

Hortonworks Cybersecurity Platform

Date of Publish: 2018-07-15

Contents

Introduction to Metron Dashboard.....	3
Functionality of Metron Dashboard.....	3
Metron Default Dashboard.....	5
Events.....	5
Enrichment.....	5
YAF.....	6
Snort.....	7
Web Request Header.....	8
DNS.....	9
 Customizing Your Metron Dashboard.....	 9
Launching the Metron Dashboard.....	9
Changing the Metron Dashboard Background Color.....	10
Adding a New Data Source.....	11
Configuring a New Data Source Index.....	11
Reviewing the New Data Source Data.....	12
Querying, Filtering, and Visualizing Data.....	12
Customizing Your Dashboard.....	13
 Sharing the Metron Dashboard.....	 13
 Triaging Alerts.....	 14
Launch the Alerts User Interface.....	14
Viewing Alerts.....	14
Using the Alerts Table.....	15
Search Alerts.....	19
Filter Alerts.....	20
Manage Alert Status.....	22
Escalate an Alert.....	25
Group Alerts.....	26
Create a Meta Alert.....	28
Save Your Searches.....	29
View Your Recent and Saved Searches.....	30
 PCAP.....	 32
Capturing pcap Data.....	32
Processing pcap Data.....	32
View pcap Data.....	34
Filtering pcap Data.....	34
Query pcap Data Using the Fixed Filter Option.....	34
Query pcap Data Using the Query Filter Option.....	35
Query pcap Data Using the REST API.....	37
Porting pcap Data to Another Application.....	37

Introduction to Metron Dashboard


The Metron dashboard is a Kibana-based dashboard designed to identify, investigate, and analyze cybersecurity data. HCP supports Kibana 4.x. Kibana is an open source analytics and visualization platform.

Functionality of Metron Dashboard

The Metron dashboard displays all of the data on a single dashboard enabling you to filter through the irrelevant data and display just the information, alerts, and context for which you are looking.

The Metron dashboard has several advantages over conventional SIEM tools, including flexibility, and the single pane of glass approach that displays all of the data on the same screen, requiring no jumping from console to console to gather the information.

Dashboard-Snort Panel



HCP supports two types of messages: metadata and alerts. By convention there should be one panel per metadata telemetry and one panel that is a "catch all" panel for alerts. The Snort panels are a good example of these two panel types. However, the Snort alerts panel only lists alerts from Snort because the default Metron dashboard contains only one data source that produces alerts.

When HCP parses the telemetry data on ingest, it extracts and normalizes different parts of the message into a standard Metron JSON. Standardizing and normalizing field names and format allows HCP to search different telemetry messages with a single query.

The first telemetry type that HCP supports is metadata messages. Metadata messages are parsed enriched messages in the JSON format.

The second telemetry type that HCP supports is alerts telemetries. Alerts telemetries come from IDS sensors like Snort or mixed telemetries like application logs that contain some metadata and some alert messages. While it is possible to set up a new panel for each alert telemetry, it is more desirable to set up a single panel that contains all of the alerts. This guarantees that the query will pull in alerts from multiple telemetries (even mixed mode telemetries).

that have some metadata and some alerts associated with them). You can then set up a detailed table containing only the alerts. To set telemetry as alert you need to set `is_alert = true`. This is already set up for HCP under the "Alerts" table.

The fields displayed for each alerts table can be customized. Ideally you want the fields of most importance (as well as the standard fields that telemetries are correlated on) to be displayed.

The following table contains a description of each of the Kibana components in the Metron dashboard.

Area Chart Panel	You can use the area chart panel for stacked timelines for which you want to see the total.
Data Table Panel	Use the data table panel to provide a detail breakdown, in tabular format, of the results of a composed aggregation. You can generate a data table from many other charts by clicking the grey bar at the bottom of the chart.
Detailed Message Panel	A detailed message panel displays the raw data from your search query.
Document Table	When you submit a search query, the 500 most recent documents that match the query are listed in the Documents table which is displayed in the center of the Discover window.
Field List	A list of all of the fields associated with a selected index pattern. This list is displayed on the left side of the Discover window.
Line Chart Panel	Use the line chart when you want to display high density time series. This chart is useful for comparing one series with another.
Mark Down Widget Panel	You can use the mark down widget panel to provide explanations or instructions for the dashboard.
Metric Panel	You can use a metric panel to display a single large number such as the number of hits or the average of a numeric field.
Pie Chart Panel	A pie chart is a circular statistical graphic that is ideal for displaying the parts of some whole.
Tile Map Panel	The tile map panel type displays a map populated with your search results. This panel type requires an Elasticsearch <code>geo_point</code> field that is mapped as type:geo_point with latitude and longitude coordinates.
Vertical Bar Chart Panel	You can use the vertical bar chart panel to display histograms. Histogram panels represent ingest rates for each individual telemetry. By convention, you should set up one for each type.

Metron Default Dashboard

The default telemetry data sources installed with HCP help highlight the useful components available in Kibana 4. The default Metron dashboard serves as a starting point for you to build your own customized dashboards. During installation, HCP sets up several telemetry data sources bundled with the platform and creates panels to display the associated data.

Events

The first panel in the dashboard highlights the variety of events being consumed by HCP. It shows the total number of events received, the variety of those events, and a histogram showing when the events were received.


Events



Enrichment

The next set of dashboard panels shows how HCP can be used to perform real-time enrichment of telemetry data. All of the IPv4 data received by HCP was cross-referenced against a geo-ip database. These locations were then used to build this set of dashboard components.


Enrichment



YAF

As part of the default sensor suite, YAF is used to generate flow records. These flow records provide significant visibility into which actors are communicating over the target network. A table panel displays the raw details of each flow record. A histogram of the duration of each flow illustrates that while most flows are relatively short-lived there are a few that are much longer in this example. Creating an index template that defined this field as numeric was required to generate the histogram.

YAF




[Related Information](#)
[Apache Metron Website](#)

Snort

Snort is a Network Intrusion Detection System (NIDS) that is being used to generate alerts identifying known bad events. Snort relies on a fixed set of rules that act as signatures for identifying abnormal events. Along with displaying the relevant details of each alert, the panel shows that there is only a single unique alert type; a test rule that creates a Snort alert on every network packet. Another table was created to show source/destination pairs that generated the most Snort alerts.


Dashboard-Snort Panel



Web Request Header

The Bro Network Security Monitor extracts application-level information from raw network packets. In this example, Bro is extracting HTTP and HTTPS requests being made over the network. The panels highlight the breakdown by request type, the total number of web requests, and raw details from each web request.


