

HCP Architecture 1

Architecture

Date of Publish: 2019-04-09



<https://docs.hortonworks.com>

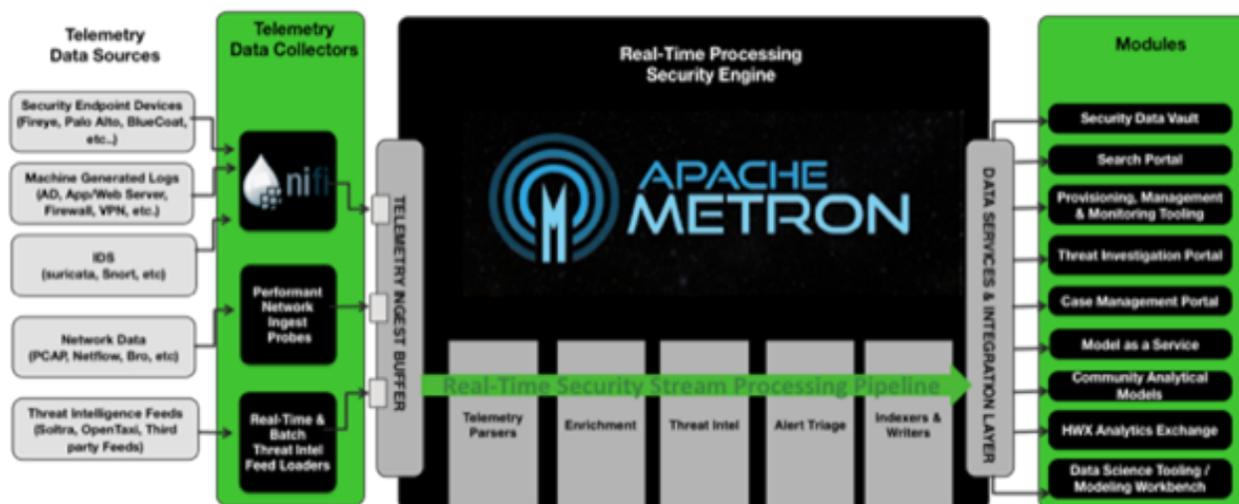
Contents

Real-Time Processing Security Engine.....	3
HCP High Level Architecture.....	3
Telemetry Data Collectors.....	4
Data Services and Integration Layer.....	4

Real-Time Processing Security Engine

The core of Hortonworks Cybersecurity Platform (HCP) architecture is the Apache Metron real-time processing security engine.

The real-time processing security engine provides the ingest buffer to capture raw events, and, in real time, parses the raw events, enriches the events with relevant contextual information, enriches the events with threat intelligence, and applies available models (such as triaging threats by using the Stellar language). The engine then writes the events to a searchable index, as well as to HDFS, for analytics.



HCP High Level Architecture

Hortonworks Cybersecurity Platform (HCP) is primarily backed by Storm and Kafka.

HCP also leverages the following components:

Zookeeper

Zookeeper provides dynamic configuration updates to running Storm topologies. This enables HCP to push updates to our Storm topologies without restarting them.

HBase

HCP uses HBase primarily for enrichments. But HBase is also used it to store user state for our UI's.

HDFS

HDFS uses HDFS for long term storage. Parsed and enriched messages land here, along with any reported exceptions or errors encountered along the way.

Solr and Elasticserach (plus Kibana)

HDP uses Solr and Elasticsearch (plus Kibana) for real-time access. HCP provides out of the box compatibility with both Solr and Elasticsearch, and custom dashboards for data exploration in Kibana.

Zeppelin

Zeppelin provides dashboards to perform custom analytics.

Kafka

Information is pushed into Metron by setting up Kafka topics for parsers to read from. There are a variety of options for setting up Kafka topics, including, but not limited to:

- Brok Kafka plugin
- Fastcapa
- NiFi

Telemetry Data Collectors

Telemetry data collectors push or stream the data source events into Apache Metron. Hortonworks Cybersecurity Platform (HCP) works with Apache NiFi to push the majority of data sources into Apache Metron.

For high-volume network data, HCP provides a performant network ingest probe. And for threat intelligence feeds, HCP supports a set of both streaming and batch loaders that enables you to push third-party intelligence feeds into Apache Metron.

Data Services and Integration Layer

The data services and integration layer is a set of three HCP modules that provides different features for different SOC personas.

HCP provides three modules for the integration layer.

Security data vault

Stores the data in HDFS.

Search portal

The Metron dashboard.

Provisioning, management, and monitoring tool

An HCP-provided management module that expedites provisioning and managing sensors. Other provisioning, management, and monitoring functions are supported through Apache Ambari.