

HDP Installing Ranger 3

Installing Apache Ranger

Date of Publish: 2018-07-15

<http://docs.hortonworks.com>

Contents

Installing Ranger Using Ambari Overview.....	3
Set Up Hadoop Group Mapping for LDAP/AD.....	3
Ranger Password Requirements.....	5
Configuring a Database Instance for Ranger.....	5
Configure a Ranger DB: MySQL/MariaDB.....	6
Configure a Ranger DB: PostgreSQL.....	7
Configure a Ranger DB: Oracle.....	8
Configure a Ranger DB: Amazon RDS.....	8
Start the Ranger Installation.....	10
Customize Services: Admin.....	11
Customize Services: Audit.....	21
Customize Services: Plugins.....	22
Customize Services: User Sync.....	23
Customize Services: Tagsync.....	32
Customize Services: Authentication.....	33
Customize Authentication: UNIX.....	33
Customize Authentication: LDAP.....	34
Customize Authentication: AD.....	36
Complete the Ranger Installation.....	39
Additional Ranger Plugin Configuration Steps for Kerberos Clusters.....	40
Additional Ranger Plugin Steps for Kerberos: HDFS.....	40
Additional Ranger Plugin Steps for Kerberos: Hive.....	41
Additional Ranger Plugin Steps for Kerberos: HBase.....	42
Additional Ranger Plugin Steps for Kerberos: Knox.....	42

Installing Ranger Using Ambari Overview

Apache Ranger can be installed either manually using the Hortonworks Data Platform (HDP) or the Ambari User Interface (UI). The Ranger service option will be made available through the Add Service wizard after the HDP cluster is installed using the installation wizard.

Once Ambari has been installed and configured, you can use the Add Service wizard to install the following components:

- Ranger Admin
- Ranger UserSync
- Ranger Key Management Service

After these components are installed and started, you can enable Ranger plugins by navigating to each individual Ranger service (HDFS, HBase, Hiveserver2, Storm, Knox, YARN, and Kafka) and modifying the configuration under advanced ranger-<service>-plugin-properties.

Note that when you enable a Ranger plugin, you will need to restart the component.

Note:

Enabling Apache Storm or Apache Kafka requires you to enable Kerberos. To enable Kerberos on your cluster, see “Configuring Authentication with Kerberos”.

Related Information

[Installing the Ranger Key Management Service](#)

Set Up Hadoop Group Mapping for LDAP/AD

To ensure that LDAP/AD group level authorization is enforced in Hadoop, you should set up Hadoop group mapping for LDAP/AD.

Before you begin

You must have access to LDAP and the connection details. Note that LDAP settings can vary depending on what LDAP implementation you are using

About this task

There are three ways to set up Hadoop group mapping:

- Using SSSD (Recommended)
- Manually create users and groups in the Linux environment
- In core-site.xml

Procedure

- Using SSSD (Recommended)

The recommended method for group mapping is to use SSSD or one of the following services to connect the Linux OS with LDAP:

- Centrify
- NSLCD
- Winbind
- SAMBA

Note that most of these services allow you to not only look up a user and enumerate their groups, but also allow you to perform other actions on the host. None of these features are required for LDAP group mapping on Hadoop -- all that is required is the ability to lookup (or "validate") a user within LDAP and enumerate their groups. Therefore, when evaluating these services, take the time to understand the difference between the NSS module (which performs user/group resolution) and the PAM module (which performs user authentication). NSS is required. PAM is not required, and may represent a security risk.

- Manually create users and groups in the Linux environment: Manually create users and groups (see link below) in your Linux environment.
- In core-site.xml, configure Hadoop to use LDAP-based group mapping:
 - a) Add the properties shown in the example below to the core-site.xml file.

You will need to provide the value for the bind user, the bind password, and other properties specific to your LDAP instance, and make sure that object class, user, and group filters match the values specified in your LDAP instance.

```
<property>
<name>hadoop.security.group.mapping</name>
<value>org.apache.hadoop.security.LdapGroupsMapping</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.bind.user</name>
<value>cn=Manager,dc=hadoop,dc=apache,dc=org</value>
</property>

<!--
<property>
<name>hadoop.security.group.mapping.ldap.bind.password.file</name>
<value>/etc/hadoop/conf/ldap-conn-pass.txt</value>
</property>
-->

<property>
<name>hadoop.security.group.mapping.ldap.bind.password</name>
<value>hadoop</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.url</name>
<value>ldap://localhost:389/</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.base</name>
<value></value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.filter.user</name>
<value>(&amp;(|(objectclass=person)(objectclass=applicationProcess))
(cn={0}))</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.filter.group</name>
<value>(objectclass=groupOfNames)</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.attr.member</name>
<value>member</value>
```

```

</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.attr.group.name</name>
<value>cn</value>
</property>

```

- b) Depending on your configuration, you may be able to refresh user and group mappings using the following HDFS and YARN commands:

```

hdfs dfsadmin -refreshUserToGroupsMappings
yarn rmadmin -refreshUserToGroupsMappings

```

- c) Verify LDAP group mapping by running the `hdfs groups` command. This command will fetch groups from LDAP for the current user. Note that with LDAP group mapping configured, the HDFS permissions can leverage groups defined in LDAP for access control.

Related Information

[Manually create users and groups](#)

[SSSD](#)

Ranger Password Requirements

This topic lists password requirements for Ranger and Ranger KMS.

Ranger user password requirements:

- Minimum of 8 characters
- Must include at least one alphabetical and one numerical character
- Must not include the following unsupported special characters: " '\ `

Ranger and Ranger KMS DB user password requirements:

- Must not include the following unsupported special characters: " '\ `

Ranger database instance password requirements:

- Refer to the password requirements for the applicable database type (MySQL, PostgreSQL, Oracle, etc.)

Configuring a Database Instance for Ranger

A database instance must be running and available to be used by Ranger. You can configure MySQL, Oracle, PostgreSQL, or Amazon RDS for this purpose.

- A MySQL, Oracle, PostgreSQL, or Amazon RDS database instance must be running and available to be used by Ranger.

The Ranger installation will create two new users (default names: `rangeradmin` and `rangerlogger`) and two new databases (default names: `ranger` and `ranger_audit`).

- Configuration of the database instance for Ranger is described in the following sections for some of the databases supported by Ranger.
 - MySQL
 - PostgreSQL
 - Oracle
 - AmazonRDS
- If you choose not to provide system Database Administrator (DBA) account details to the Ambari Ranger installer, you can use the `dba_script.py` Python script to create Ranger DB database users without exposing DBA

account information to the Ambari Ranger installer. You can then run the normal Ambari Ranger installation without specifying a DBA user name and password. For more information see “Set up Database Users Without Sharing DBA Credentials”.

Configure a Ranger DB: MySQL/MariaDB

How to configure your MySQL database instance for Ranger.

Before you begin

A MySQL/Oracle/PostgreSQL/Amazon RDS database instance must be running and available to be used by Ranger.

When using MySQL, the storage engine used for the Ranger admin policy store tables MUST support transactions. InnoDB is an example of engine that supports transactions. A storage engine that does not support transactions is not suitable as a policy store.

If you are using Amazon RDS, see “Configure an Amazon RDS Database Instance for Ranger: Prerequisites”.

Procedure

1. The MySQL database administrator should be used to create the Ranger databases. The following series of commands could be used to create the rangerdba user with password rangerdba.
 - a) Log in as the root user, then use the following commands to create the rangerdba user and grant it adequate privileges.

```
CREATE USER 'rangerdba'@'localhost' IDENTIFIED BY 'rangerdba';
GRANT ALL PRIVILEGES ON *.* TO 'rangerdba'@'localhost';

CREATE USER 'rangerdba'@'%' IDENTIFIED BY 'rangerdba';
GRANT ALL PRIVILEGES ON *.* TO 'rangerdba'@'%';

GRANT ALL PRIVILEGES ON *.* TO 'rangerdba'@'localhost' WITH GRANT
OPTION;

GRANT ALL PRIVILEGES ON *.* TO 'rangerdba'@'%' WITH GRANT OPTION;

FLUSH PRIVILEGES;
```

2. Use the exit command to exit MySQL.
3. You should now be able to reconnect to the database as rangerdba using the following command: `mysql -u rangerdba -prangerdba`.
4. After testing the rangerdba login, use the exit command to exit MySQL.
5. Confirm that the `mysql-connector-java.jar` file is in the Java share directory. This command must be run on the server where Ambari server is installed: `ls /usr/share/java/mysql-connector-java.jar`.

