- hst.server.username*: admin
- hst.server.password*: *****

Buttons for 'Cancel' and 'Save' are visible at the bottom right.

If multiple SmartSense View jars exist, choose one.

- Enter the following view instance details.

Instance Name

This is the SmartSense view instance name. This value should be unique for all SmartSense view instances you create. This value cannot contain spaces and is required.

Display Name

This is the name of the view link displayed to the user in Ambari Web.

Description	This is the description of the view displayed to the user in Ambari Web.
5. Enter the following view instance settings:	
hst.server.url	This is the HST server URL. This should be http://<HST_host>:9000/
hst.server.username	The default username is 'admin'.
hst.server.password	Unless changed after installation, the default password is 'admin'.
6. Click Save .	

Configure Workflow Manager View

Before you can access Workflow Manager, you must complete several configuration tasks.

See the following content:

Configure your cluster for Workflow Manager View

For Workflow Manager View to access HDFS, the Ambari Server daemon hosting the view must act as the proxy user for HDFS.

About this task

This allows Ambari to submit requests to HDFS on behalf of the Worklow Manager View users.

Before you begin

- Each Worklow Manager View user must have a user directory set up in HDFS.
- If the cluster is configured for Kerberos, ensure that Views is also set up for Kerberos.

Set up HDFS proxy user

To set up an HDFS proxy user for the Ambari Server daemon account, you must configure the proxy user in the HDFS configuration.

About this task

Note: If you previously set up the proxy user for another View, you can skip this task.

This configuration is determined by the account name the ambari-server daemon is running as. For example, if your ambari-server is running as root, you set up an HDFS proxy user for root.

Procedure

1. In **Ambari Web**, browse to **Services > HDFS > Configs**.
2. On the **Advanced** tab, navigate to the **Custom core-site** section.
3. Click **Add Property...**, then add the following custom properties:

```
hadoop.proxyuser.root.groups="users"
hadoop.proxyuser.root.hosts=ambari-server.hostname
```

Notice the `ambari-server` daemon account name `root` is part of the property name. Be sure to modify this property name for the account name you are running the `ambari-server` as. For example, if you were running `ambari-server` daemon under an account name of `ambariusr`, you would use the following properties instead.

```
hadoop.proxyuser.ambariusr.groups="users"
hadoop.proxyuser.ambariusr.hosts=ambari-server.hostname
```

Similarly, if you have configured Ambari Server for Kerberos, be sure to modify this property name for the primary Kerberos principal user. For example, if `ambari-server` is setup for Kerberos using principal `ambari-server@EXAMPLE.COM`, you would use the following properties instead:

```
hadoop.proxyuser.ambari-server.groups="users"
hadoop.proxyuser.ambari-server.hosts=ambari-server.hostname
```

4. Save the configuration change.

What to do next

Restart the required components as prompted by Ambari.

Set up HDFS user directory for Workflow Manager View

You must set up a directory for each user that accesses the Workflow Manager View.

About this task

Workflow Manager View stores user metadata in the user directory in HDFS. By default, the location in HDFS for the user directory is `/user/[USER_NAME]`, where `[USER_NAME]` is the user name of the currently logged in user that is accessing Workflow Manager View.

Important: Since many users leverage the default Ambari admin user for getting started with Ambari, you should create a `/user/admin` directory if one does not exist, in addition to directories for other Workflow Manager View users.

Procedure

1. Connect to a host in the cluster that includes the HDFS client.
2. Switch to the HDFS system account user.


```
su - hdfs
```
3. If working on a secure Kerberos cluster.
 - a) Destroy any existing Kerberos tickets.


```
kdestroy
```

If no ticket is found, you get an error message that you can ignore: No credentials cache file found while destroying cache.
 - b) Obtain a Kerberos ticket-granting ticket.


```
kinit -kt /etc/security/keytabs/hdfs.headless.keytab hdfs
```
4. Using the HDFS client, make an HDFS directory for the user.

For example, if your username is `wfadmin`, you would create the following directory: `hadoop fs -mkdir /user/wfadmin`
5. Set the ownership on the newly created directory.

For example, if your username is `wfadmin`, the directory owner should be `wfadmin:hadoop`.

```
hadoop fs -chown wfadmin:hadoop /user/wfadmin
```
6. Log in as root user.


```
su -
```
7. Access the Kerberos administration system.


```
kadmin.local
```

8. Create a new principal and password for the user.
`addprinc -pw [wfmadmin-password] wfmadmin@EXAMPLE.COM`
 You can use the same user name and password that you used for HDFS directory.

What to do next

Repeat steps 2 through 8 for any additional Workflow Manager View users.

Set up Kerberos for Workflow Manager View

If you install Ambari using Kerberos, the Kerberos settings for Oozie that are required for Workflow Manager are configured automatically.

About this task

The image below shows what the settings should look like.

Procedure

1. In **Ambari Web**, browse to **Services > Oozie > Configs > Advanced**.
2. On the **Advanced** tab, navigate to the **Advanced Oozie-site** section.
3. Verify that the properties match those shown in the following figure.

The screenshot shows the Ambari Web configuration interface for Oozie-site settings, divided into two sections: 'Advanced oozie-site' and 'Custom oozie-site'.

