Release Notes 1

Hortonworks Cybersecurity Platform

Date of Publish: 2018-07-15

http://docs.hortonworks.com

Contents

Hortonworks Cybersecurity Platform 1.5.1 Release Notes..... 3

Apache Component Support....3
New Features....3
Operating System Support Matrix....3
JDK Support Matrix....3
Unsupported Features....4
Community Features....4
Technical Preview Features....4
HCP 1.5.1 Repositories....4
Upgrading to HCP 1.5.1....5
Switching to Unified Enrichment Topology (Technical Preview)....5
Upgrading to Elasticsearch 5.6.2....6
Type Mapping Changes....6
Third-Party Licenses....9
Known Issues....9
Known Differences Between HCP 1.5.1 and HCP 1.5.0....10
Known Differences Between HCP 1.5.1 and Apache Metron 0.5.0....11

Hortonworks Cybersecurity Platform 1.5.1 Release Notes

This document provides you with the latest information about the Hortonworks Cybersecurity Platform (HCP) powered by Apache Metron release 1.5.1 and its product documentation.

Apache Component Support

Hortonworks Cybersecurity Platform (HCP) 1.5.1 is built on HDP 2.6.4 and HDF 3.0.1.1 and later.

The official Apache versions of all HCP 1.5.1 components are:

- Apache Metron 0.5.0
- HDP supported component versions

All components listed are official Apache releases of the most recent stable versions available.

The Hortonworks approach is to provide patches only when necessary, to ensure the interoperability of components. Unless you are explicitly directed by Hortonworks Support to take a patch update, each of the HCP components should remain at the following package version levels, to ensure a certified and supported copy of HCP 1.5.1.

Note:

For information on open source software licensing and notices, refer to the Licenses and Notices files included with the software install package.

New Features

HCP is a cybersecurity application framework that provides the ability to parse diverse security data feeds, enrich, triage, and store the data at scale, and detect cybersecurity anomalies.

HCP 1.5.1 provides the following new features:

- Performance enhanced enrichment topology
- Support for Solr 6.6 using HDP Search
- Performance improvements for Stellar

Operating System Support Matrix

HCP 1.5.1 supports a select set of operating system versions.

Unless otherwise noted, the following operating systems are validated and supported for HDP 2.6.4:

Table 1: HDP 2.6.2 Operating System Support Matrix

Operating System	Version
CentOS (64-bit)	CentOS 6.x and CentOS 7.x
Red Hat (64-bit)	RHEL 7.0†
Ubuntu	Ubuntu 14.04

†Not validated, but supported.

JDK Support Matrix

HCP 1.5.1 supports a select set of Java Development Kits (JDK) versions.

Unless otherwise noted, the following Java Development Kits (JDKs) are validated and supported for HDP 2.6.4:

Table 2: HDP 2.6.4 JDK Support Matrix

JDK	Version
Open Source	JDK8†
Oracle	JDK 8

†Not validated, but supported.

Unsupported Features

Although some features exist with HCP 1.5.1, Hortonworks does not support some community features and technical preview features.

Community Features

Some community features are developed and tested by the Hortonworks community but are not officially supported by Hortonworks. These features are excluded for a variety of reasons, including insufficient reliability or incomplete test case coverage, declaration of non-production readiness by the community at large, and feature deviation from Hortonworks best practices. Do not use these features in your production environments.

Table 3: Community Features

Feature	Description
Vagrant-based deployment	A single-node quick deployment option intended solely for development of Metron.
Docker-based deployment	A Docker-container based deployment intended solely for development of Metron.
Ansible installs	A multi-node deployment option via Ansible.

Technical Preview Features

Some features included in the HCP 1.5.1 release are not yet officially supported by Hortonworks. These technical preview features are still under development and are not recommended for a production environment.

Table 4: Technical Preview Features

Feature	Description
Meta Alerts UI	The Meta Alerts UI feature with Solr is technical preview in this release. We do not yet recommend this for production use, but please let us know about any bugs you might find. We appreciate your feedback.
Stellar in Zeppelin	The ability to run Stellar commands in Zeppelin notebook

HCP 1.5.1 Repositories

You can download HCP 1.5.1 from HCP repository locations specific to the operating system you use.

Use the following table to identify the HCP 1.5.1 repo location for your operating system and operational objectives:

Note:

When installing Elasticsearch with the management pack on Ubuntu, you must manually install the Elasticsearch repositories. The management pack does not do this, like it does on CentOS.

