

Hortonworks Data Platform

Apache Zeppelin Component Guide

(August 29, 2016)

Hortonworks Data Platform: Apache Zeppelin Component Guide

Copyright © 2012-2016 Hortonworks, Inc. Some rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, ZooKeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain free and open source.

Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [Contact Us](#) directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under **Creative Commons Attribution ShareAlike 4.0 License**.
<http://creativecommons.org/licenses/by-sa/4.0/legalcode>

Table of Contents

1. Overview	1
2. Installing Apache Zeppelin	2
2.1. Requirements for Installing Zeppelin	2
2.2. Installing Zeppelin Using Ambari	2
3. Configuring Zeppelin	5
3.1. Configuring Livy on an Ambari-Managed Cluster	5
3.2. Configuring User Impersonation for Access to Hive	7
4. Using Zeppelin	8
4.1. Launching Zeppelin	8
4.2. Working with Notes	9
4.2.1. Creating and Running a Note	10
4.2.2. Importing a Note	11
4.2.3. Exporting a Note	12
4.2.4. Using the Note Toolbar	12
4.3. Importing External Packages	13
4.4. Configuring and Using Zeppelin Interpreters	14
4.4.1. Modifying Interpreter Settings	14
4.4.2. Using Zeppelin Interpreters	14
4.4.3. Using the %jdbc Interpreter to Access Hive	16
4.4.4. Using the %livy Interpreter to Access Spark	16
5. Configuring Zeppelin Security	19
5.1. Getting Started	19
5.1.1. Prerequisites	19
5.1.2. Configuring Security on an Ambari-Managed Cluster	19
5.1.3. Configuring Security on a Cluster Not Managed by Ambari	19
5.1.4. shiro.ini Structure	20
5.2. Configure Zeppelin for Authentication: Non-Production Use	20
5.3. Configure Zeppelin for Authentication: LDAP and Active Directory	21
5.3.1. Configuring Authentication for Production Using Active Directory	21
5.3.2. Configuring Authentication for Production Using LDAP	23
5.4. Enable Access Control for Zeppelin Elements	25
5.4.1. Enabling Access Control for Interpreter, Configuration, and Credential Settings	25
5.4.2. Enabling Access Control for Notebooks	26
5.4.3. Enabling Access Control for Data	27
5.5. Configure SSL for Zeppelin	28
5.6. Configure Zeppelin for a Kerberos-Enabled Cluster	28
5.7. Shiro Settings: Reference	29
5.7.1. Active Directory Settings	29
5.7.2. LDAP Settings	30
5.7.3. General Settings	31
5.7.4. Example	31
6. Stopping the Zeppelin Server and Livy Server	33

1. Overview

Apache Zeppelin is a web-based notebook that supports interactive data exploration, visualization, and collaboration. Zeppelin supports a growing list of programming languages and interfaces, including Python, Scala, Hive, SparkSQL, shell, AngularJS, and markdown.

Apache Zeppelin is useful for working interactively with long workflows: developing, organizing, and running analytic code; and visualizing results.

The screenshot displays the Apache Zeppelin Notebook interface for an "IoT Data Analysis" workflow. At the top, there is a search bar and a user profile indicator. Below the title, a table lists various event types and their associated metrics. The table data is as follows:

Normal	N	miles	90	4,300	-89.52	40.7	0	0	0
Normal	N	miles	90	4,300	-91.05	41.72	0	0	1
Normal	N	miles	90	4,300	-91.47	41.74	0	0	0
Lane Departure	N	miles	90	4,300	-91.59	41.7	1	0	0
Unsafe following distance	N	miles	90	4,300	-88.77	40.76	0	1	1

Below the table, three interactive visualizations are shown, each with its own KSQL query and execution status:

- Bar Chart:** A grouped bar chart showing the number of occurrences for "Overspeed" and "Unsafe following distance" events. The y-axis ranges from 0 to 1,200. The "Unsafe following distance" bar is significantly higher, reaching approximately 1,100 occurrences.
- Pie Chart:** A pie chart showing the distribution of violations by certification status, with categories "N" (Not Certified) and "Y" (Certified). The "N" slice is the largest, representing about 70% of the total.
- Area Chart:** A stacked area chart showing "hoursDriven" for different event types, categorized by "isCertified" status. The y-axis ranges from 0.00 to 5.00. The chart shows a sharp spike in hours driven for the "Unsafe following distance" event.

2. Installing Apache Zeppelin

This chapter describes how to install Zeppelin on an Ambari-managed cluster.

To install Zeppelin on an HDP cluster not managed by Ambari, see [Installing and Configuring Apache Zeppelin](#) in the *Command Line Installation Guide*.

To configure Zeppelin security features, see [Configuring Zeppelin Security](#).

2.1. Requirements for Installing Zeppelin

Install Zeppelin on a node where Spark clients are already installed and running. This typically means that Zeppelin will be installed on a gateway or edge node.

Zeppelin requires the following software versions:

- HDP 2.5.0 or later.
- Apache Spark version 1.6. (Zeppelin does not currently support Spark 2.0, which is available as a technical preview with HDP 2.5.0.)
- Java 8 on the node where Zeppelin is installed.

The optional Livy server provides security features and user impersonation support for Zeppelin users. Livy is installed as part of Spark:

- For an Ambari-managed cluster, see [Installing Spark Using Ambari](#).

After installing Spark, Livy, and Zeppelin, refer to [Configuring Zeppelin](#) for post-installation steps.

- For a cluster not managed by Ambari, see [Installing and Configuring Livy](#) in the *Command Line Installation Guide*.

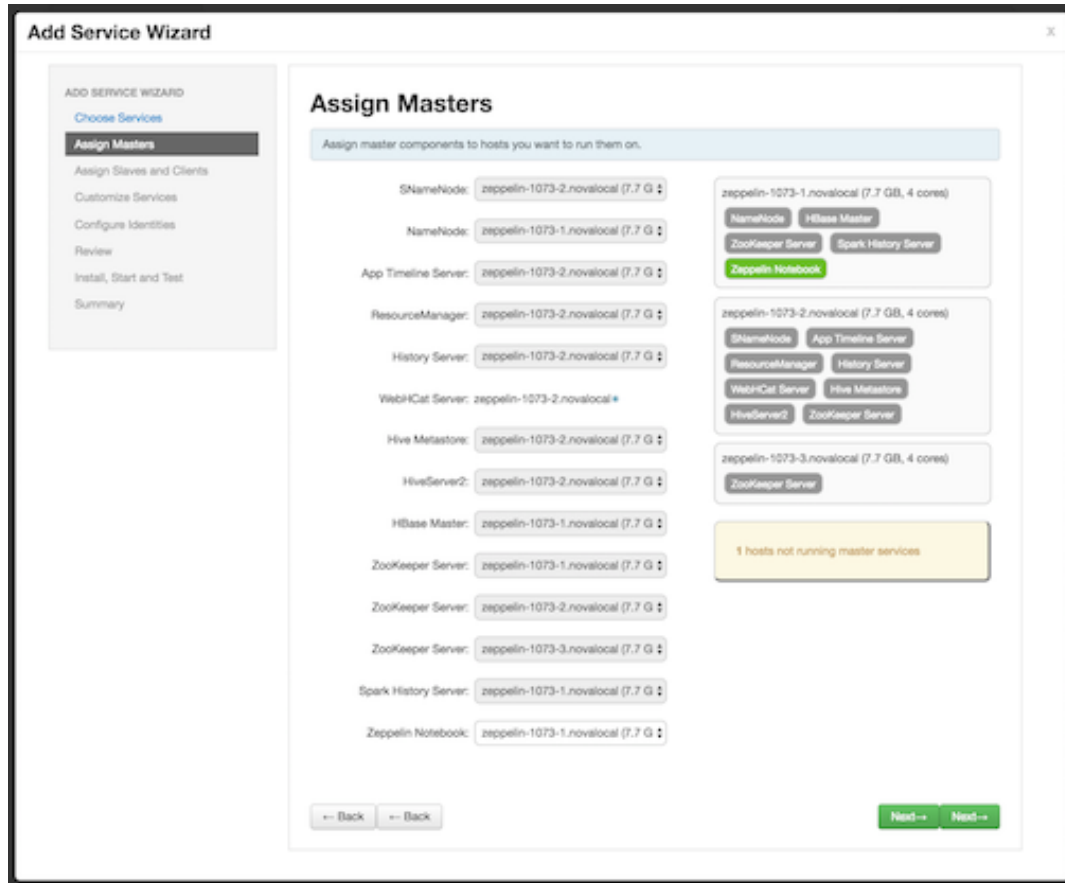
2.2. Installing Zeppelin Using Ambari

This section describes how to install Zeppelin using Ambari.