Dashboard-Bro Panel



DNS

Bro extracts DNS requests and responses being made over the network. Understanding who is making those requests, the frequency, and types can provide a deep understanding of the actors present on the network.

Dashboard-DNS Panel



Customizing Your Metron Dashboard

You can customize your Metron dashboard to display information, alerts, and the context you need to identify and analyze cybersecurity issues.

Related Information

[Kibana User Guide](#)

[Building a Dashboard](#)


Launching the Metron Dashboard

You can launch the Metron Dashboard using the Ambari UI or a the browser of your choice.

Procedure

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

Ambari Task List



- Enter the following text in a browser:

```
$KIBANA_HOST:9995
```

Changing the Metron Dashboard Background Color

You can choose to view the Metron dashboard with either a light or dark background. The dark background is sometimes preferred in a dimly lit security operations center.

Procedure

1. Click the Gear icon



(in the top right of the Metron dashboard.

You should see a check box next to **Use dark theme** near the top of the dashboard.

2. Select the check box to use the dark theme for the dashboard.

To return to the light theme, clear the check box.

Adding a New Data Source

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

Configuring a New Data Source Index

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.

Before you begin

Before you can add a new data telemetry source to the Metron dashboard, you must ensure that you've completed the following steps:

- The data telemetry source must be added to HCP.

For information on how to add a new data telemetry source, see [Adding a New Telemetry Data Source](#).

- An index template must be created for the data telemetry source.

For information on how to create an index template, see [Specifying Index Parameters Using the Management Module](#).

Procedure

1. Click the **Settings** tab on the Metron dashboard.
2. 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.

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.

3. 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-*'.
4. 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.
5. 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



).

Related Information

[Adding a New Telemetry Data Source](#)

[Creating an Index Template](#)


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.

Procedure

1. Click on the **Discover** tab and then choose the newly created data source index pattern.
2. Click any of the fields in the left column to see a representation of the variety of data for that specific field.
3. 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.

Discover Tab with Squid Elements



Querying, Filtering, and Visualizing Data

You can interactively explore your data source data using the Metron dashboard.

When HCP parses a telemetry, it extracts and normalizes different parts of the message into a standard Metron JSON object. Standardizing and normalizing field names and formats allows HCP to search different telemetry messages with a single query. You have access to every document in every index that matches your selected index patterns. The

Metron dashboard enables you to submit search queries on the data source data, filter the search results, and view the results in a number of visualizations.

In HCP, if telemetry indexing is enabled, a rotating index for every telemetry is created. By convention this index will have a name [telemetry_name]_[timestamp]. Telemetry documents indexed into this index will by convention be called [telemetry_name].doc. Queries reference the document type of the indexed telemetries.

For more information about exploring and analyzing your data, refer to the Kibana documentation:

Table 1: Querying, Filtering, and Visualizing Data

Task	Description	Where to Look
Querying your data	You can search and refine the data you receive from your data source by creating a query from the Discover page. You should create and save a query for each data source not provided by HCP. HCP includes queries for the following telemetries: • YAF • Bro • Alerts (populated by Snort) You can also add custom queries for new telemetry types.	Discovering Your Data
Filter your query results	You can use the Metron dashboard to filter your query results to further refine the information. The Metron dashboard provides two types of filters: Time Filter Restricts the search results to a specific time period. Filter by Field Filters to display only those documents that contain a particular value in a field. You can filter either from the Fields list or the Documents table.	Discover
Visualizing your data	You can filter search results to display only those documents that contain a particular value in a field. You can also create negative filters than exclude documents that contain the specified field value.	Visualize

Related Information

[Discovering Your Data](#)

[Discover](#)

[Visualize](#)

Customizing Your Dashboard

The visualizations in your Metron dashboard are stored in resizeable containers that you can arrange on the dashboard. For more information about customizing your dashboard, see [Building a Dashboard](#).

Sharing the Metron Dashboard

You might want to share the queries and visualizations you've set up with other SOC personnel.

Table 2: Sharing the Metron Dashboard

Task	Description	Where to Look
Exporting search information	You can export the contents of a query or search. This option can be very useful after you've refined your search to display only the relevant information for a cybersecurity issue and you would like to send this information to another SOC team member.	Sharing a Dashboard
Importing search information	You can import the contents of a query or search. This option can be very useful if you need to view a colleague's refined search for a cybersecurity issue.	Loading a Dashboard

Related Information[Sharing a Dashboard](#)[Loading a Dashboard](#)

Triaging Alerts

Any event that triggers your threat intelligence thresholds will trigger an alert. These alerts are how you are notified that an event needs your attention. HCP provides a graphics user interface (GUI) to view these alerts. This GUI is a standalone user interface that connects to Elasticsearch to show the alerts but also stores all other data in the browser cache.

Launch the Alerts User Interface

The Alerts user interface is bundled with HCP and installed with the Ambari management pack.

Before you begin

- Elasticsearch must be up and running and should have alerts populated by HDP topologies.
- The Alerts UI defaults to port 4201. If you are already using port 4201 for another purpose, you must change the default port for the Alerts UI to another port number.

Procedure

1. Display the **Ambari** user interface.
2. In the Services pane, select **Metron**.
3. From the **Quick Links** menu, choose **Alerts UI**.

Note: There is no login module for the Alerts UI.

Viewing Alerts

The Alerts user interface defaults to displaying the Alerts table when first opened. You can modify the alerts displayed in the Alerts table to help identify issues.

Table 3: Alerts UI Tools and Purposes

Tools	Description
Alerts table	The Alerts table displays the alerts generated by the HCP framework. The Alerts UI polls for alerts and refreshes the Alerts table at an interval that you can configure.

Tools	Description
Searches field	You can search for alerts using the search bar above the Alerts table. The search tool follows the Lucene syntax which supports a rich query language.
Filters	The Alerts UI currently provides five filters that you can apply to alerts. You can use these filters to refine the list of alerts and collect additional information on the alerts.
Alert status	You can change the status of or dismiss an alert.
Group By	You can group alerts so you can apply filters, status, etc. on multiple alerts at a time.
Meta Alerts	The meta alert feature enables you to create a system entity that contains a collection of filtered alerts.

Related Tasks

[Search Alerts](#)

[Filter Alerts](#)

[Manage Alert Status](#)

[Group Alerts](#)

[Create a Meta Alert](#)

Using the Alerts Table

The Alerts table displays the alerts generated by the HCP framework. The Alerts UI polls for alerts and refreshes the Alerts table at an interval that you can configure. This polling is paused whenever you open any configuration panels or use the **Searches** field.

