

Hortonworks Cybersecurity Package

Run Book - Technical Preview

(September 12, 2017)

Hortonworks Cybersecurity Package: Run Book - Technical Preview

Copyright © 2012-2017 Hortonworks, Inc. Some rights reserved.

Hortonworks Cybersecurity Package (HCP) is a modern data application based on Apache Metron, powered by Apache Hadoop, Apache Storm, and related technologies.

HCP provides a framework and tools to enable greater efficiency in Security Operation Centers (SOCs) along with better and faster threat detection in real-time at massive scale. It provides ingestion, parsing and normalization of fully enriched, contextualized data, threat intelligence feeds, triage and machine learning based detection. It also provides end user near real-time dashboards.

Based on a strong foundation in the Hortonworks Data Platform (HDP) and Hortonworks DataFlow (HDF) stacks, HCP provides an integrated advanced platform for security analytics.

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. Adding a New Telemetry Data Source	2
2.1. Prerequisites	2
2.2. Stream Data into HCP	3
2.2.1. OPTIONAL: Install the Squid Package	4
2.2.2. Install NiFi	4
2.2.3. Create a NiFi Flow to Stream Events to HCP	4
2.3. Parse the Squid Data Source to HCP	10
2.3.1. Parse the Squid Telemetry Event	10
2.3.2. Verify That the Events Are Indexed	18
2.3.3. Create an Index Template	19
2.4. Add New Data Source to the Metron Dashboard	20
2.4.1. Configuring a New Data Source Index in the Metron Dashboard	21
2.4.2. Reviewing the New Data Source Data	22
3. Transform the Squid Message	24
4. Enriching Telemetry Events	27
4.1. Bulk Loading Enrichment Information	28
4.1.1. OPTIONAL: Create a Mock Enrichment Source	29
4.1.2. Configuring an Extractor Configuration File	29
4.1.3. Configuring Element-to-Enrichment Mapping	31
4.1.4. Running the Enrichment Loader	32
4.2. Mapping Fields to HBase Enrichments	33
4.3. OPTIONAL: Global Configuration	36
4.4. Verify That the Events Are Enriched	36
5. Enriching Threat Intelligence Information	38
5.1. OPTIONAL: Create a Mock Threat Intel Feed Source	38
5.2. Configuring an Extractor Configuration File	38
5.3. Configuring Element-to-Threat Intel Feed Mapping	40
5.4. Run the Threat Intelligence Loader	41
5.5. Mapping Fields to HBase Enrichments	42
5.6. Verify That the Threat Intel Events Are Enriched	44
6. Prioritizing Threat Intelligence	46
6.1. Prerequisites	46
6.2. Performing Threat Triage	46
6.3. Viewing Triaged or Scored Alerts	49
6.3.1. Threat Triage Alert Panel	49
6.3.2. HCP Metron Dashboard View of Alerts	49
7. Configuring Indexing	51
7.1. Default Configuration	51
7.2. Specifying Index Parameters	52
7.3. Turning off HDFS Writer	53
8. Setting Up a Profile	55
8.1. Configuring the Profiler	55
8.2. Starting and Stopping the Profiler	58
8.3. Developing Profiles	59

List of Figures

2.1. Add Processor Dialog Box	5
2.2. New TailFile Processor	5
2.3. Configure Processor Dialog Box Settings Tab	6
2.4. Configure Processor Dialog Box Properties Tab	6
2.5. Configure Processor	7
2.6. Configure Properties	7
2.7. Create Connection Dialog Box	8
2.8. Data Flow	8
2.9. NiFi Operate Panel	9
2.10. New Sensor Panel	11
2.11. Grok Validator Panel with Sample Text	13
2.12. Grok Validator Panel with Statement Information	14
2.13. Grok Validator Panel with First Element and Test Results	15
2.14. Grok Validator Panel with Second Element and Test Results	16
2.15. mm_grok_validator_squid_complete.png	17
2.16. Elasticsearch	19
2.17. Ambari Task List	21
2.18. Configure an Index Pattern	22
2.19. Discover Tab with Squid Elements	23
3.1. Squid Schema Panel	24
3.2. New Schema Information Panel	25
3.3. Populated Schema Information Panel	26
4.1. New Schema Information Panel	35
4.2. Populated New Schema Information Panel	36
4.3. Elasticsearch *** Needs updating ***	37
5.1. New Schema Information Panel	43
5.2. Populated New Schema Information Panel	44
5.3. Elasticsearch	45
6.1. Threat Triage Rules Panel	47
6.2. Edit Rule Panel	48
6.3. Investigation Module Triaged Alert Panel	50
7.1. Management Module Advanced Panel	53
8.1. Ambari Profiler Properties	56
8.2. Enrichment Output Topic	58

List of Tables

8.1. Profiler Properties 56

1. Overview

This guide is intended for Platform Engineers and others who are responsible for adding new telemetry data sources, enriching telemetry events, triaging threat intelligence information, and ensuring telemetry events are viewable in the user interface.

This guide walks you through how to add a specific new data telemetry: Squid proxy logs. Squid is a caching proxy for the Web supporting HTTP, HTTPS, FTP, and more. It reduces bandwidth and improves response times by caching and reusing frequently-requested web pages. For more information on Squid, see squid-cache.org.

Unlike other HCP documentation, this guide provides detailed examples that are populated with information specific to the Squid data source.

Task	Description
Adding a new data source	This section describes how to add a telemetry data source to HCP.
Transforming the data source information	You can customize your sensor data to provide more meaningful data.
Adding enrichment information to the data source information	After the raw security telemetry events have been parsed and normalized, the next step is to enrich the data elements of the normalized event.
Enriching threat intelligence information	You can enrich your threat intelligence information just like you enriched your telemetry information.
Prioitizing the threat intelligence information	In HCP, you assign severity by associating possibly complex conditions with numeric scores. Then, for each message, you use a configurable aggregation function to evaluate the set of conditions and to aggregate the set of numbers for matching conditions. This aggregated score is added to the message in the <code>threat.triage.level</code> field.
Configuring indexing	The indexing topology is a topology dedicated to taking the data from a topology that has been enriched and storing the data in one or more supported indices. More specifically, the enriched data is ingested into Kafka, written in an indexing batch or bolt with a specified size, and sent to one or more specified indices. The configuration is intended to configure the indexing used for a given sensor type (for example, snort).
Setting up a Profile	A profile describes the behavior of an entity on a network. An entity can be a server, user, subnet, or application. Once you generate a profile defining what normal behavior looks like, you can build models that identify anomalous behavior.

This guide assumes that you have met all of the HCP 1.3.0 prerequisites and successfully installed HCP 1.3.0.

2. Adding a New Telemetry Data Source

This section describes how to add a telemetry data source to HCP.

For our examples, we will use the Squid data source as our new telemetry data source.

To add the Squid telemetry data source, perform the following tasks:

Task	Description
Meet Prerequisites	Before you can add a new telemetry data source you must meet the user requirements listed in this section.
Stream Data Into HCP	The first step in adding a new data source telemetry is to stream all raw events from the telemetry data source into its own Kafka topic.
Parse the Squid Data Source to HCP	Parsers transform raw data (textual or raw bytes) into JSON messages suitable for downstream enrichment and indexing by HCP. You must create a parser for each new data source.
Add New Data Source to Dashboard	After a new data telemetry source has been added to HCP, you must add it to the Metron dashboard before you can create queries and filters for it and add telemetry panels displaying its data.

2.1. Prerequisites

Before you add a new telemetry device, you must perform the following actions:

- Install HDP and HDF, and then install HCP.

For information about installing HCP, see the [HCP Installation Guide](#).

- The following use case assumes the following user requirements:
 - Proxy events from Squid logs must be ingested in real-time.
 - Proxy logs must be parsed into a standardized JSON structure suitable for analysis by Metron.
 - In real-time, the Squid proxy events must be enriched so that the domain names contain the IP information.
 - In real-time, the IP within the proxy event must be checked against for threat intelligence feeds.
 - If there is a threat intelligence hit, an alert must be raised.
 - The SOC analyst must be able to view new telemetry events and alerts from Squid.
 - HCP must profile the Squid data and incorporate statistical features of the profile into the threat triage rules.
 - HCP must be able to deploy a machine learning model that derives additional insights from the stream.

- HCP must incorporate user's machine learning model into user's threat triage rule along with threat intel, static rules, and statistical rules.
- Set HCP values.

When you install HCP, you will set up several hosts. You will need the locations of these hosts, along with port numbers, and the Metron version. These values are listed below.

- KAFKA_HOST = The host where a Kafka broker is installed.
- ZOOKEEPER_HOST = The host where a ZooKeeper server is installed.
- PROBE_HOST = The host where your sensor, probes are installed. If don't have any sensors installed, pick the host where a Storm supervisor is running.
- SQUID_HOST = The host where you want to install SQUID. If you don't care, install SQUID on the PROBE_HOST.
- NIFI_HOST = Host where you will install NIFI. This should be the same host on which you installed Squid.
- HOST_WITH_ENRICHMENT_TAG = The host in your inventory hosts file that you put under the group "enrichment."
- SEARCH_HOST = The host where you have Elastic or Solr running. This is the host in your inventory hosts file that you put under the group "search". Pick one of the search hosts.
- SEARCH_HOST_PORT = The port of the search host where indexing is configured (for example, 9300).
- METRON_UI_HOST = The host where your Metron UI web application is running. This is the host in your inventory hosts file that you put under the group "web."
- METRON_VERSION = The release of the Metron binaries you are working with (for example, 0.2.0BETA-RC2).

2.2. Stream Data into HCP

The first step in adding a new data source telemetry is to stream all raw events from the telemetry data source into its own Kafka topic.

Complete the tasks provided in the following sections:

Task	Description
OPTIONAL: Install the Squid Package [4]	In this section we install the package for our Squid telemetry data source so we can generate data. If you have already installed your data source and have a log source, you can skip this section and move to the next one.
Install NiFi	If you haven't already done so, install NiFi.
Create a NiFi Flow to Stream Events to HCP	This section provides instructions to create a data flow to capture events from Squid and push them into HCP.

2.2.1. OPTIONAL: Install the Squid Package

In this section we install the package for our Squid telemetry data source so we can generate data. If you have already installed your data source and have a log source, you can skip this section and move to the next one.



Important

Squid and NiFi must be installed on the same server. NiFi cannot read the local files on another server, so they must be colocated.

1. ssh into \$SQUID_HOST,
2. Install and start Squid:

```
sudo yum install squid  
sudo service squid start
```

2.2.2. Install NiFi

If you haven't already done so, install NiFi. For information on installing NiFi, see [NiFi Command Line Installation](#).



Important

Squid and NiFi must be installed on the same server. NiFi cannot read the local files on another server, so they must be colocated.

NiFi is built to automate the flow of data between systems. As a result, NiFi works well to collect, ingest, and push data to HCP.



Important

NiFi cannot be installed on top of HDP, so you must install NiFi manually to use it with HCP.



Note

Ensure that the NiFi web application is using port 8089.

The port number is stored in `/usr/lib/nifi<version number>/conf/nifi.properties`.

2.2.3. Create a NiFi Flow to Stream Events to HCP

This section provides instructions to create a data flow to capture events from Squid and push them into HCP. For this task we will use NiFi to create two processors, one TailFile processor that will ingest Squid events and one PutKafka processor that will stream the information to Metron. When we create and define the PutKafka processor, Kafka will create a topic for the Squid data source. We'll then connect the two processors, generate some data in the Squid log, and watch the data flow from one processor to the other.