If the file is not in the Java share directory, use the following command to install the MySQL connector .jar file:

- RHEL/CentOS/Oracle Linux: `yum install mysql-connector-java*`
 - SLES: `zypper install mysql-connector-java*`
6. Set the `jdbc/driver/path` based on the location of the MySQL JDBC driver .jar file. This command must be run on the server where Ambari server is installed: `ambari-server setup --jdbc-db={database-type} --jdbc-driver={jdbc/driver/path}`.

```
ambari-server setup --jdbc-db=mysql --jdbc-driver=/usr/share/java/mysql-
connector-java.jar
```

What to do next

“Start the Ranger Installation”

Configure a Ranger DB: PostgreSQL

How to configure your MySQL database instance for Ranger.

Before you begin

A MySQL/Oracle/PostgreSQL/Amazon RDS database instance must be running and available to be used by Ranger.

If you are using Amazon RDS, see “Configure an Amazon RDS Database Instance for Ranger: Prerequisites”.

Procedure

1. On the PostgreSQL host, install the applicable PostgreSQL connector:

Option

RHEL/CentOS/Oracle Linux

```
um install postgresql-jdbc*
```

SLES

```
zypper install -y postgresql-jdbc
```

2. Confirm that the .jar file is in the Java share directory. `ls /usr/share/java/postgresql-jdbc.jar`.
3. Change the access mode of the .jar file to 644: `chmod 644 /usr/share/java/postgresql-jdbc.jar`.
4. The PostgreSQL database administrator should be used to create the Ranger databases. The following series of commands could be used to create the rangerdba user and grant it adequate privileges:

```
echo "CREATE DATABASE $dbname;" | sudo -u $postgres psql -U postgres
echo "CREATE USER $rangerdba WITH PASSWORD '$passwd';" | sudo -u $postgres
psql -U postgres
echo "GRANT ALL PRIVILEGES ON DATABASE $dbname TO $rangerdba;" | sudo -u
$postgres psql -U postgres
```

Where:

- \$postgres is the Postgres user.
 - \$dbname is the name of your PostgreSQL database.
5. Set the jdbc/driver/path based on the location of the PostgreSQL JDBC driver .jar file. This command must be run on the server where Ambari server is installed. `ambari-server setup --jdbc-db={database-type} --jdbc-driver={/jdbc/driver/path}`.

```
ambari-server setup --jdbc-db=postgres --jdbc-driver=/usr/share/java/
postgresql-jdbc.jar
```

6. Run the following command: `export HADOOP_CLASSPATH=${HADOOP_CLASSPATH}; ${JAVA_JDBC_LIBS}:/connector.jar path`.
7. Add Allow Access details for Ranger users:
 - a) Change `listen_addresses='localhost'` to `listen_addresses='*'` (* = any) to listen from all IPs in `postgresql.conf`.
 - b) Make the following changes to the Ranger db user and Ranger audit db user in the `pg_hba.conf` file.

```
# TYPE DATABASE USER CIDR-ADDRESS METHOD
# "local" is for Unix domain socket connections only
local all postgres,rangeradmin,rangerlogger trust
# IPv4 local connections:
host all postgres,rangeradmin,rangerlogger 0.0.0.0/0 trust
# IPv6 local connections:
host all postgres,rangeradmin,rangerlogger ::/0 trust
"/var/lib/pgsql/data/pg_hba.conf" 74L, 3445C
```

8. After editing the `pg_hba.conf` file, run the following command to refresh the PostgreSQL database configuration:


```
sudo -u postgres /usr/bin/pg_ctl -D $PGDATA reload.
```

 If the `pg_hba.conf` file is located in the `/var/lib/pgsql/data` directory, the value of `$PGDATA` is `/var/lib/pgsql/data`.

What to do next

“Start the Ranger Installation”

Configure a Ranger DB: Oracle

How to configure your Oracle database instance for Ranger.

Before you begin

A MySQL/Oracle/PostgreSQL/Amazon RDS database instance must be running and available to be used by Ranger.

If you are using Amazon RDS, see “Configure an Amazon RDS Database Instance for Ranger: Prerequisites”.

Procedure

1. On the Oracle host, install the appropriate JDBC .jar file.
 - a) Download the Oracle JDBC (OJDBC) driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.
 - For Oracle Database 11g: select Oracle Database 11g Release 2 drivers > `ojdbc6.jar`.
 - For Oracle Database 12c: select Oracle Database 12c Release 1 driver > `ojdbc7.jar`.
 - b) Copy the .jar file to the Java share directory.


```
cp ojdbc7.jar /usr/share/java/
```

Make sure the .jar file has the appropriate permissions. For example: `chmod 644 /usr/share/java/ojdbc7.jar`

2. The Oracle database administrator should be used to create the Ranger databases. The following series of commands could be used to create the `RANGERDBA` user and grant it permissions using `SQL*Plus`, the Oracle database administration utility:

```
# sqlplus sys/root as sysdba
CREATE USER $RANGERDBA IDENTIFIED BY $RANGERDBAPASSWORD;
GRANT SELECT_CATALOG_ROLE TO $RANGERDBA;
GRANT CONNECT, RESOURCE TO $RANGERDBA;
QUIT;
```

3. Set the `jdbc/driver/path` based on the location of the Oracle JDBC driver .jar file. This command must be run on the server where Ambari server is installed: `ambari-server setup --jdbc-db={database-type} --jdbc-driver={/jdbc/driver/path}`.

```
ambari-server setup --jdbc-db=oracle --jdbc-driver=/usr/share/java/
ojdbc6.jar
```

What to do next

“Start the Ranger Installation”

Configure a Ranger DB: Amazon RDS

Ranger requires a relational database as its policy store. There are additional prerequisites for Amazon RDS-based databases due to how Amazon RDS is set up and managed.

About this task

Depending on your DB flavor, there are prerequisite steps to perform when using Amazon RDS. There are three ways to set up Hadoop group mapping:

- MySQL/MariaDB
- PostgreSQL
- Oracle

Procedure

- For MySQL/MariaDB, you must change the variable `log_bin_trust_function_creators` to 1 during Ranger installation. From RDS Dashboard>Parameter group (on the left side of the page):
 - a) Set the MySQL Server variable `log_bin_trust_function_creators` to 1.
 - b) (Optional) After Ranger installation is complete, reset `log_bin_trust_function_creators` to its original setting. The variable is only required to be set to 1 during Ranger installation.

- For PostgreSQL, complete the prerequisites:

The Ranger database user in Amazon RDS PostgreSQL Server should be created before installing Ranger and should be granted an existing role which must have the role `CREATEDB`.

- a) Using the master user account, log in to the Amazon RDS PostgreSQL Server from master user account (created during RDS PostgreSQL instance creation) and execute following commands:

```
CREATE USER $rangerdbuser WITH LOGIN PASSWORD 'password'
GRANT $rangerdbuser to $postgresroot
```

Where `$postgresroot` is the RDS PostgreSQL master user account (for example: `postgresroot`) and `$rangerdbuser` is the Ranger database user name (for example: `rangeradmin`).

- b) If you are using Ranger KMS, execute the following commands:

```
CREATE USER $rangerkmsuser WITH LOGIN PASSWORD 'password'
GRANT $rangerkmsuser to $postgresroot
```

Where `$postgresroot` is the RDS PostgreSQL master user account (for example: `postgresroot`) and `$rangerkmsuser` is the Ranger KMS user name (for example: `rangerkms`).

- For Oracle, due to [limitations in Amazon RDS](#), the Ranger database user and tablespace must be created manually and the required privileges must be manually granted to the Ranger database user:
 - a) Log in to the RDS Oracle Server from the master user account (created during RDS Oracle instance creation) and execute following commands:

```
create user $rangerdbuser identified by "password";
GRANT CREATE SESSION,CREATE PROCEDURE,CREATE TABLE,CREATE
VIEW,CREATE SEQUENCE,CREATE PUBLIC SYNONYM,CREATE ANY SYNONYM,CREATE
TRIGGER,UNLIMITED Tablespace TO $rangerdbuser;
create tablespace $rangerdb datafile size 10M autoextend on;
alter user $rangerdbuser DEFAULT Tablespace $rangerdb;
```

Where `$rangerdb` is a actual Ranger database name (for example: `ranger`) and `$rangerdbuser` is Ranger database username (for example: `rangeradmin`).