By default, the alerts table shows the recent alerts at the top. For example, alerts are sorted descending on timestamp. For information on modifying these configurations.

The Alerts table also provides the threat intelligence score for each alert. Next to the score is a bar that indicates the severity of the score:

Red	A score of 69 or higher
Orange	A score between 39 and 69
Yellow	A score below 39

Alerts (265379)

Filters: enrichm...country 3, host 10, ip_addr 10, ip_addr 9, sourcetype 2

Group By: 2 (sourcetype), 10 (ip_addr), 10 (host), 3 (enrichm...country), 9 (ip_addr)

Score	ID	timestamp	source type	ip_addr	enrichm...country	host	alert_status
829ed3f5-6...	a51eb0bdee	2017-08-31 11:47:55	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
de53d302-b...	2f4fcc094d	2017-08-31 11:47:55	bro	192.168.66.1		192.168.66.121	node1
06af55c9-3...	34dc1f9525	2017-08-31 11:47:55	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
13b7509a-0...	e9893696b0	2017-08-31 11:47:55	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
c4956536-0...	25cb707e67	2017-08-31 11:47:55	bro	192.168.66.1		192.168.66.121	node1
8169742b-e...	a51f275699	2017-08-31 11:47:55	bro	192.168.138.158	FR	62.75.195.236	62.75.195.236
2270eb9a-6...	ac30f9d14b	2017-08-31 11:47:55	bro	192.168.138.158	US	72.34.49.86	comarksecurity.com
1e31b227-0...	2132c0ea69	2017-08-31 11:47:55	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
b4a579fe-f...	ae8b0b0256	2017-08-31 11:47:55	bro	192.168.66.1		224.0.0.251	
395fc18f-c...	df3c95a056	2017-08-31 11:47:55	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
b73666a8-6...	219689109d	2017-08-31 11:47:58	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
17c2c850-c...	d60ed97cae	2017-08-31 11:47:58	bro	192.168.66.1		224.0.0.251	
76823098-0...	9b39380dd2	2017-08-31 11:47:58	bro	192.168.138.158	FR	62.75.195.236	62.75.195.236
9d86a47f-1...	999ed974e6	2017-08-31 11:47:58	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
c8e0d34a-b...	32b6d74b9a	2017-08-31 11:47:58	bro	192.168.66.1		224.0.0.251	
20390982-d...	4b61d3149e	2017-08-31 11:47:58	bro	192.168.138.158		192.168.138.2	
8992e07f-4...	f58e0caf19	2017-08-31 11:47:58	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
5d7a39b4-7...	eead99326f	2017-08-31 11:47:58	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
c0a5c0da-8...	11be5a243a	2017-08-31 11:47:58	bro	192.168.66.1		192.168.66.121	node1
b60c97c2-0...	635cb56726	2017-08-31 11:47:58	bro	192.168.138.158	FR	62.75.195.236	62.75.195.236
8f05f0cd-5...	bcb0004822	2017-08-31 11:47:26	bro	192.168.138.158	FR	62.75.195.236	62.75.195.236
4e31b7fe-9...	63d3f53068	2017-08-31 11:47:26	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
32c0a6d8-9...	c980b7c14	2017-08-31 11:47:26	bro	192.168.66.1		224.0.0.251	
767202af-c...	ecb8efcfe9	2017-08-31 11:47:26	bro	192.168.138.158	RJ	95.163.121.204	7qgnrzw...paysun.com
6ca1b862-6...	d4d3204765	2017-08-31 11:47:26	bro	192.168.66.1		192.168.66.121	node1

1 - 25 of 265379

Configure Table Columns

You can configure the table columns in the Alerts table to customize the type of information you display. You can modify the information that shows in each column, the title of the column, and the order in which the columns are displayed.

Procedure

1. Click



(gear icon).

The Alerts UI displays the Configure Table that lists all the columns available across all the valid search indexes.

Alerts Configure Table

Configure Table ✕

☐	Field	Short Name	Type		
☒	Score		STRING	-	-
☐	AA	rename	BOOLEAN	^	v
☐	adapter:geoadapter.begin.ts	rename	STRING	^	v
☐	adapter:geoadapter.end.ts	rename	STRING	^	v
☐	adapter:hostfromjsonlistadapter.begin.ts	rename	STRING	^	v
☐	adapter:hostfromjsonlistadapter.end.ts	rename	STRING	^	v
☐	adapter:threatinteladapter.begin.ts	rename	STRING	^	v
☐	adapter:threatinteladapter.end.ts	rename	STRING	^	v
☒	id	rename	STRING	^	v
☒	timestamp	rename	DATE	^	v
☒	source.type	rename	STRING	^	v
☒	ip_src_addr	rename	IP	^	v
☒	enrichments:geo:ip_dst_addr:country	rename	STRING	^	v
☒	ip_dst_addr	rename	IP	^	v
☒	host	rename	STRING	^	v
☒	alert_status	rename	STRING	^	v
☐	answers	rename	STRING	^	v
☐	bro_timestamp	rename	STRING	^	v
☐	comments	rename	OTHER	^	v
☐	dgmlen	rename	STRING	^	v
☐	enrichment joinbolt:joiner.ts	rename	STRING	^	v
☐	enrichments:geo:ip_dst_addr:city	rename	STRING	^	v

SAVE
CANCEL

- Select the fields you want to display and unselect the fields you do not want to display.
- You can rename the column titles by entering a new name in the **Short Name** column. For example, 'enrichments:geo:ip_dst_addr:country' can be renamed to 'Dst Country'. This is just for display convenience and the changes are not propagated to any system in HCP.

4. You can also configure the order in which the selected columns will appear in the table by using the arrow icons.
5. Click **Save** to save your changes and dismiss the **Configure Table** panel.
6. You can pause the Alerts UI polling by clicking the



(pause button).

Configure Table Row Settings

You can configure the table row settings in the Alerts table. You can use this feature to modify the appearance of the Alerts table and the refresh rate.

Procedure

1. Click the



(slides icon) at the top of the table to display the Settings dialog box.

Alerts Settings Panel

The screenshot shows the 'Settings' dialog box for the Alerts table. It has a dark background with white text. The title 'Settings' is at the top left. Below it, there are two sections: 'REFRESH RATE' and 'ROWS PER PAGE'. Each section has a row of buttons. In the 'REFRESH RATE' section, the buttons are '5s', '10s', '15s', '30s', '1m' (which is highlighted in blue), '10m', and '1h'. In the 'ROWS PER PAGE' section, the buttons are '10', '25' (which is highlighted in blue), '50', '100', '250', '500', and '1000'. At the bottom, there are two toggle switches, both of which are currently turned off. The first toggle is labeled 'HIDE Resolved Alerts' and the second is labeled 'HIDE Dismissed Alerts'.

