Managing

Date of Publish: 2019-3-18

Contents

Managing

Hortonworks Cybersecurity Platform (HCP) powered by Apache Metron provides you with several options for managing your system. Before you perform any of these asks, you should become familiar with HCP data throughput.

Update Properties

HCP configuration information is stored in Apache ZooKeeper as a series of JSON files.

You can populate your ZooKeeper configurations from multiple locations:

- $METRON_HOME/config/zookeeper
- Management UI
- Ambari
- Stellar REPL

Because Ambari explicitly manages some of these configuration properties, if you change a property explicitly managed by Ambari from a mechanism outside of Ambari, such as the Management UI, Ambari is aware of this change and overwrites it whenever the Metron topology is restarted. Therefore, you should modify Ambari-managed properties only in Ambari.

For example, the es.ip property is managed explicitly by Ambari. If you modify es.ip and change the global.json file outside Ambari, you will not see this change in Ambari. Meanwhile, the indexing topology would be using the new value stored in ZooKeeper. You will not receive any errors notifying you of the discrepancy between ZooKeeper and Ambari. However, when you restart the Metron topology component via Ambari, the es.ip property would be set back to the value stored in Ambari.

Following are the Ambari-managed properties:

Table 1: Ambari-Managed Properties

Global Configuration Property Name	Ambari Name
es.clustername	es_cluster_name
es.ip	es_hosts
es.port	es_port
es.date.format	es_date_format
profiler.period.duration	profiler_period_duration
profiler.period.duration.units	profiler_period_units
update.hbase.table	update_hbase_table
update.hbase.cf	update_hbase_cf
geo.hdfs.file	geo_hdfs_file

Understanding ZooKeeper Configurations

ZooKeeper is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services.

ZooKeeper configurations should be stored on disk in the following structure starting at $METRON_HOME/bin/zk_load_configs.sh:

global.json   The global config

sensors	The subdirectory containing the sensor enrichment configuration JSON (for example, snort.json or bro.json

By default, the sensors directory as deployed by the Ansible infrastructure is located at $METRON_HOME/config/zookeeper.

Although the configurations are stored on disk, they must be loaded into ZooKeeper to be used. You can use the utility program $METRON_HOME/bin/zk_load_config.sh to load configurations into ZooKeeper.

-f,--force	Force operation
-h,--help	Generate Help screen
-i,--input_dir <DIR>	The input directory containing the configuration files named, for example $source.json
-m,--mode <MODE>	The mode of operation: DUMP, PULL, PUSH
-o,--output_dir <DIR>	The output directory that stores the JSON configuration from ZooKeeper
-z,--zk_quorum <host:port,[host:port]*>	The ZooKeeper Quorum URL (zk1:port,zk2:port,...)

See the following list for examples of usage: Usage examples:

* To dump the existing configs from ZooKeeper on the single-ode vagrant machine:

    $METRON_HOME/bin/zk_load_configs.sh -z node1:2181 -m DUMP
* To push the configs into ZooKeeper on the single-ode vagrant machine:

    $METRON_HOME/bin/zk_load_configs.sh -z node1:2181 -m PUSH -i $METRON_HOME/config/zookeeper
* To pull the configs from ZooKeeper to the single node vagrant machine disk:

    $METRON_HOME/bin/zk_load_configs.sh -z node1:2181 -m PULL -o $METRON_HOME/config/zookeeper -f

Managing Sensors

You can manage your sensors and associated topologies using either the Hortonworks Cybersecurity Platform (HCP) Management user interface or the Apache Storm UI. The following procedures use the HCP Management UI to manage sensors. For information about using Storm to manage sensors, see the Storm documentation.