- b) If you are using Ranger KMS, execute the following commands:

```
create user $rangerdbuser identified by "password";
GRANT CREATE SESSION,CREATE PROCEDURE,CREATE TABLE,CREATE
VIEW,CREATE SEQUENCE,CREATE PUBLIC SYNONYM,CREATE ANY SYNONYM,CREATE
TRIGGER,UNLIMITED Tablespace TO $rangerkmsuser;
create tablespace $rangerkmsdb datafile size 10M autoextend on;
alter user $rangerkmsuser DEFAULT Tablespace $rangerkmsdb;
```

Where \$rangerkmsdb is a actual Ranger database name (for example: rangerkms) and \$rangerkmsuser is Ranger database username (for example: rangerkms).

What to do next

“Start the Ranger Installation”

Start the Ranger Installation

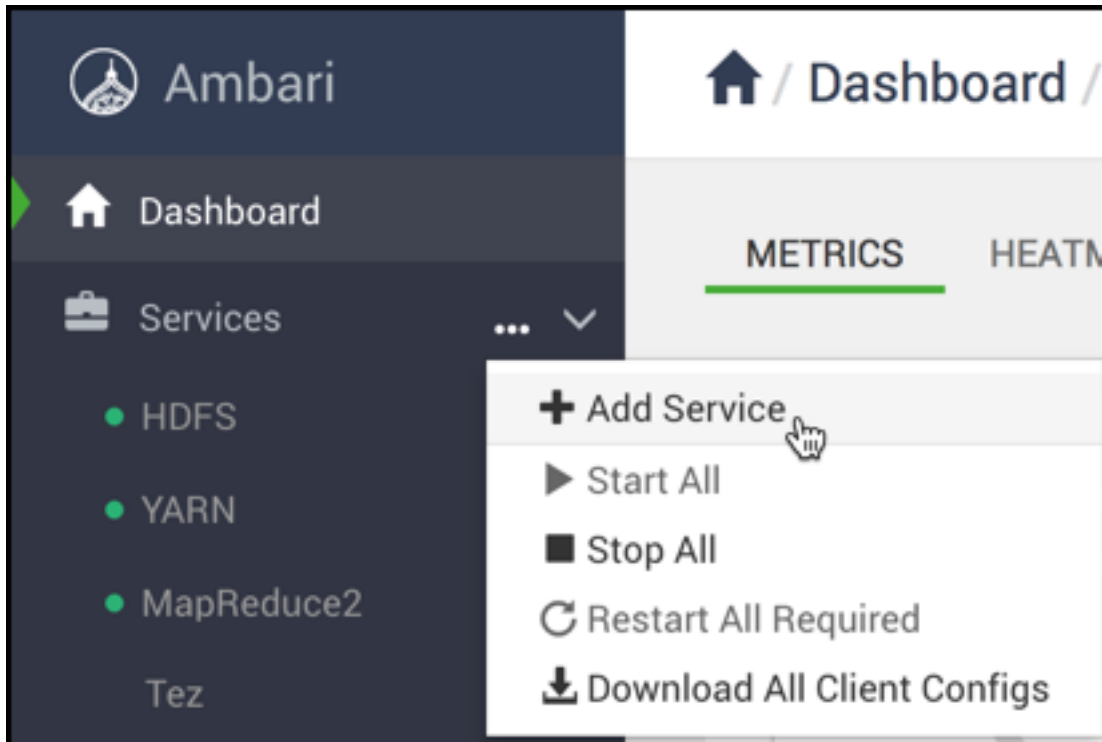
How to begin installing Ranger via Ambari.

Before you begin

You must have configured a database instance for Ranger.

Procedure

1. Log into your Ambari cluster with your designated user credentials.
The main Ambari Dashboard page will be displayed.
2. In the left navigation menu, click Actions, then select Add Service.



3. On the Choose Services page, select Ranger, then click Next.

<input checked="" type="checkbox"/>	Ambari Metrics	0.1.0	A system for metrics collection that provides storage and retrieval capability for metrics collected from the cluster
<input checked="" type="checkbox"/>	Atlas	1.0.0	Atlas Metadata and Governance platform
<input checked="" type="checkbox"/>	Kafka	1.0.1	A high-throughput distributed messaging system
<input checked="" type="checkbox"/>	Knox	1.0.0	Provides a single point of authentication and access for Apache Hadoop services in a cluster
<input type="checkbox"/>	Log Search	0.5.0	Log aggregation, analysis, and visualization for Ambari managed services. This service is Technical Preview .
<input checked="" type="checkbox"/>	Ranger	1.0.0	Comprehensive security for Hadoop
<input checked="" type="checkbox"/>	Ranger KMS	1.0.0	Key Management Server
<input type="checkbox"/>	SmartSense	1.5.0.2.7.0.0-703	SmartSense - Hortonworks SmartSense Tool (HST) helps quickly gather configuration, metrics, logs from common HDP services that aids to quickly troubleshoot support cases and receive cluster-specific recommendations.
<input type="checkbox"/>	Spark2	2.3.1	Apache Spark 2.3 is a fast and general engine for large-scale data processing.
<input type="checkbox"/>	Zeppelin Notebook	0.8.0	A web-based notebook that enables interactive data analytics. It enables you to make beautiful data-driven, interactive and collaborative documents with SQL, Scala and more.
<input type="checkbox"/>	Druid	0.12.1	A fast column-oriented distributed data store.
<input type="checkbox"/>	Superset	0.15.0	Superset is a data exploration platform designed to be visual, intuitive and interactive. This service is Technical Preview .

The Ranger Requirements page appears.

- If you have not already done so, run `ambari-server setup --jdbc-db=$database-type --jdbc-driver=$jdbc/driver/path`
`ambari-server setup --jdbc-db=mysql --jdbc-driver=/usr/share/java/mysql-connector-java.jar`
- You are prompted to Assign Masters. Make a note of the Ranger Admin host for use in subsequent installation steps. Click Next when finished to continue with the installation.

The screenshot shows the 'Add Service Wizard' interface. On the left, a sidebar lists the steps: 1. Choose Services, 2. Assign Masters (current step), 3. Assign Slaves and Clients, 4. Customize Services, 5. Review, 6. Install, Start and Test, 7. Summary. The main area is titled 'Assign Masters' and contains a list of services with dropdown menus for assigning master components. The services listed are: NameNode, SNameNode, ResourceManager, Timeline Service V2.0 Reader, Timeline Service V1.5, YARN Registry DNS, and History Server. All dropdowns are set to 'dw-weekly.field.hortonworks.com (29.3 GB, 16 ...)'. On the right, a preview window shows a grid of service icons, with 'Ranger Admin' and 'Ranger Usersync' highlighted in green.

The Customize Services page appears. These settings are described in the Ranger Installation: Customize Services section.

- On the **Assign Slaves and Clients** page, click **Next**.

What to do next

“Customize Services: Admin”

Related Information

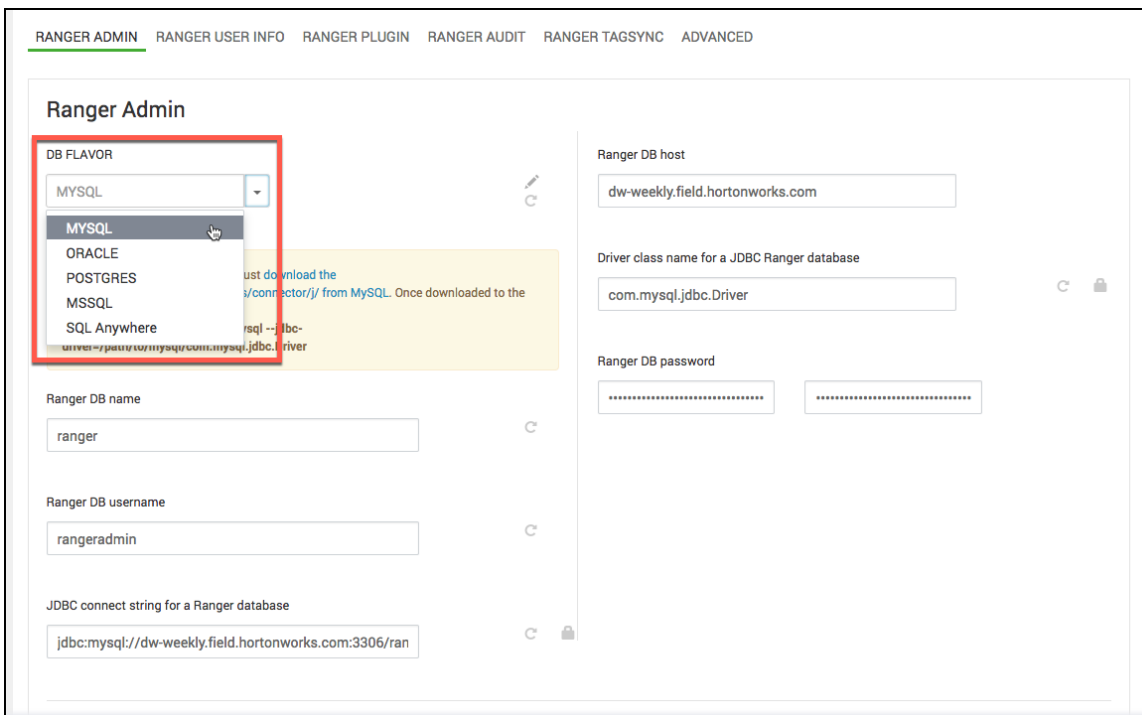
[Configuring a Database Instance for Ranger](#)

Customize Services: Admin

How to customize the Ranger Admin service when installing Ranger via Ambari.