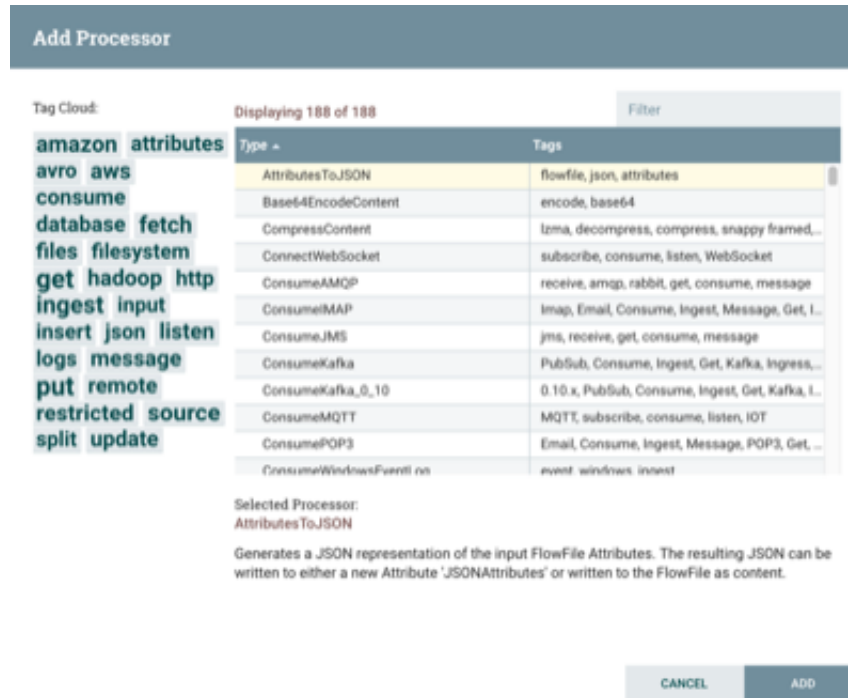
1. Drag the first icon on the toolbar  (the processor icon) to your workspace.
Nifi displays the Add Processor dialog box.

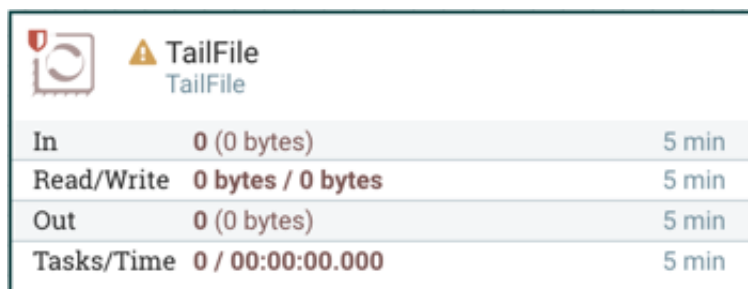
Figure 2.1. Add Processor Dialog Box



2. Select the TailFile type of processor and click **Add**.

NiFi displays a new TailFile processor.

Figure 2.2. New TailFile Processor



3. Right-click the processor icon and select **Configure** to display the **Configure Processor** dialog box.

- In the **Settings** tab, change the name to **Ingest Squid Events**.

Figure 2.3. Configure Processor Dialog Box Settings Tab

Configure Processor

SETTINGS SCHEDULING PROPERTIES COMMENTS

Name: Ingest Squid Events Enabled success Automatically Terminate Relationships
All FlowFiles are routed to this Relationship.

Id: 13a1a081-015c-1000-7972-7dc6816628b0

Type: TailFile

Penalty Duration: 30 sec Yield Duration: 1 sec

Bulletin Level: WARN

CANCEL APPLY

- In the **Properties** tab, enter the path to the squid `access.log` file in the **Value** column for the **File(s) to Tail** property.

Figure 2.4. Configure Processor Dialog Box Properties Tab

Configure Processor

SETTINGS SCHEDULING PROPERTIES COMMENTS

Required field +

Property	Value
Tailing mode	Single file
File(s) to Tail	/usr/log/squid/access.log
Rolling Filename Pattern	No value set
Base directory	No value set
Initial Start Position	Beginning of File
State Location	Local
Recursive lookup	false
Rolling Strategy	Fixed name
Lookup frequency	10 minutes
Maximum age	24 hours

CANCEL APPLY

4. Click **Apply** to save your changes and dismiss the **Configure Processor** dialog box.
5. Add another processor by dragging the **Processor** icon to the main window.
6. Select the **PutKafka** type of processor and click **Add**.
7. Right-click the processor and select **Configure**.

8. In the Settings tab, change the name to **Stream to Metron** and then select the Automatically Terminate Relationships check boxes for **Failure** and **Success**.

Figure 2.5. Configure Processor




The screenshot shows the 'Configure Processor' interface with the 'SETTINGS' tab selected. The 'Name' field is 'Stream to Metron' and the 'Enabled' checkbox is checked. The 'Id' is '13ada490-015c-1000-1805-77a61f75a5ab' and the 'Type' is 'PutKafka'. The 'Penalty Duration' is '30 sec' and the 'Yield Duration' is '1 sec'. The 'Bulletin Level' is set to 'WARN'. The 'Automatically Terminate Relationships' section has 'failure' and 'success' checked. The 'failure' checkbox is checked and the text below it says 'Any FlowFile that cannot be sent to Kafka will be routed to this Relationship'. The 'success' checkbox is checked and the text below it says 'Any FlowFile that is successfully sent to Kafka will be routed to this Relationship'. At the bottom right, there are 'CANCEL' and 'APPLY' buttons.

9. In the Properties tab, set the following three properties:

- Known Brokers: \$KAFKA_HOST:6667
- Topic Name: squid
- Client Name: nifi-squid

Figure 2.6. Configure Properties



The screenshot shows the 'Configure Processor' interface with the 'PROPERTIES' tab selected. A table lists properties and their values. The 'Client Name' property is highlighted in yellow.

Property	Value
Known Brokers	\$KAFKA_HOST:6667
Topic Name	squid
Partition Strategy	Round Robin
Partition	No value set
Kafka Key	No value set
Delivery Guarantee	Best Effort
Message Delimiter	No value set
Max Buffer Size	5 MB
Max Record Size	1 MB
Communications Timeout	30 secs
Batch Size	16384
Queue Buffering Max Time	No value set
Compression Codec	None
Client Name	nifi-squid

At the bottom right, there are 'CANCEL' and 'APPLY' buttons.

10. Click **Apply** to save your changes and dismiss the **Configure Processor** dialog box.

11. Create a connection by dragging the arrow from the Ingest Squid Events processor to the Stream to Metron processor.

NiFi displays a **Create Connection** dialog box.

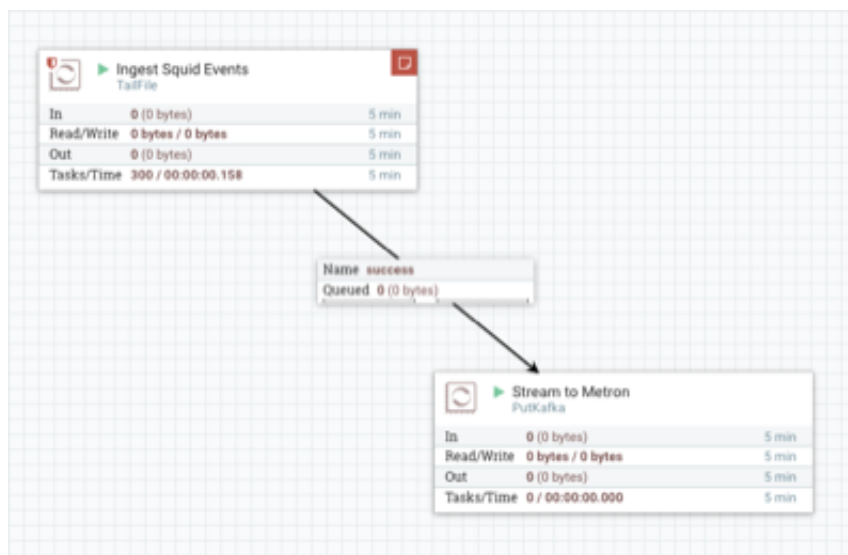
Figure 2.7. Create Connection Dialog Box



12. Click **Add** to accept the default settings for the connection.

13. Press the Shift key and draw a box around both parsers to select the entire flow.

Figure 2.8. Data Flow




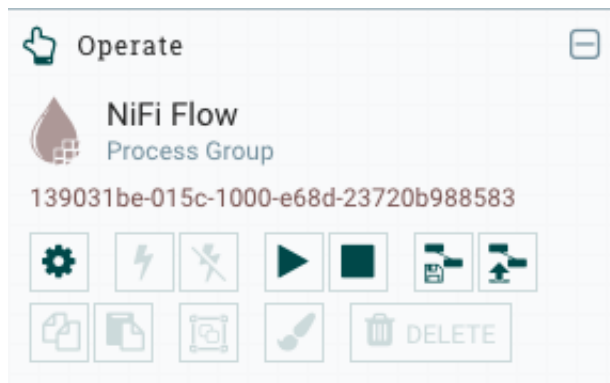
14. Click  (Start button) in the **Operate** panel.

Figure 2.9. NiFi Operate Panel



15. Add NiFi to the Squid group:

- a. Use ssh to access the host for the new data source.
- b. Add NiFi to the Squid group:

```
sudo usermod -G squid nifi
```

16. Generate some data using the new data processor client.

- a. With Squid started, look at the different log files that get created:

```
sudo su -  
cd /var/log/squid  
ls
```

The file you want for Squid is the `access.log`, but another data source might use a different name.

- b. Generate entries for the log so you can see the format of the entries.

```
squidclient -h 127.0.0.1 "http://www.atmape.ru"
```

You will see the following data in the `access.log` file.

```
1481143984.330 1111 127.0.0.1 TCP_MISS/301 714 GET http://www.atmape.  
ru/ - DIRECT/212.109.217.71 text/html
```

- c. Using the Squid log entries, you can determine the format of the log entries is:

```
timestamp | time elapsed | remotehost | code/status | bytes | method |  
URL rfc931 peerstatus/peerhost | type
```

You should see metrics on the processor of data being pushed into Metron.

17. Look at the Storm UI for the parser topology and you should see tuples coming in.

- Go to Ambari UI.
- From the **Quick Links** pull-down menu in the upper center of the window, select **Storm UI**.

18. Before leaving this section, run the following commands to fill the `access.log` with data. You'll need this data when you enrich the telemetry data.

```

squidclient -h 127.0.0.1 "http://www.aliexpress.com/af/shoes.html?ltype=
wholesale&d=y&origin=n&isViewCP=y&catId=0&initiative_id=SB_20160622082445&
SearchText=shoes"
squidclient -h 127.0.0.1 "http://www.help.land1.co.uk/domains-c40986/
transfer-domains-c79878"
squidclient -h 127.0.0.1 "http://www.pravda.ru/science/"
squidclient -h 127.0.0.1 "http://www.brightsideofthesun.com/2016/6/25/
12027078/anatomy-of-a-deal-phoenix-suns-pick-bender-chriss"
squidclient -h 127.0.0.1 "https://www.microsoftstore.com/store/msusa/en_US/
pdp/Microsoft-Band-2-Charging-Stand/productID.329506400"
squidclient -h 127.0.0.1 "https://tfl.gov.uk/plan-a-journey/"
squidclient -h 127.0.0.1 "https://www.facebook.com/Africa-Bike-
Week-1550200608567001/"
squidclient -h 127.0.0.1 "http://www.ebay.com/itm/02-Infiniti-QX4-Rear-
spoiler-Air-deflector-Nissan-Pathfinder-/172240020293?fits=Make%3AInfiniti
%7CModel%3AQX4&hash=item281a4e2345:g:iMkAAOSwoBtW4Iwx&vxp=mtr"
squidclient -h 127.0.0.1 "http://www.recruit.jp/corporate/english/company/
index.html"
squidclient -h 127.0.0.1 "http://www.lada.ru/en/cars/4x4/3dv/about.html"
squidclient -h 127.0.0.1 "http://www.help.land1.co.uk/domains-c40986/
transfer-domains-c79878"
squidclient -h 127.0.0.1 "http://www.aliexpress.com/af/shoes.html?ltype=
wholesale&d=y&origin=n&isViewCP=y&catId=0&initiative_id=SB_20160622082445&
SearchText=shoes"

```

For more information about creating a NiFi data flow, see the NiFi documentation.

