

CCP Architecture 2.0.0

Architecture

Date of publish: 2017-11-06

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2019. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 ("ASLv2"), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER'S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Real-Time Processing Security Engine.....4

CCP High Level Architecture..... 4

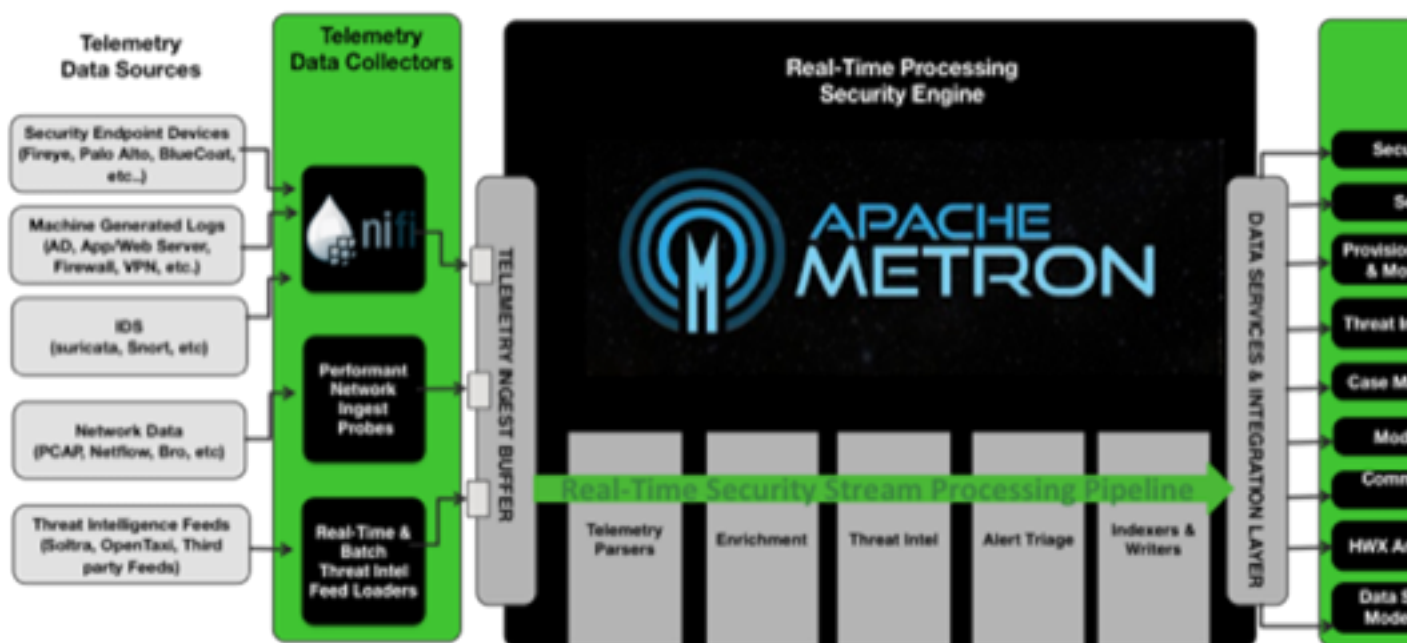
Telemetry Data Collectors..... 5

Data Services and Integration Layer..... 5

Real-Time Processing Security Engine

The core of Cloudera Cybersecurity Platform (CCP) 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.



CCP High Level Architecture

Cloudera Cybersecurity Platform (CCP) is primarily backed by Storm and Kafka.

CCP also leverages the following components:

ZooKeeper

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

HBase

CCP uses HBase primarily for enrichments. But HBase is also used to store user state for our UIs.

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 Elasticsearch (plus Kibana)

HDP uses Solr and Elasticsearch (plus Kibana) for real-time access. CCP 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:

- Grok Kafka plugin
- Fastcapa
- NiFi

Telemetry Data Collectors

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

For high-volume network data, CCP provides a performant network ingest probe. And for threat intelligence feeds, CCP 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 CCP modules that provides different features for different SOC personas.

CCP 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

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