Procedure

1. On the Customize Services page, select the Ranger Admin tab, then use the DB Flavor drop-down to select the database type that you are using with Ranger.



2. Enter the database server address in the Ranger DB Host box:

Table 1: Ranger DB Host

DB Flavor	Host	Example
MySQL	<HOST[:PORT]>	c6401.ambari.apache.org or c6401.ambari.apache.org:3306
Oracle	<HOST:PORT:SID>	c6401.ambari.apache.org:1521:ORCL
	<HOST:PORT/Service>	c6401.ambari.apache.org:1521/XE
PostgreSQL	<HOST[:PORT]>	c6401.ambari.apache.org or c6401.ambari.apache.org:5432
MS SQL	<HOST[:PORT]>	c6401.ambari.apache.org or c6401.ambari.apache.org:1433
SQLA	<HOST[:PORT]>	c6401.ambari.apache.org or c6401.ambari.apache.org:2638

3. Ranger DB name -- The name of the Ranger Policy database, e.g. ranger_db or ranger. Please note that if you are using Oracle, you must specify the Oracle tablespace name here.
4. Driver class name for a JDBC Ranger database -- the driver class name is automatically generated based on the selected DB Flavor. The table below lists the default driver class settings. Currently Ranger does not support any third party JDBC driver.

Table 2: Driver Class Name

DB Flavor	Driver class name for a JDBC Ranger database
MySQL	com.mysql.jdbc.Driver
Oracle	oracle.jdbc.driver.OracleDriver
PostgreSQL	org.postgresql.Driver
MS SQL	com.microsoft.sqlserver.jdbc.SQLServerDriver
SQLA	sap.jdbc4.sqlanywhere.IDriver

5. Ranger DB username and Ranger DB Password -- Enter the user name and passwords for your Ranger database server. The following table describes these settings in more detail. You can use the MySQL database that was installed with Ambari, or an external MySQL, Oracle, PostgreSQL, MS SQL or SQL Anywhere database.

Table 3: Ranger DB Username Settings

Property	Description	Default Value	Example Value	Required?
Ranger DB username	The username for the Policy database.	rangeradmin	rangeradmin	Yes
Ranger DB password	The password for the Ranger Policy database user.		PassWORD!	Yes

6. JDBC connect string

Note:

Currently the Ambari installer generates the JDBC connect string using the jdbc:oracle:thin:@//host:port/db_name format. You must replace the connection string as described in the following table:

Table 4: JDBC Connect String

DB Flavor	Syntax	Example Value
MySQL	jdbc:mysql://DB_HOST:PORT/db_name	jdbc:mysql://c6401.ambari.apache.org:3306/ranger_db
Oracle	For Oracle SID: jdbc:oracle:thin:@DB_HOST:PORT:SID	jdbc:oracle:thin:@c6401.ambari.apache.org:1521:ORCL
	For Oracle Service Name: jdbc:oracle:thin:@//DB_HOST[:PORT][//ServiceName]	jdbc:oracle:thin:@//c6401.ambari.apache.org:1521/XE
PostgreSQL	jdbc:postgresql://DB_HOST/db_name	jdbc:postgresql://c6401.ambari.apache.org:5432/ranger_db
MS SQL	jdbc:sqlserver://DB_HOST;databaseName=db_name	jdbc:sqlserver://c6401.ambari.apache.org:1433;databaseName=ranger_db
SQLA	jdbc:sqlanywhere:host=DB_HOST;database=db_name	jdbc:sqlanywhere:host=c6401.ambari.apache.org:2638;database=ranger_db

7. Setup Database and Database User:

Option**Yes**

If set to Yes -- The Database Administrator (DBA) user name and password will need to be provided as described in the next step.

Note:

Option

Ranger does not store the DBA user name and password after setup. Therefore, you can clear these values in the Ambari UI after the Ranger setup is complete.

No

If set to No -- A No indicates that you do not wish to provide Database Administrator (DBA) account details to the Ambari Ranger installer. Setting this to No continues the Ranger installation process without providing DBA account details. In this case, you must perform the system database user setup as described in “Set up Database Users Without Sharing DBA Credentials”, and then proceed with the installation.

Note:

If No is selected and the UI still requires you to enter a user name and password in order to proceed, you can enter any value -- the values do not need to be the actual DBA user name and password.

8. Database Administrator (DBA) username and Database Administrator (DBA) password -- The DBA username and password are set when the database server is installed. If you do not have this information, contact the database administrator who installed the database server.

Table 5: DBA Credential Settings

Property	Description	Default Value	Example Value	Required?
Database Administrator (DBA) username	The Ranger database user that has administrative privileges to create database schemas and users.	root	root	Yes
Database Administrator (DBA) password	The root password for the Ranger database user.		root	Yes

If the Oracle DB root user Role is SYSDBA, you must also specify that in the Database Administrator (DBA) username parameter. For example, if the DBA user name is orcl_root you must specify orcl_root AS SYSDBA.

Note:

As mentioned in the note in the previous step, if Setup Database and Database User is set to No, a placeholder DBA username and password may still be required in order to continue with the Ranger installation.

The following images show examples of the DB settings for each Ranger database type.

MySQL

Ranger Admin [Ranger User Info](#) [Ranger Plugin](#) [Ranger Audit](#) [Ranger Tagsync](#) [Advanced](#)

Ranger Admin

DB FLAVOR
MYSQL

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
jdbc:mysql://c6402.ambari.apache.org:330

Ranger DB host
c6402.ambari.apache.org

Driver class name for a JDBC Ranger database
com.mysql.jdbc.Driver

Ranger DB password

Setup Database and Database User
 Yes

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password

JDBC connect string for root user
jdbc:mysql://c6402.ambari.apache.org:330

Oracle -- if the Oracle instance is running with a Service name.

Ranger Admin [Ranger User Info](#) [Ranger Plugin](#) [Ranger Audit](#) [Ranger Tagsync](#) [Advanced](#)

Ranger Admin

DB FLAVOR
ORACLE ▾

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
//c6402.ambari.apache.org:1521/XE/ranger

Ranger DB host
c6402.ambari.apache.org:1521/XE

Driver class name for a JDBC Ranger database
oracle.jdbc.driver.OracleDriver

Ranger DB password

Setup Database and Database User
 Yes

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password

JDBC connect string for root user
cthin:@//c6402.ambari.apache.org:1521/XE

Oracle -- if the Oracle instance is running with a SID.