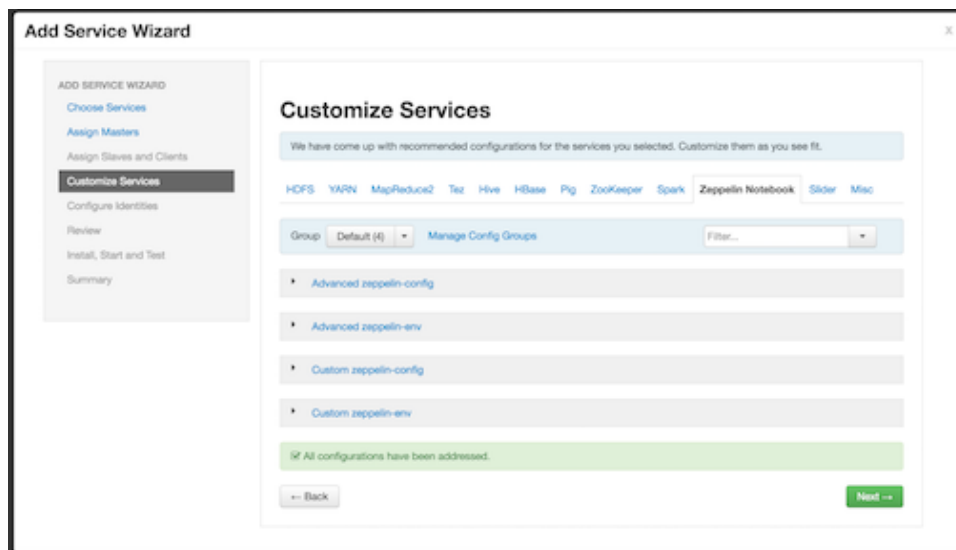
The Ambari installation wizard sets default values for Zeppelin configuration settings. At first you should accept the default settings. Later, when you are more familiar with Zeppelin, consider customizing Zeppelin configuration settings.

To install Zeppelin using Ambari, add the Zeppelin service:

1. In the "Actions" menu at the bottom left of the Ambari dashboard, select "Add Service". This starts the Add Service wizard, which displays the "Choose Services" page.
2. On the Choose Services page, select the Zeppelin service.
3. Click "Next" to continue.
4. On the Assign Masters page, review the node assignment for Zeppelin, and modify as needed:



5. Click "Next" to continue.
6. On the Customize Services page, review default values and click "Next" to continue.



7. On the Review page, confirm the node selected to run the Zeppelin service:

```
Please review the configuration before installation

http://s3.amazonaws.com/dev.hortonworks.com/HDP/centos7/2.x/BUILDS/2.4.0.0-169/
redhat7 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos7

suse11 (HDP-2.4):
http://s3.amazonaws.com/dev.hortonworks.com/HDP/suse11sp3/2.x/BUILDS/2.4.0.0-169/
suse11 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/suse11sp3

ubuntu12 (HDP-2.4):
http://s3.amazonaws.com/dev.hortonworks.com/HDP/ubuntu12/2.x/BUILDS/2.4.0.0-169/
ubuntu12 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/ubuntu12

ubuntu14 (HDP-2.4):
http://s3.amazonaws.com/dev.hortonworks.com/HDP/ubuntu14/2.x/BUILDS/2.4.0.0-169/
ubuntu14 (HDP-UTILS-1.1.0.20):
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/ubuntu12

Services:
Zeppelin Notebook
Notebook : c6401.ambari.apache.org
```

8. Click "Deploy" to complete the installation process.

To validate your Zeppelin installation, open the Zeppelin Web UI in a browser window. Use the port number configured for Zeppelin (9995 by default); for example:

```
http://<zeppelin-host>:9995
```

To check the Zeppelin version number, type the following command on the command line:

```
/usr/hdp/current/zeppelin-server/bin/zeppelin-daemon.sh --version
```

Zeppelin stores configuration settings in the `/etc/zeppelin/conf` directory. Note, however, that if your cluster is managed by Ambari you should not modify configuration settings directly. Instead, use the Ambari web UI.

Zeppelin stores log files in `/var/log/zeppelin` on the node where Zeppelin is installed.

3. Configuring Zeppelin

This chapter describes how to configure the following optional features:

- Configuring Livy on an Ambari-managed cluster. (For a cluster not managed by Ambari, see [Installing and Configuring Livy](#) in the *Command Line Installation Guide*.)
- Configuring Zeppelin to access Hive using identity propagation, on Kerberos-enabled clusters.

3.1. Configuring Livy on an Ambari-Managed Cluster

Livy is a proxy service for Apache Spark. On HDP 2.5, Livy offers two main capabilities:

- Zeppelin users can launch a Spark session on a cluster, submit code, and retrieve job results, all over a secure connection.
- When Zeppelin runs with authentication enabled, Livy propagates user information when a session is created. Livy user impersonation offers an extended multi-tenant experience, allowing users to share RDDs and cluster resources. Multiple users can access their own private data and session, and collaborate on a notebook.




Note

Livy can be accessed through the `%livy` interpreter in Zeppelin notebooks, but direct use of Livy REST APIs is not supported.

Livy supports Kerberos, but does not require it.

The following graphic shows process communication among Zeppelin, Livy, and Spark:



On an Ambari-managed cluster, Livy is available as an option when you install Spark. After installing Livy, complete the following configuration steps.

1. (*Optional*, for Kerberos-enabled clusters) Ensure that access is enabled only for groups and hosts where Livy runs.
2. Specify the Livy host URL:

- a. Navigate to the Interpreter configuration page in the Zeppelin Web UI.
- b. In the livy interpreter section, make sure that the `zeppelin.livy.url` property contains the full Livy host name; replace `localhost` if necessary.
- c. Scroll down to the "Save" button, and click "Save".

Note: on an Ambari-managed cluster you can find the Livy host from the Ambari dashboard: navigate to Spark > Summary > Livy Server.

3. (Optional) Configure Livy impersonation:


- a. From the Ambari dashboard, navigate to Spark > Configs.
- b. Open the "Custom livy-conf" category.
- c. Ensure that `livy.superusers` is listed; if not, add the property.
- d. Set `livy.superusers` to the user account associated with Zeppelin, `zeppelin.livy.principal`.

For example, if `zeppelin.livy.principal` is `zeppelin-sr1@example.com`, set `livy.superusers` to the same account, `zeppelin-sr1@example.com`.

4. (Optional) Specify a timeout value for Livy sessions.


By default, Livy preserves cluster resources by recycling sessions after one hour of session inactivity. When a Livy session times out, the Livy interpreter must be restarted.