2. To modify the rate at which the Alerts table is refreshed with new alert information, choose a value under **Refresh Rate**.
3. To modify the number of rows displayed in the Alerts table, choose a value under **Rows Per Page**.

Note: The number of rows that are visible in the Alerts table is restricted by the size of your browser window.

4. To hide resolved alerts or dismissed alerts, click the slide button next to the appropriate action.
HIDE Resolved Alerts and HIDE Dismissed Alerts are non-functional features in this release.

Display Additional Alerts Information

In addition to displaying alert information in the Alerts table, you can display all the information about the alert in Elasticsearch in a separate panel.

Procedure

1. Select an alert by clicking on empty space in the alert row.

The Alerts UI displays a panel listing all available data in Elasticsearch about the alert.

Alerts Information Panel

AVuKz1_n1LEanKS6qbtb ✕

Status: **NEW** | ESCALATE | OPEN | DISMISS | RESOLVE

alert_status	OPEN
dgmlen	40
enrichments:geo:ip_sr c_addr:city	Phoenix
enrichments:geo:ip_sr c_addr:country	US
enrichments:geo:ip_sr c_addr:dmaCode	753
enrichments:geo:ip_sr c_addr:latitude	33.4499
enrichments:geo:ip_sr c_addr:locID	5308655
enrichments:geo:ip_sr c_addr:location_point	33.4499,-112.0712
enrichments:geo:ip_sr c_addr:longitude	-112.0712
enrichments:geo:ip_sr c_addr:postalCode	85004
ethdst	00:00:00:00:00:00
ethlen	0x3C
ethsrc	00:00:00:00:00:00
guid	5eba8dec-278f-4f9f-b655-e98f6f4c1983
id	1906
ip_dst_addr	192.168.138.158
ip_dst_port	49197

2. The Status states at the top of the panel display the current status of the alert.

Search Alerts

You can search for alerts using the search bar above the Alerts table. The search tool follows the Lucene syntax which supports a rich query language.

Procedure

1. To search on an item that is displayed in the Alerts table, simply click on the item and it will display in the **Searches** field.

Searches Field



2. You can also directly type in the **Searches** field to enter search criteria.
For example, you can enter source:type:snort.
3. To remove an item in the **Searches** field, mouse over the information in the **Searches** field until an **x** appears at the end of the text. Click on the **x** to remove the search filter and the operator following or preceding it.
4. To clear the entire **Searches** field, click the **x** at the end of the field.
5. You can specify the time range of your search by using the time range selector on the far right of the **Searches** field.



Note:

The time-range selector is not available if you put a timestamp in the **Searches** field.

The time-range button defaults to **All time** which displays all alerts corresponding to the Searches parameters. To customize the time range, click the time-range drop-down menu and select one of the following:


Time Range

Enables you to choose the start and end dates and times for your search.

Quick Ranges

Provides a list of pre-specified time ranges that you can choose.

Time Selector Dialog Box



After you make your choice, the time-selector label will reflect your selection.



Related reference

[Apache Lucene - Query Parser Syntax](#)

Filter Alerts

The Alerts UI currently provides five filters that you can apply to alerts. You can use these filters to refine the list of alerts and collect additional information on the alerts. These filters are listed in the **Filters** panel on the left of the **Alerts** window.

Procedure

1. Click one of the filters in the **Filters** panel on the left of the window.

The Filter expands to list all of the facet values contained in the filter. For example, in the following figure, the **enrichments:geo_dst_addr:country** filter contain the countries Russia, France, and USA.

The screenshot shows the Apache Metron Alerts interface. On the left, there are filters for 'enrich...country' (3), 'host' (10), 'ip_dst_addr' (10), 'ip_src_addr' (9), and 'sourcetype' (2). The main area displays a list of alerts grouped by 'source.type' (2), 'ip_dst_addr' (10), 'host' (10), and 'enrich...country' (3). The table columns are: Score, id, timestamp, source.type, ip_src_addr, enrich...country, and ip_dst_addr. The right sidebar shows the status of the selected alert (829ed3f6-6034-4969-91c7-87...) with buttons for 'ESCALATE', 'NEW', 'OPEN', 'DISMISS', and 'RESOLVE'. Below the status is a comments section with an 'ADD COMMENT' button.

Score	id	timestamp	source.type	ip_src_addr	enrich...country	ip_dst_addr
-	06af55c9-3...34dc1f9525	2017-08-31 11:47:55	bro	192.168.138.158	RU	95.163.121.20
-	829ed3f6-6...a51eb0bdee	2017-08-31 11:47:55	bro	192.168.138.158	RU	95.163.121.20
-	de53d302-b...2fafcc094d	2017-08-31 11:47:55	bro	192.168.66.1		192.168.66.12
10	f6f44a88-8...c938c01be6	2017-08-30 08:07:59	snort	192.168.66.1		192.168.66.12
-	10a83fa3-8...41395e6201	2017-08-30 12:43:58	bro	192.168.138.158	FR	62.75.195.238
-	d80006c0-4...afcc68715	2017-08-30 12:43:58	bro	192.168.138.158	FR	62.75.195.238
-	f63236fe-3...2208f599f4	2017-08-30 12:43:58	bro	192.168.66.1		192.168.66.12
-	9ef44ca9-c...a991447655	2017-08-30 12:43:58	bro	192.168.138.158	US	204.152.354.5
-	0b94d0d6-7...fb113e0bce	2017-08-30 12:43:58	bro	192.168.138.158	FR	62.75.195.238
-	578f8cb-a...7420f2bac2	2017-08-30 12:43:58	bro	192.168.66.1		192.168.66.12
-	3964e90f-b...aeb9439c39	2017-08-30 12:43:58	bro	192.168.138.158	RU	95.163.121.20
-	90403c32-0...2ecf46b0d8	2017-08-30 12:43:58	bro	192.168.66.1		224.0.0.251
-	88950cee-5...359faa0a78	2017-08-30 12:43:58	bro	192.168.138.158	RU	95.163.121.20
-	605fe112-c...c2619a8a36	2017-08-30 12:43:58	bro	192.168.138.158		192.168.138.5
-	327e845e-1...3c21a1a799	2017-08-30 12:44:06	bro	192.168.138.158	RU	95.163.121.20
-	9ebf8599-5...5ec9069f53	2017-08-30 12:44:06	bro	192.168.138.158	FR	62.75.195.238
-	ca0d424e-6...40146306a	2017-08-30 12:44:06	bro	192.168.138.158	FR	62.75.195.238
-	779b607e-2...f1365e19a0	2017-08-30 12:44:06	bro	192.168.138.158	US	72.34.49.86
-	3b22460c-a...b8a14515cb	2017-08-30 12:44:06	bro	192.168.138.158		192.168.138.5
-	e56c2754-b...737a2ecb5e	2017-08-30 12:44:06	bro	192.168.66.1		224.0.0.251

Note:

The UI displays the number of alerts corresponding to each facet next to the facet.