Start a Sensor

After you install a sensor, you can start it using Management userinterface.
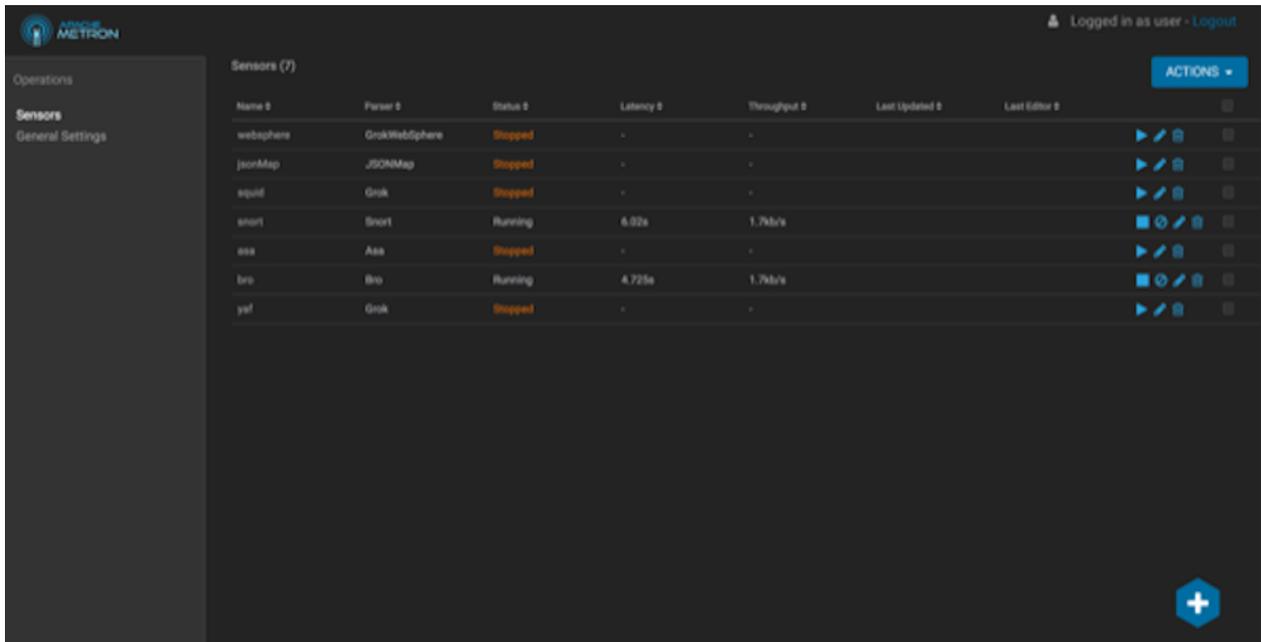
Procedure

From the main window, click

(start) in the

(tool bar) on the right side of the window.

Starting the sensor might take a few minutes. When the operation completes successfully, the Status value for the sensor changes to Running.

Stop a Sensor

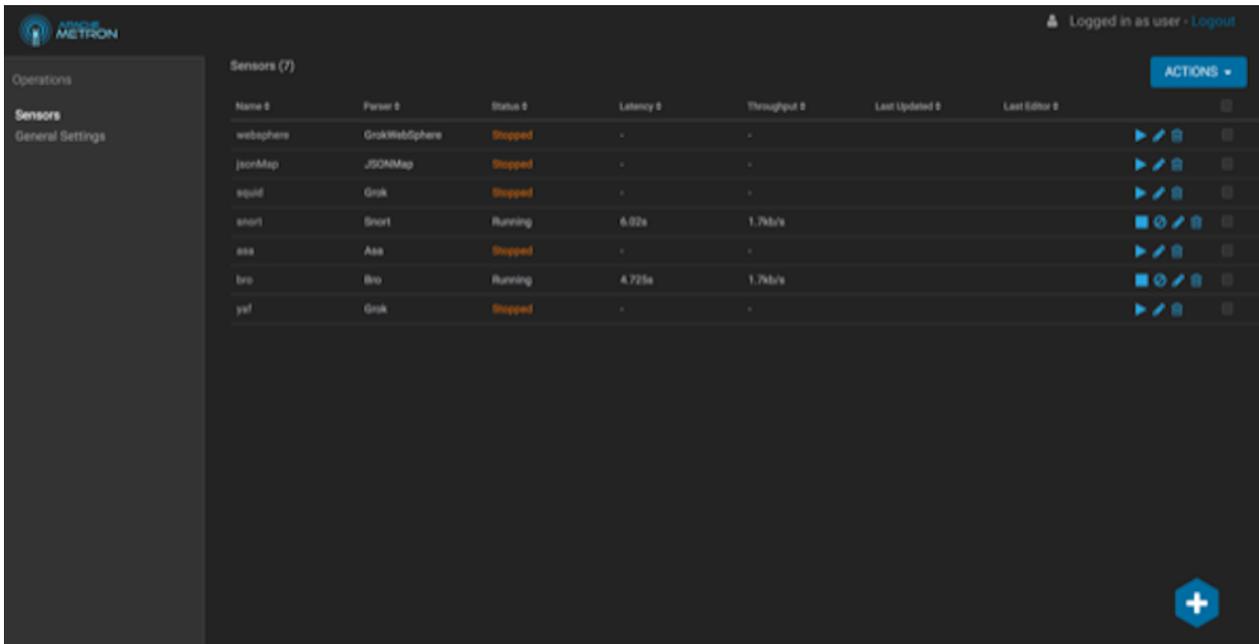After you install a sensor, you can stop it using the Management user interface.

Procedure

From the main window, click

(stop) in

(tool bar) on the right side of the window.