To specify a larger or smaller value using Ambari, navigate to the `livy.server.session.timeout` property in the "Advanced livy-conf" section of the Spark service. Specify the timeout in milliseconds (the default, one hour, is 3600000):



5. If you changed any Livy interpreter settings, restart the Livy interpreter.

Navigate to the Interpreter configuration page in the Zeppelin Web UI. Locate the Livy interpreter and click "restart":



To verify that the Livy server is running, access the Livy web UI in a browser window. The default port is 8998:

```
http://<livy-hostname>:8998/
```

3.2. Configuring User Impersonation for Access to Hive

For clusters with Kerberos enabled: if you want Hive queries to run under the user ID associated with the Zeppelin session, enable user impersonation for the `%jdbc(hive)` interpreter.

1. Enable authentication via the Shiro configuration.

In the Zeppelin UI, navigate to the JDBC interpreter section of the Interpreter page, and then specifying authorization type, keytab, and principal:

- a. Set `zeppelin.jdbc.auth.type` to `KERBEROS`.
- b. Set `zeppelin.jdbc.principal` to the value of the principal.
- c. Set `zeppelin.jdbc.keytab.location` to the keytab location.

2. The JDBC interpreter automatically adds the end user as proxy user and sends the string to `HiveServer2`; for example:

```
jdbc:hive2://HiveHost:10000/default;principal=hive/_HOST@HOST1.COM;hive.  
server2.proxy.user=testuser
```

4. Using Zeppelin

A Zeppelin notebook is a browser-based GUI for interactive data exploration, modeling, and visualization.

As a notebook author or collaborator, you write code in a browser window. When you run the code from the browser, Zeppelin sends the code to backend processors such as Spark. The processor or service and returns results; you can then use Zeppelin to review and visualize results in the browser.

Note: Zeppelin on HDP does not support sharing a note by sharing its URL, due to lack of proper access control over how and with whom a note can be shared.

Apache Zeppelin is supported on the following browsers:

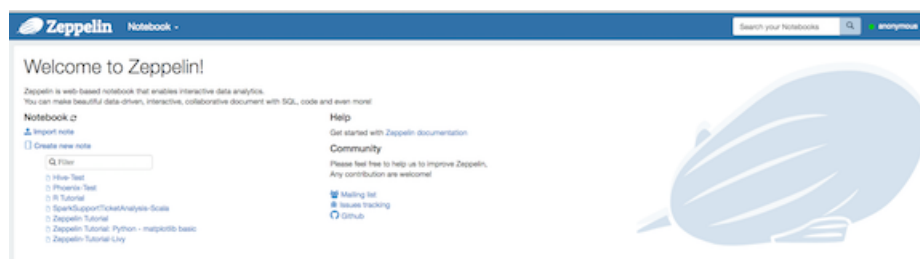
- Internet Explorer, latest supported releases. (Zeppelin is not supported on versions 8 or 9, due to lack of native support for WebSockets.)
- Google Chrome, latest stable release.
- Mozilla Firefox, latest stable release.
- Apple Safari, latest stable release. Note that when SSL is enabled for Zeppelin, Safari requires a Certificate Authority-signed certificate to access the Zeppelin UI.

4.1. Launching Zeppelin

To launch Zeppelin in your browser, access the host and port associated with the Zeppelin server. The default port is 9995:


```
http://<zeppelinhost>:9995.
```

When you first connect to Zeppelin, you will see the home page:



The following menus are available in the top banner of all Zeppelin pages:

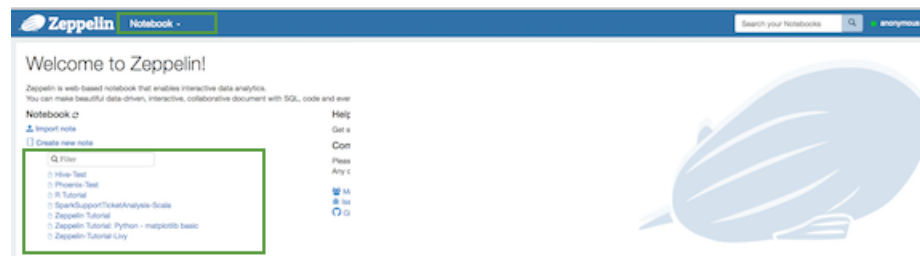
- Notebook: open a note, filter the list of notes by name, or create a note
- User settings: listed as your username, or "anonymous" if your Zeppelin deployment uses the default Shiro configuration.



- List version information about Zeppelin
- Review interpreter settings and configure, add, or remove interpreter instances
- Save credentials for data sources
- Display configuration settings

Each instance of Zeppelin contains "notes", which are the fundamental elements of a Zeppelin notebook.

Zeppelin lists available notes on the left side of the welcome screen and in the "Notebook" menu. Zeppelin ships with several sample notes, including tutorials:



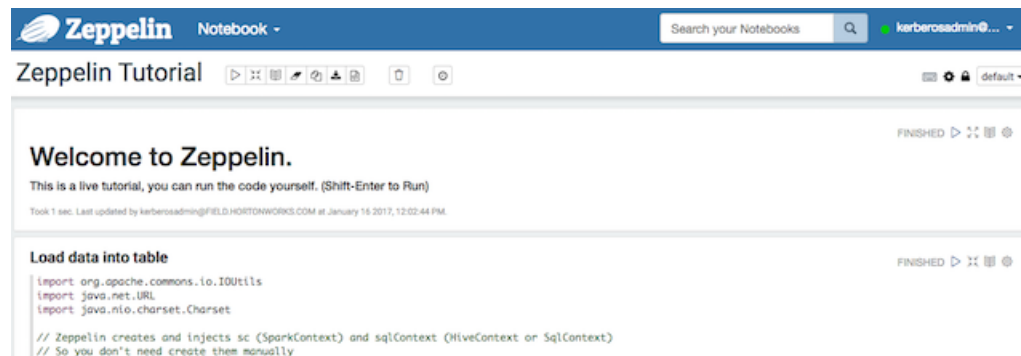
4.2. Working with Notes

A note consists of one or more paragraphs of code, which you can use to define and run snippets of code in a flexible manner.

A paragraph contains code to access services, run jobs, and display results. A paragraph consists of two main sections: an interactive box for code, and a box that displays results. To the right is a set of paragraph commands. The following graphic shows paragraph layout.



Zeppelin ships with several sample notes, including tutorials that demonstrate how to run Spark scala code, Spark SQL code, and create visualizations.



To run a tutorial:

1. Navigate to the tutorial: click one of the Zeppelin tutorial links on the left side of the welcome page, or use the Notebook pull-down menu.
2. Zeppelin presents the tutorial, a sequence of paragraphs prepopulated with code and text.
3. Starting with the first paragraph, click the triangle button at the upper right of the paragraph. The status changes to PENDING, RUNNING, and then FINISHED when done.
4. When the first cell finishes execution, results appear in the box underneath your code. Review the results.
5. Step through each cell, running the code and reviewing results.

4.2.1. Creating and Running a Note

To create a note:

1. Click "Create new note" on the welcome page, or click the "Notebook" menu and choose "+ Create new note."
2. Type your commands into the blank paragraph in the new note.

When you create a note, it appears in the list of notes on the left side of the home page and in the Notebook menu. By default, Zeppelin stores notes in the `$ZEPPELIN_HOME/notebook` folder.

To run your code:

1. Click the triangle button in the cell that contains your code:




2. Zeppelin displays status near the triangle button: PENDING, RUNNING, ERROR, or FINISHED.

3. When finished, results appear in the result section below your code.

The settings icon (outlined in red) offers several additional commands:



These commands allow you to perform several note operations, such as showing and hiding line numbers, clearing the results section, and deleting the paragraph.



