Installing HCP with Ambari

Date of Publish: 2018-11-15

http://docs.hortonworks.com

Contents

Installing HCP Using Ambari

Installing Hortonworks Cybersecurity Platform (HCP) using Apache Ambari uses both the graphic user interface of Ambari and the Metron user interface. Both of these tools promote a faster installation that preinstalls much of the configuration you need.

Prerequisites for an Existing Cluster

You can install HCP on an Ambari-managed cluster running HDP 2.5.x or 2.6.x and Ambari 2.4.2 (or later). However, the cluster must meet requirements for both the Hadoop cluster and the Metron nodes.

Specifications for Hadoop Cluster

All Hadoop-related nodes running HCP must meet operating system, HDP, and cluster requirements.

All Hadoop-related nodes must meet the following specifications:

- All cluster nodes must be running CentOS 6.x, CentOS 7.x, or Ubuntu 14.04
- The cluster must be running HDP 2.5.x or HDP 2.6.x managed by Ambari 2.4.2 (or later)
- The cluster must have a minimum of the following nodes:
  - Two Hadoop master nodes
  - Four Hadoop slaves nodes
  - One node for Ambari
- Each of the Hadoop Slave and Master nodes must meet the minimum specifications.
- The following services must be installed across the Hadoop Master and Slave nodes:
  - HDFS
  - HBase
  - ZooKeeper
  - Kafka
  - Storm
  - YARN
  - Spark 2.3.0 or later

  To determine the supported version for each service, refer to Ambari, and choose Admin > Stacks and Versions.
- Each of the following components must be installed on at least one node. The YARN ATS must installed on the master node. All other services in the list should be installed on multiple nodes.

  

  Note:

  For security reasons, no other workloads should be running on the cluster.

  Ambari Component

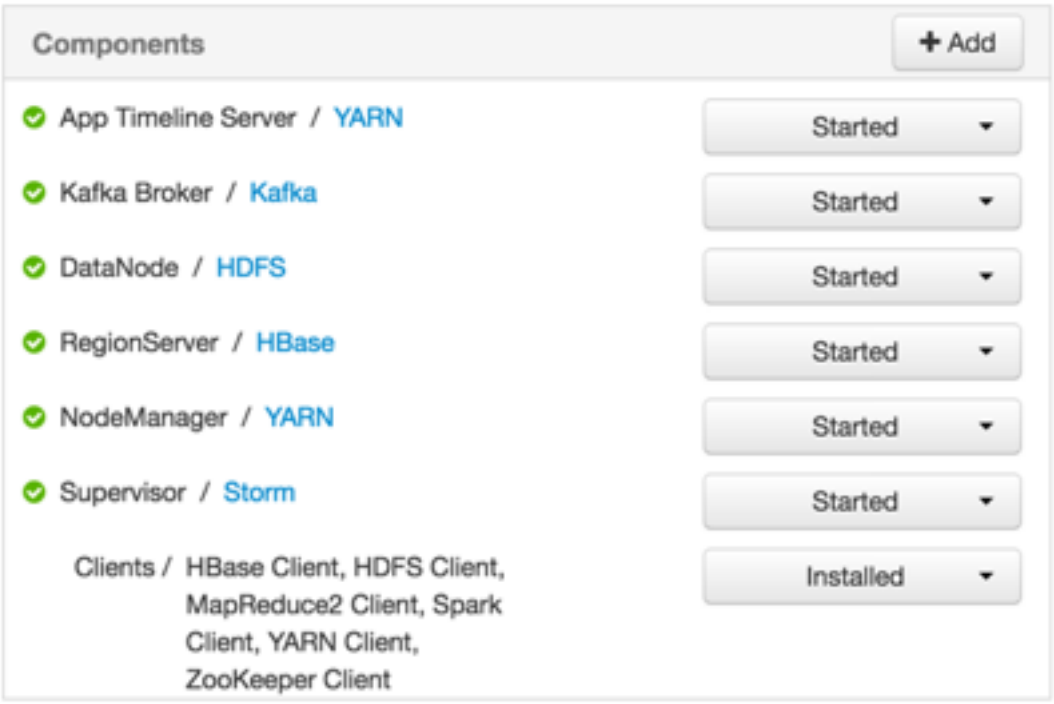

Specifications for Metron Nodes

All Metron nodes must meet specifications for the number of nodes dedicated for Metron-specific components and the ability to access the nodes.

The Metron nodes must meet the following specifications:

- At least three nodes must be dedicated for Metron-specific components.
- You must have root access on all Metron nodes.

The following figure illustrates a sample deployment architecture based on the previous specifications:

Sample Deployment Architecture

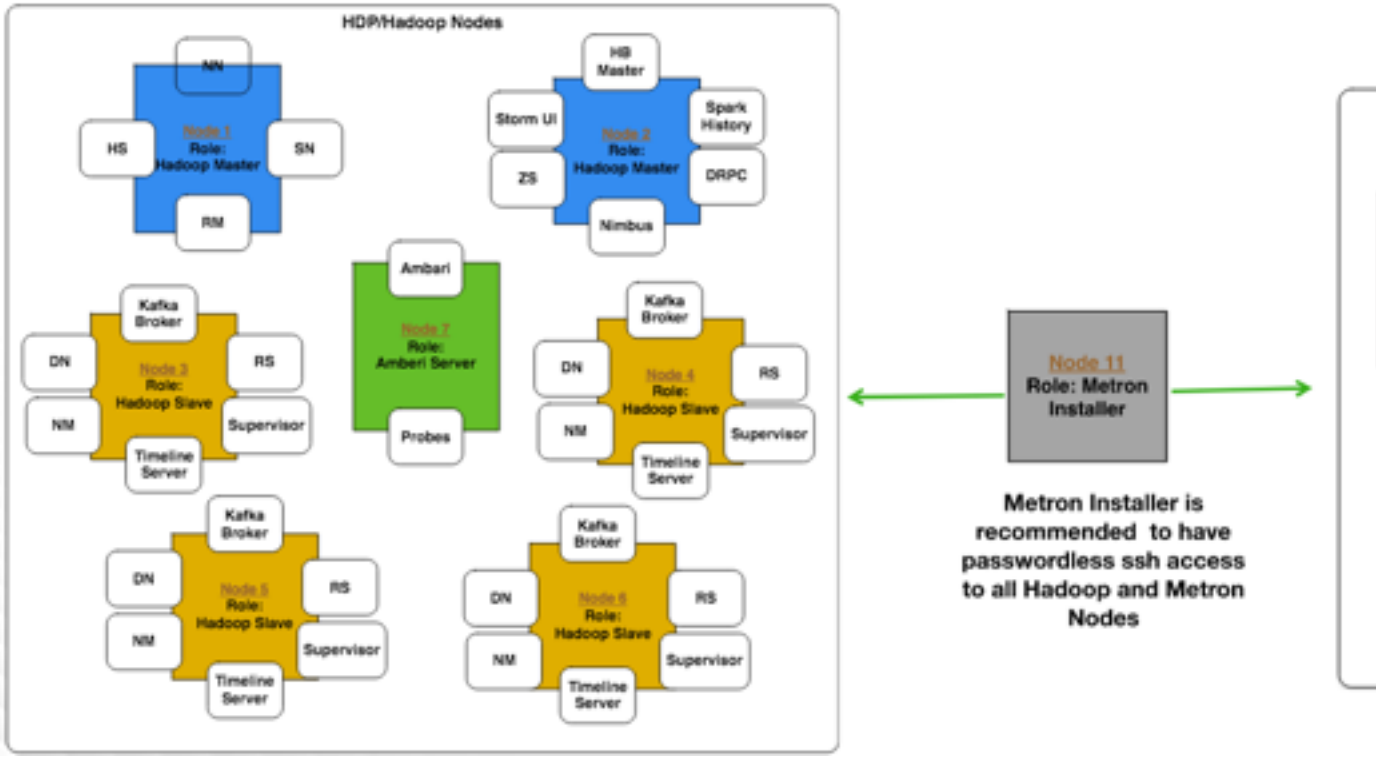

Set up the REST Application Database

Prior to installing HCP, you must set up the REST application database.