2.3. Parse the Squid Data Source to HCP

Parsers transform raw data (textual or raw bytes) into JSON messages suitable for downstream enrichment and indexing by HCP. There is one parser for each data source and the information is piped to the Enrichment/Threat Intelligence topology.

The following sections guide you through how to add the Squid telemetry to HCP. The following procedures use the Management module UI whenever possible. If you would like to see these steps performed using the CLI, see the [Hortonworks Cybersecurity Package Administration Guide](#).

Task	Description
Parse the Squid Telemetry Event [10]	This section uses the Management module to create a parser for the new data source.
Verify that the Events are Indexed	After you finish parsing your new data source, you should verify that the data source events are indexed and the output matches any Stellar transformation functions you used.
Create an Index Template	To work with a new data source data in the Metron dashboard, you need to ensure that the data is landing in the search index (Elasticsearch) with the correct data types. You can achieve this by defining an index template.


2.3.1. Parse the Squid Telemetry Event

1. Launch the HCP Management module:

- a. From the Ambari Dashboard panel, click **Metron**.
- b. Make sure you have the **Summary** tab selected.
- c. Select the Metron Management UI from the Summary list.


The **Management UI** tool should display in a separate browser tab.

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

2. Click **Sensors** on the left side of the window, under **Operations**.
3. Click  (the add button) in the lower right corner of the screen.

The Management module displays a panel used to create the new sensor.

Figure 2.10. New Sensor Panel



4. In the **NAME** field, enter the name of the sensor.

For our example, we use the name **squid**.

If a Kafka topic already exists for the sensor name, the module displays a message similar to **Kafka Topic Exists. Emitting** and displays the name of the **Parser Type**. You can skip to the next section, [Verify Events are Indexed](#).

If no matching Kafka topic is found, the module displays **No Matching Kafka Topic**.

5. In the **Parser Type** field, choose the type of parser for the new sensor.

If you chose a Grok parser type, the module prompts for a Grok statement.

6. If no Kafka topic exists for squid, create a Grok statement for the new parser:

- a. In the Grok Statement box, click the  (expand window button) to display the Grok Validator panel.

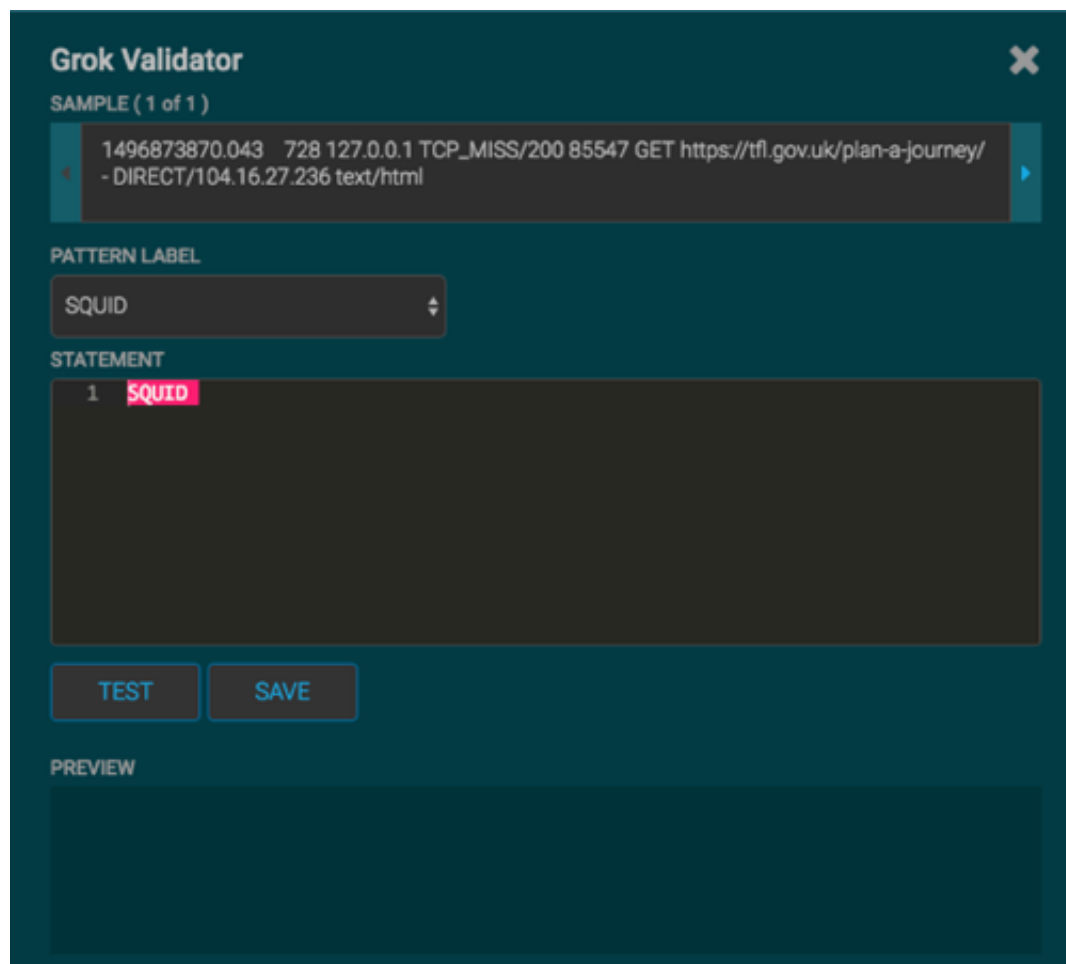
- b. Choose or enter a name for the Grok statement in the **PATTERN LABEL** field.

For our example, we chose SQUID.

- c. If there is no data in the **SAMPLE** text field, enter a sample log entry for the data source.

To create sample log entries, see [Step 15](#).

Figure 2.11. Grok Validator Panel with Sample Text



- d. Refer to the format of the log entries you determined in [Step 11d](#).

For example, the log entry format for Squid is:

```
timestamp | time elapsed | remotehost | code/status | bytes | method |  
URL rfc931 peerstatus/peerhost | type
```

- e. In the **STATEMENT** field, enter the first element in the sensor log file.

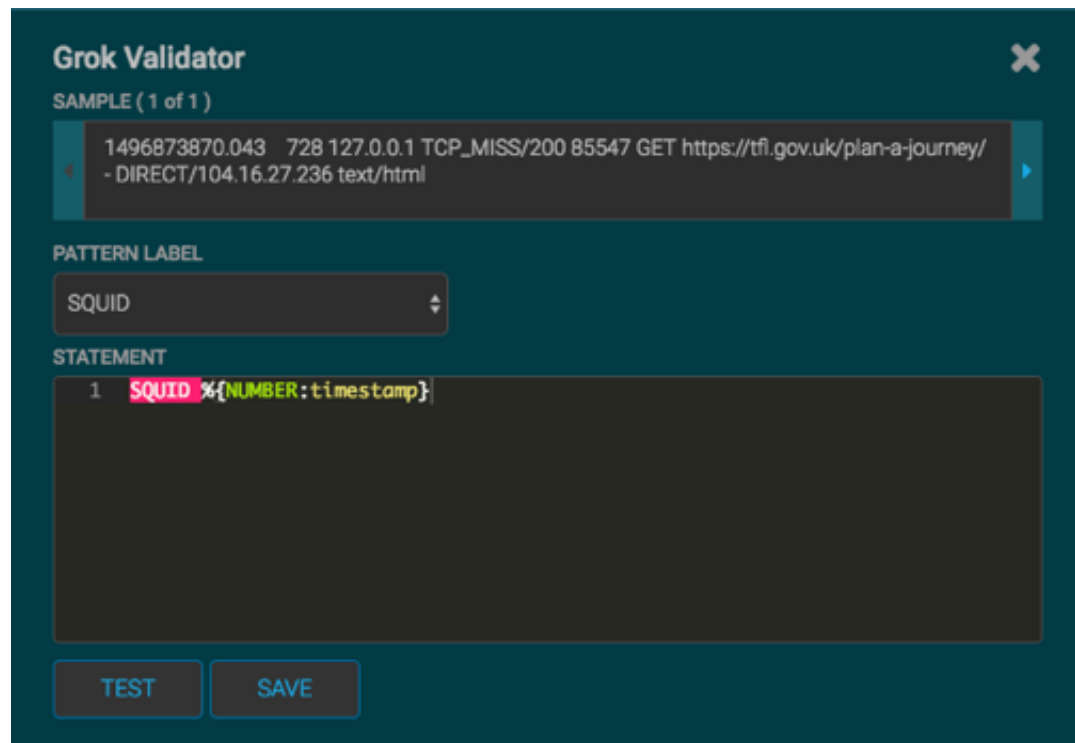
For Squid, the first element in the sensor log file is `timestamp` which is a number, so we enter `%{NUMBER:timestamp}`.



Note

The Management module automatically completes partial words in your Grok statement as you enter them.

Figure 2.12. Grok Validator Panel with Statement Information



f. Click **TEST**.

If the validator finds an error, it displays the error information. If the validation succeeds, it displays the valid mapping in the **PREVIEW** field.

Because you entered the timestamp element, the validator parses the timestamp correctly and leaves the rest of the information as random data.

Figure 2.13. Grok Validator Panel with First Element and Test Results



g. Refer to the log entry format and enter the second element.

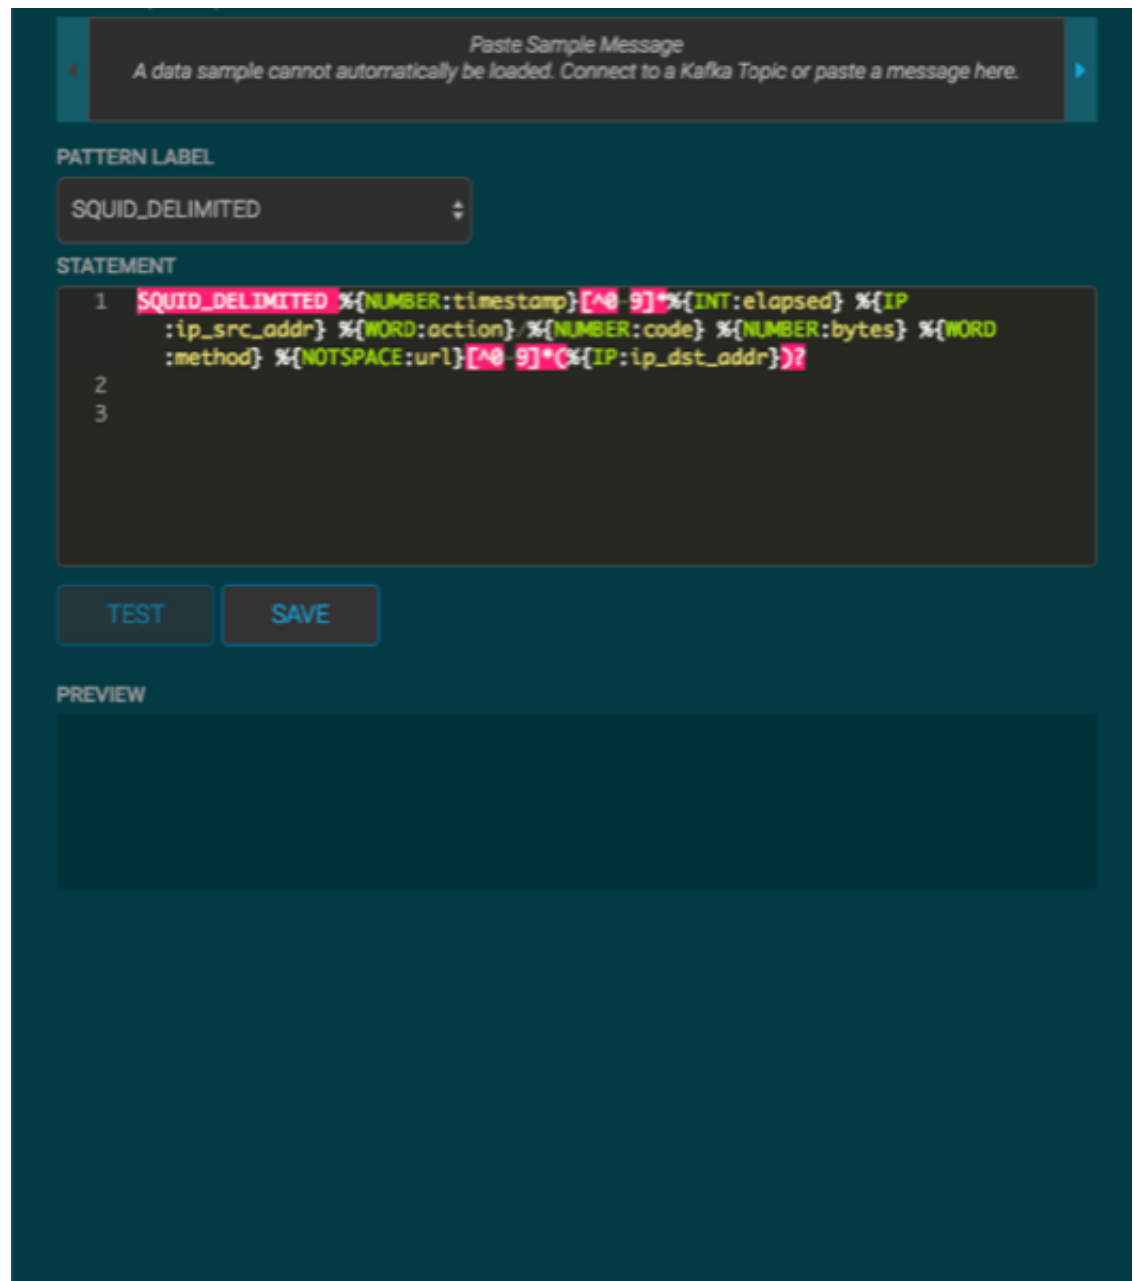
For Squid, the second element is `elapsed`, so we enter `%{INT:elapsed}` and click **TEST** again.