Ranger Admin [Ranger User Info](#) [Ranger Plugin](#) [Ranger Audit](#) [Ranger Tagsync](#) [Advanced](#)

Ranger Admin

DB FLAVOR
ORACLE

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
jdbc:oracle:thin:@//c6402.ambari.apache.org:1521:ORCL

Ranger DB host
c6402.ambari.apache.org:1521:ORCL

Driver class name for a JDBC Ranger database
oracle.jdbc.driver.OracleDriver

Ranger DB password

Setup Database and Database User
 Yes

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password

JDBC connect string for root user
jdbc:oracle:thin:@//c6402.ambari.apache.org:1521:ORCL

PostgreSQL

Ranger Admin [Ranger User Info](#) [Ranger Plugin](#) [Ranger Audit](#) [Ranger Tagsync](#) [Advanced](#)

Ranger Admin

DB FLAVOR
POSTGRES ▾

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
sql://c6402.ambari.apache.org:5432/ranger

Ranger DB host
c6402.ambari.apache.org:5432

Driver class name for a JDBC Ranger database
org.postgresql.Driver

Ranger DB password

Setup Database and Database User
 Yes

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password

JDBC connect string for root user
t://c6402.ambari.apache.org:5432/postgres

MS SQL

Ranger Admin Ranger User Info Ranger Plugin Ranger Audit Ranger Tagsync Advanced

Ranger Admin

DB FLAVOR
MSSQL

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
ari.apache.org:1433;databaseName=ranger

Ranger DB host
c6402.ambari.apache.org:1433

Driver class name for a JDBC Ranger database
com.microsoft.sqlserver.jdbc.SQLServerC

Ranger DB password

Setup Database and Database User
 Yes

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password

JDBC connect string for root user
sqlserver://c6402.ambari.apache.org:1433;

SQL Anywhere

Ranger Admin [Ranger User Info](#) [Ranger Plugin](#) [Ranger Audit](#) [Ranger Tagsync](#) [Advanced](#)

Ranger Admin

DB FLAVOR
SQL Anywhere ▾

Ranger DB name
ranger

Ranger DB username
rangeradmin

JDBC connect string for a Ranger database
!.ambari.apache.org:2638;database=ranger

Ranger DB host
c6402.ambari.apache.org:2638

Driver class name for a JDBC Ranger database
sap.jdbc4.sqlanywhere.IDriver

Ranger DB password
***** *****

Setup Database and Database User
 Yes ✎ 🔒

Database Administrator (DBA) username
rangerdba

Database Administrator (DBA) password
***** *****

JDBC connect string for root user
/here:host=c6402.ambari.apache.org:2638;

9. To test the DB settings, click Test Connection. If a Ranger database has not been pre-installed, Test Connection will fail even for a valid configuration.

What to do next

“Customize Services: Audit”

Related Information

[Set up Database Users Without Sharing DBA Credentials](#)

Customize Services: Audit

How to customize the Ranger Audit service when installing Ranger via Ambari.

About this task

Apache Ranger uses Apache Solr to store audit logs and provides UI searching through the audit logs. Solr must be installed and configured before installing Ranger Admin or any of the Ranger component plugins. The default configuration for Ranger Audits to Solr uses the shared Solr instance provided under the Ambari Infra service. Solr is both memory and CPU intensive. If your production system has high volume of access requests, make sure that the Solr host has adequate memory, CPU, and disk space.

SolrCloud is the preferred setup for production usage of Ranger. SolrCloud, which is deployed with the Ambari Infra service, is a scalable architecture that can run as a single node or multi-node cluster. It has additional features such as replication and sharding, which is useful for high availability (HA) and scalability. You should plan your deployment based on your cluster size. Because audit records can grow dramatically, plan to have at least 1 TB of free space in the volume on which Solr will store the index data. Solr works well with a minimum of 32 GB of RAM. You should provide as much memory as possible to the Solr process. It is highly recommended to use SolrCloud with at least two Solr nodes running on different servers with replication (CCDR) enabled. SolrCloud also requires Apache ZooKeeper.

It is recommended that you store audits in both HDFS and Solr. The default configuration for Ranger Audits to Solr uses the shared Solr instance provided under the Ambari Infra service. For more information about Audits to Solr, see [Using Apache Solr for Ranger Audits](#).

Procedure

1. On the Customize Services page, select the Ranger Audit tab.

It is recommended that you store audits in Solr and HDFS. Both of these options are set to ON by default. Solr provides the capability to index and search on the most recent logs while HDFS is used as the more permanent or longer term store. By default, Solr is used to index the preceding 30 days of audit logs.

2. Under Audit to Solr, turn ON SolrCloud.

The SolrCloud configuration settings will be loaded automatically.

Audit to Solr

Audit to Solr

ON

SolrCloud

ON

External SolrCloud

OFF

ranger.audit.solr.zookeepers

ranger.audit.solr.username

ranger.audit.solr.password

Audit to HDFS

Audit to HDFS

ON

Destination HDFS Directory

What to do next

“Customize Services: Plugins”

Related Information

[Using Ambari Core Services>Ambari Infra](#)

[SolrCloud](#)

[Cross Data Center Replication \(CDCR\)](#)

Customize Services: Plugins

How to enable Ranger Plugins when installing Ranger via Ambari.

About this task

If you are using a Kerberos-enabled cluster, there are a number of additional steps you must follow to ensure that you can use the Ranger plugins on a Kerberos cluster.

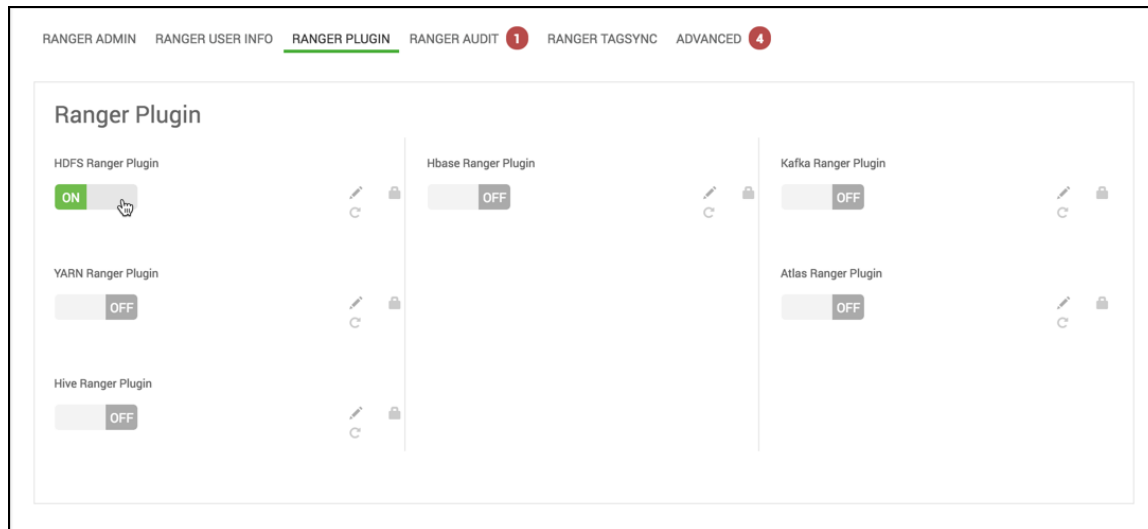
The following Ranger plugins are available:

- HDFS
- Hive
- HBase
- Kafka

- Knox
- YARN
- Storm
- Atlas
- Solr

Procedure

1. From the **Ranger Plugin** tab, turn **On** the plugins you want.



For every plugin you enable, you will have to restart the associated component. E.G., if you enable the HDFS Ranger plugin, you will have to restart HDFS.

2. Click **Next**.

What to do next

“Customize Services: User Sync”

Customize Services: User Sync

How to customize Ranger User Sync for either UNIX or LDAP/AD when installing Ranger via Ambari.

About this task

You can customize User Sync for either Unix or LDAP/AD.

Before you begin

Test Run User Sync

Before committing to usersync changes, it is recommended that you test-run that users and groups are being retrieved as intended.

To test-run loading User and Group data into Ranger before committing to the changes:

1. Set `ranger.usersync.policymanager.mockrun=true`. This parameter can be found in Ambari> Ranger> Configs> Advanced> Advanced ranger-ugsync-site.
2. View the Users and Groups that will be loaded into Ranger: `tail -f /var/log/ranger/usersync/usersync.log`.
3. After confirming that the users and groups are retrieved as intended, set `ranger.usersync.policymanager.mockrun=false` and restart Ranger Usersync.

This will sync the users shown in the usersync log to the Ranger database.

Procedure

- Customize User Sync for UNIX:
 - a) On the Customize Services page, select the Ranger User Info tab.
 - b) Click Yes under Enable User Sync.
 - c) Use the Sync Source drop-down to select UNIX, then set the following properties.

Table 6: UNIX User Sync Properties

Property	Description	Default Value
Sync Source	Only sync users above this user ID.	500
Password File	The location of the password file on the Linux server.	/etc/passwd
Group File	The location of the groups file on the Linux server.	/etc/group

The screenshot shows the 'Add Service Wizard' interface. On the left is a sidebar with navigation options: 'Choose Services', 'Assign Masters', 'Assign Slaves and Clients', 'Customize Services' (highlighted), 'Configure Identities', 'Review', 'Install, Start and Test', and 'Summary'. The main area is titled 'Customize Services' and contains a list of services: HDFS, YARN, MapReduce2, Tez, Hive, HBase, Pig, Oozie, ZooKeeper, Falcon, Storm, Ambari Metrics, Atlas, Kafka, Knox, Ranger (selected), Slider, SmartSense, Spark, and Misc. Below this, it states 'There are 43 configuration changes in 8 services' with a 'Show Details' link. A 'Group' dropdown is set to 'Default (1)' and a 'Filter...' input is present. The 'Ranger User Info' tab is active, showing 'Enable User Sync' as 'Yes', 'Sync Source' as 'UNIX', 'Minimum User ID' as 500, 'Password File' as /etc/passwd, and 'Group File' as /etc/group.