- You can continue to apply filters to the alerts displayed in the **Alerts** window to further refine the alerts list. As you select filters and facets, they are displayed in the **Searches** field.

For example, in the following figure, we've applied the source.type filter with the bro facet and then the ip_dst_addr filter with the IP address 95.163.121.204.

- To clear filters that have been populated to the **Searches** field, click



(delete icon) at the end of the **Searches** field.

Manage Alert Status

You can manage one or more alerts at a time using the **ACTIONS** menu. You can use the **ACTIONS** to change the status of or dismiss an alert.

Procedure

- Select an alert by clicking on empty space in the alert row.

The Alerts UI displays a panel listing the status of the alert all available data in Elasticsearch about the alert.

Alerts Information Panel

AVuKz1_n1LEanKS6qbtb

Status

NEW

ESCALATE

OPEN

DISMISS

RESOLVE

alert_status	OPEN
dgmlen	40
enrichments:geo:ip_sr c_addr:city	Phoenix
enrichments:geo:ip_sr c_addr:country	US
enrichments:geo:ip_sr c_addr:dmaCode	753
enrichments:geo:ip_sr c_addr:latitude	33.4499
enrichments:geo:ip_sr c_addr:locID	5308655
enrichments:geo:ip_sr c_addr:location_point	33.4499,-112.0712
enrichments:geo:ip_sr c_addr:longitude	-112.0712
enrichments:geo:ip_sr c_addr:postalCode	85004
ethdst	00:00:00:00:00:00
ethlen	0x3C
ethsrc	00:00:00:00:00:00
guid	5eba8dec-278f-4f9f- b655-e98f6f4c1983
id	1906
ip_dst_addr	192.168.138.158
ip_dst_port	49197

The current alert status is highlighted.


Note:

To manage more than one alert at a time, click the check boxes at the end of alert rows, then select the action you want to perform from the ACTIONS menu.

- Click the new status you want to apply to the alert, then dismiss the panel.
- You can also add a comment to this action by clicking



(Comment button), entering your comment in the **Comments** field, and clicking **ADD COMMENT**.



The screenshot shows a dark-themed UI for managing alerts. At the top, there is a header bar with a yellow vertical bar on the left, a truncated alert ID '829ed3f6-6034-4969-91c7-87...' in the center, and a close button (X) on the right. Below the header, there is a 'Status' section with a grid of buttons: 'NEW', 'OPEN', 'DISMISS', and 'RESOLVE'. The 'ESCALATE' button is highlighted in blue. Below the status section is a 'Comments' section with a large text input area and an 'ADD COMMENT' button.

The Alerts UI indicates that an alert has one or more comments by displaying



(comment icon) next to the alert status in the **Alerts** window.

Note:

You cannot add a comment to an alert contained in a meta alert. You can only add comments to the meta alert.

4. To delete a comment, click the comment to delete, then click the trash can icon. Click OK in the **Confirmation** dialog box.


Escalate an Alert

You can escalate one or more alerts at a time to create an event that can be tracked by an external ticketing system.

Procedure

1. Select an alert by clicking on empty space in the alert row.

The Alerts UI displays a panel listing the status of the alert all available data in Elasticsearch about the alert.




The current alert status is highlighted.

Note:

To manage more than one alert at a time, click the check boxes at the end of alert rows, then select the action you want to perform from the **ACTIONS** menu.

2. Click **Escalate**.



829ed3f6-6034-4969-91c7-87...

Status

ESCALATE

NEW OPEN DISMISS

RESOLVE

Comments

ADD COMMENT

HCP writes the event to a Kafka escalation topic. An external orchestration software can pick up the event from the topic and use the API to create an incident or append to an existing incident.

3. You can also add a comment to this action by clicking



(Comment button), entering your comment in the **Comments** field, and clicking **ADD COMMENT**.

Group Alerts


You can group alerts so you can apply filters, status, etc. to multiple alerts at a time.

Procedure

1. Click one of the groups listed by **Group By**.


The **Alerts** table view changes to a tree view listing the values of the groups.

In the following example, the group is source.type and the values are Snort and Bro.



Note: The icon to the left of the value provides the cumulative severity score for all the alerts in the value. If the score exceeds 999, then the value displays as 999+.

2. Click one of the values to list the alerts for that value.



3. You can click an alert to add it to the Searches field.

Note: Searches will search through all the groups, not just the group containing the alert.

4. All features that are available for the Alerts table are available for the tree view.

For example, if you apply an action, such as Escalate, to an alert, it will apply to all alerts within the group. Similarly, if you search for a parameter, it will search all alerts within the group.

5. You can continue to refine your alerts by applying additional groups.

You can change the order in which the groups are applied to the alerts by clicking and dragging the groups on the **Groups By** line.

6. To ungroup your alerts and return to the Alerts window, click Ungroup which is located on the far right of the list of groups.




Create a Meta Alert

The meta alert feature enables you to create a save a group of filtered alerts. Like the group feature, you can group filtered alerts that pertain to an incident. However, with meta alert, you can save your grouping, creating a system entity, to view it later. Also, when you filter alerts, if a relevant alert is contained in a meta alert, the entire meta alert will be included in the filter results.

Procedure

1. Click one of the groups listed by **Group By**.

The **Alerts** table view changes to a tree view listing the values of the groups.



2. Use the **Search** and **GroupBy** options to create one or more groups containing alerts on which you want to focus.
3. When you have selected a group of alerts that you want to focus on, click



(meta alert icon), then confirm that you wish to create a meta alert with the selected alerts.

The meta alert disappears from the tree view. You can still see the meta alert in the alerts table view.

4. You can rename your meta alert by completing the following steps:
 - a) Display the Alerts UI display panel by clicking on empty space in the meta alert row.