Procedure

1. Connect to MySQL and create a Metron REST database:
   mysql -uroot -p -e "CREATE DATABASE IF NOT EXISTS metronrest;"
2. Create a Metron user in MySQL with a password, then apply database access permission to the Metron user:

```
CREATE USER 'metron'@'$REST_HOST' IDENTIFIED BY 'Myp@ssw0rd';
GRANT ALL PRIVILEGES ON metronrest.* TO 'metron'@'$REST_HOST';
```

3. Create user and authorities tables:

```
use metronrest;
create table if not exists users(
  username varchar(50) not null primary key,
  password varchar(50) not null,
  enabled boolean not null
);
create table authorities (
  username varchar(50) not null,
  authority varchar(50) not null,
  constraint fk_authorities_users foreign key(username) references
  users(username)
);
create unique index ix_auth_username on authorities (username,authority);
```

4. Add one or more users to the REST application:

```
use metronrest;
```

```
insert into users (username, password, enabled) values ('your_username',
 'your_password',1);
insert into authorities (username, authority) values ('your_username',
 'ROLE_USER');
```

5. Exit MySQL:

```
quit
```

6. Install the appropriate MySQL client library for your version of MySQL. For example:

```
cd $METRON_HOME/lib
wget https://dev.mysql.com/get/Downloads/Connector-J/mysql-connector-
java-5. 1.41.tar.gz
tar xf mysql-connector-java-5.1.41.tar.gz
```

7. To add additional users:

```
use metronrest;
insert into users (username, password, enabled) values ('your_username',
 'your_password',1);
insert into authorities (username, authority) values ('your_username',
 'ROLE_USER');
commit;
```

Install HCP on an Ambari Cluster

Prior to installing the HCP Ambari management pack, you must meet HCP's requirements for the cluster, Metron node, and Ambari server.

Before you begin

Prior to installing the HCP Ambari management pack, you must complete the following:

Procedure

- Meet all of the cluster specifications listed in Specifications for Hadoop Cluster.
- Meet all of the metron node specifications listed in Specifications for Metron Nodes.
- Download and install Ambari.
- Set up the Ambari server.

Install HCP Ambari Management Pack

An HCP Ambari management pack bundles service definitions, stack definitions, and stack add-on service definitions so they do not need to be included with the Ambari core functionality and can be updated in between major releases. You can use the HCP management pack to install Metron, plus the parser topologies, indexing topologies, and enrichment topologies.

About this task

You can find the management pack repositories for each of the operating systems supported by HCP in the HCP Release Notes. The following is an example of installing the HCP Ambari management pack on CentosOS 7.

Procedure

1. Download the HCP management pack tar file from the HCP repo location:

```
wget -nv http://public-repo-1.hortonworks.com/HCP/centos7/1.x/
updates/1.7.1.0/tars/metron/hcp-ambari-mpack-1.7.1.0-24.tar.gz
```

You can find the management pack repositories for each of the operating systems supported by HCP at HCP Repositories.

Note: When installing Elasticsearch with the HCP management pack on Ubuntu, you must manually install the Elasticsearch repositories. You also do not need to download and install the elasticsearch_mpack.

2. If you are using Elasticsearch, download the Elasticsearch management pack tar file from the HCP repo location:

```
wget -nv http://public-repo-1.hortonworks.com/HCP/centos7/1.x/
updates/1.7.1.0/tars/metron/hcp-ambari-mpack-1.7.1.0-24.tar.gz
```

3. Install the HCP management packs onto the Ambari server:

   Install the elasticsearch_mpack only if you are using Elasticsearch.

```
ambari-server install-mpack --mpack=/${MPACK_DOWNLOAD_DIRECTORY}/hcp-
ambari-mpack-1.7.1.0-24.tar.gz --verbose ambari-server install-mpack --
mpack=/${MPACK_DOWNLOAD_DIRECTORY}/elasticsearch_mpack-1.7.1.0-24.tar.gz
 --verbose
```

   You should see a message saying that the management pack completed successfully.

Install Solr

If you are using Apache Solr, install it using the Ambari HDP Search management pack.

Procedure

1. From Ambari, stop the following:
   - Metron
   - Kibana
   - Elasticsearch
2. Install the Ambari HDP Search Management pack.

   For instructions on downloading and using the Ambari HDP Search management pack, see Apache Solr Search Installation.

   The Meta Alerts UI feature with Solr is technical preview in this release. We do not yet recommend this for production use, but please let us know about any bugs you might find. We appreciate your feedback.

   Important: Ensure the Java threat stack size parameter is set to greater than 320kb. The default setting for SOLR_JAVA_STACK_SIZE is not sufficient to start the Solr service.

   Ambari automatically creates collections for the following:
   - bro
   - snort
   - yaf
   - metaalert
   - error
3. If you want to create a collection for a schema not supplied by HCP, perform the following steps:

   a) Set Solr environmental variables in ZooKeeper.

```
# Path to the zookeeper node used by Solr
export ZOOKEEPER=node1:2181/solr
# Define SOLR_HOME
export SOLR_HOME=/opt/lucidworks-hdpsearch/solr/
# Set to true if Kerberos is enabled
export SECURITY_ENABLED=true
```

b) Create a collection.

For example:

```
su $SOLR_USER -c "$SOLR_HOME/bin/solr create -c bro -d $METRON_HOME/
config/schema/bro/"
```

c) Pull all configurations from ZooKeeper to the Metron config directory:

```
$METRON_HOME/bin/zk_load_configs.sh -m PULL -z $ZOOKEEPER -o
 $METRON_HOME/config/zookeeper -f
```

4. Add "source.type.field" : "source.type" and threat.triage.score.field" : "threat.triage.score" to the global.json file located at$METRON_HOME/config/zookeeper/global.json:

```
$METRON_HOME/bin/zk_load_configs.sh -m PUSH -z $ZOOKEEPER -i $METRON_HOME/
config/zookeeper
```

The global.json file should look similar to:

```
{
  "es.clustername" : "metron",
  "es.ip" : "blah:9300",
  "es.date.format" : "yyyy.MM.dd.HH",
  "parser.error.topic" : "indexing",
  "update.hbase.table" : "metron_update",
  "update.hbase.cf" : "t",
  "es.client.settings" : {
    "client.transport.ping_timeout" : "500s"
  },
  "profiler.client.period.duration" : "15",
  "profiler.client.period.duration.units" : "MINUTES",
  "source.type.field" : "source.type",
  "threat.triage.score.field" : "threat:triage:score",
  "user.settings.hbase.table" : "user_settings",
  "user.settings.hbase.cf" : "cf",
  "geo.hdfs.file" : "/apps/metron/geo/default/GeoLite2-City.mmdb.gz"
}
```

5. Push the configuration to ZooKeeper:

```
$METRON_HOME/bin/zk_load_configs.sh -m PUSH -z $ZOOKEEPER -i $METRON_HOME/
config/zookeeper
```

6. Restart Metron.
7. Start Solr.
8. From Ambari, select **Metron** in the components panel.
9. Click the **Configs** tab, then click the **Indexing** tab.
10. Choose Solr in the **Index Writer - Random Access** pull down menu.

Index Updates

Indexing Update Table

metron_update

Indexing Update Column Family

t

Index Writer - Random Access

Random Access Search Engine

Solr

Elasticsearch

Solr

m Access

1

Enrichment Ackers for Random Access

1

Indexing childopts

11. Click **Save**.
12. From Ambari, stop and restart the Metron Alerts user interface.
13. From Ambari, stop and restart Metron REST.

What to do next

You can access Solr by choosing **Solr UI** from the **Quick Links** pull down menu in Ambari.

Start the Ambari Server

After you install the HCP Ambari management pack, you need to start or restart the Ambari server, depending on whether your are installing HCP on a new or existing cluster.

