# Hortonworks Cybersecurity Package

## Release Notes

(January 26, 2018)

# Hortonworks Cybersecurity Package: Release Notes



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.

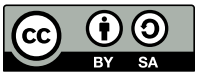

# Table of Contents

# List of Tables

# 1. Hortonworks Cybersecurity Suite 1.4.1 Release Notes

This document provides you with the latest information about the Hortonworks Cybersecurity Package (HCP) powered by Apache Metron release 1.4.1 and its product documentation.



## 1.1. Apache Component Support

**Component Versions**

HCP is built on HDP 2.5.6 and HDF 3.0.1.1 and later. The official Apache versions of all HCP 1.4.1 components are:

- Apache Metron 0.4.2
- 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 1.4.1.

### Note

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

## 1.2. 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 1.4.1 provides the following new features:

- Support for Elasticsearch 5.6.2.

  Elasticsearch 2.x is no longer supported.

- Support for Kibana 5.6.2 including updated dashboards.
- Support for Curator utility provided by Elasticsearch.

  Data Pruner is no longer supported.
- Alerts user interface
  - Displaying alerts
  - Searching alerts
  - Saving searches
  - Viewing your recent and saved searches
  - Configuring Alerts table
  - The ability to group alerts into meta-alerts
  - Adding comments
  - Alert status based workflow
  - Ability to escalate alerts to external ticketing systems
- Significant performance improvement for parsing
- M-pack based installation and configuration for Profiling
- Performance improvement for Indexing
- Additional geospatial and hash functions in Stellar
- Short circuit evaluation and multi-line Stellar statements

# 1.3. Platform Support Matrices

This section outlines the platform support matrices for HCP 1.4.1.

- Operating System Support Matrix [2]
- JDK Support Matrix [3]

## 1.3.1. Operating System Support Matrix

Unless otherwise noted, the following operating systems are validated and supported for HDP 2.6.2:

### Table 1.1. 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.

## 1.3.2. JDK Support Matrix

Unless otherwise noted, the following Java Development Kits (JDKs) are validated and supported for HDP 2.6.0:

### Table 1.2. HDP 2.6.0 JDK Support Matrix

| JDK | Version |
|---|---|
| Open Source | JDK8† |
| Oracle | JDK 8 |

†Not validated, but supported.

# 1.4. Unsupported Features

Although the following features exist within HCP 1.4.1, Hortonworks does not currently support these specific capabilities:

- Community Features [3]

## 1.4.1. Community Features

The following 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 1.3. 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. |

# 1.5. HCP 1.4.1 Repositories

Use the following table to identify the HCP 1.4.1 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 1.4. 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/1.4.1.0/hcp.repo |
| | HCP Management Pack | http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.4.1.0/tars/metron/hcp-ambari-mpack-1.4.1.0-18.tar.gz |
| RedHat Enterprise Linux / CentOS 7 (64-bit) | Repo | http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.4.1.0/hcp.repo |
| | HCP Management Pack | http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.4.1.0/tars/metron/hcp-ambari-mpack-1.4.1.0-18.tar.gz |
| Ubuntu 14 | Repo | http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.4.1.0/hcp.list |
| | HCP Management Pack | http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.4.1.0/tars/metron/hcp-ambari-mpack-1.4.1.0-18.tar.gz |

# 1.6. Upgrading to HCP 1.4.1

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

- Updating Elasticsearch Templates
- Updating Existing Indexes

For information on how to upgrade to HCP 1.4.1 from a previous release, see Apache Metron Upgrading.

# 1.7. Upgrading to Elasticsearch 5.6.2

There are a number of template changes in Elasticsearch 5.2, most notably around string type handling, that may cause issues when upgrading. If you are upgrading from Elasticsearch 2.x to Elasticsearch 5.6.2, you will need to re-index. For information on the type mapping changes, see Section 1.7.1, “Type Mapping Changes” [4].

For more information, see Upgrading Elasticsearch.

## 1.7.1. 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" }`<br><br>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

# 1.8. Third-Party Licenses

Global: Apache 2.0

| Apache Component | Subcomponents | License |
|---|---|---|
| Storm | Logback | EPL |

# 1.9. Known Issues

The HCP 1.4.1 release has the following known issue.

- During HCP installation, some versions of Zeppelin might fail to install. If the Zeppelin notebooks are not installed, import the Apache Zeppelin Notebook manually. See https://docs.hortonworks.com/HDPDocuments/HCP1/HCP-1.4.1/bk_installation/content/importing_zeppelin_notebook_manually.html for more information.

For a list of known differences between HCP 1.4.1 and Apache Metron 0.4.2, see Section 1.9.1, "Known Differences Between HCP 1.4.1 and Apache Metron 0.4.2" [7].

## 1.9.1. Known Differences Between HCP 1.4.1 and Apache Metron 0.4.2

The following Apache bugs identify known differences between HCP 1.4.1 and Apache Metron 0.4.2.

### Table 1.5. Known Differences Between HCP 1.4.1 and Apache Metron 0.4.2

| Feature | Description |
|---|---|
| https://issues.apache.org/jira/browse/METRON-1091 | Shell: Stand Alone installation |
| https://issues.apache.org/jira/browse/METRON-1427 | Add support for storm 1.1 and hdp 2.6 |
| https://issues.apache.org/jira/browse/METRON-1391 | Typos in Documentation/Examples within metron-management/README.md |
| https://issues.apache.org/jira/browse/METRON-1389 | Zeppelin notebook import does not work with Ambari 2.6 |
| https://issues.apache.org/jira/browse/METRON-1432 | JDK Install Fails on Ubuntu Development Environment |
| https://issues.apache.org/jira/browse/METRON-1431 | Add REGEXP_REPLACE function to Stellar |
| https://issues.apache.org/jira/browse/METRON-1410 | Some more upgrade fallout... Can't restart Metron Indexing |