- Customize User Sync for LDAP/AD:
 - a) On the Customize Services page, select the Ranger User Info tab.
 - b) Click Yes under Enable User Sync.
 - c) Use the Sync Source drop-down to select LDAP/AD.
 - d) Set the following properties on the Common Configs tab.

Table 7: LDAP/AD Common Configs

Property	Description	Default Value	Sample Values
LDAP/AD URL	Add URL depending upon LDAP/AD sync source	ldap://{host}:{port}	ldap://ldap.example.com:389 or ldaps://ldap.example.com:636
Bind Anonymous	If Yes is selected, the Bind User and Bind User Password are not required.	NO	
Bind User	The location of the groups file on the Linux server.	The full distinguished name (DN), including common name (CN), of an LDAP/AD user account that has privileges to search for users. The LDAP bind DN is used to connect to LDAP and query for users and groups.	cn=admin,dc=example,dc=com or admin@example.com
Bind User Password	The password of the Bind User.		
Incremental Sync	<p>If Yes is selected, Ranger Usersync saves the latest timestamp of all the objects that are synced previously and uses that timestamp to perform the next sync. Usersync then uses a polling mechanism to perform incremental sync by using LDAP attributes uSNChanged (for AD) or modifytimestamp (for LDAP).</p> <p>Enabling Incremental Sync for the first time will initiate a full sync; subsequent sync operations will be incremental.</p> <p>When Incremental Sync is enabled, Group Sync (under the Group Configs tab) is mandatory.</p> <p>Recommended for large deployments.</p>	<p>For upgrade installations: No</p> <p>For new installations: Yes</p>	Yes

Ranger Admin **Ranger User Info** Ranger Plugin Ranger Audit Ranger Tagsync Advanced

Ranger User Info

Enable User Sync
 Yes

Sync Source
 LDAP/AD

Common Configs **User Configs** Group Configs

LDAP/AD URL

Bind User

Bind User Password

Incremental Sync
 Yes

e) Set the following properties on the User Configs tab.

Table 8: LDAP/AD User Configs

Property	Description	Default Value	Sample Values
Group User Map Sync	Sync specific groups for users.	Yes	Yes
Username Attribute	The LDAP user name attribute.		sAMAccountName for AD, uid or cn for OpenLDAP
User Object Class	Object class to identify user entries.	person	top, person, organizationalPerson, user, or posixAccount

Property	Description	Default Value	Sample Values
User Search Base	Search base for users. Ranger can search multiple OUs in AD. Ranger UserSync module performs a user search on each configured OU and adds all the users into single list. Once all the OUs are processed, a user's group membership is computed based on the group search.		cn=users,dc=example,dc=com;ou=example1,ou=example2
User Search Filter	Optional additional filter constraining the users selected for syncing.		Sample filter to retrieve all the users: cn=* Sample filter to retrieve all the users who are members of groupA or groupB: (&((memberof=CN=GroupA,OU=groups,DC=example)(memberof=CN=GroupB,OU=groups,DC=example))
User Search Scope	This value is used to limit user search to the depth from search base.	sub	base, one, or sub
User Group Name Attribute	Attribute from user entry whose values would be treated as group values to be pushed into the Access Manager database. You can provide multiple attribute names separated by commas.	memberof,ismemberof	memberof, ismemberof, or gidNumber
Enable User Search	This option is available only when the "Enable Group Search First" option is selected.	No	Yes

Ranger User Info

Enable User Sync
 Yes

Sync Source
LDAP/AD

[Common Configs](#) [User Configs](#) [Group Configs](#)

Username Attribute
sAMAccountName

User Object Class
person

User Search Base
cn=users,dc=example,dc=com;ou=example1,ou=example2

User Search Filter

User Search Scope
sub

User Group Name Attribute
memberof, ismemberof

Group User Map Sync
 Yes

f) Set the following properties on the Group Configs tab.


Table 9: LDAP/AD Group Configs


Property	Description	Default Value	Sample Values
Enable Group Sync	<p>If Enable Group Sync is set to No, the group names the users belong to are derived from "User Group Name Attribute". In this case no additional group filters are applied.</p> <p>If Enable Group Sync is set to Yes, the groups the users belong to are retrieved from LDAP/AD using the following group-related attributes.</p> <p>Enabled by default if "Incremental Sync" is enabled under the Common Configs tab.</p>	No	Yes
Group Member Attribute	The LDAP group member attribute name.		member
Group Name Attribute	The LDAP group name attribute.		distinguishedName for AD, cn for OpenLdap
Group Object Class	LDAP Group object class.		group, groupofnames, or posixGroup
Group Search Base	<p>Search base for groups.</p> <p>Ranger can search multiple OUs in AD. Ranger UserSync module performs a user search on each configured OU and adds all the users into single list. Once all the OUs are processed, a user's group membership is computed based on the group search configuration. Each OU segment needs to be separated by a ; (semi-colon).</p>		ou=groups,DC=example,DC=com;ou=group1;ou=group2
Group Search Filter	Optional additional filter constraining the groups selected for syncing.		<p>Sample filter to retrieve all groups: cn=*</p> <p>Sample filter to retrieve only the groups whose cn is Engineering or Sales: ((cn=Engineering) (cn=Sales))</p>
Enable Group Search First	<p>When Enable Group Search First is selected, there are two possible ways of retrieving users:</p> <ul style="list-style-type: none"> If Enable User Search is not selected: users are retrieved from the "member" attribute of the group. If Enable User Search is selected: user membership is computed by performing an LDAP search based on the user configuration. 	No	Yes


Property	Description	Default Value	Sample Values
<p>state: anchor=nested_ldapad_syncSync Nested Groups</p>	<p>Enables nested group memberships in Ranger so that the policies configured for parent groups are applied for all the members in the subgroups.</p> <p>If a group itself is a member of another group, the users belonging to the member group are part of the parent group as well.</p> <p>Group Hierarchy Levels determines evaluated nested group depth.</p> <p>If you do not see the Sync Nested Groups flag, upgrade to Ambari 2.6.0+.</p>	No	Yes, No
Group Hierarchy Levels	<p>Determines how many nested groups to evaluate in support of Sync Nested Groups.</p> <p>If Group Hierarchy Levels is greyed out, enable Sync Nested Groups.</p> <p>Set to any integer >0.</p>	0	2


[Common Configs](#) [User Configs](#) [Group Configs](#)


Enable Group Sync
 Yes


Group Member Attribute
 


Group Name Attribute
 


Group Object Class
 

Group Search Base
 

Group Search Filter
 

Enable Group Search First
 Yes 

Sync Nested Groups
 Yes 

Group Hierarchy Levels
 

What to do next

“Customize Services: Tagsync”


Related Information

[Set Up Hadoop Group Mapping for LDAP/AD](#)

Customize Services: Tagsync

How to customize the Ranger Tagsync service when installing Ranger via Ambari.

About this task



Procedure

1. To configure Ranger Tagsync, select Ranger Tagsync on the Customize Services page, then specify a Tagsync source.

It is recommended that you only configure Atlas Tag Source. Configuring File Tag Source or Atlas REST Tag Source is generally not required, and should only be attempted by advanced users.

2. Configure Atlas Tag Source Properties:

Table 10: Atlas Tag Source Properties

Property	Description
Atlas Source: Kafka endpoint	The Kafka endpoint: <kafka_server_url>:6667
Atlas Source: ZooKeeper endpoint	The ZooKeeper endpoint: <zookeeper_server_url>:2181

Property	Description
Atlas Source: Kafka consumer group	The Ranger entities consumer.

What to do next

“Customize Services: Authentication”

Customize Services: Authentication

This section describes how to configure Ranger authentication for UNIX, LDAP, and AD.

Customize Authentication: UNIX

How to customize the Ranger UNIX Authentication service when installing Ranger via Ambari.