Procedure

1. To start the Ambari server, enter the following:

   ambari-server start
2. To restart the Ambari server, enter the following:

   ambari-server restart

Install, Configure, and Deploy a HDP Cluster with HCP

You can use the Ambari Install wizard running in your browser to install, configure, and deploy your cluster.

About this task

To keep your changes to the indices writer, you must stop or restart the indexing topology only through Ambari. If you start or stop the indices writer through REST, the writer resets its settings to the Elasticsearch default settings.

Procedure

1. Open Ambari Web using a web browser.
   a) Point your browser to http://<your.ambari.server>:8080, where <your.ambari.server> is the name of your ambari server host.

      For example, a default Ambari server host is located at http://c6401.ambari.apache.org:8080.
   b) Log in to the Ambari Server using the default user name/password: admin/admin.

      You can change these credentials later.

      For a new cluster, the Ambari install wizard displays a Welcome page from which you launch the Ambari Install wizard.
2. For an existing cluster, select **Choose Services** from the **Actions/Add Service Wizard** menu and skip to Step 7.
3. From the Ambari Welcome page, choose **Launch Install Wizard**.
4. In **Name your cluster**, type a name for the cluster you want to create, and then choose **Next**.

   Avoid white spaces or special characters in the name.
5. Select the HDP stack you want to run.
6. Enter the set up information for which the install wizard prompts you.

   You need to supply the fully qualified domain name (FQDN) of each of your hosts. The wizard also needs to access the private key file you created in Set Up Password-less SSH. Using the host names and key file information, the wizard can locate, access, and interact securely with all hosts in the cluster.
   a) Use the **Target Hosts** text box to enter your list of host names, one per line.

      You can use ranges inside brackets to indicate larger sets of hosts. For example, for host01.domain through host10.domain use host[01-10].domain

Note: If you are deploying on EC2, use the internal Private DNS host names.

b) If you want to let Ambari automatically install the Ambari Agent on all your hosts using SSH, select **Provide your SSH Private Key** and either use the **Choose File** button in the **Host Registration Information** section to find the private key file that matches the public key you installed earlier on all your hosts or cut and paste the key into the text box manually.

Note: If you are using Internet Explorer 9, the Choose File button might not appear. Use the text box to cut and paste your private key manually. Fill in the user name for the SSH key you have selected. If you do not want to use root , you must provide the user name for an account that can execute sudo without entering a password.

c) Click **Register** and **Confirm** to continue.

Ambari displays the **Choose Services** dialog box that lists the services that Ambari can install into the cluster.

7. Choose the services to install onto the cluster, and then click **Next**.

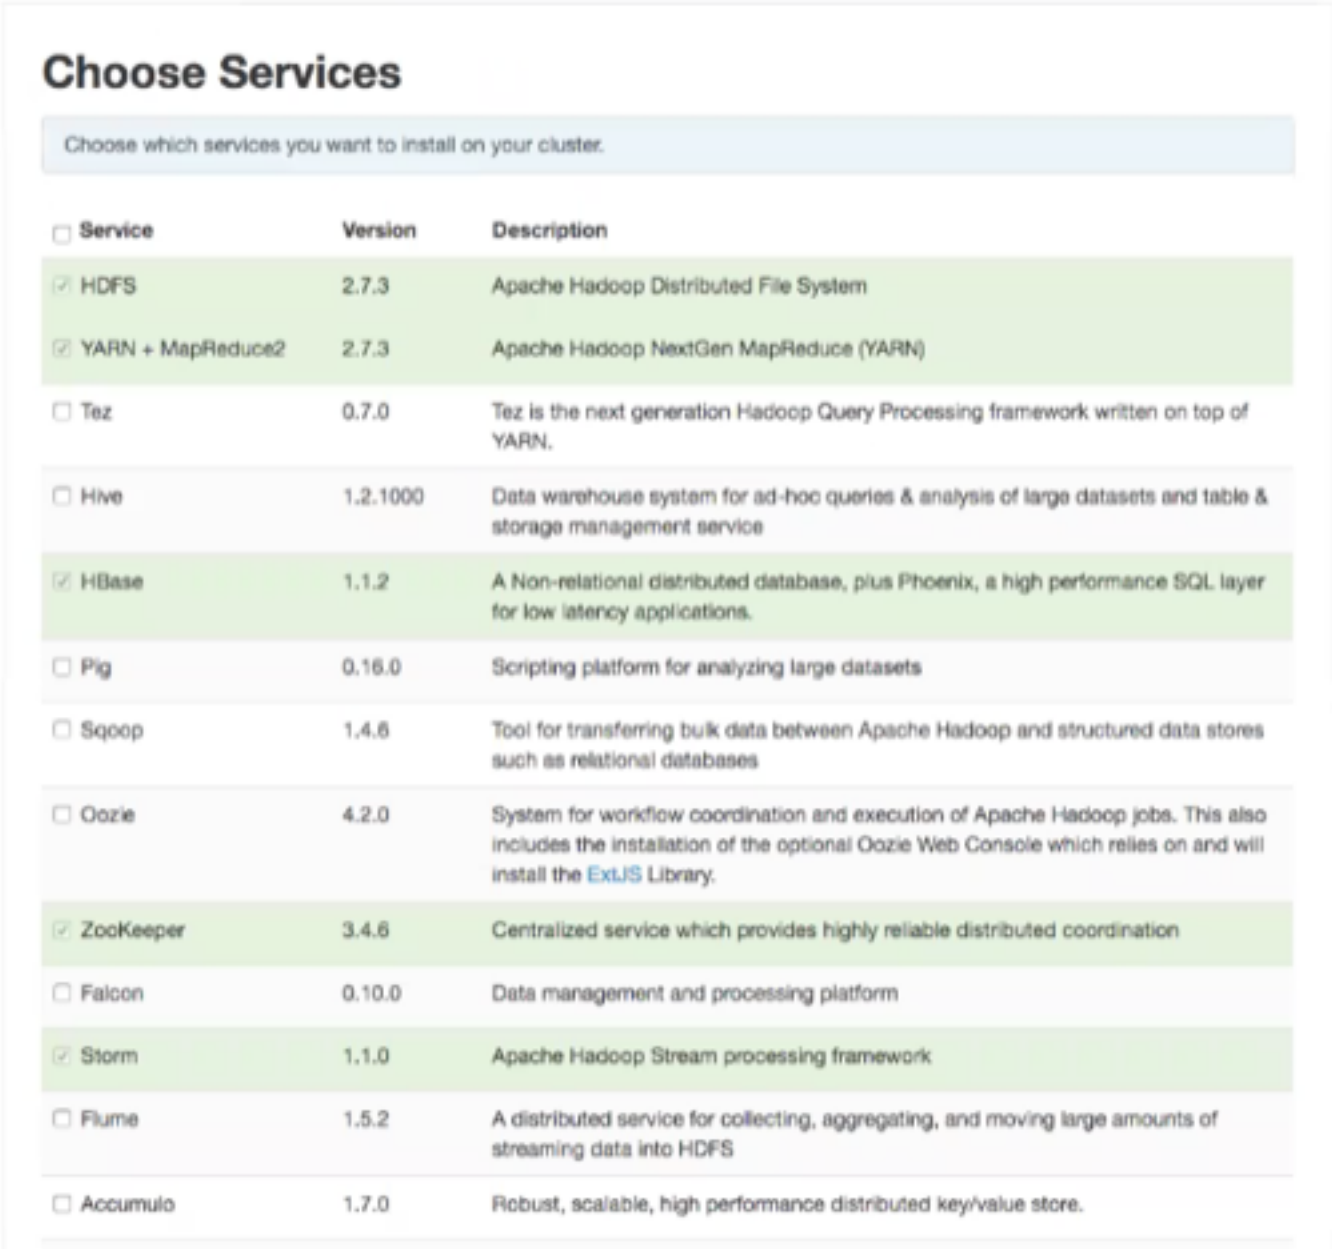

HCP requires the following services:

- HDFS
- HBase
- ZooKeeper
- Storm
- Kafka
- Ambari Metric Service
- Metron
- Elasticsearch (Can be installed either manually or by Ambari. Hortonworks recommends installing Elasticsearch by Ambari.)
- Kibana (Can be installed either manually or by Ambari. Hortonworks recommends installing Kibana by Ambari.)
- Zeppelin Notebook
- Spark
- Hive
- Tez
- Yarn

Ambari displays the **Assign Masters** window.