Advanced oozie-site

- `oozie.authentication.kerberos.name.rules`: `RULE:[1:$1@$0](ambari-qa-cl1@EXAMPLE.COM)s/.*ambari-qa/RULE:[1:$1@$0](cstm-`
- `oozie.authentication.type`: `kerberos`
- `oozie.service.HadoopAccessorService.kerberos.enabled`: `true`

Custom oozie-site

- `oozie.authentication.kerberos.keytab`: `/etc/security/keytabs/spnego.service.keytab`
- `oozie.authentication.kerberos.principal`: `HTTP/_HOST@EXAMPLE.COM`
- `oozie.service.HadoopAccessorService.kerberos.principal`: `oozie/_HOST@EXAMPLE.COM`
- `oozie.service.HadoopAccessorService.keytab.file`: `/etc/security/keytabs/oozie.service.keytab`

At the bottom of the 'Custom oozie-site' section, there is a link labeled 'Add Property ...'.

Set up proxy user for Oozie

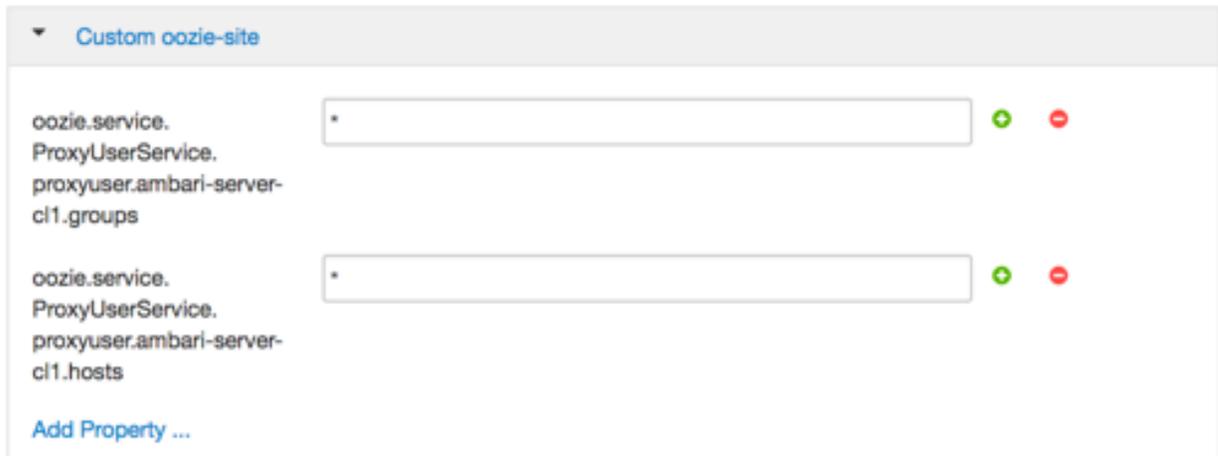
If you are using Kerberos, you must configure the proxy user for Oozie.

About this task

Workflow Manager uses Oozie as its scheduling engine.

Procedure

1. In **Ambari Web**, browse to **Services > Oozie > Configs**.
2. Expand the **Custom oozie-site** section.
3. Click **Add Property...**, then add the following custom properties:
`oozie.service.ProxyUserService.proxyuser.[AMBARI_SERVER_CL1].groups=*`
 Replace [AMBARI_SERVER_CL1] with the server principal name used when configuring Kerberos.
`oozie.service.ProxyUserService.proxyuser.[AMBARI_SERVER_CL1].hosts=*`
 Replace [AMBARI_SERVER_CL1] with the server principal name used when configuring Kerberos.
 For example:



4. Save the configuration change.

What to do next

Restart the required components as prompted by Ambari.

Create and configure a Workflow Manager View instance

Use Ambari Admin to create a Workflow Manager view instance.

About this task

You can configure multiple WFM View instances. You might want to have multiples instances if you want to assign different users and permissions for each instance. You can also have instances that run locally and others that run remotely.

Procedure

1. Click **Manage Ambari** to open the **Ambari Admin** user interface.
2. Click **Views**, expand the **Workflow_Manager View**, and click **Create Instance**.
3. On the Create Instance page, select the **Version**.
 If multiple View versions exist, choose one.
4. Enter the following view instance Details:

Table 6: Workflow Manager View Instance Details

Property	Description	Example Value
Instance Name	This is the Workflow Manager View instance name and must be unique for all instances you create. The instance name cannot be modified after the instance is created.	wfm_local_instance This value cannot contain spaces or special characters other than an underscore.
Display Name	This is the name of the view link displayed to the user in Ambari Web. The display name can be modified after the instance is created.	WFM View This value can contain spaces and underscores, but no other special characters.
Description	This is the description of the view displayed to the user in Ambari Web.	Local instance of WFM
Visible	This checkbox determines whether the view is displayed to users in Ambari Web.	Visible or Not Visible

- In **Cluster Configuration**, select the **Local Cluster** or **Remote Cluster** instance type and select the cluster name.
- Review the remaining settings and make changes as desired.
The Settings and Cluster Configuration options depend on a few cluster and deployment factors in your environment. Typically, you can accept the default Settings unless you are using the Workflow Manager View with a Kerberos-enabled cluster.
- Click **Save**.
The instance is created and a success message displays.
- Scroll to the **Permissions** section at the bottom of the Views configuration form.
- Grant permission on the Workflow Manager View for the set of users and groups who can access the view instance.

The screenshot shows the 'Permissions' section of the Ambari configuration interface. It is divided into three main areas:

- Permission:** A dropdown menu currently set to 'Use'.
- Grant permission to these users:** A text input field containing the user 'admin'.
- Grant permission to these groups:** A text input field with an 'Add Group' button.
- Local Cluster Permissions:** A section titled 'Grant Use permission for the following mycluster Roles:' with a list of roles, each with a checked checkbox:
 - Cluster Administrator
 - Cluster Operator
 - Service Operator
 - Service Administrator
 - Cluster User

At the bottom of the Local Cluster Permissions section, there are links for 'Check All' and 'Clear All'.

Related Information[Workflow Manager Guide](#)