Procedure

1. Select the Advanced tab on the Customize Services page.
2. Under Ranger Settings, specify the Ranger Access Manager/Service Manager host address in the External URL box in the format `http://<your_ranger_host>:6080`.
3. Under Ranger Settings, select UNIX.
HTTP is enabled by default -- if you disable HTTP, only HTTPS is allowed.
4. Under UNIX Authentication Settings, set the following properties.

Table 11: UNIX Authentication Settings

Property	Description	Default Value	Example Value
Allow remote Login	Flag to enable/disable remote login. Only applies to UNIX authentication.	true	true
ranger.unixauth.service.hostname	The address of the host where the UNIX authentication service is running.	{{ugsync_hos}}	{{ugsync_host}}
ranger.unixauth.service.port	The port number on which the UNIX authentication service is running.	5151	5151

Note:

Properties with value `{{xyz}}` are macro variables that are derived from other specified values in order to streamline the configuration process. Macro variables can be edited if required -- if you need to restore the original value, click the Set Recommended symbol at the right of the property box.

The screenshot shows the 'Add Service Wizard' window with the following settings:

- Ranger Settings:**
 - External URL:
 - Authentication method:
 - LDAP
 - ACTIVE_DIRECTORY
 - UNIX
 - NONE
 - HTTP enabled:
- Unix Authentication Settings:**
 - Allow remote Login:
 - ranger.unixauth.service.hostname:
 - ranger.unixauth.service.port:
- Knox SSO Settings:** (Collapsed)
- Advanced ranger-admin-site:** (Collapsed)

What to do next

“Complete the Ranger Installation”

Customize Authentication: LDAP

How to customize the Ranger LDAP Authentication service when installing Ranger via Ambari.

Procedure

1. Select the Advanced tab on the Customize Services page.
2. Under Ranger Settings, specify the Ranger Access Manager/Service Manager host address in the External URL box in the format `http://<your_ranger_host>:6080`.
3. Under Ranger Settings, select LDAP.
4. Under LDAP Settings, set the following properties.

Table 12: LDAP Authentication Settings

Property	Description	Default Value	Example Value
ranger ldap.base.dn	The Distinguished Name (DN) of the starting point for directory server searches.	dc=example,dc=com	dc=example,dc=com

Property	Description	Default Value	Example Value
Bind User	The full Distinguished Name (DN), including Common Name (CN) of an LDAP user account that has privileges to search for users. This is a macro variable value that is derived from the Bind User value from Ranger User Info > Common Configs.	{{ranger_ug_ldap_bind_dn}}	{{ranger_ug_ldap_bind_dn}}
Bind User Password	Password for the Bind User. This is a macro variable value that is derived from the Bind User Password value from Ranger User Info > Common Configs.		
ranger.ldap.group.roleattribute	The LDAP group role attribute.	cn	cn
ranger.ldap.referral	See description below.	ignore	follow ignore throw
LDAP URL	The LDAP server URL. This is a macro variable value that is derived from the LDAP/AD URL value from Ranger User Info > Common Configs.	{{ranger_ug_ldap_url}}	{{ranger_ug_ldap_url}}
ranger.ldap.user.dnpattern	The user DN pattern is expanded when a user is being logged in. For example, if the user "ldapadmin" attempted to log in, the LDAP Server would attempt to bind against the DN "uid=ldapadmin,ou=users,dc=example,dc=com" using the password the user provided>	uid={0},ou=users,dc=xasecure,dc=net	cn=ldapadmin,ou=Users,dc=example,dc=com
User Search Filter	The search filter used for Bind Authentication. This is a macro variable value that is derived from the User Search Filter value from Ranger User Info > User Configs.	{{ranger_ug_ldap_user_searchfilter}}	{{ranger_ug_ldap_user_searchfilter}}

Note:

Properties with value {{xyz}} are macro variables that are derived from other specified values in order to streamline the configuration process. Macro variables can be edited if required -- if you need to restore the original value, click the Set Recommended symbol at the right of the property box.

There are three possible values for `ranger.ldap.referral`: follow, throw, and ignore. The recommended setting is follow.

When searching a directory, the server might return several search results, along with a few continuation references that show where to obtain further results. These results and references might be interleaved at the protocol level.

- When this property is set to follow, the LDAP service provider processes all of the normal entries first, and then follows the continuation references.
- When this property is set to throw, all of the normal entries are returned in the enumeration first, before the `ReferralException` is thrown. By contrast, a "referral" error response is processed immediately when this property is set to follow or throw.
- When this property is set to ignore, it indicates that the server should return referral entries as ordinary entries (or plain text). This might return partial results for the search.

Add Service Wizard

Ranger Settings

External URL:

Authentication method: LDAP
 ACTIVE_DIRECTORY
 UNIX
 NONE

HTTP enabled:

LDAP Settings

ranger.ldap.base.dn:

Bind User:

Bind User Password:

ranger.ldap.group.roleattribute:

ranger.ldap.referral:

LDAP URL:

ranger.ldap.user.dnpattern:

User Search Filter:

▸ Knox SSO Settings

▸ Advanced ranger-admin-site

▸ Advanced ranger-env

▸ Advanced ranger-ugsync-site

▸ Custom admin-properties

What to do next

“Complete the Ranger Installation”

Customize Authentication: AD

How to customize the Ranger AD Authentication service when installing Ranger via Ambari.

Procedure

1. Select the Advanced tab on the Customize Services page.
2. Under Ranger Settings, specify the Ranger Access Manager/Service Manager host address in the External URL box in the format `http://<your_ranger_host>:6080`.
3. Under Ranger Settings, select `ACTIVE_DIRECTORY`.
4. Under AD Settings, set the following properties.

Table 13: AD Settings

Property	Description	Default Value	Example Value
<code>ranger.ldap.ad.base.dn</code>	The Distinguished Name (DN) of the starting point for directory server searches.	<code>dc=example,dc=com</code>	<code>dc=example,dc=com</code>
<code>ranger.ldap.ad.bind.dn</code>	The full Distinguished Name (DN), including Common Name (CN) of an LDAP user account that has privileges to search for users. This is a macro variable value that is derived from the Bind User value from Ranger User Info > Common Configs.	<code>{{ranger_ug_ldap_bind_dn}}</code>	<code>{{ranger_ug_ldap_bind_dn}}</code>
<code>ranger.ldap.ad.bind.password</code>	Password for the bind.dn. This is a macro variable value that is derived from the Bind User Password value from Ranger User Info > Common Configs.		
Domain Name (Only for AD)	The domain name of the AD Authentication service.		<code>dc=example,dc=com</code>
<code>ranger.ldap.ad.referral</code>	See description below.	<code>ignore</code>	<code>follow ignore throw</code>
<code>ranger.ldap.ad.url</code>	The AD server URL. This is a macro variable value that is derived from the LDAP/AD URL value from Ranger User Info > Common Configs.	<code>{{ranger_ug_ldap_url}}</code>	<code>{{ranger_ug_ldap_url}}</code>
<code>ranger.ldap.ad.user.searchfilter</code>	The search filter used for Bind Authentication. This is a macro variable value that is derived from the User Search Filter value from Ranger User Info > User Configs.	<code>{{ranger_ug_ldap_user_searchfilter}}</code>	<code>{{ranger_ug_ldap_user_searchfilter}}</code>

Note:

Properties with value `{{xyz}}` are macro variables that are derived from other specified values in order to streamline the configuration process. Macro variables can be edited if required -- if you need to restore the original value, click the Set Recommended symbol at the right of the property box.

There are three possible values for `ranger.ldap.ad.referral`: `follow`, `throw`, and `ignore`. The recommended setting is `follow`.

When searching a directory, the server might return several search results, along with a few continuation references that show where to obtain further results. These results and references might be interleaved at the protocol level.

- When this property is set to `follow`, the AD service provider processes all of the normal entries first, and then follows the continuation references.
- When this property is set to `throw`, all of the normal entries are returned in the enumeration first, before the `ReferralException` is thrown. By contrast, a "referral" error response is processed immediately when this property is set to `follow` or `throw`.

- When this property is set to ignore, it indicates that the server should return referral entries as ordinary entries (or plain text). This might return partial results for the search. In the case of AD, a PartialResultException is returned when referrals are encountered while search results are processed.

When you have finished configuring all of the Customize Services Settings, click Next at the bottom of the page to continue with the installation.

- When you save the authentication method as Active Directory, a Dependent Configurations pop-up may appear recommending that you set the authentication method as LDAP. This recommended configuration should not be applied for AD, so you should clear (un-check) the `ranger.authentication.method` check box, then click OK.