8. Use the **Assign Masters** window to assign the Master components to the appropriate hosts in your cluster.

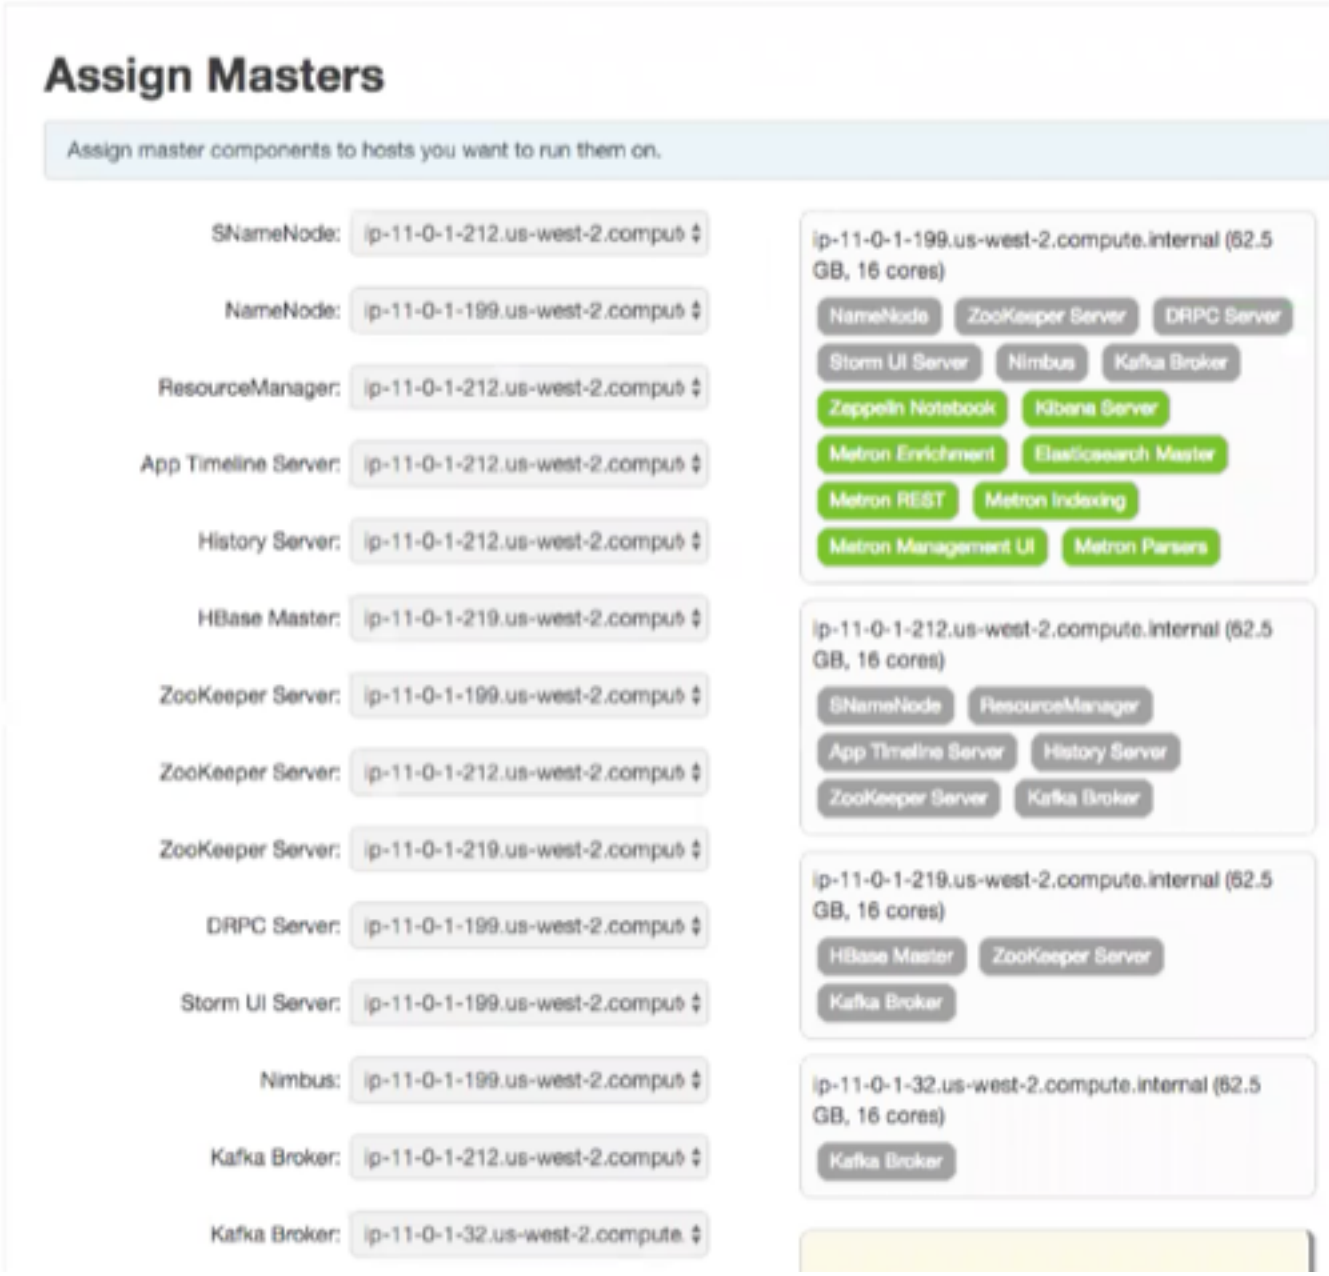

If Ambari detects any errors in your master component assignments, it will indicate the error in red.

a) To change the host assignment for a service, select a host name from the drop-down menu for that service.
b) To remove a ZooKeeper instance, click the green minus icon next to the host address you want to remove.
c) When you are satisfied with the assignments, click **Next**.

   Ambar displays the **Assign Slaves and Clients** window.

9. Use the **Assign Slaves and Clients** window to assign cluster nodes (DataNodes, NodeManagers, and RegionServers) to run with worker processes such as Elasticsearch.

a) Use all or none to select all of the hosts in the column or none of the hosts, respectively.

   If a host has an asterisk next to it, that host is also running one or more master components. Hover your mouse over the asterisk to see which master components are on that host.
b) Select a minimum of one Elasticsearch data node. The data node cannot be on same host as the master.
c) Fine-tune your selections by using the check boxes next to specific hosts.
d) Check the **Client** checkbox for any components that have the **Supervisor** checkbox checked.