Stopping the sensor might take a few minutes. When the operation completes successfully, the Status value for the sensor changes to Stopped.

Modify a Sensor

You can modify any sensor listed in Hortonworks Cybersecurity Platform (HCP) Management user interface.
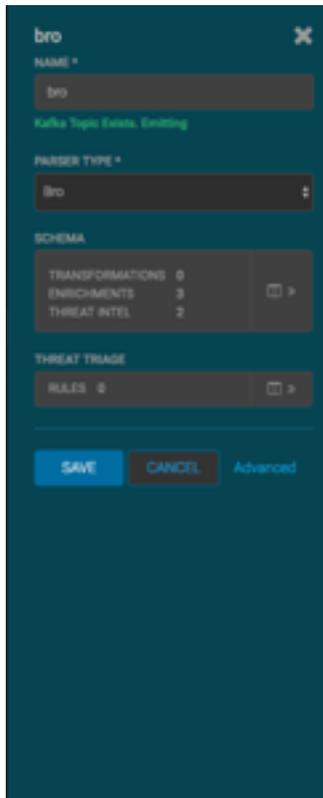
Procedure

1. From the **Operations** panel of the main window, select **Sensors**. click

(edit) for the sensor you want to modify.

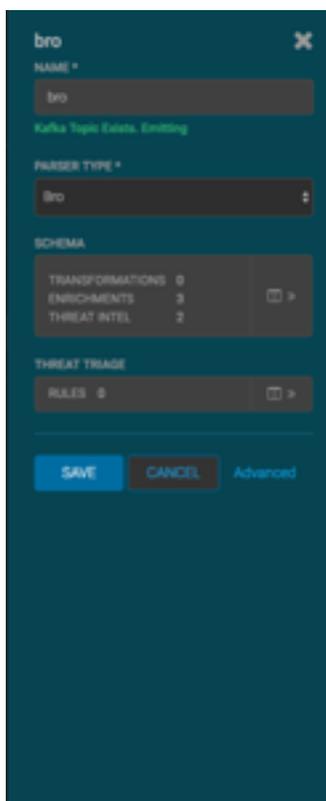The Management UI displays a panel populated with the sensor configuration information:

---

2. Click

(edit) for the sensor you want to modify.

The Management UI displays a panel populated with the sensor configuration information:

3. Modify the following information for the sensor, as necessary:

- Sensor name
- Parser type
- Schema information
- Threat triage information

4. Click **Save**.

Delete a Sensor

You can delete a sensor if you don't need it.

Before you begin

You must take the sensor offline before deleting it.

Procedure

1. In the Ambari user interface, click the **Services** tab.

2. Click **Metron** from the list of services.

3. Click **Configs** and then click **Parsers**.

4. Delete the name of the parser you want to delete from the **Metron Parsers** field.

5. Display the Management module.

6. Select the check box next to the appropriate sensor in the Sensors table.

   You can delete more than one sensor by clicking multiple check boxes.

7. From the **Actions** menu, select **Delete**.

   The Management module deletes the sensor from ZooKeeper.

8. Finally, delete the json file for the sensor on the Ambari master node:

```
ssh $AMBARI_MASTER_NODE
cd $METRON_HOME/config/zookeeper/parser
rm $DATASOURCE.json
```

Start and Stop Parsers

You might want to stop or restart parsers as you refine your cybersecurity monitoring. You can easily stop and start parsers by using Ambari.

Procedure

1. Display the Ambari UI and navigate to **Services** > **Metron** > **Summary**:

2. Click **Metron Parsers** to display the **Components** window.

The Components window displays a list of Metron hosts and which components reside on each host.

3. Click **Started/Stopped** to change the status of the parser; then click **OK** in the **Confirmation** dialog box.

Ambari displays the **Background Operation Running** dialog box which provides the status of the operation.

4. Click OK to exit the **Background Operation Running** dialog box.

Start and Stop Enrichments

You might want to stop or start enrichments as you refine or focus your cybersecurity monitoring. You can easily stop and start enrichments by using Ambari.

Procedure

1. Display the Ambari tool and navigate to **Services** > **Metron** > **Summary**.

   Ambari Metron Summary Window

2. Click **Metron Enrichments** to display the **Components** window.

   This window displays a list of HCP hosts and which components reside on each host.

   Components Window

3. Click the **Started/Stopped** button by **Metron Enrichments** to change the status of the Enrichments then click the **Confirmation** button to verify that you want to start or stop the enrichments.

   Ambari displays the **Background Operation Running** dialog box.

4. Click **Stop Metron Enrichments**.

Ambari displays the **Stop Metron Enrichments** dialog box.