Alerts Information Panel

AVuKz1_n1LEanKS6qbtb

×

Status

NEW

ESCALATE

OPEN

DISMISS

RESOLVE

alert_status	OPEN
dgmilen	40
enrichments:geo:ip_sr c_addr:city	Phoenix
enrichments:geo:ip_sr c_addr:country	US
enrichments:geo:ip_sr c_addr:dmaCode	753
enrichments:geo:ip_sr c_addr:latitude	33.4499
enrichments:geo:ip_sr c_addr:locID	5308655
enrichments:geo:ip_sr c_addr:location_point	33.4499,-112.0712
enrichments:geo:ip_sr c_addr:longitude	-112.0712
enrichments:geo:ip_sr c_addr:postalCode	85004
ethdst	00:00:00:00:00:00
ethlen	0x3C
ethsrc	00:00:00:00:00:00
guid	5eba8dec-278f-4f9f- b655-e98f6f4c1983
id	1906
ip_dst_addr	192.168.138.158
ip_dst_port	49197

- b) Click the current meta alert name at the top of the panel and enter your new meta alert name.
- c) Dismiss the panel by clicking the X in the upper right corner of the panel.

Save Your Searches

You can save your Alert searches for future reuse.

Procedure

1. To save a search, click the



(save button) next to the **Searches** field.


2. When prompted, enter a name for the saved search parameters, then click **Save**.
This will save both the search parameters and the column configurations.

View Your Recent and Saved Searches

You can view both your recent searches and saved searches in the Alerts UI.

Procedure


Click the



button to the left of the **Searches** field.

The Alerts UI displays the Searches panel.

Searches Panel



The **Searches** panel lists two types of searches:

Recent Searches

This is a list of your most recent searches.

To display the saved search, simply click on the search name.

The Alerts UI saves a maximum of ten of your most recent searches.

Saved Searches

This is a list of your saved searches.

To display the saved search, simply click on the search name.

You can delete any of these saved searches by clicking the trash can icon that becomes visible when you mouse over each saved search.

PCAP

The pcap data source can rapidly ingest raw data directly into HDFS from Kafka. As a result, you can store all of the raw packet capture data in HDFS and review or query it at a later date.

The pcap data is not displayed in the Metron dashboard, but you can query, view, or retrieve the data in order to port it to another application like Wireshark.

Capturing pcap Data

In your production environment there is likely to be one or more hosts configured with one or more span ports that receives raw packet data from a packet aggregation device. You can use one of HCP's packet capture programs to capture the pcap data; pycapa and DPDK. These programs are responsible for capturing the raw packet data off the wire and sending that data to Kafka where it can be ingested by HCP.

The following example uses Pycapa.

```
service pycapa start
```

If everything worked correctly, the raw packet data can be consumed from a Kafka topic called pcap. The data is binary.

```
$ /usr/hdp/current/kafka-broker/bin/kafka-console-consumer.sh -z
  zookeeper1:2181 --topic pcap
E)###>K#####P#"ssLQlJ
      P##0
E(  #@##x#####>K####"PQlJ
      ssLPF#
```

Processing pcap Data

After you capture some pcap data, the next step is to have HCP process the pcap data and store it in HDFS. Start the PCAP topology to begin this process. A Storm topology called 'pcap' is launched that consumes the raw pcap data from the Kafka topic and writes this data into sequence files in HDFS.

```
$ $METRON_HOME/bin/start_pcap_topology.sh
Running: /usr/jdk64/jdk1.8.0_77/bin/java -server -Ddaemon.name= -
Dstorm.options= -Dstorm.home=/usr/hdp/2.5.0.0-1245/storm -Dstorm.log.dir=/
var/log/storm -Djava.library.path=/usr/local/lib:/opt/local/lib:/usr/lib -
Dstorm.conf.file= -cp /usr/hdp/2.5.0.0-1245/storm/lib/log4j-core-2.1.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/storm-core-1.0.1.2.5.0.0-1245.jar:/usr/
hdp/2.5.0.0-1245/storm/lib/minlog-1.3.0.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/objenesis-2.1.jar:/usr/hdp/2.5.0.0-1245/storm/lib/ring-cors-0.1.5.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/storm-rename-hack-1.0.1.2.5.0.0-1245.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/disruptor-3.3.2.jar:/usr/hdp/2.5.0.0-1245/
storm/lib/kryo-3.0.3.jar:/usr/hdp/2.5.0.0-1245/storm/lib/log4j-over-
slf4j-1.6.6.jar:/usr/hdp/2.5.0.0-1245/storm/lib/reflectasm-1.10.1.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/log4j-slf4j-impl-2.1.jar:/usr/
```

```

hdp/2.5.0.0-1245/storm/lib/log4j-api-2.1.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/clojure-1.7.0.jar:/usr/hdp/2.5.0.0-1245/storm/lib/zookeeper.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/servlet-api-2.5.jar:/usr/hdp/2.5.0.0-1245/
storm/lib/slf4j-api-1.7.7.jar:/usr/hdp/2.5.0.0-1245/storm/lib/
asm-5.0.3.jar org.apache.storm.daemon.ClientJarTransformerRunner
org.apache.storm.hack.StormShadeTransformer /usr/metron/0.3.0/lib/metron-
pcap-backend-0.3.0.jar /tmp/d5f844e8b1a611e6a6d10a0a570e5f4d.jar
Running: /usr/jdk64/jdk1.8.0_77/bin/java -client -Ddaemon.name= -
Dstorm.options= -Dstorm.home=/usr/hdp/2.5.0.0-1245/storm -Dstorm.log.dir=/
var/log/storm -Djava.library.path=/usr/local/lib:/opt/local/lib:/usr/
lib:/usr/hdp/current/storm-client/lib -Dstorm.conf.file= -cp /usr/
hdp/2.5.0.0-1245/storm/lib/log4j-core-2.1.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/storm-core-1.0.1.2.5.0.0-1245.jar:/usr/hdp/2.5.0.0-1245/storm/lib/
minlog-1.3.0.jar:/usr/hdp/2.5.0.0-1245/storm/lib/objenesis-2.1.jar:/usr/
hdp/2.5.0.0-1245/storm/lib/ring-cors-0.1.5.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/storm-rename-hack-1.0.1.2.5.0.0-1245.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/disruptor-3.3.2.jar:/usr/hdp/2.5.0.0-1245/storm/lib/kryo-3.0.3.jar:/usr/
hdp/2.5.0.0-1245/storm/lib/log4j-over-slf4j-1.6.6.jar:/usr/hdp/2.5.0.0-1245/
storm/lib/reflectasm-1.10.1.jar:/usr/hdp/2.5.0.0-1245/storm/lib/log4j-
slf4j-impl-2.1.jar:/usr/hdp/2.5.0.0-1245/storm/lib/log4j-api-2.1.jar:/usr/
hdp/2.5.0.0-1245/storm/lib/clojure-1.7.0.jar:/usr/hdp/2.5.0.0-1245/storm/
lib/zookeeper.jar:/usr/hdp/2.5.0.0-1245/storm/lib/servlet-api-2.5.jar:/
usr/hdp/2.5.0.0-1245/storm/lib/slf4j-api-1.7.7.jar:/usr/hdp/2.5.0.0-1245/
storm/lib/asm-5.0.3.jar:/tmp/d5f844e8b1a611e6a6d10a0a570e5f4d.jar:/usr/hdp/
current/storm-supervisor/conf:/usr/hdp/2.5.0.0-1245/storm/bin -Dstorm.jar=/
tmp/d5f844e8b1a611e6a6d10a0a570e5f4d.jar org.apache.storm.flux.Flux --
remote /usr/metron/0.3.0/flux/pcap/remote.yaml --filter /usr/metron/0.3.0/
config/pcap.properties
#####      ###      #####      ###
#####      ###      #####
#####      ###      ###      #####
#####      ###      ###      #####
###      #####
###      #####      #####      ###      ###
+-      Apache Storm      +-
+- data Flow User eXperience      +-
Version: 1.0.1
Parsing file: /usr/metron/0.3.0/flux/pcap/remote.yaml
636 [main] INFO o.a.s.f.p.FluxParser - loading YAML from input stream...
638 [main] INFO o.a.s.f.p.FluxParser - Performing property substitution.
639 [main] INFO o.a.s.f.p.FluxParser - Not performing environment variable
substitution.
907 [main] WARN o.a.s.f.FluxBuilder - Found multiple invocable methods
for class org.apache.metron.spout.pcap.SpoutConfig, method from, given
arguments [END]. Using the last one found.
976 [main] INFO o.a.s.f.FluxBuilder - Detected DSL topology...
----- TOPOLOGY DETAILS -----
Topology Name: pcap
----- SPOUTS -----
kafkaSpout [1] (org.apache.metron.spout.pcap.KafkaToHDFSspout)
----- BOLTS -----
----- STREAMS -----
1157 [main] INFO o.a.s.f.Flux - Running remotely...
1157 [main] INFO o.a.s.f.Flux - Deploying topology in an ACTIVE state...
1194 [main] INFO o.a.s.StormSubmitter - Generated ZooKeeper secret payload
for MD5-digest: -8340121339010421700:-4824301672672404920
1268 [main] INFO o.a.s.s.a.AuthUtils - Got AutoCreds []
1343 [main] INFO o.a.s.StormSubmitter - Uploading topology jar /tmp/
d5f844e8b1a611e6a6d10a0a570e5f4d.jar to assigned location: /data1/hadoop/
storm/nimbus/inbox/stormjar-49aedc3d-a259-409d-a96b-4b615ce07076.jar
1810 [main] INFO o.a.s.StormSubmitter - Successfully uploaded topology jar
to assigned location: /data1/hadoop/storm/nimbus/inbox/stormjar-49aedc3d-
a259-409d-a96b-4b615ce07076.jar