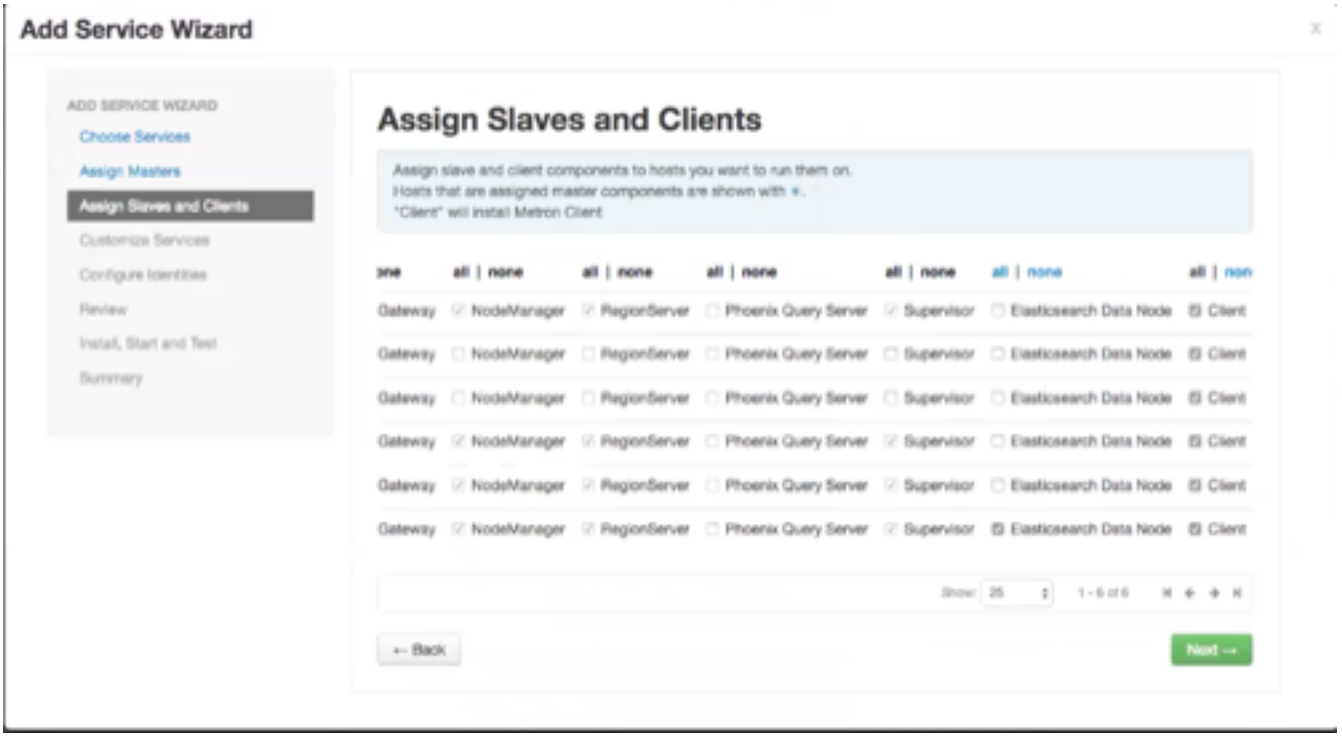

e) When you are satisfied with your assignments, click **Next**.

10. Use the **Customize Services** window to configure or customize cluster service property settings.

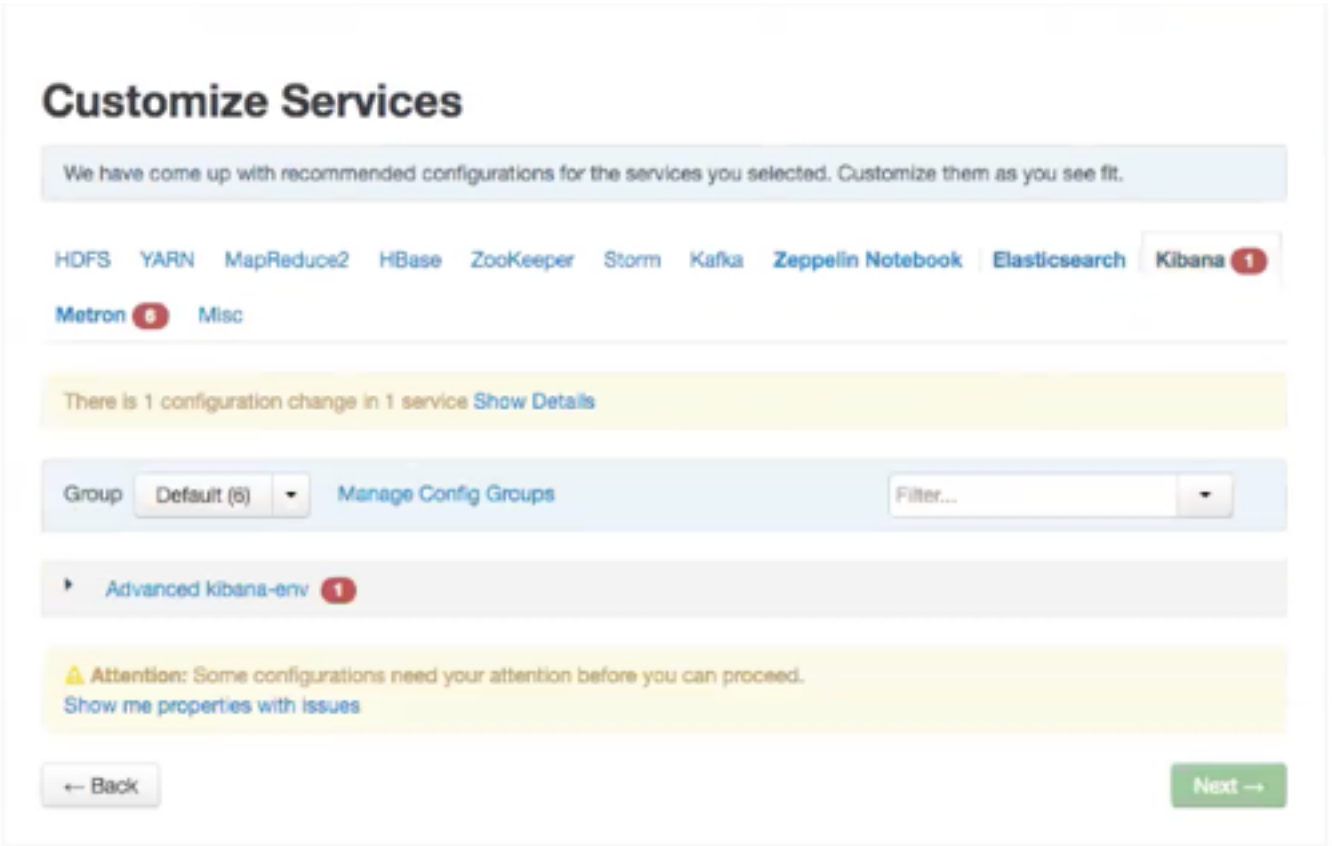

a) Browse through each service tab.

Any tab that requires input displays a red badge with the number of properties that need attention. Select each service tab that displays a red badge number and enter the appropriate information.

By hovering your cursor over each of the properties, you can see a brief description of what the property does. The number of service tabs shown depends on the services you decided to install in your cluster.

The following is a list of service tabs for which you need to provide information:

Kibana

kibana_es_url

Set to the fully-qualified url for the Elasticsearch master: http://es-master-host:9200.

Metron

The Metron tab contains a few tabs that contain information that is critical to HCP set up.

- Index Settings

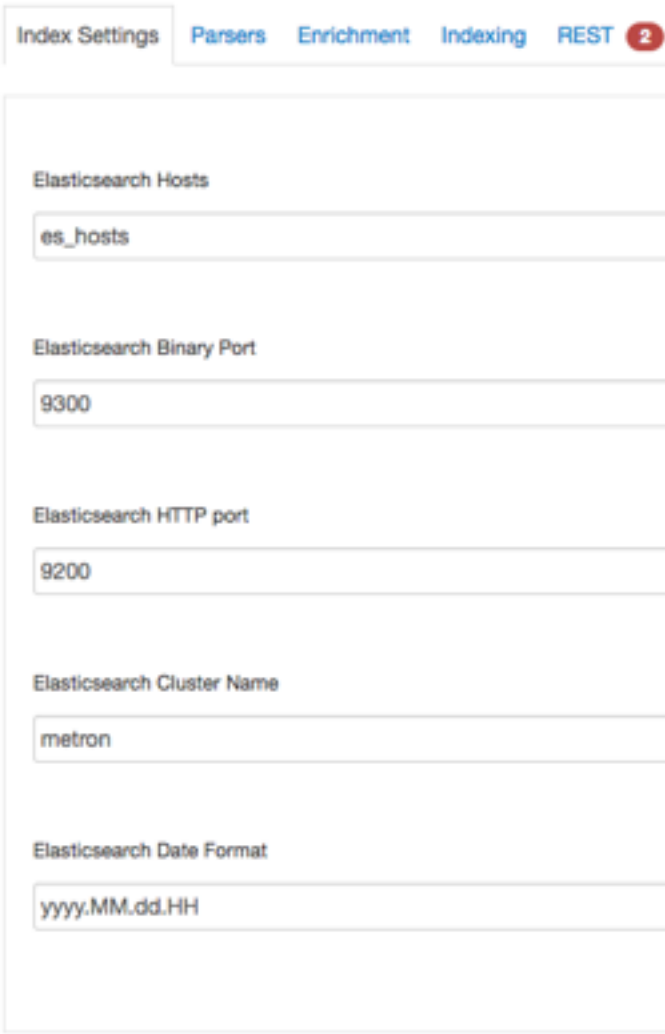

  tab
- Elasticsearch Hosts
- A comma separated list of Elasticsearch data nodes that you identified in Step 10.