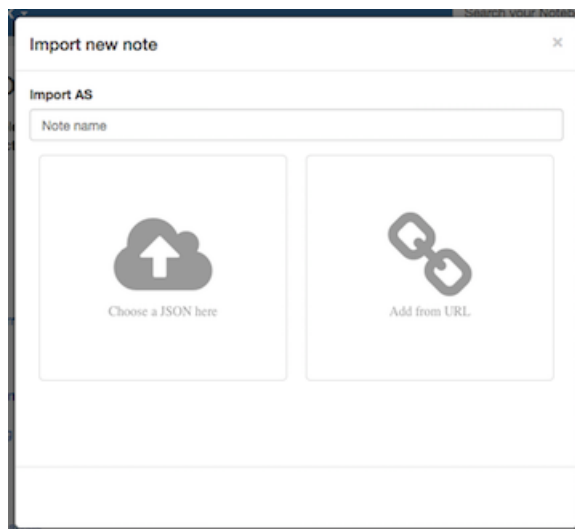
4.2.2. Importing a Note

To import a note from a URL or from a JSON file in your local file system:

1. Click "Import note" on the Zeppelin home page:



2. Zeppelin displays an import dialog box:



3. To upload the file or specify the URL, click the associated box.

By default, the name of the imported note is the same as the original note. You can rename it by providing a new name in the "Import AS" field.

4.2.3. Exporting a Note

To export a note to a local JSON file, use the export note icon in the note toolbar:



Zeppelin downloads the note to the local file system.

Note: Zeppelin exports code and results sections in all paragraphs. If you have a lot of data in your results sections, consider trimming results before exporting them.

4.2.4. Using the Note Toolbar

At the top of each note there is a toolbar with buttons for running code in paragraphs and for setting configuration, security, and display options:

There are several buttons in the middle of the toolbar:




These buttons perform the following operations:

- Execute all paragraphs in the note sequentially, in the order in which they are displayed in the note.
- Hide or show the code section of all paragraphs.
- Hide or show the result sections in all paragraphs.
- Clear the result section in all paragraphs.
- Clone the current note.
- Export the current note to a JSON file.

Note that the code and result sections in all paragraphs are exported. If you have extra data in some of your result sections, trim the data before exporting it.

- Commit the current note content.
- Delete the note.
- Schedule the execution of all paragraphs using CRON syntax. This feature is not currently operational. If you need to schedule Spark jobs, consider using Oozie Spark action.

There are additional buttons on the right side of the toolbar:



These buttons perform the following operations (from left to right):

- Display all keyboard shortcuts.
- Configure interpreters that are bound to the current note.
- Configure note permissions.
- Switch display mode:
 - Default: the notebook can be shared with (and edited by) anyone who has access to the notebook.
 - Simple: similar to default, with available options shown only when your cursor is over the cell.
 - Report: only your results are visible, and are read-only (no editing).

Note: Zeppelin on HDP does not support sharing a note by sharing its URL, due to lack of proper access control over who and how a note can be shared.

4.3. Importing External Packages

If you want to use an external package within a Zeppelin note, you can follow one of these approaches:

- Specify the dependency for the associated interpreter on the Interpreter page. For more information, see [Dependency Management for Interpreter](http://zeppelin.apache.org/Dependency%20Management%20for%20Interpreter) at zeppelin.apache.org.
- For Spark jobs, you can pass a jar, package, or list of files to `spark-submit` using `SPARK_SUBMIT_OPTIONS`; for example:
 - `SPARK_SUBMIT_OPTIONS` in `conf/zeppelin-env.sh`

```
export SPARKSUBMITOPTIONS="--packages com.databricks:spark-csv_2.10:1.2.0
--jars /path/mylib1.jar,/path/mylib2.jar --files /path/mylib1.py,/path/
mylib2.zip,/path/mylib3.egg"
```

- In `SPARK_HOME/conf/spark-defaults.conf`

```
spark.jars /path/mylib1.jar,/path/mylib2.jar spark.jars.packages com.
databricks:spark-csv_2.10:1.2.0 spark.files /path/mylib1.py,/path/mylib2.
egg,/path/mylib3.zip
```

If you want to import a library for a note that uses the Livy interpreter, see [Using the %livy Interpreter to Access Spark](#).

4.4. Configuring and Using Zeppelin Interpreters

An interpreter is a plugin that enables you to access processing engines and data sources from the Zeppelin UI. For example, if you want to use scala code in your Zeppelin notebook, you need a scala interpreter. Each interpreter runs in its own JVM on the same node as the Zeppelin server. The Zeppelin server communicates with interpreters through the use of Thrift.

Apache Zeppelin on HDP supports the following interpreters:

- Spark
- Hive (supported via JDBC)
- Shell
- Markdown
- Livy (supports Spark, SparkSQL, PySpark, and SparkR)
- AngularJS

4.4.1. Modifying Interpreter Settings

Before using an interpreter, you might want to modify default settings such as the home directory for the `spark` interpreter, the name of the Hive JDBC driver for the JDBC interpreter, or the keytab and principal name for a secure installation.

To set custom configuration values for an interpreter:


1. Click the user settings menu and navigate to the Interpreter page.
2. Scroll down to the Properties section for the interpreter of interest, and click "edit".
3. Make your changes.
4. Scroll to the end of the list of settings and click "Save".

Note: The Interpreter page is subject to access control settings. If the page does not list a set of interpreters, check with your system administrator.

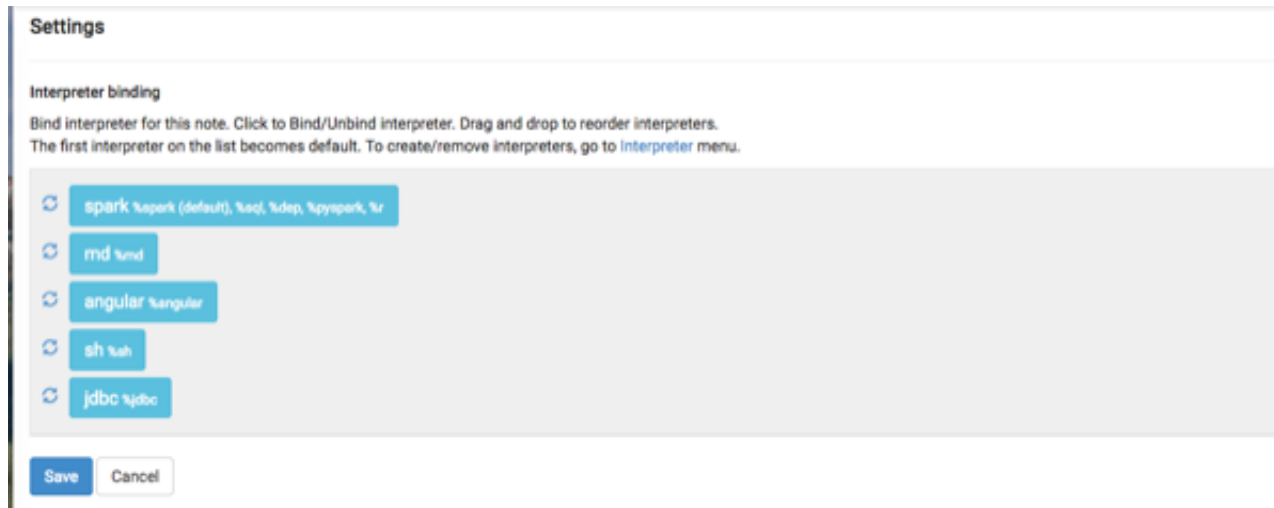
4.4.2. Using Zeppelin Interpreters

Before using an interpreter, ensure that the interpreter is available for use in your note:

1. Navigate to your note.
2. Click on "interpreter binding":



- Under "Settings", make sure that the interpreter you want to use is selected (in blue text). Unselected interpreters appear in white text:



- To select an interpreter, click on the interpreter name to select the interpreter. Each click operates as a toggle.
- You should unselect interpreters that will not be used. This makes your choices clearer. For example, if you plan to use `%livy` to access Spark, unselect the `%spark` interpreter.