Property	Service	Config Group	File Name	Current Value	Recommended Value
<input type="checkbox"/> ranger.authentication.method	Ranger	Default	ranger-admin-site	UNIX	LDAP

- Update the Ranger admin truststore configuration:
 - In **Ambari > Ranger > Configs > Advanced > Advanced ranger-admin-site**, add the following parameters:

```
ranger.truststore.file=/etc/ranger/admin/truststore
ranger.truststore.password=password
```

- Restart Ranger.

What to do next

“Complete Ranger Installation”

Complete the Ranger Installation

How to finish installing Ranger via Ambari, after customizing services.

Procedure

1. On the Review page, carefully review all of your settings and configurations. If everything looks good, click Deploy to install Ranger on the Ambari server.

Add Service Wizard

ADD SERVICE WIZARD

- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Configure Identities
- Review**
- Install, Start and Test
- Summary

Review

Please review the configuration before installation

Admin Name : admin
 Cluster Name : Thomas1
 Total Hosts : 3 (0 new)

Repositories:

```

redhat5 (HDP-2.2):
http://public-repo-1.hortonworks.com/HDP/centos5/2.x/updates/2.2.6.0
redhat5 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos5
redhat6 (HDP-2.2):
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.2.6.0
redhat6 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos6
suse11 (HDP-2.2):
http://public-repo-1.hortonworks.com/HDP/suse11sp3/2.x/updates/2.2.6.0
suse11 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/suse11sp3
ubuntu12 (HDP-2.2):
http://public-repo-1.hortonworks.com/HDP/ubuntu12/2.x/updates/2.2.6.3
  
```

← Back Print Deploy →

2. When you click Deploy, Ranger is installed on the specified host on your Ambari server. A progress bar displays the installation progress.

Add Service Wizard

ADD SERVICE WIZARD

- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Configure Identities
- Review
- Install, Start and Test**
- Summary

Install, Start and Test

Please wait while the selected services are installed and started.

24 % overall

Show: All (3) | In Progress (3) | Warning (0) | Success (0) | Fail (0)

Host	Status	Message
c6401.ambari.apache.org	8%	Installing Ranger Admin
c6402.ambari.apache.org	33%	Install complete (Waiting to start)
c6403.ambari.apache.org	33%	Install complete (Waiting to start)

3 of 3 hosts showing - Show All Show: 25 | 1 - 3 of 3

Next →

- When the installation is complete, a Summary page displays the installation details. You may need to restart services for cluster components after installing Ranger.
If the installation fails, you should complete the installation process, then reconfigure and reinstall Ranger.

Additional Ranger Plugin Configuration Steps for Kerberos Clusters

If you are using a Kerberos-enabled cluster, there are a number of steps you need to follow in order use the following Ranger plugins on a Kerberos cluster

If you are using a Kerberos-enabled cluster, there are a number of steps you need to follow in order use the following Ranger plugins on a Kerberos cluster.

Note:

These procedures assume that you have already enabled the component Ranger plugins.

Additional Ranger Plugin Steps for Kerberos: HDFS

How to enable the Ranger HDFS plugin on a Kerberos cluster.

Before you begin

This procedure assumes that you have already completed “Customize Services: Plugins”.

Procedure

- Create the system (OS) user rangerhdfslookup. Make sure this user is synced to Ranger Admin (under Settings>Users/Groups tab in the Ranger Admin User Interface).
- Create a Kerberos principal for rangerhdfslookup: `kadmin.local -q 'addprinc -pw rangerhdfslookup rangerhdfslookup@example.com'`.

Note:

A single user/principal (e.g., rangerrepouser) can also be created and used across services.

- Navigate to the HDFS service.
- Click the Config tab.
- Navigate to advanced ranger-hdfs-plugin-properties and update the properties listed in the table shown below.

Table 14: HDFS Plugin Properties

Configuration Property Name	Value
Ranger repository config user	rangerhdfslookup@example.com
Ranger repository config password	rangerhdfslookup
common.name.for.certificate	blank

6. After updating these properties, click Save and restart the HDFS service.

Additional Ranger Plugin Steps for Kerberos: Hive

How to enable the Ranger Hive plugin on a Kerberos cluster.

Before you begin

This procedure assumes that you have already completed “Customize Services: Plugins”.

Procedure

1. Create the system (OS) user rangerhivelookup. Make sure this user is synced to Ranger Admin (under Settings>Users/Groups tab in the Ranger Admin UI).
2. Create a Kerberos principal for rangerhivelookup: `kadmin.local -q 'addprinc -pw rangerhivelookup rangerhivelookup@example.com'`.
3. Navigate to the Hive service.
4. Click the Config tab and navigate to advanced ranger-hive-plugin-properties.
5. Update the following properties with the values listed in the table below.

Table 15: Hive Plugin Properties

Configuration Property Name	Value
Ranger service config user	rangerhivelookup@example.com
Ranger service config password	rangerhivelookup
common.name.for.certificate	blank

6. After updating these properties, click Save and then restart the Hive service.

Additional Ranger Plugin Steps for Kerberos: HBase

How to enable the Ranger HBase plugin on a Kerberos cluster.

Before you begin

This procedure assumes that you have already completed “Customize Services: Plugins”.

Procedure

1. Create the system (OS) user rangerhbaselookup. Make sure this user is synced to Ranger Admin (under users/groups tab in the Ranger Admin UI).
2. Create a Kerberos principal for rangerhbaselookup: `kadmin.local -q 'addprinc -pw rangerhbaselookup rangerhbaselookup@example.com'`.
3. Navigate to the HBase service.
4. Click the Config tab and go to advanced ranger-hbase-plugin-properties.
5. Update the following properties with the values listed in the table below.

Table 16: HBase Plugin Properties

Configuration Property Name	Value
Ranger service config user	rangerhbaselookup@example.com
Ranger service config password	rangerhbaselookup
common.name.for.certificate	blank

6. After updating these properties, click Save and then restart the HBase service.

Additional Ranger Plugin Steps for Kerberos: Knox

How to enable the Ranger Knox plugin on a Kerberos cluster.

Before you begin

This procedure assumes that you have already completed “Customize Services: Plugins”.

Procedure

1. Create the system (OS) user rangerknoxlookup. Make sure this user is synced to Ranger Admin (under Settings>Users/Groups tab in the Ranger Admin UI).
2. Create a Kerberos principal for rangerknoxlookup: `kadmin.local -q 'addprinc -pw rangerknoxlookup rangerknoxlookup@example.com'`.
3. Navigate to the Knox service.
4. Click the Config tab and navigate to advanced ranger-knox-plugin-properties.
5. Update the following properties with the values listed in the table below.

Table 17: Knox Plugin Properties

Configuration Property Name	Value
Ranger service config user	rangerknoxlookup@example.com
Ranger service config password	rangerknoxlookup
common.name.for.certificate	blank

6. After updating these properties, click Save and then restart the Knox service.

7. Open the Ranger Admin UI by entering the following information:
 - `http://ranger-host>:6080`
 - `username/password - admin/admin`. or use username as shown in advanced ranger-env under the Config tab of the Ranger service, and password as shown in Admin Settings.
8. After you have successfully logged into the system, you will be redirected to the Access Manager page.
9. Click the repository (clusterName_hadoop) Edit option under the HDFS box.

The screenshot shows the Ranger Admin UI 'Config Properties' page. The form includes the following fields:

- `username`: `hadoop`
- `password`: `----`
- `fs.default.name`: `hdfs://rangeramban-maradkar-17mar03-rhel6-6-cs1cloud.internal` (highlighted with a red box)
- `hadoop.security.authorization`: `false`
- `hadoop.security.authentication`: `simple`
- `hadoop.security.auth_to_local`: `DEFAULT`
- `dfs.datanode.kerberos.principal`:
- `dfs.namenode.kerberos.principal`:
- `dfs.secondary.namenode.kerberos.principal`:
- `hadoop.rpc.protection`: (highlighted with a red box)
- `Common Name For Certificate`: (highlighted with a red box)

A 'Test Connection' button is located at the bottom of the form.

10. Update the following properties listed in the table below under the Config Properties section:

Table 18: Knox Configuration Properties

Configuration Property Name	Value
<code>fs.default.name</code>	<code>hdfs</code>
<code>hadoop.rpc.protection</code>	<code>blank</code>
<code>common.name.for.certificate</code>	<code>blank</code>

11. Click Named Test Connection. You should see a Connected Successfully dialog box appears.
12. Click Save.