5. Click the entry for your Metron cluster, then click **Metron Enrichments Stop** again.
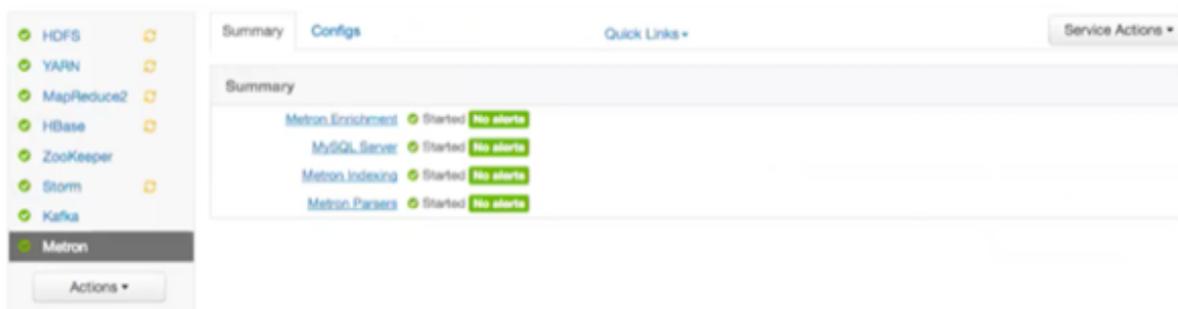
Ambari displays a dialog box for your Metron cluster which lists the actions as is stops the enrichments.

Start and Stop Indexing

You might want to stop or start indexing as you refine or focus your cybersecurity monitoring. You can easily stop and start indexing by using Ambari.

Procedure

1. Display the Ambari tool and navigate to **Services** > **Metron** > **Summary**.

2. Click **Metron Indexing**.

This window displays a list of HCP hosts and which components reside on each host.

3. Click **Started/Stopped** by **Metron Indexing** to change the status of the Indexing then .

   Ambari displays the **Background Operation Running** dialog box.

4. Click the **Confirmation** button to verify that you want to start or stop the indexing.

5. Click **Stop Metron Indexing**.

   Ambari displays the **Stop Metron Indexing** dialog box.

6. Click the entry for your Metron cluster, then click **Metron Indexing Stop** again.

   Ambari displays a dialog box for your Metron cluster which lists the actions as it stops the indexing.

Prune Data from Elasticsearch

Elasticsearch provides tooling to prune index data through its Curator utility.

Procedure

1. Use the following command to prune the Elasticsearch data:

   The following is a sample invocation that you can configure through Cron to prune indexes based on the timestamp in the index name.

```
/opt/elasticsearch-curator/curator_cli --host localhost delete_indices --
filter_list '
    {
      "filtertype": "age",
```

```
        "source": "name",
        "timestring": "%Y.%m.%d",
        "unit": "days",
        "unit_count": 10,
        "direction": "older"
    }'
```

Using name as the source value causes Curator to look for a timestring value within the index or snapshot name, and to convert that into an epoch timestamp (epoch implies UTC).

2. For finer-grained control over indexes pruning, provide multiple filters as an array of JSON objects to filter_list. Chaining multiple filters implies logical AND.

```
--filter_list
 '[{"filtertype":"age","source":"creation_date","direction":"older","unit":"days","un
{"filtertype":"pattern","kind":"prefix","value":"logstash"}]'
```

For finer-grained control over the indexes pruning that will be pruned, you can also provide multiple filters as an array of JSON objects to filter_list. Chaining multiple filters implies there is an implicit logical AND when chaining multiple filters.

```
--filter_list
 '[{"filtertype":"age","source":"creation_date","direction":"older","unit":"days","un
{"filtertype":"pattern","kind":"prefix","value":"logstash"}]'
```

Tune Apache Solr

To tune and customize Apache Solr, refer to the *Apache Solr Reference Guide*.

Back Up the Metron Dashboard

You can back up your Metron dashboard to avoid losing your customizations:

Procedure

To back up your Metron dashboard use the following command:

```
python packaging/ambari/metron-mpack/src/main/resources/common-services/
KIBANA/5.6.x/package/scripts/dashboard/dashboardindex.py \
  $ES_HOST 9200 \
 $SOME_PATH/dashboard.p -s
```

Restore Your Metron Dashboard Backup

You can restore a back up of your Metron dashboard by writing the Kibana dashboard to Solr or Elasticsearch.

Procedure

To restore a back up of your Metron dashboard, you can write the Kibana dashboard to Solr or Elasticsearch.

For example:

```
python packaging/ambari/metron-mpack/src/main/resources/common-services/
KIBANA/5.6.x/package/scripts/dashboard/dashboardindex.py \
  $ES_HOST 9200 \
 $SOME_PATH/dashboard.p
```

Note that this overwrites the .kibana index.