REST tab

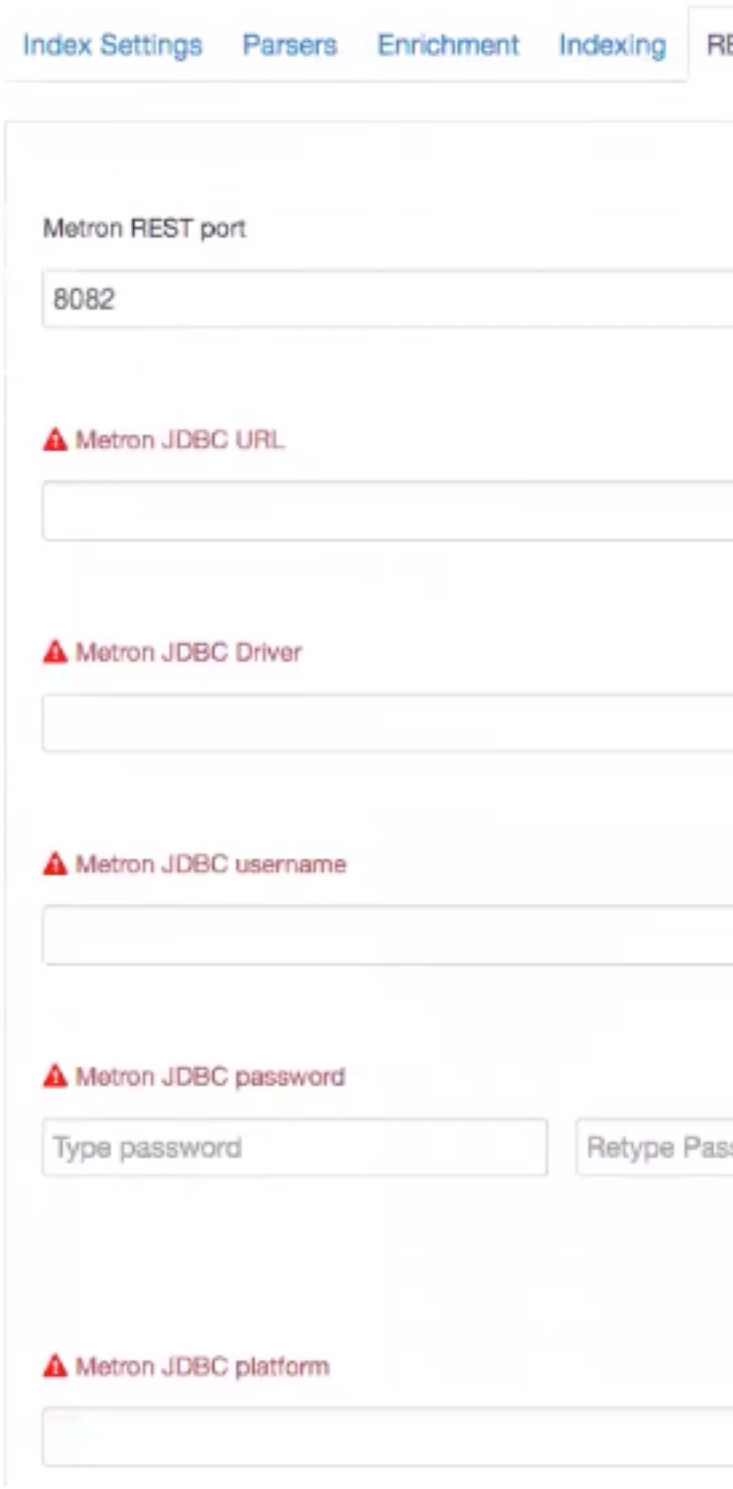

	Metron REST port Use 8082.
JDBC URL	jdbc:mysql://mysql_host:3306/metronrest
JDBC Driver	com.mysql.jdbc.Driver
	You can choose between the following databases for the REST configuration.

	• PostgreSQL • MySQL • H2 • Oracle
JDBC Username	Metron REST user name
JDBC Password	Metron REST password
Metron JDBC client path	<MYSQL_JAVA_CONNECTOR_PATH>/mysql-connector-java-5.1.41-bin.jar
Advanced Tab (Metron)	Most of the fields in the Advanced tab are auto populated and should not be modified.
Misc tab	The service account users and groups are available under the **Misc** tab. These are the operating system accounts the service components will run as. If these users do not exist on your hosts, Ambari will automatically create the users and groups locally on the hosts. If these users already exist, Ambari will use those accounts. Depending on how your environment is configured, you might not allow groupmod or usermod operations. If this is the case, you must be sure all users and groups are already created and be sure to select the **Skip group modifications** option on the **Misc** tab. This tells Ambari to not modify group membership for the service users.

11. OPTIONAL: Switch to using LDAP to define access privileges.

HCP defaults to Java Database Connectivity (JDBC) to define access privileges. You can easily switch to using LDAP.

a) Start your LDAP tool.
b) In Ambari, click **Metron** in the **Actions** menu and then click the **Config** tab.
c) Click the **Security** tab, and set **LDAP Enables** to **On**.
d) Modify each of the fields to match your LDAP configuration.

LDAP Trustore and **LDAP Truststore Password** are only used with LDAP enabled SSL Ensure proper certificcates are imported into a trustore for use by HCP.

e) Set the **Bind user password** to match the admin user's password.
f) Click the **Summary** tab to display all of the Metron components.
g) Restart **Metron REST**.

Now, when you go to Swagger or the UIs, you should be able to view your assigned roles and permissions.