Table 5: HCP Repo Locations

OS	Format	Download Location
RedHat Enterprise Linux / CentOS 6 (64-bit)	Repo	http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.5.1.0/hcp.repo
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.5.1.0/tars/metron/hcp-ambari-mpack-1.5.1.0-16.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.5.1.0/tars/metron/elasticsearch_mpack-1.5.1.0-16.tar.gz
RedHat Enterprise Linux / CentOS 7 (64-bit)	Repo	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.1.0/hcp.repo
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.1.0/tars/metron/hcp-ambari-mpack-1.5.1.0-16.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.1.0/tars/metron/elasticsearch_mpack-1.5.1.0-16.tar.gz
Ubuntu 14.04	Repo	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.1.0/hcp.list
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.1.0/tars/metron/hcp-ambari-mpack-1.5.1.0-16.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.1.0/tars/metron/elasticsearch_mpack-1.5.1.0-16.tar.gz

Upgrading to HCP 1.5.1

For information on how to upgrade to HCP 1.5.1 from a previous release, see Hortonworks Cybersecurity Platform Upgrade Guide.

Switching to Unified Enrichment Topology (Technical Preview)

Switching from the current split-join enrichment topology to the new unified enrichment topology can reduce the latency of enrichment messages and avoid overloading the enrichment cache during times of heavy traffic.

Procedure

1. Stop the Metron enrichment topology in Ambari.
   a) Click **Metron Enrichment** in the **Summary** list.
   b) Choose **Stop** from the menu next to **Metron Enrichment / Metron**.
2. In the **Enrichment** tab, choose **Unified** from the **Enrichment Topology** menu.

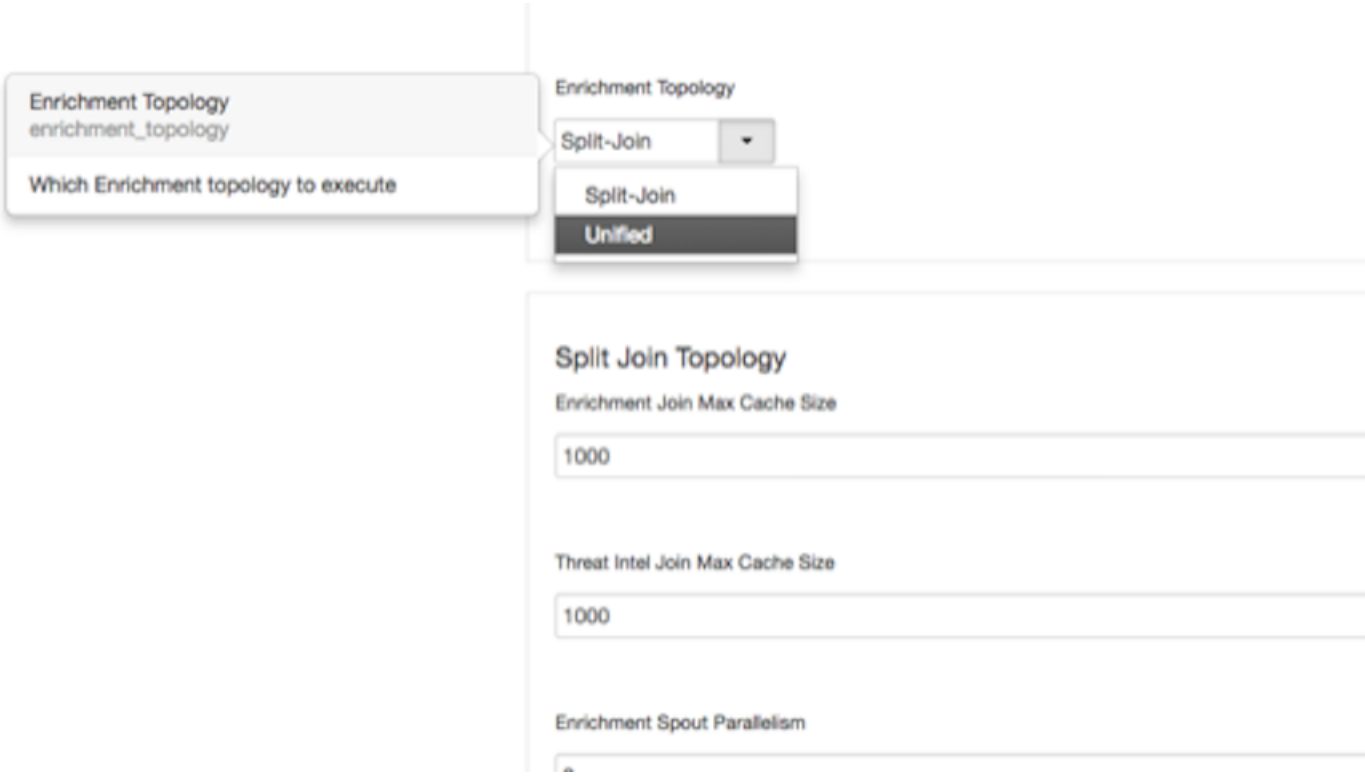

Where appropriate, the unified topology reuses the same settings from the split-join topology.