Figure 2.14. Grok Validator Panel with Second Element and Test Results



- h. Continue to build and test the Grok statement until you have entries for each element in the log entry.

Figure 2.15. mm_grok_validator_squid_complete.png



Note

You should perform the Grok validation using several different sensor log entries to ensure that the Grok statement is valid for all sensor logs. To display additional sensor log entries, click the forward or backward arrow icon on the side of the **SAMPLE** text box.

- i. When your Grok statement is complete and valid, click **SAVE** to save the Grok statement for the sensor.

Your Grok full statement should look like the following:

```
SQUID_DELIMITED %{NUMBER:timestamp}[^0-9]*%{INT:elapsed}  
%{IP:ip_src_addr} %{WORD:action}/%{NUMBER:code}  
%{NUMBER:bytes} %{WORD:method} %{NOTSPACE:url}[^0-9]*(%{IP:ip_dst_addr})?
```

7. Click **SAVE** to save the sensor information and add it to the list of Sensors.

This new data source processor topology ingests from the \$Kafka topic and then parses the event with the HCP Grok framework using the Grok pattern. The result is a standard JSON Metron structure that then is added to the "enrichment" Kafka topic for further processing.

8. Test that a Kafka topic has been created for the Squid parser:

- a. Navigate to the following directory:

```
/usr/hdp/current/kafka-broker/bin
```

- b. List all of the Kafka topics:

```
./kafka-topics.sh --zookeeper localhost:2181 --list
```

You should see the following list of Kafka topics:

- amari_kafka_service_check
- bro
- enrichments
- escalation
- indexing
- snort
- squid
- yaf

2.3.2. Verify That the Events Are Indexed

After you finish adding your new data source, you should verify that the data source events are indexed and the output matches any Stellar transformation functions you used.

By convention, the index where the new messages are indexed is called `squid_index_[timestamp]` and the document type is `squid_doc`.

Use the Elasticsearch Head plug-in to verify that the messages were indexed correctly:

1. Log in to \$SEARCH_HOST host:

```
ssh into Host $SEARCH_HOST
```

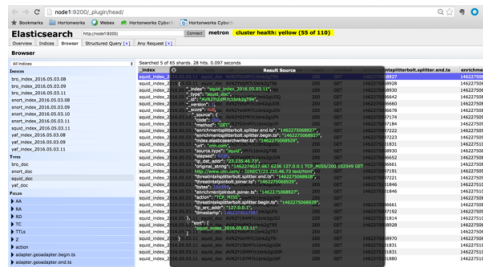
2. Install the head plug-in:

```
/usr/share/elasticsearch/bin/plugin install mobz/elasticsearch-head/1.x
```

3. Navigate to ElasticSearch Head UI: `http://$SEARCH_HOST:9200/_plugin/head/`.
4. Click the Browser tab and select the squid document in the left panel; then select one of the sample docs.

You should see something like the following:

Figure 2.16. Elasticsearch



5. Review the output to ensure it reflects the Stellar transformation functions you used.

2.3.3. Create an Index Template

To work with a new data source data in the Metron dashboard, you need to ensure that the data is landing in the search index (Elasticsearch) with the correct data types. You can achieve this by defining an index template.



Note

You will need to update the Index template after you add or change enrichments for a data source.

1. Run the following command to create an index template for the Squid.

```
curl -XPOST $SEARCH_HOST:$SEARCH_PORT/_template/squid_index -d '{
  "template": "squid_index*",
  "mappings": {
    "squid_doc": {
      "_timestamp": {
        "enabled": true
      },
      "properties": {
        "timestamp": {
          "type": "date",
          "format": "epoch_millis"
        },
        "ip_src_addr": {
          "type": "ip"
        },
        "ip_src_port": {
          "type": "integer"
        }
      }
    }
  }
}
```



```
"ip_dst_addr": {  
  "type": "ip"  
},  
"ip_dst_port": {  
  "type": "integer"  
}  
}  
}'
```

The example assumes the following:

- The template applies to any indices that are named `squid_index*`.
- The index has one document type that must be named `squid_doc`.
- The index contains timestamps.
- The properties section defines the types of each field. This example defines the five common fields that most sensors contain.
- Additional fields can be added following the five that are already defined.

By default, Elasticsearch will attempt to analyze all fields of type string. This means that Elasticsearch will tokenize the string and perform additional processing to enable free-form text search. In many cases, and all cases for the Squid data, you want to treat each of the string fields as enumerations. This is why most fields in the index template are ``not_analyzed``.

2. [OPTIONAL] An index template will only apply for indices that are created after the template is created. If you want to delete **existing** indices for the new data source so that new ones can be generated with the index template, run the following command.



Important

Running this command will delete all existing indices.

```
curl -XDELETE $SEARCH_HOST:9200/squid*
```

3. Wait for the Squid index to be re-created.

This might take a minute or two based on how fast the new data source data is being consumed in your environment.

```
curl -XGET $SEARCH_HOST:9200/$DATASOURCE*
```

2.4. Add New Data Source to the Metron Dashboard

After a new data telemetry source has been added to HCP, you will need to add it to the Metron dashboard before you can create queries and filters for it and add telemetry panels displaying its data. Complete the steps in the following sections to add a new telemetry source to the Metron dashboard:

Task	Description
Configuring a New Data Source Index in the Metron Dashboard [21]	Now that you have an index for the new data source with all of the right data types, you need to tell the Metron dashboard about this index.
Reviewing the New Data Source Data [22]	Now that the Metron dashboard is aware of the new data source index, you can look at the data.

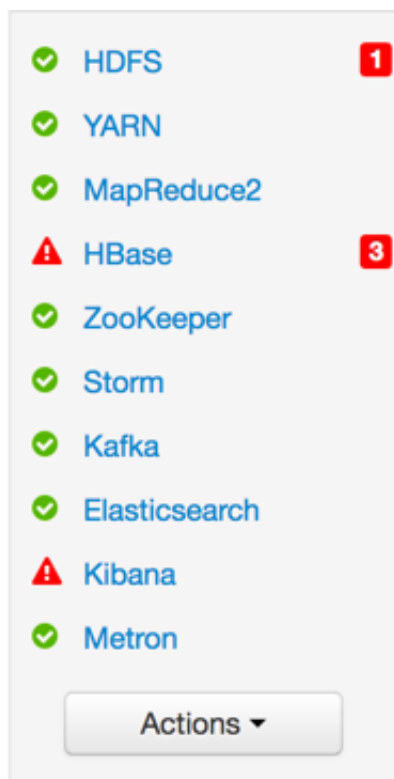
2.4.1. Configuring a New Data Source Index in the Metron Dashboard

Now that you have an index for the new data source with all of the right data types, you need to load the Metron dashboard (also known as Kibana) and tell the dashboard about this index.

1. Launch the Metron dashboard if you have not already done so:

1. From Ambari, click Kibana in the list of quick tasks.

Figure 2.17. Ambari Task List



2. Select **Metron UI** from the Quick Links menu in the top center of the window.
2. Click the **Settings** tab on the Metron dashboard.
3. Make sure you have the **Indices** tab selected, then click **+Add New**.

Kibana displays the **Configure an index pattern** window. Use the index pattern window to identify your telemetry source.

Figure 2.18. Configure an Index Pattern

Configure an index pattern

In order to use Kibana you must configure at least one index pattern. Index patterns are used to identify the Elasticsearch index to run search and analytics against. They are also used to configure fields.

Index contains time-based events
 Use event times to create index names [DEPRECATED]

Index name or pattern
Patterns allow you to define dynamic index names using * as a wildcard. Example: logstash-*

logstash-*

Do not expand index pattern when searching (not recommended)
By default, searches against any time-based index pattern that contains a wildcard will automatically be expanded to query only the indices that contain data within the currently selected time range.

Searching against the index pattern logstash-* will actually query elasticsearch for the specific matching indices (e.g. logstash-2015.12.21) that fall within the current time range.

Unable to fetch mapping. Do you have indices matching the pattern?


4. In the **Index name or pattern** field, enter the name of the index pattern of your data telemetry source.

In most cases the name of the index pattern will match the sensor name. For example, the 'bro' sensor has an index pattern of 'bro-*'.

5. If your data telemetry source **does not** contain time-based events, clear the **Index contains time-based events** check box.

If your data telemetry source does contain time-based events, leave the check box as is. Most of your data telemetry sources will contain time-based events.

6. Click **Create** to add the index pattern for your new data telemetry source.

If you would like this new index pattern to be the default, click the Green Star icon ().

2.4.2. Reviewing the New Data Source Data

Now that the Metron dashboard is aware of the new data source index, you can look at the data.

1. Display the Metron dashboard.
2. Click on the **Discover** tab and then choose the newly created data source index pattern.
3. Click any of the fields in the left column to see a representation of the variety of data for that specific field.
4. Click the Right Facing Arrow icon next to a specific record in the center of the window (the **Document** table) to expand the record and display the available data.

3. Transform the Squid Message

You can customize your sensor data to provide more meaningful data. For example, you can choose to transform a url to provide the domain name of the outbound connection or the IP address. To do this, you need to add transformation information:




Note

Your sensor must be running and producing data before you can add transformation information.

1. In the Management module, click  (edit button) for your sensor.

The Management module displays the schema panel.

Figure 3.1. Squid Schema Panel



The screenshot shows a configuration window titled "Squid" with a close button (X) in the top right corner. The window contains several sections:

- NAME ***: A text input field containing "Squid".
- No Matching Kafka Topic**: A warning message in orange text.
- PARSER TYPE ***: A dropdown menu with "Grok" selected.
- GROK STATEMENT**: A text input field with a right-pointing arrow icon.
- SCHEMA**: A section containing a table with three rows: "TRANSFORMATIONS 0", "ENRICHMENTS 0", and "THREAT INTEL 0". To the right of the table is a right-pointing arrow icon.
- THREAT TRIAGE**: A section containing a text input field with "RULES 0" and a right-pointing arrow icon.

At the bottom of the window, there are three buttons: "SAVE" (blue), "CANCEL" (grey), and "Advanced" (blue).