12. OPTIONAL: Configure the PCAP topology by setting your PCAP properties in the **PCAP** tab.

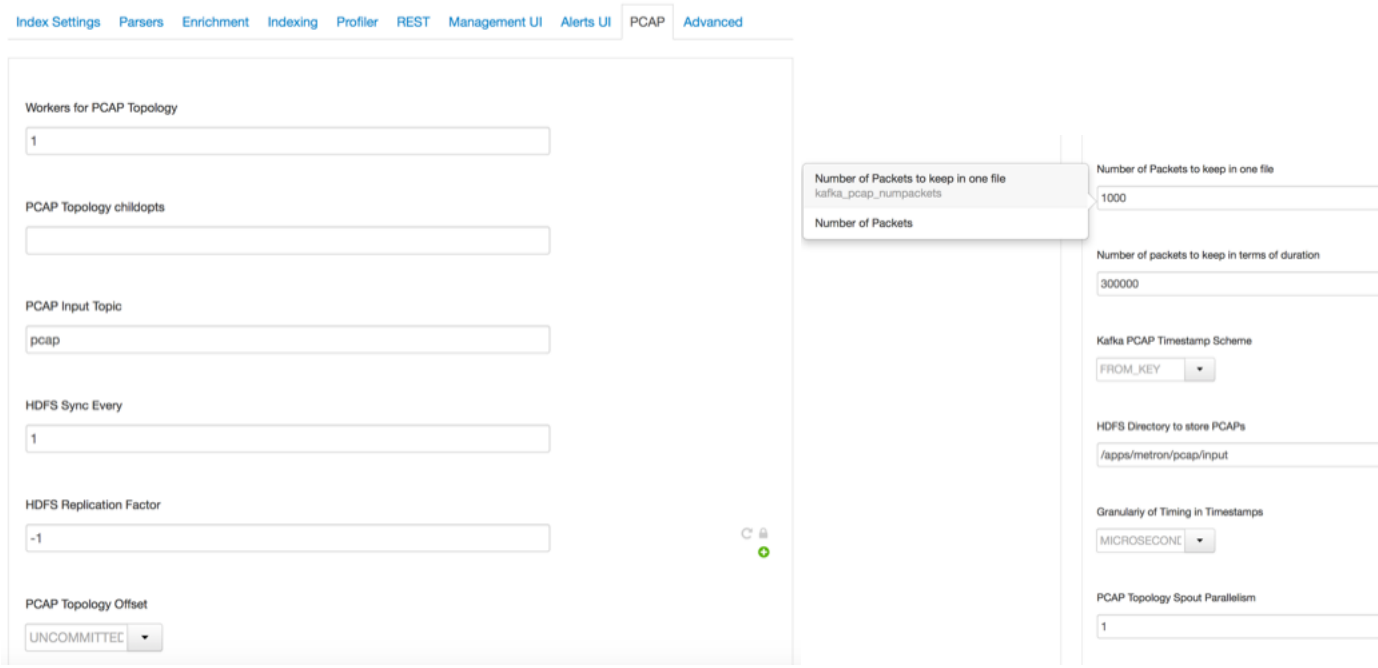

13. Check the assignments displayed by Ambari to ensure that everything is correct, and then click **Deploy**.

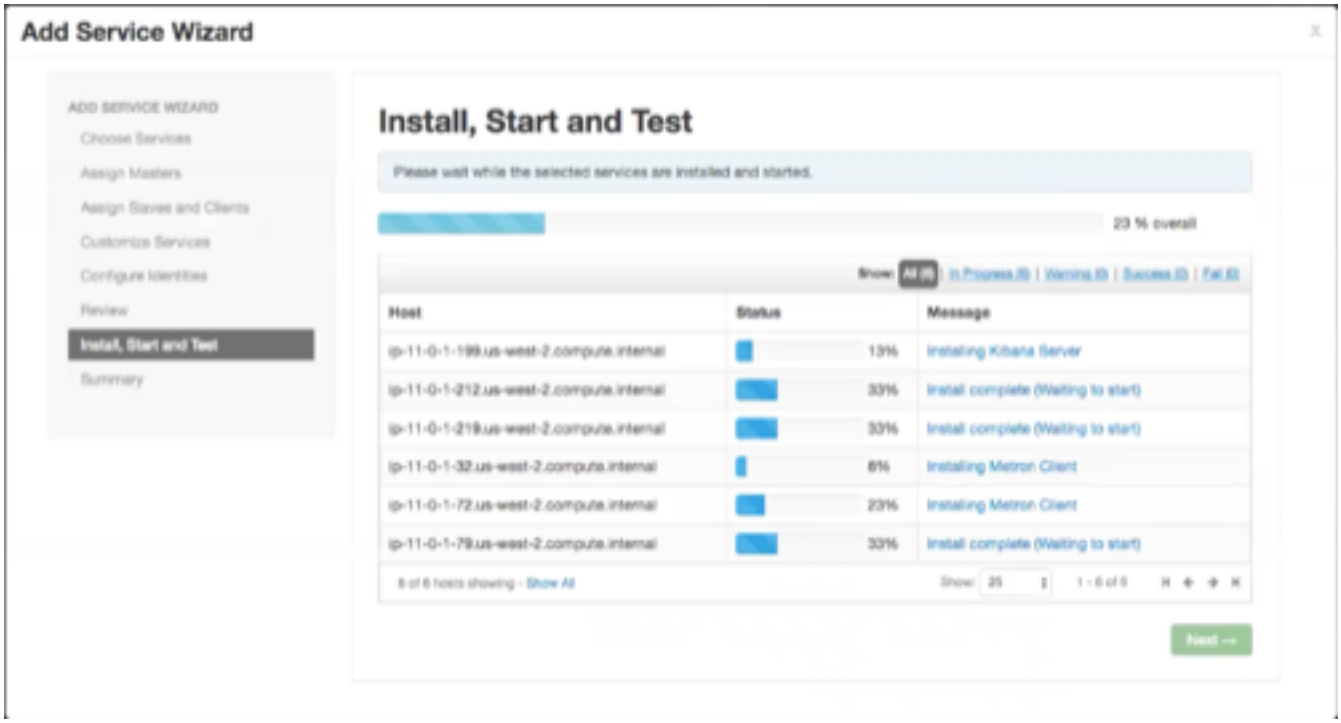

14. If you need to make changes, use the left navigation bar to return to the appropriate screen.

    The progress of the install displays on the screen. Ambari installs, starts, and runs a simple test on each component. Overall status of the process displays in a progress bar at the top of the screen and host-by-host status displays in the main section. Do not refresh your browser during this process. Refreshing the browser might interrupt the progress indicators.

15. OPTIONAL: To see specific information on what tasks have been completed per host, click the link in the **Message** column for the appropriate host. In the **Tasks** pop-up, click the individual task to see the related log files. You can select filter conditions by using the **Show** drop-down list. To see a larger version of the log contents, click the **Open** icon or, to copy the contents to the clipboard, use the **Copy** icon.

16. When Successfully installed and started the services appears, click **Next**.

Import Apache Zeppelin Notebook Using Ambari

If you would like to install Apache Zeppelin, complete the following steps after you have successfully installed HCP. You can use the Apache Zeppelin dashboard to view and analyze telemetry data provided by HCP.

Procedure

1. Login to Ambari at http://$AMBARI_HOST:8080.
2. In Ambari, click **Metron>Service Actions>Zeppelin Notebook Import**.
   Ambari imports the Zeppelin Notebook.
3. Login to Zeppelin at http://$ZEPPELIN_HOST:9995.
4. Search for the notebook named **Metron - YAF Telemetry**.

Streaming Data into HCP

To prepare for HCP to ingest data source data into HCP, you must stream each raw event stream from the telemetry data source into its own individual Kafka topic. This applies to the telemetry data sources for which HCP includes parsers (for example, Bro, Snort, and YAF). Even though HCP includes parsers for these data sources, HCP does not install these data sources or ingest the raw data. This is something that you must do.

Depending on the type of data you are streaming into HCP, you can use one of the following methods:

NiFi	This type of streaming method works for most types of data sources.  **Note:** Ensure that the NiFi web application is using port 8089.
Performant network ingestion probes	This type of streaming method is ideal for streaming high volume packet data.
Real-time and batch threat intelligence feed loaders	This type of streaming method is used for real-time and batch threat intelligence feed loadNiFiers.

Verify That HCP Deployed Successfully for Ambari Install

After you install Hortonwork Cybersecurity Platform, you need to verify that your services are displayed in Ambari and that you can access the Metron Dashboard.

Procedure

1. Verify that the topologies bundled with HCP are deployed.
   From Ambari, navigate to **Storm** > **Quick Links** > **Storm UI**.
   You should see the following topologies listed:
   - Snort
   - pcap
   - YAF (Yet Another Flowmeter)
   - Bro Network Security Monitor
   - Indexing topology

2. Check that the enrichment topology has emitted some data.

   This could take a few minutes to show up in the Storm UI. The Storm enrichment topology UI should look something like the following:

   Storm UI with Enrichment Details

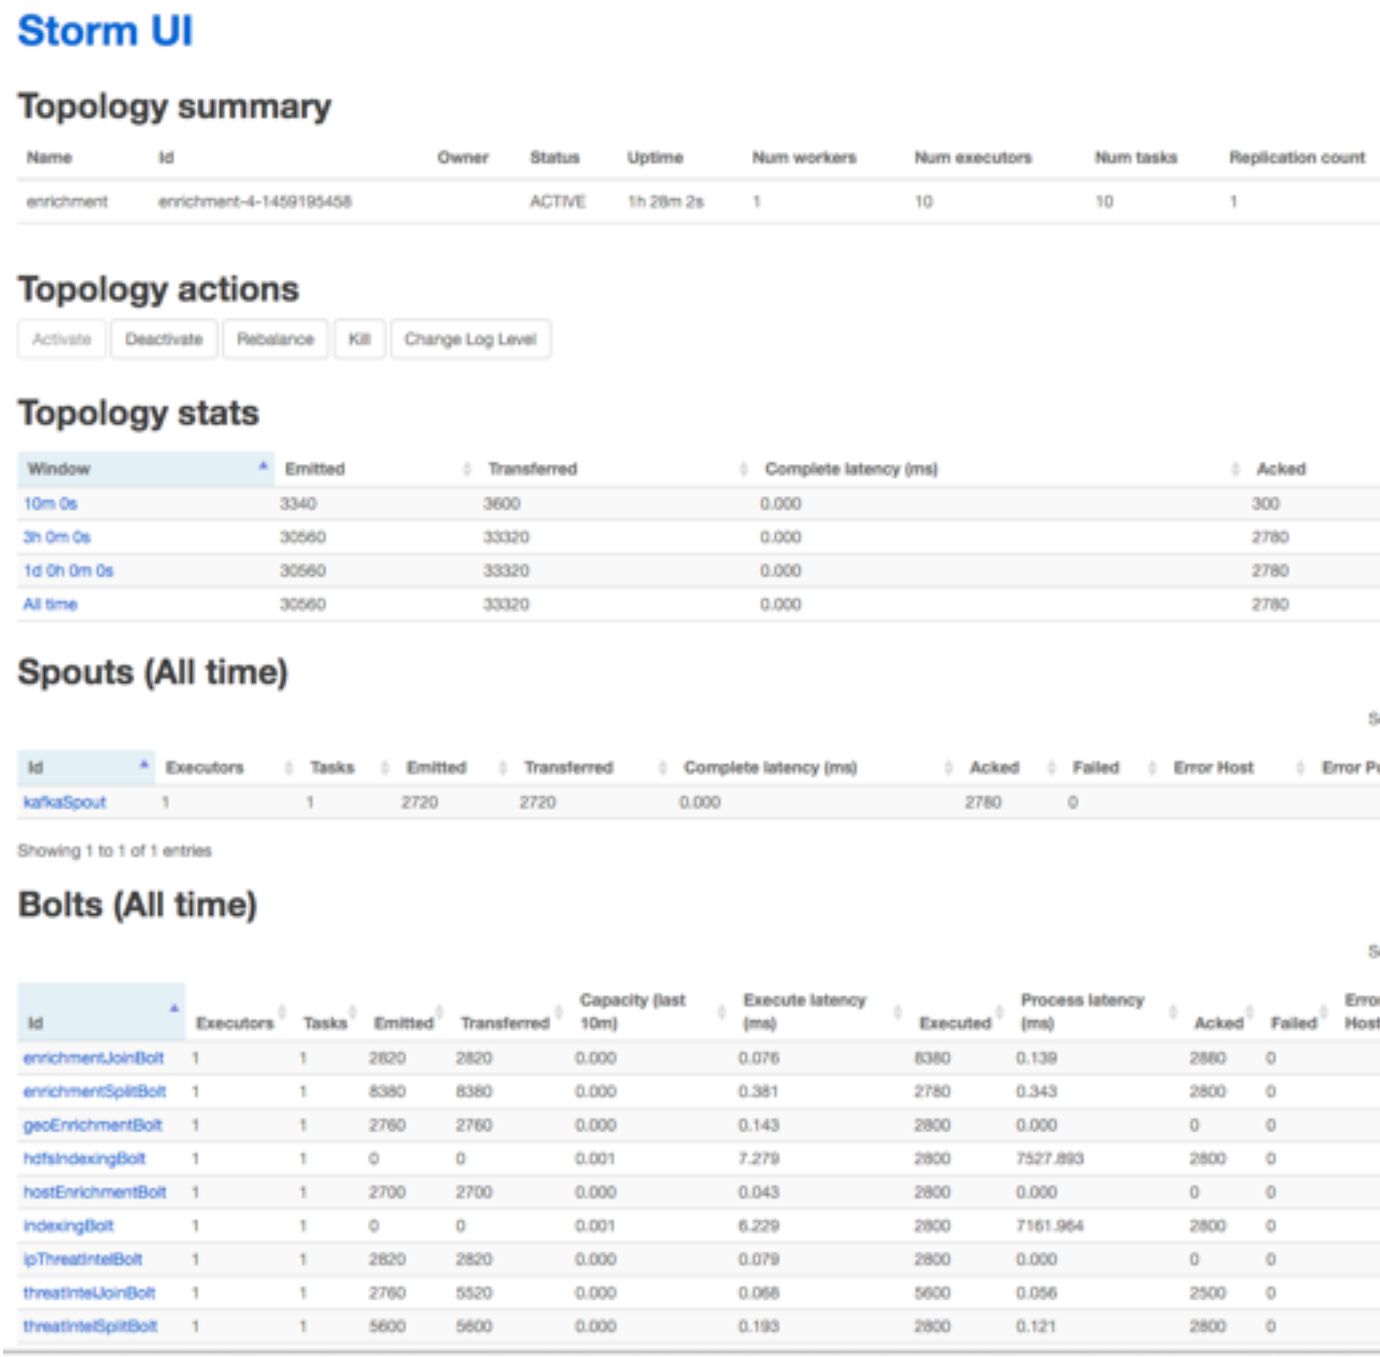

3. Ensure that the Metron dashboard is available and receiving data by displaying the dashboard at $METRON_UI_HOST:5000.

   Check to ensure that the indexing is done correctly and the data is visualized.

4. Check to ensure that some data is written into HDFS at /apps/metron for at least one of the data sources.

What to do next

Customize HCP to meet your own needs.

Open the Metron Dashboard

After you install and configure HCP, you can load and launch the Metron dashboard. The Metron dashboard enables you to identify, investigate, and analyze cybersecurity data.

Procedure

1. Ensure that you have selected **Metron** in the left navigation panel in Ambari.
2. From the **Service Action** menu, select **Kibana Dashboard Install**.
3. After the dashboard installs, click **Kibana** in the left navigation panel.
4. From the **Quick Links** pull-down menu, select **Metron UI**.

   The Metron dashboard should display in a separate browser tab.

What to do next

If you have already installed the Metron dashboard, reloading the dashboard will not overwrite your customizations to the dashboard. If you want to overwrite your customizations to the dashboard, you must delete the .kibana index from Elasticsearch and reload the Metron dashboard again from Ambari.

Opening the Management User Interface

You can use the HCP Management user interface to add and configure telemetry parsers to Hortonworks Cybsersecurity Platform (HCP). You can launch the UI either from Ambari or from a broswer.

Procedure

1. From the Ambari Dashboard navigation panel, click **Metron**.
2. Verify that the **Summary** tab is selected.
3. From the **Quick Links** menu, select **Management UI**.

   The Metron Management UI tool displays in a separate browser tab.

   Alternatively, you can launch the module from $METRON_MANAGEMENT_UI_HOST:4200 in a browser.

Opening the Alerts User Interface

You can use the Alerts user interface to display, filter, and sort events and their associated fields. You can also use the UI to escalate, add comments to, and group events.

Procedure

1. From the Ambari Dashboard navigation panel, click **Metron**.
2. Verify that the **Summary** tab is selected.
3. From the **Quick Links** menu, select **Alerts UI**.

   The Alerts UI tool displays in a separate browser tab.

Optimization Guidelines

In any Storm-based platform, there are many parameters that control the system’s performance. The values of these parameters vary greatly with differences in cluster size and data velocity. You will need to ensure that you have a properly tuned index is key to overall system performance. See the Storm user guide for detailed discussion.

- num.workers

- num.ackers
- max.spout.pending
- topology.worker.childopts – increase heap size (-XmxNNNNm –XmsNNNNm)
- topology.workers