3. Verify that the unified topology settings are appropriate for your system.
4. Restart the enrichment topology in Ambari.

Upgrading to Elasticsearch 5.6.2

There are a number of template changes in Elasticsearch 5.6.2, most notably around string type handling, that may cause issues when upgrading.

For Elasticsearch 5.x, the existing indexes and templates need to upgraded. For more information, see:

- Updating Elasticsearch Templates to Work with Elasticsearch 5.x
- Updating Existing Indexes to Work with Elasticsearch 5.x

If you are upgrading from Elasticsearch 2.x to Elasticsearch 5.6.2, you will need to re-index.

Related Information

Upgrade Elasticsearch

Type Mapping Changes

Type mappings in Elasticsearch 5.6.2 have changed from ES 2.x. This section provides an overview of the most significant changes.

The following is a list of the major changes in Elasticsearch 5.6.2:

- String fields replaced by text/keyword type
- Strings have new default mappings as follows:

```
{
  "type": "text",
  "fields": {
    "keyword": {
      "type": "keyword",
```

```
        "ignore_above": 256
      }
    }
  }
```

- There is no longer a _timestamp field that you can set "enabled" on.

  This field now causes an exception on templates. The Metron model has a timestamp field that is sufficient.

The semantics for string types have changed. In 2.x, index settings are either "analyzed" or "not_analyzed" which means "full text" and "keyword", respectively. Analyzed text means the indexer will split the text using a text analyzer, thus allowing you to search on substrings within the original text. "New York" is split and indexed as two buckets, "New" and "York", so you can search or query for aggregate counts for those terms independently and match against the individual terms "New" or "York." "Keyword" means that the original text will not be split/analyzed during indexing and instead treated as a whole unit. For example, "New" or "York" will not match in searches against the document containing "New York", but searching on "New York" as the full city name will match. In Elasticsearch 5.6 language, instead of using the "index" setting, you now set the "type" to either "text" for full text, or "keyword" for keywords.

Below is a table listing the changes to how String types are now handled.

sort, aggregate, or access values	Elasticsearch 2.x	Elasticsearch 5.x	Example
no	`"my_property" : { "type": "string", "index": "analyzed" }`	`"my_property" : { "type": "text" }` Additional defaults: "index": "true", "fielddata": "false"	"New York" handled via in-mem search as "New" and "York" buckets. No aggregation or sort.
yes	`"my_property": { "type": "string", "index": "analyzed" }`	`"my_property": { "type": "text", "fielddata": "true" }`	"New York" handled via in-mem search as "New" and "York" buckets. Can aggregate and sort.
yes	`"my_property": { "type": "string", "index": "not_analyzed" }`	`"my_property" : { "type": "keyword" }`	"New York" searchable as single value. Can aggregate and sort. A search for "New" or "York" will not match against the whole value.
yes	`"my_property": { "type": "string", "index": "analyzed" }`	`"my_property": { "type": "text", "fields": { "keyword": { "type": "keyword", "ignore_above": 256 } } }`	"New York" searchable as single value or as text document, can aggregate and sort on the sub term "keyword."

If you want to set default string behavior for all strings for a given index and type, you can do so with a mapping similar to the following (replace ${your_type_here} accordingly):

```
# curl -XPUT 'http://${ES_HOST}:${ES_PORT}/_template/
default_string_template' -d '
{
    "template": "*",
    "mappings" : {
        "${your_type_here}": {
            "dynamic_templates": [
                {
                    "strings": {
                        "match_mapping_type": "string",
                        "mapping": {
                            "type": "text"
                            "fielddata": "true"
                        }
                    }
                }
            ]
        }
    }
}
```

By specifying the template property with value *, the template will apply to all indexes that have documents indexed of the specified type (${your_type_here}).

The following are other settings for types in ES:

- doc_values
  - On-disk data structure
  - Provides access for sorting, aggregation, and field values
  - Stores same values as _source, but in column-oriented fashion better for sorting and aggregating
  - Not supported on text fields
  - Enabled by default
- fielddata
  - In-memory data structure
  - Provides access for sorting, aggregation, and field values
  - Primarily for text fields
  - Disabled by default because the heap space required can be large

Third-Party Licenses

HCP deploys numerous third-party licenses and dependencies, all of which are compatible with the Apache software license. For complete third-party license information, see the licenses and notice files contained within the distribution.

Related Information

Apache 2.0

Known Issues

The HCP 1.5.1 release has the following known issue:

- Queries in the Alerts UI against source.type:metaalert will not work in Solr. Specifically, filtering by source.type:metaalert will not return any results. Searches against other types of metaalerts, for example ip.src.addr:192.168.1.1 will produce results.