2. In the Schema box, click  (expand window button).

The Management module displays the Schema panel and populates it with message, field, and value information.

The Sample field, at the top of the panel, displays a parsed version of a sample message from the sensor. The Management module will test your transformations against this parsed message.

You can use the right and left arrow buttons in the Sample field to view the parsed version of each sample message available from the sensor.

You can apply transformations to an existing field or create a new field. Typically users choose to create and transform a new field, rather than transforming an existing field.


3. To add a new transformation, either click the  next to a field or click the



(plus sign) at the bottom of the **Schema** panel.

The module displays a new dialog box for your transformations.

Figure 3.2. New Schema Information Panel



4. Choose the field you want to transform from the **INPUT FIELD** box, enter the name of the new field in the **NAME** field, and then choose a function with the appropriate parameters in the **TRANSFORMATIONS** box. You can apply more than transformation to the input field.

Figure 3.3. Populated Schema Information Panel

The screenshot shows a configuration panel for a field named 'ip_dst_addr_copy'. It is divided into several sections: 'INPUT FIELD' with a dropdown menu showing 'ip_dst_addr'; 'NAME' with a text input field containing 'ip_dst_addr_copy'; 'TRANSFORMATIONS' with two dropdown menus, the first showing 'DOMAIN_REMOVE_SUBDOMAINS' and the second showing 'DOMAIN_REMOVE_TLD', each with a minus sign button to its right; a third dropdown menu showing the composed transformation 'DOMAIN_REMOVE_TLD(DOMAIN_REMOVE_SUBDOMAINS(ip_dst_addr))'; 'ENRICHMENTS' with an empty dropdown menu; and 'THREAT INTEL' with an empty dropdown menu. A 'SAVE' button is located at the bottom left of the panel.

5. Click **SAVE** to save your additions.

The Management module populates the Transforms field with the number of transformations applied to the sensor.

If you change your mind and want to remove a transformation, click "-" next to the field.

6. Click **SAVE** in the parser panel to save the transformation information.

4. Enriching Telemetry Events

After the raw security telemetry events have been parsed and normalized, the next step is to enrich the data elements of the normalized event.

Enrichments add external data from data stores (such as HBase). HCP uses a combination of HBase, Storm, and the telemetry messages in json format to enrich the data in real time to make it relevant and consumable. You can use this enriched information immediately rather than needing to hunt in different silos for the relevant information.

HCP supports two types of configurations: global and sensor specific. The sensor specific configuration configures the individual enrichments and threat intelligence enrichments for a given sensor type (for example, squid). This section describes sensor specific configurations. For more information on global configuration, see [Global Configuration](#).

HCP provides two types of enrichment:

- Telemetry events
- Threat intelligence information

For information about threat intelligence see [Enriching Threat Intelligence Information](#).

HCP provides the following telemetry enrichment sources but you can add your own enrichment sources to suit your needs:

- Asset
- GeolP
- User



Note

The telemetry data sources for which HCP includes parsers (for example, Bro, Snort, and YAF) already include enrichment topologies. These topologies will become effective when you start the data sources in HCP.

Prior to enabling an enrichment capability within HCP, the enrichment store (which for HCP is primarily HBase) must be loaded with enrichment data. The dataload utilities convert raw data sources to a primitive key (type, indicator) and value and place it in HBase.

HCP supports three types of enrichment loaders:

- Bulk load from HDFS via MapReduce
- Taxii Loader
- Flat File ingestion

For simplicity's sake, we use the bulk loader to load enrichments:

Task	Description
Bulk Loading Enrichment Information	Bulk loading is used to load information that does not change frequently.
Mapping Fields to HBase Enrichments	Now that you have data flowing into the HBase table, you need to ensure that the enrichment topology can be used to enrich the data flowing past.
OPTIONAL:Global Configuration	Global enrichments are applied to all data sources as opposed to other enrichments that are applied at the field level to individual sensors. This type of enrichment can save you time by applying common enrichments to all of your sensors.
Verify Events are Enriched	After you finish enriching your new data source, you should verify that the output matches your enrichment information.

4.1. Bulk Loading Enrichment Information

Bulk loading is used to load information that does not change frequently. For example, bulk loading is ideal for loading from an asset database on a daily or even weekly basis because you don't typically change the number of assets on a network very often.

Enrichment data can be bulk loaded from the local file system, HDFS. The enrichment loader transforms the enrichment into a JSON format that is understandable to Metron. The loading framework has additional capabilities for aging data out of the enrichment stores based on time. Once the stores are loaded, an enrichment bolt that can interact with the enrichment store can be incorporated into the enrichment topology.

You can bulk load enrichment information from the following sources:

- CSV Flat File Ingestion
- HDFS via MapReduce
- Taxii Loader

For our example, we will use CSV flat file ingestion. For more information about Taxii Loader, see [Bulk Loading Threat Intelligence Information](#). For more information about HDFS via MapReduce, see [Bulk Loading Threat Intelligence Information](#)

Bulk loading information consists of the following major steps:

Task	Description
OPTIONAL: Create a Mock Enrichment Source [29]	For our runbook demonstration we create a mock enrichment source. In your production environment you will want to use a genuine enrichment source.
Configuring an Extractor Configuration File [29]	The extractor configuration file is used to bulk load the enrichment store into HBase.
Configuring Element-to-Enrichment Mapping [31]	This section configures what element of a tuple should be enriched with what enrichment type.
Running the Enrichment Loader [32]	After the enrichment source and enrichment configuration are defined, you must run the loader to move the data from the enrichment source to the HCP enrichment store (HBase) and store the enrichment configuration in ZooKeeper.

4.1.1. OPTIONAL: Create a Mock Enrichment Source

For our runbook demonstration, we create a mock enrichment source. In your production environment you will want to use a genuine enrichment source.

To create a mock enrichment source, complete the following steps:

1. As root user, log into `$HOST_WITH_ENRICHMENT_TAG`.

```
sudo -s $HOST_WITH_ENRICHMENT_TAG
```

2. Copy and paste the following data into a file called `whois_ref.csv` in `$METRON_HOME/config`. This CSV file represents our enrichment source.

```
google.com, "Google Inc.", "US", "Dns Admin",874306800000  
work.net, "", "US", "PERFECT PRIVACY, LLC",788706000000  
capitalone.com, "Capital One Services, Inc.", "US", "Domain Manager",  
795081600000  
cisco.com, "Cisco Technology Inc.", "US", "Info Sec",547988400000  
cnn.com, "Turner Broadcasting System, Inc.", "US", "Domain Name Manager",  
748695600000  
news.com, "CBS Interactive Inc.", "US", "Domain Admin",833353200000  
nba.com, "NBA Media Ventures, LLC", "US", "C/O Domain Administrator",  
786027600000  
espn.com, "ESPN, Inc.", "US", "ESPN, Inc.",781268400000  
pravda.com, "Internet Invest, Ltd. dba Imena.ua", "UA", "Whois privacy  
protection service",806583600000  
hortonworks.com, "Hortonworks, Inc.", "US", "Domain Administrator",  
1303427404000  
microsoft.com, "Microsoft Corporation", "US", "Domain Administrator",  
673156800000  
yahoo.com, "Yahoo! Inc.", "US", "Domain Administrator",790416000000  
rackspace.com, "Rackspace US, Inc.", "US", "Domain Admin",903092400000  
landl.co.uk, "1 & 1 Internet Ltd", "UK", "Domain Admin",943315200000
```

Make sure you don't have an empty newline character as the last line of the CSV file, as that will result in a null pointer exception.

4.1.2. Configuring an Extractor Configuration File

The extractor configuration file is used to bulk load the enrichment store into HBase.

Complete the following steps to configure the extractor configuration file:

1. Log in as root to the host on which Metron is installed.

```
sudo -s $METRON_HOME
```

2. Determine the schema of the enrichment source.

The schema of our mock enrichment source is `domain|owner|registeredCountry|registeredTimestamp`.

3. Create an extractor configuration file called `extractor_config_temp.json` at `$METRON_HOME/config` and populate it with the enrichment source schema.

For example:

```
{
  "config" : {
    "columns" : {
      "domain" : 0
      ,"owner" : 1
      ,"home_country" : 2
      ,"registrar": 3
      ,"domain_created_timestamp": 4
    }
    ,"indicator_column" : "domain"
    ,"type" : "whois"
    ,"separator" : ","
  }
  ,"extractor" : "CSV"
}
```

4. You can transform and filter the enrichment data as it is loaded into HBase by using Stellar extractor properties in the extractor configuration file. HCP supports the following Stellar extractor properties:

Extractor Property	Description	Example
value_transform	Transforms fields defined in the columns mapping with Stellar transformations. New keys introduced in the transform are added to the key metadata.	"value_transform" : { "domain" : "DOMAIN_REMOVE_TLD(domain)"
value_filter	Allows additional filtering with Stellar predicates based on results from the value transformations. In the following example, records whose domain property is empty after removing the TLD are omitted.	"value_filter" : "LENGTH(domain) > 0", "indicator_column" : "domain",
indicator_transform	Transforms the indicator column independent of the value transformations. You can refer to the original indicator value by using indicator as the variable name, as shown in the following example. In addition, if you prefer to piggyback your transformations, you can refer to the variable domain, which allows your indicator transforms to inherit transformations done to this value during the value transformations.	"indicator_transform" : { "indicator" : "DOMAIN_REMOVE_TLD(indicator)"
indicator_filter	Allows additional filtering with Stellar predicates based on results from the value transformations. In the following example, records whose indicator value is empty after removing the TLD are omitted.	"indicator_filter" : "LENGTH(indicator) > 0", "type" : "top_domains",

If you include all of the supported Stellar extractor properties in the extractor configuration file, it will look similar to the following:

```
{
  "config" : {
    "zk_quorum" : "$ZOOKEEPER_HOST:2181",
    "columns" : {
      "rank" : 0,
```

```
"domain" : 1
},
"value_transform" : {
"domain" : "DOMAIN_REMOVE_TLD(domain)"
},
"value_filter" : "LENGTH(domain) > 0",
"indicator_column" : "domain",
"indicator_transform" : {
"indicator" : "DOMAIN_REMOVE_TLD(indicator)"
},
"indicator_filter" : "LENGTH(indicator) > 0",
"type" : "top_domains",
"separator" : ",",
},
"extractor" : "CSV"
}
```

Running a file import with the above data and extractor configuration will result in the following two extracted data records:

Indicator	Type	Value
google	top_domains	{ "rank": "1", "domain": "google" }
yahoo	top_domains	{ "rank": "2", "domain": "yahoo" }

- Remove any non-ASCII invisible characters that might have been included if you copy and pasted:

```
iconv -c -f utf-8 -t ascii extractor_config_temp.json -o extractor_config.json
```



Note

The `extractor_config.json` file is not stored anywhere by the loader. This file is used once by the bulk loader to parse the enrichment dataset. If you would like to keep a copy of this file, be sure to save a copy to another location.

4.1.3. Configuring Element-to-Enrichment Mapping

We now need to configure what element of a tuple should be enriched with what enrichment type.

This configuration is stored in ZooKeeper.

- Log in as root user to the host that has Metron installed.

```
sudo -s $METRON_HOME
```

- Copy and paste the following into a file called `enrichment_config_temp.json` at `$METRON_HOME/config`.

```
{
  "zkQuorum" : "$ZOOKEEPER_HOST:2181"
, "sensorToFieldList" : {
  "squid" : {
    "type" : "ENRICHMENT"
  , "fieldToEnrichmentTypes" : {
    "domain_without_subdomains" : [ "whois" ]
  }
}
}
```

```
}
}
}
```

3. Remove any non-ASCII invisible characters that might have been included if you copy and pasted:

```
iconv -c -f utf-8 -t ascii enrichment_config_temp.json -o enrichment_config.json
```

4.1.4. Running the Enrichment Loader

After the enrichment source and enrichment configuration are defined, you must run the loader to move the data from the enrichment source to the HCP enrichment store (HBase) and store the enrichment configuration in ZooKeeper.



Note

There is a special configuration parameter to the Extractor config that is only considered during this loader:

`inputFormatHandler` This specifies how to consider the data. The two implementations are `BY_LINE` and `org.apache.metron.dataloads.extractor.inputformat.`

The default is `BY_LINE`, which makes sense for a list of CSVs where each line indicates a unit of information which can be imported. However, if you are importing a set of STIX documents, then you want each document to be considered as input to the Extractor.

1. Use the loader to move the enrichment source to the enrichment store in ZooKeeper.

Perform the following from the location containing your extractor and enrichment configuration files and your enrichment source. In our example, this information is located at `$METRON_HOME/config`.

```
$METRON_HOME/bin/flatfile_loader.sh -n enrichment_config.json -i whois_ref.csv -t enrichment -c t -e extractor_config.json
```

The parameters for the utility are as follows:

Short Code	Long Code	Required	Description
-h		No	Generate the help screen/set of options
-e	-extractor_config	Yes	JSON document describing the extractor for this input data source
-t	-hbase_table	Yes	The HBase table to import into
-c	-hbase_cf	Yes	The HBase table column family to import into
-i	-input	Yes	The input data location on local disk. If this is a

Short Code	Long Code	Required	Description
			file, then that file will be loaded. If this is a directory, then the files will be loaded recursively under that directory.
-l	-log4j	No	The log4j properties file to load
-n	-enrichment_config	No	The JSON document describing the enrichments to configure. Unlike other loaders, this is run first if specified.

HCP loads the enrichment data into Apache HBase and establishes a ZooKeeper mapping. The data is extracted using the extractor and configuration defined in the `extractor_config.json` file and populated into an HBase table called `enrichment`.

2. Verify that the logs were properly ingested into HBase:

```
hbase shell
scan 'enrichment'
```

3. Verify that the ZooKeeper enrichment tag was properly populated:

```
$METRON_HOME/bin/zk_load_configs.sh -m DUMP -z $ZOOKEEPER_HOST:2181
```

4. Generate some data by using the Squid client to execute requests.

- a. Use ssh to access the host for Squid.

- b. Start Squid and navigate to `/var/log/squid`:

```
sudo service squid start
sudo su -
cd /var/log/squid
```

- c. Generate some data by entering the following:

```
squidclient http://www.cnn.com
```

4.2. Mapping Fields to HBase Enrichments

Now that you have data flowing into the HBase table, you need to ensure that the enrichment topology can be used to enrich the data flowing past.

You can refine the parser output in three ways:

- Transformations
- Enrichments
- Threat Intel

Each of the parser outputs is added or modified in the **Schema** field. To modify any of the parser outputs, complete the following steps:



Note

To load sample data from your sensor, the sensor must be running and producing data.

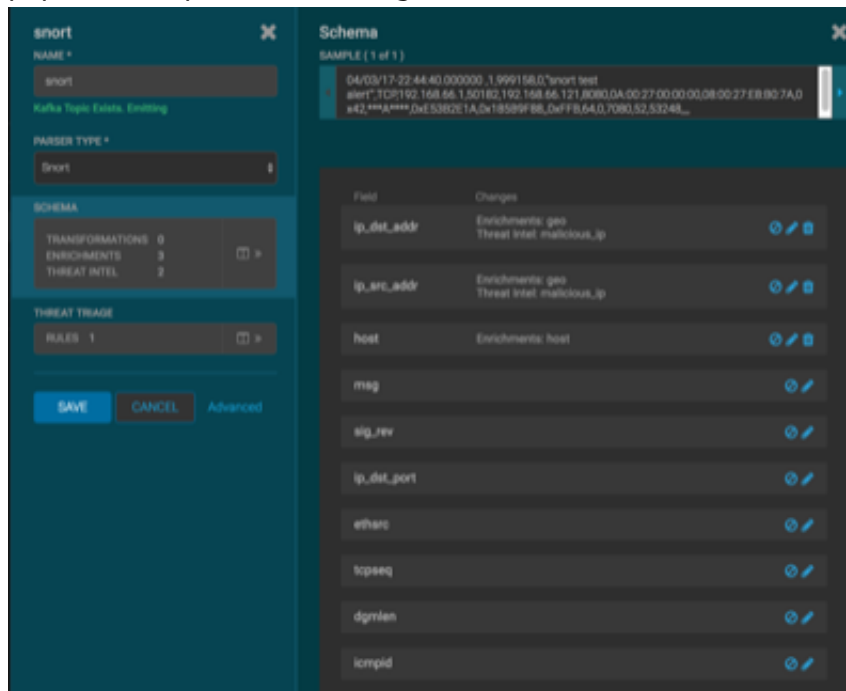
1. Display the Management module UI.
2. Select the new sensor from the list of sensors on the main window.

3. Click the pencil icon in the list of tool icons  for the new sensor.

The Management Module displays the sensor panel for the new sensor.

4. In the Schema box, click  (expand window button).


The Management module displays a second panel and populates the panel with message, field, and value information.



The Sample field, at the top of the panel, displays a parsed version of a sample message from the sensor. The Management module will test your transformations against these parsed messages.

You can use the right and left arrow buttons in the Sample field to view the parsed version of each sample message available from the sensor.

5. You can apply transformations to an existing field or create a new field. Click

the  (edit icon) next to a field to apply transformations to that field. Or click



(plus sign) at the bottom of the Schema panel to create new fields.

Typically users store transformations in a new field rather than overriding existing fields.

For both options, the Management module expands the panel with a dialog box containing fields in which you can enter field information.

Figure 4.1. New Schema Information Panel

6. In the dialog box, enter the name of the new field in the **NAME** field, choose an input field from the **INPUT FIELD** box, and choose your transformation from the **TRANSFORMATIONS** field or enrichment from the **ENRICHMENTS** field.


For example, to create a new field showing the lower case version of the method field, do the following:

- Enter method-uppercase in the **NAME** field.
- Choose `method` from the **INPUT FIELD**.
- Choose `TO_UPPER` in the **TRANSFORMATIONS** field.

Your new schema information panel should look like this:

Figure 4.2. Populated New Schema Information Panel

The screenshot shows a configuration panel for a schema named 'method_uppercase'. It is organized into several sections: 'INPUT FIELD' with a dropdown menu set to 'method'; 'NAME' with a text input field containing 'method_uppercase'; 'TRANSFORMATIONS' with a dropdown menu set to 'TO_UPPER' and a minus sign button to its right, and a preview area below it showing 'TO_UPPER(method)'; 'ENRICHMENTS' with an empty dropdown menu; and 'THREAT INTEL' with an empty dropdown menu. A 'SAVE' button is located at the bottom left of the panel.

7. Click **SAVE** to save your changes.
8. You can suppress fields from showing in the Index by clicking  (suppress icon).
9. Click **SAVE** to save the changed information.

The Management module updates the Schema field with the number of changes applied to the sensor.

4.3. OPTIONAL: Global Configuration

Global enrichments are applied to all data sources as opposed to other enrichments that are applied at the field level to individual sensors. This type of enrichment can save you time by applying common enrichments to all of your sensors. For more information on configuring global configurations, see [Global Configuration](#).

4.4. Verify That the Events Are Enriched

After you finish enriching your new data source, you should verify that the output matches your enrichment information.

5. Enriching Threat Intelligence Information

You can enrich your threat intelligence information just like you enriched your telemetry information.

You can choose to skip this section and come back to it later if you don't want to enrich your threat intelligence information at this time.

Metron provides an extensible framework to plug in threat intel sources. Each threat intel source has two components: an enrichment data source and an enrichment bolt. The threat intelligence feeds are loaded into a threat intelligence store similar to how the enrichment feeds are loaded. The keys are loaded in a key-value format. The key is the indicator and the value is the JSON formatted description of what the indicator is.

We recommend using a threat feed aggregator such as [Soltra](#) to dedup and normalize the feeds via STIX/Taxii. Metron provides an adapter that is able to read Soltra-produced STIX/Taxii feeds and stream them into HBase, which is the preferred data store to back high-speed threat intel lookups on HCP. HCP additionally provides a flat file and STIX bulk loader that can normalize, dedup, and bulk load or poll threat intel data into HBase even without the use of a threat feed aggregator.

5.1. OPTIONAL: Create a Mock Threat Intel Feed Source

Metron is designed to work with STIX/Taxii threat feeds, but can also be bulk loaded with threat data from a CSV file. In this example, we will explore the CSV example. The same loader framework that is used for enrichment here is used for threat intelligence. Similar to enrichments, we need to set up a data.csv file, the extractor config JSON, and the enrichment config JSON.

For this example, we will use a Zeus malware tracker list located here: <https://zeustracker.abuse.ch/blocklist.php?download=domainblocklist>.

1. Log into the \$HOST_WITH_ENRICHMENT_TAG as root.
2. Copy the contents from the Zeus malware tracker list link to a file called domainblocklist.csv.

```
curl https://zeustracker.abuse.ch/blocklist.php?download=domainblocklist  
| grep -v "^#" | grep -v "^$" | grep -v "^https" | awk '{print $1",abuse.  
ch"}' > domainblocklist.csv
```

5.2. Configuring an Extractor Configuration File

Complete the following steps to configure the extractor configuration file:

1. Log in as root to the host on which Metron is installed.

```
sudo -s $METRON_HOME
```

2. Determine the schema of the enrichment source.

The schema of our mock enrichment source is domain|owner|registeredCountry|registeredTimestamp.

3. Create an extractor configuration file called `extractor_config_temp.json` at `$METRON_HOME/config` and populate it with the threat intel source schema.

HCP supports a subset of STIX messages for importation:

STIX Type	Specific Type	Enrichment Type Name
Address	IPV_4_ADDR	address:IPV_4_ADDR
Address	IPV_6_ADDR	address:IPV_6_ADDR
Address	E_MAIL	address:E_MAIL
Address	MAC	address:MAC
Domain	FQDN	domain:FQDN
Hostname		hostname

The following example configures the STIX extractor to load from a series of STIX files, however we only want to bring in IPv4 addresses from the set of all possible addresses. Note that if no categories are specified for import, all are assumed. Also, only address and domain types allow filtering via `stix_address_categories` and `stix_domain_categories` config parameters.

For example:

```
{
  "config": {
    "columns": {
      "malicious_host": 0, "source": 1
    },
    "indicator_column": "malicious_host", "type": "malicious_ip",
    "separator": ",",
    "extractor": "CSV"
  }
}
```

4. Remove any non-ASCII invisible characters that might have been included if you copy and pasted:

```
iconv -c -f utf-8 -t ascii extractor_config_temp.json -o
extractor_config_temp.json
```

5. OPTIONAL: You also have the ability to transform and threat intel data using Stellar as it is loaded into HBase. This feature is available to all extractor types.

As an example, we will be providing a CSV list of top domains as an enrichment and filtering the value metadata, as well as the indicator column, with Stellar expressions.

```
{
  "config" : {
    "zk_quorum" : "node1:2181",
    "columns" : {
      "rank" : 0,
      "domain" : 1
    }
  }
}
```

```
    },  
    "value_transform" : {  
      "domain" : "DOMAIN_REMOVE_TLD(domain)"  
    },  
    "value_filter" : "LENGTH(domain) > 0",  
    "indicator_column" : "domain",  
    "indicator_transform" : {  
      "indicator" : "DOMAIN_REMOVE_TLD(indicator)"  
    },  
    "indicator_filter" : "LENGTH(indicator) > 0",  
    "type" : "top_domains",  
    "separator" : ",",  
  },  
  "extractor" : "CSV"  
}
```

5.3. Configuring Element-to-Threat Intel Feed Mapping

We now need to configure what element of a tuple should be enriched with what enrichment type.

This configuration is stored in ZooKeeper.

1. Log in as root user to the host that has Metron installed.

```
sudo -s $METRON_HOME
```

2. Copy and paste the following into a file called `enrichment_config_temp.json` at `$METRON_HOME/config`.

```
{  
  "zkQuorum": "ip-10-0-163-99.us-west-1.compute.internal:2181",  
  "sensorToFieldList":{  
    "squid":{  
      "type": "THREAT_INTEL", "fieldToEnrichmentTypes":{  
        "domain_without_subdomains":["malicious_ip"]  
      }  
    }  
  }  
}
```

You must specify the following:

- The zookeeper quorum which holds the cluster configuration
- The mapping between the fields in the enriched documents and the enrichment types.

This configuration allows the ingestion tools to update ZooKeeper post-ingestion so that the enrichment topology can take advantage immediately of the new type.

3. Remove any non-ASCII invisible characters that might have been included if you copy and pasted:

```
iconv -c -f utf-8 -t ascii enrichment_config_temp.json -o enrichment_config.json
```

5.4. Run the Threat Intelligence Loader

Now that you have the threat intel source, threat intel extractor, and threat intel mapping config defined, you can run the loader to move the data from the threat intel source to the Metron threat intel Store and store the enrichment config in ZooKeeper.

1. Log into the \$HOST_WITH_ENRICHMENT_TAG as root.
2. Run the loader:

```
/usr/metron/$METRON_RELEASE/bin/flatfile_loader.sh -n enrichment_config.json  
-i domainblocklist.csv -t threatintel -c t -e extractor_config_temp.json
```

The previous command adds the threat intel data into HBase and establishes a ZooKeeper mapping. The data is populated into an HBase table called threatintel.

3. To verify that the logs were properly ingested into HBase, run the following command:

```
hbase shell  
scan 'threatintel'
```

4. Now check if the ZooKeeper enrichment tag was properly populated:

```
/usr/metron/$METRON_RELEASE/bin/zk_load_configs.sh -m DUMP -z  
$ZOOKEEPER_HOST:2181
```

You should see a configuration for the Squid sensor something like the following

```
{  
  "index" : "squid",  
  "batchSize" : 1,  
  "enrichment" : {  
    "fieldMap" : {  
      "hbaseEnrichment" : [ "ip_src_addr" ]  
    },  
    "fieldToTypeMap" : {  
      "ip_src_addr" : [ "user" ]  
    },  
    "config" : { }  
  },  
  "threatIntel" : {  
    "fieldMap" : { },  
    "fieldToTypeMap" : { },  
    "config" : { },  
    "triageConfig" : {  
      "riskLevelRules" : { },  
      "aggregator" : "MAX",  
      "aggregationConfig" : { }  
    }  
  },  
  "configuration" : { }  
}
```

5. Generate some data by using the Squid client to execute http requests.

```
squidclient http://www.actdhaka.com
```

5.5. Mapping Fields to HBase Enrichments

Now that you have data flowing into the HBase table, you need to ensure that the enrichment topology can be used to enrich the data flowing past.

You can refine the parser output in three ways:

- Transformations
- Enrichments
- Threat Intel

Each of the parser outputs is added or modified in the **Schema** field. To modify any of the parser outputs, complete the following steps:



Note

To load sample data from your sensor, the sensor must be running and producing data.


1. Select the new sensor from the list of sensors on the main window.

2. Click the pencil icon in the list of tool icons  for the new sensor.

The Management Module displays the sensor panel for the new sensor.

3. In the Schema box, click  (expand window button).

The Management module displays a second panel and populates the panel with message, field, and value information.



The Sample field, at the top of the panel, displays a parsed version of a sample message from the sensor. The Management module will test your transformations against these parsed messages.

You can use the right and left arrow buttons in the Sample field to view the parsed version of each sample message available from the sensor.

4. You can apply transformations to an existing field or create a new field. Click

the  (edit icon) next to a field to apply transformations to that field. Or click




(plus sign) at the bottom of the Schema panel to create new fields.

Typically users store transformations in a new field rather than overriding existing fields.

For both options, the Management module expands the panel with a dialog box containing fields in which you can enter field information.

Figure 5.1. New Schema Information Panel




5. In the dialog box, enter the name of the new field in the **NAME** field, choose an input field from the **INPUT FIELD** box, and choose your transformation from the **TRANSFORMATIONS** field or enrichment from the **ENRICHMENTS** field.


For example, to create a new field showing the lower case version of the method field, do the following:

- Enter `method-uppercase` in the **NAME** field.
- Choose `method` from the **INPUT FIELD**.
- Choose `TO_UPPER` in the **TRANSFORMATIONS** field.

Your new schema information panel should look like this:

Figure 5.2. Populated New Schema Information Panel



6. Click **SAVE** to save your changes.
7. You can suppress fields from showing in the Index by clicking  (suppress icon).
8. Click **SAVE** to save the changed information.

The Management module updates the Schema field with the number of changes applied to the sensor.

5.6. Verify That the Threat Intel Events Are Enriched

After you finish enriching your new data source, you should verify that the output matches your enrichment information.

By convention, the index where the new messages are indexed is called `squid_index_[timestamp]` and the document type is `squid_doc`.

Use the Elasticsearch Head plug-in to verify that the messages display your enrichment information:

1. Log in to `$SEARCH_HOST` host:

```
ssh into Host $SEARCH_HOST
```

2. Install the head plug-in:

If you are connected to the internet, you can use the following command:

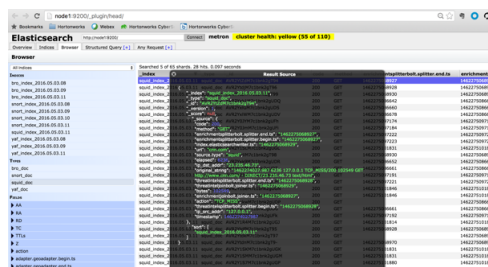
```
/usr/share/elasticsearch/bin/plugin install mobz/elasticsearch-head/1.x
```

If you are not connected to the internet, complete the following steps:

- a. Download the tar.gz by hand.
 - b. untar into the `plugins` directory.
 - c. At this point, the URL `http://:9200/_plugin/head/` should work (don't forget the trailing `/`).
3. Navigate to ElasticSearch Head UI: `http://$SEARCH_HOST:9200/_plugin/head/`.
4. Click the Browser tab and select the squid document in the left panel; then select one of the sample docs.

You should see something like the following:

Figure 5.3. Elasticsearch



5. Review the output to ensure it reflects the enrichment functions you used.

6. Prioritizing Threat Intelligence

Not all threat intelligence indicators are equal. Some require immediate response, while others can be dealt with or investigated as time and availability permits. As a result you need to triage and rank threats by severity.

In HCP, you assign severity by associating possibly complex conditions with numeric scores. Then, for each message, you use a configurable aggregation function to evaluate the set of conditions and to aggregate the set of numbers for matching conditions. This aggregated score is added to the message in the `threat.triage.level` field. For more information about Stellar and threat triage configurations, see [Using Stellar to Set up Threat Triage Configurations](#).

This section details the steps to understand and create severity rules, configure them in ZooKeeper, and view the resulting alerts in the HCP Investigation module:

- [Prerequisites \[46\]](#)
- [Performing Threat Triage \[46\]](#)
- [Viewing Triageed or Scored Alerts \[49\]](#)

6.1. Prerequisites

Before you can prioritize a threat intelligence enrichment, you must ensure that the enrichment is working properly. See [Enriching Telemetry Events](#) and [Enriching Threat Intelligence Information](#) for more information.

6.2. Performing Threat Triage

To create a threat triage rule configuration, you must first define your rules. These rules identify the conditions in the data source data flow and associate alert scores with those conditions. Following are some examples:





Important

In order to apply rules to a sensor, the sensor json file must contain `"is_alert": "true"`.

- Rule 1 If a threat intelligence enrichment type is alerted, imagine that you want to receive an alert score of 5.
- Rule 2 If the URL ends with neither `.com` nor `.net`, then imagine that you want to receive an alert score of 10.

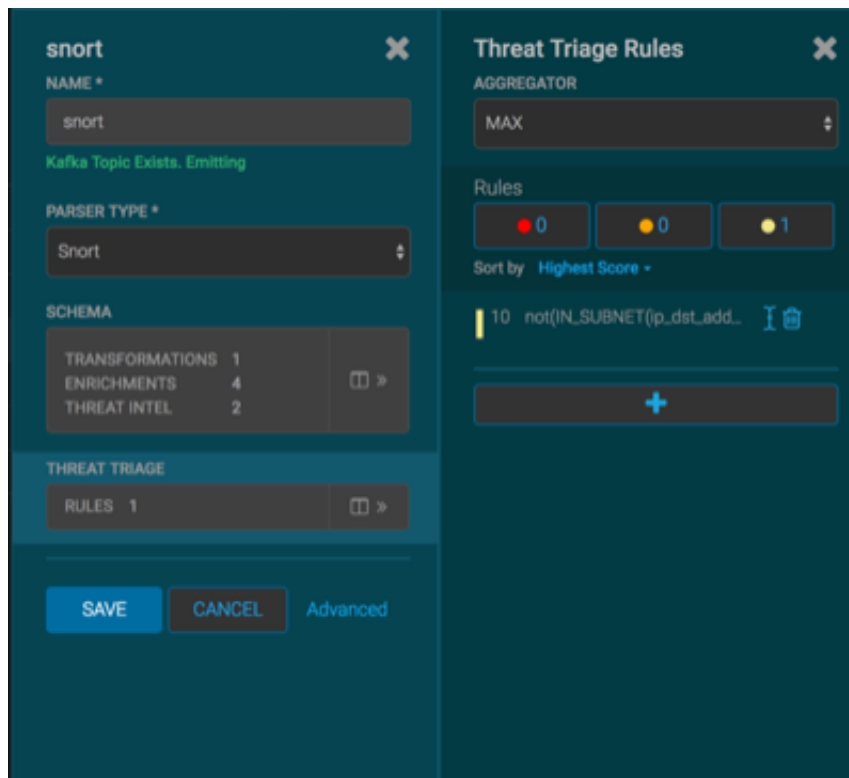
To create these rules, complete the following steps:

1.  Click the (edit button) for your sensor.

2. In the Threat Triage field, click the  icon (expand window).

The module displays the Threat Triage Rules panel.


Figure 6.1. Threat Triage Rules Panel



3. Click the + button to add a rule.

The module displays the **Edit Rule** panel.

Figure 6.2. Edit Rule Panel



4. Assign a name to the new rule by entering the name in the NAME field.
5. In the Text field, enter the syntax for the new rule.

For example:

```
EXISTS(is_alert)
```

6. Use the **SCORE ADJUSTMENT** slider to choose the threat score for the rule.
7. Click **SAVE** to save the new rule.

The new rule is listed in the Threat Triage Rules panel.

8. Choose how you want to aggregate your rules by choosing a value from the Aggregator menu.

You can choose between:

- | | |
|---------------|--|
| MAX | The maximum of all of the associated values for matching queries. |
| MIN | The minimum of all of the associated values for matching queries. |
| MEAN | the mean of all of the associated values for matching queries. |
| POSITIVE_MEAN | The mean of the positive associated values for the matching queries. |

9. You can use the **Rules** section and the **Sort by** pull down menu below the **Rules** section to filter how threat triages display.

For example, to display only high levels alerts, click the box containing the red indicator. To sort the high level alerts from highest to lowest, choose **Highest Score** from the **Sort by** pull down menu.

10. Click **SAVE** on the Sensor panel to save your changes.

6.3. Viewing Triaged or Scored Alerts

You can view triaged alerts in the indexing topic in Kafka or in the triaged alert panel in the HCP Metron dashboard.

- [Threat Triage Alert Panel \[49\]](#)
- [HCP Metron Dashboard View of Alerts \[49\]](#)

6.3.1. Threat Triage Alert Panel

To view the Threat Triage Alert Panel in Kafka, complete the following steps. An alert in the indexing topic in Kafka will appear similar to the following:

1. List the Kafka topics to find the threat triage alert panel:

```
/usr/hdp/current/kafka-broker/bin/kafka-topics.sh --zookeeper  
$ZOOKEEPER_HOST:2181 --list
```

2. View the threat triage alert Kafka topic:

```
cd $METRON_HOME/bin/.stellar  
THREAT_TRIAGE_PRINT(conf)
```

The topic should appear similar to the following:

```
> THREAT_TRIAGE_PRINT(conf)  
#####  
# Name                # Comment # Triage Rule  
#                   # Score # Reason  
  
#  
#####  
# Abnormal DNS Port #          # source.type == "bro" and protocol == "dns" and  
ip_dst_port != 53 # 10      # FORMAT("Abnormal DNS Port: expected: 53, found:  
%s:%d", ip_dst_addr, ip_dst_port) #  
#####
```

6.3.2. HCP Metron Dashboard View of Alerts

The following figure shows you an example of a triaged alert panel in the HCP Metron dashboard. For URLs from `cnn.com`, no threat alert is shown, so no triage level is set. Notice the lack of a `threat.triage.level` field:

Figure 6.3. Investigation Module Triaged Alert Panel



The screenshot displays a table titled "Triaged Alerts" with the following columns: Time, sourceType, threat triage level, full_hostname, ip_src_addr, and ip_dst_addr. A red box highlights the "threat triage level" column, and a red arrow points to it from a label "New Triage Level Column".

Time	sourceType	threat triage level	full_hostname	ip_src_addr	ip_dst_addr
* June 29th 2016, 17:14:30.483	squld	5	www.acthaka.com	127.0.0.1	198.50.239.7
* June 29th 2016, 17:14:29.196	squld	5	www.acthaka.com	127.0.0.1	198.50.239.7
* June 29th 2016, 17:14:28.025	squld	5	www.acthaka.com	127.0.0.1	198.50.239.7

7. Configuring Indexing

The indexing topology is a topology dedicated to taking the data from a topology that has been enriched and storing the data in one or more supported indices. More specifically, the enriched data is ingested into Kafka, written in an indexing batch or bolt with a specified size, and sent to one or more specified indices. The configuration is intended to configure the indexing used for a given sensor type (for example, snort).

Currently, HCP supports the following indices:

- Elasticsearch
- Solr
- HDFS under `/apps/metron/indexing/indexed`

Depending on how you start the indexing topology, it can have HDFS and either elasticsearch or SOLR writers running.

Just like the Global Configuration file, the Indexing Configuration file format is a JSON file stored in ZooKeeper and on disk at `$METRON_HOME/config/zookeeper/indexing`.



Note

Errors during indexing are sent to a Kafka queue called `index_errors`

Within the sensor-specific configuration, you can configure the individual writers. The parameters currently supported are:

<code>index</code>	The name of the index to write to (defaulted to the name of the sensor).
<code>batchSize</code>	The size of the batch that is written to the indices at once (defaulted to 1).
<code>enabled</code>	Whether the index or writer is enabled (default true).

7.1. Default Configuration

If you do not configure the individual writers, the sensor-specific configuration will use the default values. You can choose to use this default configuration by either not creating the Indexing Configuration file or by entering the following in the file. You can name the file anything you like, for example `index_config.json`, but it must be located at `$METRON_HOME/config/zookeeper/indexing`.

```
{  
}
```


If a writer configuration is unspecified, then a warning is indicated in the Storm console. For example, `WARNING: Default and (likely) unoptimized writer config used for hdfs writer and sensor squid`. You can ignore this warning message if you intend to use the default configuration.

This default configuration uses the following configuration:

- elasticsearch writer
 - index name the same as the sensor
 - batch size of 1
 - enabled
- hdfs writer
 - index name the same as the sensor
 - batch size of 1
 - enabled


7.2. Specifying Index Parameters

You can to specify the parameters for the writers rather than using the default values using the HCP Management Module..

1. Edit your sensor by clicking  (the edit button) next your sensor in the Management Module.
2. Click the **Advanced** button next to **Save** and **Cancel**.

The Management Module expands the panel to display the Advanced fields.

Figure 7.1. Management Module Advanced Panel



3. Enter index configuration information for your sensor.
4. Click **Save** to save your changes and push your configuration to ZooKeeper.

7.3. Turning off HDFS Writer

You can also turn off the HDFS index or writer using the following syntax in the index.json file.

```
{
```

```
"elasticsearch": {  
  "index": "foo",  
  "enabled" : true  
},  
"hdfs": {  
  "index": "foo",  
  "batchSize": 100,  
  "enabled" : false  
}  
}
```

8. Setting Up a Profile

A profile describes the behavior of an entity on a network. An entity can be a server, user, subnet, or application. Once you generate a profile defining what normal behavior looks like, you can build models that identify anomalous behavior.

Any field contained within a message can be used to generate a profile. A profile can even be produced by combining fields that originate in different data sources. You can transform the data used in a profile by leveraging the Stellar language.

The Profiler is automatically installed and started when you install HCP through Ambari install.

- [Configuring the Profiler \[55\]](#)
- [Starting and Stopping the Profiler \[58\]](#)
- [Developing Profiles \[59\]](#)

8.1. Configuring the Profiler

The configuration for the Profiler topology is stored in ZooKeeper at `/metron/topology/profiler`. These properties also exist in the default installation of HCP at `$METRON_HOME/config/zookeeper/profiler.json`. You can change these values two ways: with Ambari or on disk and then uploaded to ZooKeeper using `$METRON_HOME/bin/zk_load_configs.sh`. The following task uses Ambari to configure the Profiler.

You might need to work with your Platform Engineer to modify Profiler values.



Note

The Profiler can persist any serializable object, not just numeric values.

To configure the Profiler, complete the following steps:

1. Display the Ambari user interface and click the **Services** tab.
2. Click **Metron** in the list of services, then click the **Configs** tab.
3. Click the **Profiler** tab.

Ambari displays a list of Profiler properties that you can use to configure Profiler.

Figure 8.1. Ambari Profiler Properties

4. Use these properties to configure the Profiler, then click the Save button near the top of the window.

The following table provides a list of the configurable properties and descriptions.

Table 8.1. Profiler Properties

Ambari Configs Field.	Settings.	Description
-----------------------	-----------	-------------

Kafka Input Topic Start	profiler_kafka_start	One of EARLIEST, LATEST, UNCOMMITTED_EARLIEST, UNCOMMITTED_LATEST
Profiler Setup Period Duration	profiler_period_duration	The duration of each profile period. This value should be defined along with profiler_period_duration_units.
Period Units	profiler_period_units	The units is used to specify the profiler_period_duration. This value should be define along with profiler_period_duration.
Time to Live	profiler_ttl	If a message has not been applied to a Profile in this period of time, the Profile will be terminated and its resources will be cleaned up. This value should be defined along with profiler_ttl_units. This time-to-live does not affect the persisted Profile data in HBase. It only affects the state stored in memory during the execution of the latest profile period. This state will be deleted if the time_to_live is exceeded.
Time to Live Units	profiler_ttl_units	The units used to specify the profiler_ttl.
HBase Table	profiler_hbase_table	The name of the HBase table the profiler is written to. The profiler expects that the table exists and is writable.
HBase Table Column Family	profiler_hbase_cf	The column family used to store profile data in HBase.
HBase Batch Size	profiler_hbase_batch	The number of puts that are written to HBase in a single batch.
HBase Flush Interval	profiler_hbase_flush_interval	The maximum number of seconds between batch writes to HBase.
Storm topology.worker.childopts	profiler_topology_worker_childopts	Extra topology child opts for the storm opts.
Number of Workers	profiler_topology_workers	The profiler storm topology storm workers
Number of Acker Executors	profiler_acker_executors	The profiler storm topology acker executors



Note

The profiler input topic is bound to the enrichment output topic. If that enrichment output topic is changed, then the profiler will restart as well as the enrichment topology.

Figure 8.2. Enrichment Output Topic

The screenshot displays the configuration interface for the Enrichment service in Ambari. At the top, there are navigation tabs: Index Settings, Parsers, Enrichment (selected), Indexing, Profiler, REST, Management UI, and Advanced. Below the tabs, the configuration is organized into two main sections: Adapters and Kafka.

Adapters

- GEOIP Load Datafile URL:** A text input field containing the URL `http://geolite.maxmind.com/download/geoip/database/GeoLite2-City.mmdb.gz`.
- Host Enrichment:** A text input field containing a JSON array: `[{"ip": "10.1.128.236", "local": "YES", "type": "webserver", "asset_value": "important"}, {"ip": "10.1.128.236"}]`.

Kafka

- Enrichment Offset:** A dropdown menu currently set to `UNCOMMITTED`.
- Enrichment Input Topic:** A text input field containing `enrichments`.
- Enrichment Output Topic:** A text input field containing `indexing`.
- Enrichment Error Topic:** A text input field containing `indexing`.
- Threat Intel Error Topic:** A text input field containing `indexing`.

8.2. Starting and Stopping the Profiler

The Profiler is automatically started when you install HCP. However, if you want to stop or restart the Profiler, complete the following steps:

1. Display the Ambari user interface.
2. Select the **Services** tab, then select **Metron** from the list of services.

Ambari displays a list of Metron components.

3. Select **Metron Profiler**.

Metron displays a complete list of Metron components.

4. Select the pull down menu next to **Metron Profiler / Metron** and select the appropriate status:

- Restart
- Stop
- Turn on Maintenance Mode

8.3. Developing Profiles

Troubleshooting issues when programming against a live stream of data can be difficult. The Stellar REPL (an interactive top level or language shell) is a powerful tool to help work out the kinds of enrichments and transformations that are needed. The Stellar REPL can also be used to help when developing profiles for the Profiler.

Follow these steps in the Stellar REPL to see how it can be used to help create profiles.

1. Take a first pass at defining your profile.

As an example, in the editor copy/paste the basic "Hello, World" profile below.

```
[Stellar]>>> conf := SHELL_EDIT()
[Stellar]>>> conf
{
  "profiles": [
    {
      "profile": "hello-world",
      "onlyif": "exists(ip_src_addr)",
      "foreach": "ip_src_addr",
      "init": { "count": "0" },
      "update": { "count": "count + 1" },
      "result": "count"
    }
  ]
}
```

2. Initialize the Profiler.

```
[Stellar]>>> profiler := PROFILER_INIT(conf)
[Stellar]>>> profiler
org.apache.metron.profiler.StandAloneProfiler@4f8ef473
```

3. Create a message to simulate the type of telemetry that you expect to be profiled.

As an example, in the editor copy/paste the JSON below.


```
[Stellar]>>> message := SHELL_EDIT()
[Stellar]>>> message
{
  "ip_src_addr": "10.0.0.1",
  "protocol": "HTTPS",
  "length": "10",
  "bytes_in": "234"
}
```

4. Apply some telemetry messages to your profiles. The following applies the same message 3 times.

```
[Stellar]>>> PROFILER_APPLY(message, profiler)
org.apache.metron.profiler.StandAloneProfiler@4f8ef473

[Stellar]>>> PROFILER_APPLY(message, profiler)
org.apache.metron.profiler.StandAloneProfiler@4f8ef473

[Stellar]>>> PROFILER_APPLY(message, profiler)
org.apache.metron.profiler.StandAloneProfiler@4f8ef473
```

5. Flush the Profiler to see what has been calculated.

A flush is what occurs at the end of each 15 minute period in the Profiler. The result is a list of profile measurements. Each measurement is a map containing detailed information about the profile data that has been generated.

```
[Stellar]>>> values := PROFILER_FLUSH(profiler)
[Stellar]>>> values
[{"period":{"duration=900000, period=1669628, start=1502665200000, end=1502666100000},
  profile=hello-world, groups=[], value=3, entity=10.0.0.1}]
```

This profile counts the number of messages by IP source address. Notice that the value is '3' for the entity '10.0.0.1' as we applied 3 messages with an 'ip_src_addr' of '10.0.0.1'. There will always be one measurement for each [profile, entity] pair.

6. If you are unhappy with the data that has been generated, then 'wash, rinse and repeat' this process. After you are satisfied with the data being generated by the profile, then follow the [Getting Started](#) guide to use the profile against your live, streaming data in a Metron cluster.