```

```
1820 [main] INFO o.a.s.StormSubmitter - Submitting
      topology pcap in distributed mode with conf
      {"topology.workers":1,"storm.zookeeper.topology.auth.scheme":"digest","storm.zookeeper
2004 [main] INFO o.a.s.StormSubmitter - Finished submitting topology: pcap
```

View pcap Data

To view the pcap data, use the pcap inspector utility, `$METRON_HOME/bin/pcap_inspector.sh`. This utility enables you to retrieve and view portions of the sequence files which store the pcap data in HDFS.

Procedure

To view pcap data, use the following command:

```
usage: PcapInspector
-h, --help          Generate Help screen
-i, --input <SEQ_FILE> Input sequence file on HDFS
-n, --num_packets <N> Number of packets to dump
```

Filtering pcap Data

You can search or filter the pcap data using either a command line tool or a REST API.

Query pcap Data Using the Fixed Filter Option

You can search or filter the PCAP data by the packet header with the fixed filter command line tool.

The packet header filter is specified via the `-pf` or `--packet_filter` options.

The fixed filter option tool is executed by `${metron_home}/bin/pcap_query.sh [fixed|query]`

You can filter or query for the following fields in the PCAP data:

- `ip_scr_addr`
- `ip_dst_addr`
- `ip_src_port`
- `ip_dst_port`
- `protocol`
- `timestamp`

Fixed filter options:

```
-bop, --base_output_path <arg> Query result output path. Default is
                               '/tmp'.
-bp, --base_path <arg>         Base PCAP data path. Default is
                               '/apps/metron/pcap'.
-da, --ip_dst_addr <arg>       Destination IP address.
-df, --date_format <arg>       Date format to use for parsing start_time
                               and end_time. Default is to use time in
                               millis since the epoch.
-dp, --ip_dst_port <arg>       Destination port.
-pf, --packet_filter <arg>     Packet filter regex
-et, --end_time <arg>          Packet end time range. Default is current
                               system time.
-nr, --num_reducers <arg>      The number of reducers to use. Default
                               is 10.
-h, --help                     Display help.
-ir, --include_reverse         Indicates if filter should check swapped
```

-p, --protocol <arg>	src/dest addresses and IPs.
-rpf	IP Protocol.
-sa, --ip_src_addr <arg>	Maximum number of records per file.
-sp, --ip_src_port <arg>	Source IP address.
-st, --start_time <arg>	Source port.
	(required) Packet start time range.

Fixed filter examples:

```
$METRON_HOME/bin/pcap_query.sh fixed \
    -st "20160617" \
    -df "yyyyMMdd" \
    -sa 192.168.138.158 \
    -da 123.456.789.012 \
    -sp 49197 \
    -dp 80 \
    -p 6
    -rpf 500
```

To search for every packet that has an ip_dst_port of 8080 and contains the text "persist", run:

```
$METRON_HOME/bin/pcap_query.sh fixed \
    --ip_dst_port 8080 \
    --packet_filter \
    "\persist\" \
    -st "20170425" \
    -df "yyyyMMdd"
```

Query pcap Data Using the Query Filter Option

You can search or filter the PCAP data using a binary regular expression which can be run on the packet payload itself. This query filter option can produce a very large output and create multiple files populating them with the specified number of records and titling them with timestamps.