- During HCP installation, some versions of Zeppelin might fail to install. If the Zeppelin notebooks are not installed, import the Apache Zeppelin Notebook manually.

Related Information

Importing the Apache Zeppelin Notebook Manually

Known Differences Between HCP 1.5.1 and HCP 1.5.0

The following bugs identify known differences between HCP 1.5.1 and HCP 1.5.0.

Table 6: Known Differences Between HCP 1.5.1 and HCP 1.4.2

Feature	Description
METRON-1421	Create a SolrMetaAlertDao
METRON-1489	Retrofit UI tests to run reliably during nightly QE runs
METRON-1532	Getting started documentation improvements
METRON-1533	Create KAFKA_FIND Stellar Function
METRON-1544	Flaky test: org.apache.metron.stellar.common.CachingStellarProcessorTest#testCaching
METRON-1547	Solr Comment Fields
METRON-1553	Validate JIRA Script Error
METRON-1565	Metaalerts fix denormalization after moving to active status
METRON-1566	Alert updates are not propagated to metaalert child alerts
METRON-1568	Stellar should have a _ special variable which returns the message in map form
METRON-1569	Allow user to change field name conversion when indexing to Elasticsearch
METRON-1571	Correct KAFKA_TAIL Seek to End Logic
METRON-1572	Enhance KAFKA_PUT function
METRON-1573	Enhance KAFKA_* functions to return partition and offset details
METRON-1574	Update version to 0.5.0
METRON-1575	Add leet gpg public key to the KEYS file
METRON-1576	bundle.css RAT failure for metron-interface/metron-alerts
METRON-1577	Solr searches don't include the index of the result
METRON-1579	Stellar should return the expression that failed in the exception
METRON-1580	Release candidate check script requires Bro Plugin
METRON-1584	Fix multivalue field errors in Bro Solr schema
METRON-1585	SolrRetrieveLatestDao does not use the collection lookup
METRON-1586	Defaulting for the source type field in alerts UI does not work
METRON-1589	'/api/v1/search/search' fails when 'Solr Zookeeper Urls' has comma separated multiple zookeeper urls
METRON-1592	Unable to use third party parser with Storm versions >= 1.1.0
METRON-1593	Setting Metron rest additional classpath removes HBase and Hadoop configs from classpath
METRON-1594	Indexing Topology Crashes with Invalid Message
METRON-1598	NoClassDefFoundError when running with Elasticsearch X-Pack
METRON-1599	Allow user to define global property 'source.type.field' in Ambari

Feature	Description
METRON-1601	Rename metaalert alert nested field to metron_alert to avoid collision
METRON-1603	Fix multivalue field errors in Bro Solr schema
METRON-1607	Update public web site to point at 0.5.0 new release
METRON-1608	Add configuration for threat.triage.field name
METRON-1609	Elasticsearch settings in Ambari should not be required if Solr is the indexer
METRON-1611	Increment master version number to 0.5.1 for on-going development
METRON-1612	Fix website download links
METRON-1613	Metaalerts status update broken in Alerts UI
METRON-1616	Changing alert status fails if no metaalerts have been created yet
METRON-1622	Allow user to define global property 'threat.triage.score.field' in Ambari
METRON-1624	Set Profiler and Enrichment batch parameters in Ambari
METRON-1625	Merge master into Solr feature branch
METRON-1626	Set Profiler and Enrichment batch parameters in Ambari
METRON-1627	Alerts UI: Metaalert details missing in details panel when trying to add alert to existing metaalert
METRON-1629	Update Solr documentation
METRON-1630	Add threat.triage.score.field to READMEs
METRON-1633	Incorrect instructions when merging PR into feature branch
METRON-1637	Wrong path to escalate alert REST endpoint

Known Differences Between HCP 1.5.1 and Apache Metron 0.5.0

The following bugs identify known differences between HCP 1.5.1 and Apache Metron 0.5.0.

Table 7: Known Differences Between HCP 1.5.1 and HCP 1.4.2

Feature	Description
METRON-508	Expand Elasticsearch templates to support the standard bro logs
METRON-986	Enhance Fastcapa to Support Intel X520
METRON-990	Clean up and organize flux properties
METRON-1007	Ship the metron-management jar as part of the mpack install
METRON-1012	Update Metron public web site for 0.4.0 release
METRON-1014	StellarShell class name typo
METRON-1021	increment metron version number to 0.4.1 in poms etc
METRON-1489	Retrofit UI tests to run reliably during nightly QE runs
METRON-1624	Set Profiler and Enrichment batch parameters in Ambari
METRON-1634	Alerts UI add comment doesn't immediately show up
METRON-1637	Wrong path to escalate alert REST endpoint