Whenever one or more interpreters could be used to access the same underlying service, you can specify the precedence of interpreters within a note:

- Drag and drop interpreters into the desired positions in the list.
- When finished, click "Save".

Using an Interpreter in a Paragraph

To use an interpreter, specify the interpreter directive at the beginning of a paragraph, using the format `% [INTERPRETER_NAME]`. The directive must appear before any code that uses the interpreter.

The following paragraph uses the `%sh` interpreter to access the system shell and list the current working directory:

```
%sh
pwd

home/zeppelin
```

Some interpreters support more than one form of the directive. For example, the `%livy` interpreter supports directives for PySpark, SparkR, and Spark SQL.

To view interpreter directives and settings, navigate to the Interpreter page and scroll through the list of interpreters or search for the interpreter name. The list of directives is at the start of the associated interpreter section of the notebook, followed by options and property settings.

livy %livy, %livy.pyspark, %livy.sparkr, %livy.sql

Note: The Interpreter page is subject to access control settings. If the Interpreters page does not list settings, check with your system administrator for more information.

Using Interpreter Groups

Each interpreter belongs to an interpreter group. Interpreters in the same group can reference each other. For example, if the Spark SQL interpreter and the Spark interpreter are in the same group, the Spark SQL interpreter can reference the Spark interpreter to access its SparkContext.

For more information about interpreters, see [Interpreters in Apache Zeppelin](http://zeppelin.apache.org) at zeppelin.apache.org.

4.4.3. Using the %jdbc Interpreter to Access Hive

The %jdbc interpreter supports access to Apache Hive data.

To use the JDBC interpreter for Hive access, specify the following directive:

```
%jdbc(hive)
```

Next, add the statement that accesses Hive; here is a sample paragraph:

```
%jdbc(hive)
SELECT * FROM db_name;
```

If you receive an error, you might need to complete the following two additional steps:

1. Copy Hive jar files to `/usr/hdp/current/zeppelin-server/interpreter/jdbc` (or create a soft link).
2. On the Interpreters page in the Zeppelin Web UI, restart the JDBC interpreter.

The JDBC interpreter connects via Thrift.


4.4.4. Using the %livy Interpreter to Access Spark

The Livy interpreter offers several advantages over the default Spark interpreter (%spark):

- Sharing of Spark context across multiple Zeppelin instances.
- Reduced resource use, by recycling resources after 60 minutes of activity (by default). The default Spark interpreter runs jobs—and retains job resources—indefinitely.
- User impersonation. When the Zeppelin server runs with authentication enabled, the Livy interpreter propagates user identity to the Spark job so that the job runs as the originating user. This is especially useful when multiple users are expected to connect to the same set of data repositories within an enterprise. (The default Spark interpreter runs jobs as the default Zeppelin user.)

- The ability to run Spark in yarn-cluster mode.

Prerequisite: Before using Livy in a note, check the Interpreter page to ensure that the Livy interpreter is configured properly for your cluster:



Note: The Interpreter page is subject to access control settings. If the Interpreters page does not list settings, check with your system administrator for more information.

To use access Spark using Livy, specify the corresponding interpreter directive in a paragraph, before the code that accesses Spark; for example:

- PySpark:

```
%livy.pyspark
print "1"
```

- SparkR:

```
%livy.sparkr
hello <- function( name ) {
  sprintf( "Hello, %s", name );
}

hello("livy")
```



Important

To use SQLContext with Livy, do not create SQLContext explicitly. Zeppelin creates SQLContext by default.

If necessary, remove the following lines from the SparkSQL declaration area of your note:

```
//val sqlContext = new org.apache.spark.sql.SQLContext(sc)
//import sqlContext.implicits._
```

Livy sessions are recycled after a specified period of session inactivity. The default is one hour. If your session times out, you will need to restart the interpreter. On the Interpreters configuration page, locate the Livy interpreter and click "restart".

Importing External Packages

To import an external package for use in a note that runs with Livy:

1. Navigate to the interpreter settings.
2. If you are running the Livy interpreter in local mode (as specified by `livy.spark.master`), add jar files to the `/usr/hdp/<version>/livy/repl-jars` directory.

3. If you are running the Livy interpreter in yarn-cluster mode, either complete step 2 or edit the Livy configuration on the Interpreters page as follows:
 - a. Add a new key, `livy.spark.jars.packages`.
 - b. Set its value to `<group>:<id>:<version>`.

Here is an example for the `spray-json` library, which implements JSON in Scala:

```
io.spray:spray-json_2.10:1.3.1
```

5. Configuring Zeppelin Security

Zeppelin uses Apache Shiro to provide authentication and authorization (access control).

This chapter describes how to configure and enable several Zeppelin security features:

1. Configure authentication. Zeppelin supports a Shiro-based identity source for testing and informal use, as well as LDAP and Active Directory identity sources for production use. After authentication is enabled, when users connect to Apache Zeppelin they are prompted for login credentials.
2. Optionally, limit who can configure Zeppelin interpreter, credential, and configuration settings; notebooks; and data.
3. Optionally, configure the Zeppelin UI to run over SSL (HTTPS).
4. Optionally, configure Zeppelin to run on a Kerberos-enabled cluster.

5.1. Getting Started

5.1.1. Prerequisites

To use LDAP or Active Directory (AD) as the identity store, LDAP or AD must be installed and running on your cluster. You will need LDAP or AD coordinates to configure them for use with Zeppelin. In addition, the associated user accounts must be defined on your Zeppelin nodes.

5.1.2. Configuring Security on an Ambari-Managed Cluster

If your cluster is managed by Ambari, navigate to the Configs tab. Edit settings in the "Advanced zeppelin-env" and "Advanced zeppelin-config" categories, as described in following subsections.

Changes to `shiro_ini_content` require restarting the Zeppelin server. Ambari indicates this with a warning, and offers a menu option to restart Zeppelin.

5.1.3. Configuring Security on a Cluster Not Managed by Ambari

If your cluster is not managed by Ambari:

1. Locate the `shiro.ini` template file in the Zeppelin `/conf` folder:

```
/usr/hdp/current/zeppelin-server/conf/shiro.ini.template.
```

2. Copy the template file as `shiro.ini`:

```
/usr/hdp/current/zeppelin-server/conf/shiro.ini
```

3. Edit `shiro.ini` file contents, as described in the following subsections.

4. After editing the `shiro.ini` file, restart the Zeppelin server:

```
./bin/zeppelin-daemon.sh restart
```

5.1.4. shiro.ini Structure

The `shiro_ini_content` property (Ambari) and `shiro.ini` file (non-Ambari) contain several sections for configuring authentication:

- `[main]`, which contains definitions for LDAP or Active Directory objects and properties.
- `[users]`, which can be used to specify user accounts and passwords for simple deployments that do not require secure passwords, and require only a small number of statically-defined accounts.
- `[roles]`, for defining roles associated with access control.
- `[urls]`, for configuring URL-based security. For Zeppelin, the `[urls]` section is used to specify authentication method and define access control filters.

For more information about Shiro configuration and processing, see [Apache Shiro Configuration](#).

5.2. Configure Zeppelin for Authentication: Non-Production Use

The following steps provide a quick, basic form of authentication. This approach is not for production use; usernames and passwords are exposed in clear text. For production use, you should use LDAP or Active Directory as the identity source.

To configure authentication for informal use or testing:

1. Populate the `[urls]` section as follows:
 - a. Specify `authc` as the authentication method in the URL section of `shiro.ini` contents, and make sure that the `authc` line is not commented out.
 - b. To disable anonymous access to Zeppelin, add a comment character (`#`) at the start of the line containing `/** = anon`.

Here is an example:

```
[urls]
#/api/version = anon
#/** = anon
/** = authc
```

2. Populate the `[users]` section as follows:

Specify authorized accounts and associated passwords in `shiro_ini` settings: for clusters managed by Ambari, update `shiro_ini_content`; for non-Ambari clusters, update the `shiro.ini` file.

The following example configures authentication for users `admin`, `user1`, and `user2`, with passwords `password1`, `password2`, and `password3`, respectively:

```
[users]
admin = password1
user1 = password2
user2 = password3
```

3. Restart the Zeppelin server using Ambari or, for a cluster not managed by Ambari, follow the instructions in [Installing and Configuring Apache Zeppelin](#) in the *Non-Ambari Cluster Installation Guide*.
4. After completing these steps, Zeppelin requires authentication of user credentials before allowing access to the Zeppelin UI.

5.3. Configure Zeppelin for Authentication: LDAP and Active Directory

Zeppelin supports LDAP and Active Directory (AD) as identity stores for authentication. Because Active Directory is based on LDAP requirements, the configuration process is similar; however, the properties differ.

Before configuring LDAP or AD, user accounts must exist on Zeppelin nodes; and users, groups, and domain information must be stored in your LDAP or AD directory. You will need user, group, and domain information to configure LDAP or AD for Zeppelin.



Important

You should [configure and enable SSL](#) to the Zeppelin Web server whenever you enable LDAP or AD authentication. Without SSL, network communication is visible.

5.3.1. Configuring Authentication for Production Using Active Directory

To enable AD authentication, complete the following steps.

Note: Zeppelin currently uses Bind requests to authenticate end users; it does not support the LDAP `compare` operation.

The following steps describe basic settings. For more information about these and other settings, see [Shiro Settings: Reference](#) and [Shiro authentication for Apache Zeppelin](#).

1. Secure the HTTP channel.

In the `[urls]` section of `shiro.ini` contents, uncomment the line `/** = authc` and comment out the line `/** = anon` (to disable anonymous access):

```
[urls]
/api/version = anon
#/** = anon
/** = authc
```


Note: The [urls] section is processed from top to bottom; earlier statements have precedence. If you have two conflicting lines, the first is honored.

2. In the [main] section of shiro.ini contents, enable `activeDirectoryRealm` and modify the following settings for your operating environment. For clusters managed by Ambari, update `shiro_ini_content`; for non-Ambari clusters, update the `shiro.ini` file.

Note that there are two types of directory references, those that refer to the AD database, and those that refer to user accounts and groups. Domain information can differ between the two.

```
[main]

# authentication settings
activeDirectoryRealm = org.apache.zepelin.server.ActiveDirectoryGroupRealm
activeDirectoryRealm.url = ldap://<ldap-domain>:389
activeDirectoryRealm.searchBase = DC=<user-org-level-domain>,DC=<user-
second-level-domain>,DC=<user-top-level-domain>

# general settings
sessionManager = org.apache.shiro.web.session.mgt.DefaultWebSessionManager
cacheManager = org.apache.shiro.cache.MemoryConstrainedCacheManager
securityManager.cacheManager = $cacheManager
securityManager.sessionManager = $sessionManager
securityManager.sessionManager.globalSessionTimeout = 86400000

shiro.loginUrl = /api/login
```

3. *Optional:* Zeppelin supports connections to AD over SSL. To force Zeppelin to make an SSL connection to AD, change the value of `activeDirectoryRealm.url` from `ldap` to `ldaps` and specify the AD SSL port; for example:

```
activeDirectoryRealm.url = ldaps://hdp.example.com:636
```

If LDAP is using a self-signed certificate, import the certificate into the truststore of JVM running Zeppelin:

```
echo -n | openssl s_client -connect ldap.example.com:389 | sed -ne '/--BEGIN
CERTIFICATE-/,/--END CERTIFICATE-/p' > /tmp/examplecert.crt
keytool -import -keystore $JAVA_HOME/jre/lib/security/cacerts -storepass
changeit -noprompt -alias mycert -file /tmp/examplecert.crt
```

4. Secure the Websocket channel.

On an Ambari-managed cluster, navigate to the "Advanced zepelin-config" section and set `zepelin.anonymous.allowed` to `false`. HDP 2.6: check category

On a cluster not managed by Ambari, edit the `conf/zepelin-site.xml` file. Set `zepelin.anonymous.allowed` to `false`. (If the file does not exist, rename `conf/zepelin-site.xml.template` to `conf/zepelin-site.xml`, and then edit `zepelin-site.xml`.)

5. *Optional:* If you want to keep clear passwords from appearing in `shiro.ini`, complete the following steps:

- a. At your OS command line interface, use the Hadoop `credential` command to create an entry for the Active Directory credential:

```
> hadoop credential create activeDirectoryRealm.systemPassword -provider
jceks:///etc/zepelin/conf/credentials.jceks
Enter password:
Enter password again:

activeDirectoryRealm.systemPassword has been successfully created.
org.apache.hadoop.security.alias.JavaKeyStoreProvider has been updated.
```

- b. Using `chmod 400`, make the `credentials.jceks` file readable and writable only by the Zeppelin system user.
- c. Add the following line to `shiro.ini` contents:

```
activeDirectoryRealm.systemPassword -provider jceks:///etc/zepelin/conf/
credentials.jceks
```

6. Restart the Zeppelin server using Ambari or, for a cluster not managed by Ambari, manually restart the Zeppelin server:

```
./bin/zeppelin-daemon.sh restart
```

To use Zeppelin after configuring AD authentication:

1. Open the Zeppelin UI in a new browser tab:

```
http://<zeppelin-hostname>:9995
```

2. Click the Login button at the top right corner of the page.
3. In the login dialog box, specify a valid Active Directory username and password.

Unless the `activeDirectoryRealm.principalSuffix` property was set during AD configuration, users must fully qualify their account name; for example:

```
ad-username@AD.DOMAIN.COM
```

4. If Zeppelin access to Active Directory is configured accurately, Zeppelin presents its home page.

5.3.2. Configuring Authentication for Production Using LDAP

To use any form of LDAP other than AD, complete the steps in this section.

Note: Zeppelin currently uses LDAP Bind requests to authenticate end users; it does not support the LDAP `compare` operation.

The following steps describe basic settings. For more information about these and additional settings, see [Shiro Settings: Reference](#) and [Shiro authentication for Apache Zeppelin](#).

1. Secure the HTTP channel.

In the `[urls]` section of `shiro.ini` contents, uncomment the line `/** = authc`, and comment out the line `/** = anon` (to disable anonymous access):

```
[urls]
/api/version = anon
#/** = anon
/** = authc
```

Note: The `[urls]` section is processed from top to bottom; earlier statements have precedence. If you have two conflicting lines, the first is honored.

2. In the `[main]` section of `shiro.ini` contents, enable `ldapRealm` and modify the following settings for your operating environment. For clusters managed by Ambari, update `shiro_ini_content`; for non-Ambari clusters, update the `shiro.ini` file.

Note that there are two types of directory references: those that refer to the LDAP database, and those that refer to user accounts and groups. The domain information can differ between the two.

```
[main]

# authentication settings
ldapRealm = org.apache.zeppelin.server.LdapGroupRealm
ldapRealm.contextFactory.environment[ldap.searchBase] = DC=<user-second-level-domain>,DC=<user-top-level-domain>
ldapRealm.userDnTemplate = uid={0},OU=<user-account>,DC=<user-second-level-domain>,DC=<user-top-level-domain>
ldapRealm.contextFactory.url = ldap://<ldap-domain>:389
ldapRealm.contextFactory.authenticationMechanism = simple

# general settings
sessionManager = org.apache.shiro.web.session.mgt.DefaultWebSessionManager
cacheManager = org.apache.shiro.cache.MemoryConstrainedCacheManager
securityManager.cacheManager = $cacheManager
securityManager.sessionManager = $sessionManager
securityManager.sessionManager.globalSessionTimeout = 86400000

shiro.loginUrl = /api/login
```

3. Optional: Zeppelin supports connections to LDAP over SSL. To force Zeppelin to make an SSL connection to LDAP, change the `contextFactory.url` value from `ldap` to `ldaps` and specify the LDAP SSL port; for example:

```
ldapRealm.contextFactory.url = ldaps://hdp.example.com:636
```

If LDAP is using a self-signed certificate, import the certificate into the truststore of JVM running Zeppelin:

```
echo -n | openssl s_client -connect ldap.example.com:389 | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/examplecert.crt
keytool -import -keystore $JAVA_HOME/jre/lib/security/cacerts -storepass changeit -noprompt -alias mycert -file /tmp/examplecert.crt
```

4. Secure the Websocket channel.

On an Ambari-managed cluster, navigate to the "Advanced zeppelin-config" section and set `zeppelin.anonymous.allowed` to `false`.

On a cluster not managed by Ambari, edit the `conf/zeppelin-site.xml` file. Set `zeppelin.anonymous.allowed` to `false`. (If the file does not exist, rename `conf/zeppelin-site.xml.template` to `conf/zeppelin-site.xml`.)

- Restart the Zeppelin server using Ambari or, for a cluster not managed by Ambari, manually restart the Zeppelin server:

```
./bin/zeppelin-daemon.sh restart
```

- To use Zeppelin after configuring LDAP authentication:

- Open the Zeppelin UI in a new browser tab:

```
http://<zeppelin-hostname>:9995
```

- Click the Login button at the top right corner of the page.
- In the login dialog box, specify a valid LDAP username and password.
- If Zeppelin access to LDAP is configured accurately, Zeppelin presents its home page.

5.4. Enable Access Control for Zeppelin Elements

After configuring authentication, you might want to restrict access control to configure Zeppelin interpreters, access notes, or access data. You can authorize access at three levels within Zeppelin:

- UI authorization* restricts access to Zeppelin Interpreter, Credential, and Configuration pages based on administrative privileges.
- Note-level authorization* restricts access to notes based on permissions (owner, reader, or writer) granted to users and groups.
- Data-level authorization* restricts access to specific data sets.

5.4.1. Enabling Access Control for Interpreter, Configuration, and Credential Settings

By default, any authenticated account can access Zeppelin interpreter, credential, and configuration settings. When access control is enabled, unauthorized users can see the page heading, but no settings. There are two steps: defining roles, and specifying which roles have access to which settings.

Prerequisite: Users and groups must be defined on all Zeppelin nodes and in the associated identity store.

To enable access control for the Zeppelin interpreter, credential, or configuration pages, complete the following steps:

- Define a `[roles]` section in `shiro.ini` contents, and specify permissions for defined groups.

The following example grants all permissions ("*") to users in group `admin`:

```
[roles]
admin = *
```

2. In the `[urls]` section of the `shiro.ini` contents, uncomment the interpreter, configurations, or credential line(s) to enable access to the interpreter, configuration, or credential page(s), respectively. (If the `[urls]` section is not defined, add the section. Include the three `/api` lines listed in the following example.)

The following example specifies access to interpreter, configurations, and credential settings for role "admin":

```
[urls]
/api/version = anon
/api/interpreter/** = authc, roles[admin]
/api/configurations/** = authc, roles[admin]
/api/credential/** = authc, roles[admin]
#/** = anon
/** = authc
```

To add more roles, separate role identifiers with commas inside the square brackets.

Note: The sequence of lines in the `[urls]` section is important. The `/api/version` line must be the first line in the `[urls]` section:

```
/api/version = anon
```

Next, specify the three `/api` lines in any order:

```
/api/interpreter/** = authc, roles[admin]
/api/configurations/** = authc, roles[admin]
/api/credential/** = authc, roles[admin]
```

The `authc` line must be last in the `[urls]` section:

```
/** = authc
```

3. When unauthorized users attempt to access the interpreter, configurations, or credential page, they see the page heading but not settings.

5.4.2. Enabling Access Control for Notebooks

The next optional security step is to add access control for Zeppelin notes, granting permissions to specific users and groups. There are two main steps to this process: defining the `searchBase` property in the Zeppelin Shiro configuration, and then specifying permissions.

To restrict notebook access to authorized users:

1. Administrator:

In Zeppelin configuration settings, specify `activeDirectoryRealm.searchBase` or `ldapRealm.searchBase`, depending on whether you use AD or LDAP for authentication. The value of `searchBase` controls where Zeppelin looks for users and groups.

For more information, refer to [Shiro Settings: Reference](#). For an example, see [Configure Authentication for Authentication: LDAP and Active Directory](#).

2. Owner of a notebook:
 - a. Navigate to the note.
 - b. Click the lock icon on the notebook:



- c. Zeppelin presents a popup menu. Enter the user and groups that should have access to the note. To search for an account, start typing the name.

Note: If you are using Shiro as the identity store, users should be listed in the [user] section. If you are using AD or LDAP users and groups should be stored in the realm associated with your Shiro configuration.

Note Permissions (Only note owners can change)

Enter comma separated users and groups in the fields.
Empty field (*) implies anyone can do the operation.

Owners	<input type="text" value="search for users"/>	Owners can change permissions,read and write the note.
Readers	<input type="text" value="search for users"/>	Readers can only read the note.
Writers	<input type="text" value="search for users"/>	Writers can read and write the note.

For more information, see Apache [Zeppelin Notebook Authorization](#).

5.4.3. Enabling Access Control for Data

Access control for data brought into Zeppelin depends on the underlying data source:

- To configure access control for Spark data, Zeppelin must be running as an end user ("identity propagation"). Zeppelin implements access control using Livy. When identity propagation is enabled via Livy, data access is controlled by the type of data source being accessed. For example, when you access HDFS data, access is controlled by HDFS permissions.
- To configure access control for Hive data, use the JDBC interpreter.
- To configure access control for the Spark shell, define permissions for end users running the shell.

5.5. Configure SSL for Zeppelin

To configure the Zeppelin UI for access over SSL:

1. In Ambari, access the "Advanced zeppelin-config" section of the Zeppelin configuration settings:
2. Set `zeppelin.ssl` to `true`.
3. Configure key manager and key store settings with the correct values for your system:
 - a. Set `zeppelin.ssl.key.manager.password` to the password associated with the key manager.
 - b. Set `zeppelin.ssl.keystore.password` to the password associated with the key store.
 - c. Set `zeppelin.ssl.keystore.path` to the path associated with the key store.
 - d. Set `zeppelin.ssl.keystore.type` to the type of key store configured on the cluster (for example, JKS).
4. If you wish to use client-side certificate authentication, enable client-side authentication and configure the associated trust store settings:
 - a. Set `zeppelin.ssl.client.auth` to `true`
 - b. Set `zeppelin.ssl.truststore.path` to the path associated with your trust store.
 - c. Set `zeppelin.ssl.truststore.password` to the password associated with your trust store.
 - d. Set `zeppelin.ssl.truststore.type` to the type of trust store configured on the cluster (for example, JKS).
5. Check to make sure that all settings are valid.
6. Connect using HTTPS.

Note: When SSL is enabled for Zeppelin, the Safari browser requires a Certificate Authority-signed certificate to access the Zeppelin UI.

5.6. Configure Zeppelin for a Kerberos-Enabled Cluster

The Zeppelin daemon needs a Kerberos account and keytab to run in a Kerberized cluster.

- When you enable Kerberos on an Ambari-managed cluster, Ambari configures Kerberos for Zeppelin and automatically creates a Kerberos account and keytab for it. For more information, see [Configuring Ambari and Hadoop for Kerberos](#).

- If your cluster is not managed with Ambari and you plan to enable Kerberos for the Zeppelin server, see [Creating Service Principals and Keytab Files for HDP](#) in the Hadoop Security Guide.

5.7. Shiro Settings: Reference

This section lists more information about Shiro settings described in the authentication and authorization configuration subsections, plus a few more settings that might be useful.

5.7.1. Active Directory Settings

Active Directory (AD) stores users and groups in a hierarchical tree structure, built from containers including the organizational unit (ou), organization (o), and domain controller (dc). The path to each entry is a Distinguished Name (DN) that uniquely identifies a user or group.

User and group names typically have attributes such as a common name (cn) or unique ID (uid).

Specify the DN as a string, for example `cn=admin,dc=example,dc=com`. White space is ignored.

`activeDirectoryRealm.url`

specifies the host and port where Active Directory is set up. For more information, see [Apache Shiro Realms](#).

If the protocol element is specified as `ldap`, SSL is not used. If the protocol is specified as `ldaps`, access is over SSL.

Note: If Active Directory uses a self-signed certificate, import the certificate into the truststore of the JVM running Zeppelin; for example:

```
echo -n | openssl s_client -connect ldap.
example.com:389 | sed -ne '/-BEGIN CERTIFICATE-/
,-END CERTIFICATE-/p' > /tmp/examplecert.crt

keytool -import -keystore $JAVA_HOME/jre/lib/
security/cacerts -storepass changeit -noprompt -
alias mycert -file /tmp/examplecert.crt
```

`activeDirectoryRealm.principalSuffix`

simplifies the logon information that users must use to log in. Otherwise, AD requires a username fully qualified with domain information. For example, if a fully-qualified user account is `user@hdpqa.example.com`, you can specify a shorter suffix such as `user@hdpqa`.

```
activeDirectoryRealm.principalSuffix = @<user-
org-level-domain>
```

`activeDirectoryRealm.searchBase`

defines the base distinguished name from which the directory search starts. A distinguished name defines

each entry; "dc" entries define a hierarchical directory tree.

`activeDirectoryRealm.systemUsername,`
`activeDirectoryRealm.systemPassword`
 defines the username and password that Zeppelin uses to connect to Active Directory when it searches for users and groups. These two settings are used for controlling access to UI features, not for authentication. The Bind method does not require a valid user password.

`activeDirectoryRealm.groupRolesMap`
 a comma-separated list that maps groups to roles. These settings are used by Zeppelin to restrict UI features to specific AD groups. The following example maps group `hdpdv_admin` at `hdp3.example.com` to the "admin" role:

```
CN=hdpdv_admin,DC=hdp3,DC=example,DC=com:admin
```

`activeDirectoryRealm.authorizationCachingEnabled`
 specifies whether to use caching to improve performance. To enable caching, set this property to `true`.

5.7.2. LDAP Settings

LDAP stores users and groups in a hierarchical tree structure, built from containers including the organizational unit (ou), organization (o), and domain controller (dc). The path to each entry is a Distinguished Name (DN) that uniquely identifies a user or group.

User and group names typically have attributes such as a common name (cn) or unique ID (uid).

Specify the DN as a string, for example `cn=admin,dc=example,dc=com`. White space is ignored.

Zeppelin LDAP authentication uses templates for user DNs. See [LDAP Realm Settings](#) for options you can set for an LDAP realm.

`LdapRealm`
 specifies the class name to use for LDAP authentication. You should set this to `org.apache.zeppelin.server.LdapGroupRealm` unless you are familiar with LDAP and prefer to use `org.apache.shiro.realm.ldap.JndiLdapRealm`. For more information, see [Apache Shiro Realms](#).

`LdapRealm.contextFactory.environment[ldap.searchBase]`
 defines the base distinguished name from which the LDAP search starts. Shiro searches for `userDnTemplate` at this address.

If the protocol is specified as `ldap`, SSL is not used. If the protocol is specified as `ldaps`, access is over SSL.

- `ldapRealm.userDnTemplate` specifies the search location where the user is to be found. Shiro replaces `{0}` with the username acquired from the Zeppelin UI. Zeppelin uses User DN templates to configure associated realms.
- `ldapRealm.contextFactory.url` specifies the host and port on which LDAP is running.
- If the protocol element is specified as `ldap`, SSL will not be used. If the protocol is specified as `ldaps`, access will be over SSL.
- Note:** If LDAP is using a self-signed certificate, import the certificate into the truststore of JVM running Zeppelin; for example:
- ```
echo -n | openssl s_client -connect ldap.
example.com:389 | sed -ne '/-BEGIN CERTIFICATE-/,
/-END CERTIFICATE-/p' > /tmp/examplecert.crt
keytool -import -keystore $JAVA_HOME/jre/lib/
security/cacerts -storepass changeit -noprompt -
alias mycert -file /tmp/examplecert.crt
```
- `ldapRealm.contextFactory.systemUsername,`  
`ldapRealm.contextFactory.systemPassword` define the username and password that Zeppelin uses to connect to LDAP, to search for users and groups. These two settings are used for controlling access to UI features, not for authentication. The Bind method does not require a valid user password.
- `ldapRealm.authorizationCachingEnabled` specifies whether to use caching to improve performance. To enable caching, set this property to `true`.

### 5.7.3. General Settings

- `securityManager.sessionManager.globalSessionTimeout` specifies how long to wait (in milliseconds) before logging out a user, if they are logged in and are not moving the cursor.

The default is 86,400,000 milliseconds, which equals 24 hours.

### 5.7.4. Example

The following example shows a minimum set of `shiro.ini` settings for authentication and access control, assuming that the corresponding account information is configured in Active Directory (at `adhost.field.hortonworks.com`) and on Zeppelin nodes.

```
[main]
AD authentication settings
```

```
activeDirectoryRealm = org.apache.zeppelin.server.ActiveDirectoryGroupRealm
activeDirectoryRealm.url = ldap://adhost.org.hortonworks.com:389
activeDirectoryRealm.searchBase = DC=org,DC=hortonworks,DC=com

general settings
sessionManager = org.apache.shiro.web.session.mgt.DefaultWebSessionManager
cacheManager = org.apache.shiro.cache.MemoryConstrainedCacheManager
securityManager.cacheManager = $cacheManager
securityManager.sessionManager = $sessionManager
securityManager.sessionManager.globalSessionTimeout = 86400000
shiro.loginUrl = /api/login

[roles]
admin = *

[urls]
authentication method and access control filters
/api/version = anon
/api/interpreter/** = authc, roles[admin]
/api/configurations/** = authc, roles[admin]
/api/credential/** = authc, roles[admin]
#/** = anon
/** = authc
```

## 6. Stopping the Zeppelin Server and Livy Server

To stop the Zeppelin server on a cluster managed by Ambari, use the Ambari Web UI.

To stop the Zeppelin server on a cluster that is not managed by Ambari, issue the following command from user `zeppelin`:

```
/usr/hdp/current/zeppelin-server/bin/zeppelin-daemon.sh stop
```

To stop the Livy server, issue the following command from user `livy`:

```
/usr/hdp/<hdp_version>/livy/bin/livy-server stop
```

For example:

```
/usr/hdp/2.5.0.0-1245/livy/bin
```