The query filter option is specified via the `BYTEARRAY_MATCHER(pattern, data)` Stellar function. The first argument is the regex pattern and the second argument is the data. The packet data will be exposed via the packet variable in Stellar.

The query filter option tool is executed by `${metron_home}/bin/pcap_query.sh [fixed|query]`.

You can filter or query for the following fields in the PCAP data:

- ip_scr_addr
- ip_dst_addr
- ip_src_port
- ip_dst_port
- protocol
- timestamp

Query filter options:

-bop, --base_output_path <arg>	Query result output path. Default is '/tmp'.
-bp, --base_path <arg>	Base PCAP data path. Default is '/apps/metron/pcap'.
-df, --date_format <arg>	Date format to use for parsing start_time and end_time. Default is to use time in millis since the epoch.
-et, --end_time <arg>	Packet end time range. Default is current system time.
-nr, --num_reducers <arg>	The number of reducers to use. Default is 10.

-h, --help	Display help.
-q, --query <arg>	Query string to use as a filter.
-rpf	Maximum number of records per file.
-st, --start_time <arg>	(required) Packet start time range.

The Query filter's --query argument specifies the Stellar expression to execute on each packet. To interact with the packet, a few variables are exposed:

- packet : The packet data (a byte[])
- ip_src_addr : The source address for the packet (a String)
- ip_src_port : The source port for the packet (an Integer)
- ip_dst_addr : The destination address for the packet (a String)
- ip_dst_port : The destination port for the packet (an Integer)
- BYTEARRAY_MATCHER : The first argument is the regex pattern and the second argument is the data. The packet data will be exposed via the packet variable in Stellar.

Query filter examples:

```
$METRON_HOME/bin/pcap_query.sh query \
    -st "20160617" \
    -df "yyyyMMdd" \
    --query "ip_src_addr ==
'192.168.138.158' and ip_src_port == '49197' \
    and ip_dst_addr ==
'123.456.789.012' and ip_dst_port == '80' \
    and protocol == '6'"
    -rpf 500
```

To search for every packet that has an ip_dst_port of 8080 and contains the text "persist", run:

```
$METRON_HOME/bin/pcap_query.sh query \
    --query "ip_dst_port == 8080 &&
    BYTEARRAY_MATCHER('\`persist\`', packet)" \
    -st "20170425" \
    -df "yyyyMMdd"
```

You can also do proper binary regexes that look for packets containing the text "persist" and the 2 byte sequence 0x1F909 (in hex):

```
$METRON_HOME/bin/pcap_query.sh query \
    --query "BYTEARRAY_MATCHER('1F90', packet) &&
    BYTEARRAY_MATCHER('\`persist\`', packet)" \
    -st "20170425" \
    -df "yyyyMMdd"
```

Other examples:

```
$METRON_HOME/bin/pcap_query.sh query \
    -st "1466136000000" \
    --query "IN_SUBNET(ip_src_addr,
'192.168.0.0/24') and ip_src_port == '49197' \
    and ip_dst_addr ==
'123.456.789.012' and ip_dst_port == '80' \
    and protocol == '6'"
    -rpf 500
```

```
# subnet function checks IP is in specified subnet
--query "IN_SUBNET(ip_src_addr, '192.168.0.0/24') \
    and ip_src_port == '49197' \
    and ip_dst_addr == '123.456.789.012' \
```

```

        and ip_dst_port == '80' \
        and protocol == '6'"

# range queries on ports
--query "ip_src_port <= 50000 and ip_dst_port >= 30000"

# range queries with conditionals and parens
--query "(ip_src_port < 50000 and ip_src_port > 40000) \
        or (ip_src_port < 20000 and ip_src_port > 10000)"

# in/not in list of values
--query "ip_src_port < 10000 and ip_dst_port in ['54056', '54057',
        '8080']"

```

Query pcap Data Using the REST API

The REST API is an alternative to using the CLI and Stellar to query the pcap data. The purpose of this service is to provide a middle tier to negotiate retrieving packet capture data that flows into HCP. This packet data is in a form that libpcap-based tools can read. The REST API exposes the query functionality via http.

This fixed filter is less expressive than the CLI/Stellar filter option. The fixed filter gives you the following values that you can filter on:

- srcIp
- srcPort
- dstIp
- dstPort
- startTime
- endTime

The fixed filter filters on explicit matches only so you cannot use any specialized functions or comparison operators. The query filter is far more expressive because it exposes Stellar, which enables the user to invoke specialized functions, use comparison operators, and perform boolean logic on the parameters.

You can start the fixed filter service either via the init.d script installed, /etc/init.d/pcapservice or directly via the yarn jar command:

```

yarn jar $METRON_HOME/lib/metron-api-$METRON_VERSION.jar
org.apache.metron.pcapservice.rest.PcapService -port $SERVICE_PORT -
query_hdfs_path $QUERY_PATH -pcap_hdfs_path $PCAP_PATH

```

where

- METRON_HOME is the location of the metron installation
- METRON_VERSION is the version of the metron installation
- SERVICE_PORT is the port to bind the REST service to
- QUERY_PATH is the temporary location to store query results. They are deleted after the service reads them.
- PCAP_PATH is the path to the packet data on HDFS

Porting pcap Data to Another Application

You can port pcap data to another application using the libpcap-compliant pcap file.

When you use the pcap query utility to extract pcap data, the utility creates a libpcap-compliant pcap file in the current working directory.

```


[root@ip-10-0-0-53 0.3.0]# ls -l

```

```
total 72
drwxr-xr-x. 2 livy games 4096 Nov 22 22:36 bin
drwxr-xr-x. 3 livy games 4096 Nov 23 17:10 config
drwxr-xr-x. 2 livy games 4096 Sep 29 17:44 ddl
drwxr-xr-x. 6 livy games 4096 Aug 22 14:54 flux
drwxr-xr-x. 2 root root 4096 Nov 23 17:07 lib
drwxr-xr-x. 2 livy games 4096 Nov 22 22:36 patterns
-rw-r--r--. 1 root root 48506 Nov 23 18:06 pcap-
data-20161123180607184+0000.pcap
```

```
[root@ip-10-0-0-53 0.3.0]# file pcap-data-20161123180607184+0000.pcap
pcap-data-20161123180607184+0000.pcap: tcpdump capture file (little-endian)
- version 2.4 (Ethernet, capture length 65535)
```

You can open the libpcap-compliant pcap file with any third-party tool that supports the file type. For example, you can load Wireshark and choose File > Open. Wireshark will load the pcap file.



The content of the file will be similar to the following:

