Hortonworks Cybersecurity Platform

Release Notes

(May 18, 2018)

Hortonworks Cybersecurity Platform: Release Notes

Hortonworks Cybersecurity Platform (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 dashboarding.

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 Platform 1.5.0 Release Notes

This document provides you with the latest information about the Hortonworks Cybersecurity Platform (HCP) powered by Apache Metron release 1.5.0 and its product documentation.

1.1. Apache Component Support

Component Versions

HCP is built on HDP 2.6.4 and HDF 3.0.1.1 and later. The official Apache versions of all HCP 1.5.0 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 1.5.0.

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 1.5.0 provides the following new features:

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

- Performance improvements for Stellar

1.3. Platform Support Matrices

This section outlines the platform support matrices for HCP 1.5.0.

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

1.3.1. Operating System Support Matrix

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

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 2.6.4:

Table 1.2. HDP 2.6.4 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 1.5.0, Hortonworks does not currently support these specific capabilities:

- Community Features [2]

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.

1.4.2. Technical Preview Features

Table 1.4. 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

1.5. HCP 1.5.0 Repositories

Use the following table to identify the HCP 1.5.0 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 1.5. 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/1.5.0.0/hcp.repo
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.5.0.0/tars/metron/hcp-ambari-mpack-1.5.0.0-9.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/centos6/1.x/updates/1.5.0.0/tars/metron/elasticsearch_mpack-1.5.0.0-9.tar.gz
RedHat Enterprise Linux / CentOS 7 (64-bit)	Repo	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.0.0/hcp.repo
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.0.0/tars/metron/hcp-ambari-mpack-1.5.0.0-9.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/centos7/1.x/updates/1.5.0.0/tars/metron/elasticsearch_mpack-1.5.0.0-9.tar.gz
Ubuntu 14.04	Repo	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.0.0/hcp.list
	HCP Management Pack	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.0.0/tars/metron/hcp-ambari-mpack-1.5.0.0-9.tar.gz
	Elasticsearch Management Pack	http://public-repo-1.hortonworks.com/HCP/ubuntu14/1.x/updates/1.5.0.0/tars/metron/elasticsearch_mpack-1.5.0.0-9.tar.gz

1.6. Upgrading to HCP 1.5.0

For information on how to upgrade to HCP 1.5.0 from a previous release, see Hortonworks Cybersecurity Platform Upgrade Guide.

1.7. 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.

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.

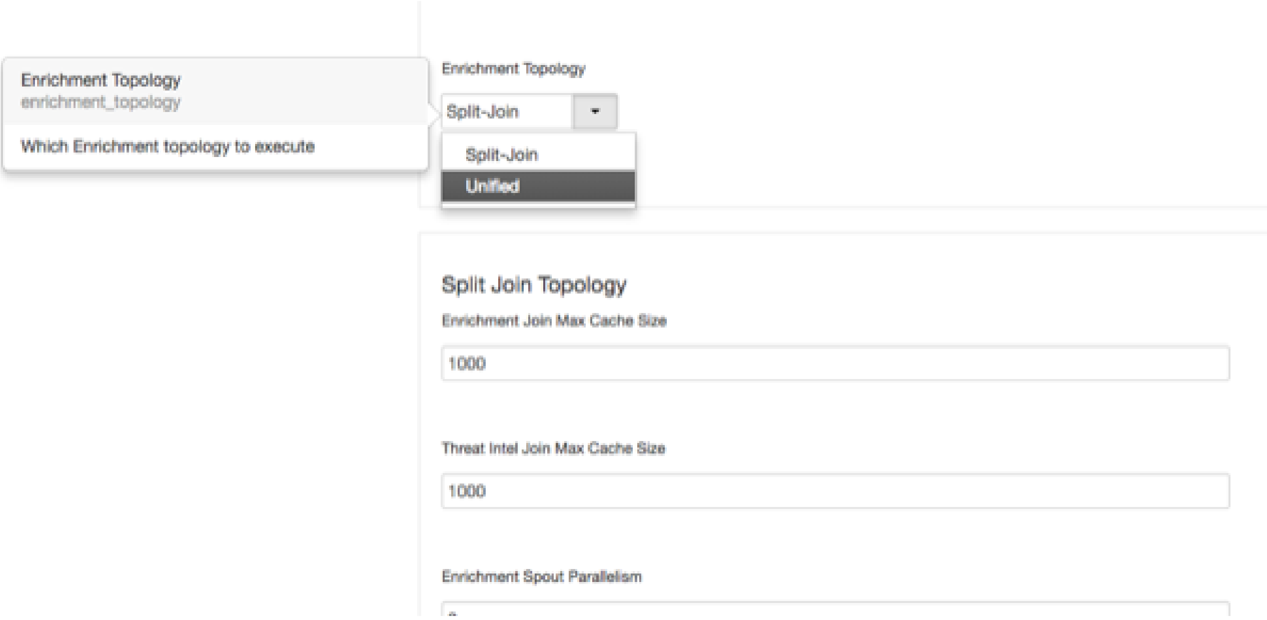

Enrichment Topology
enrichment_topology

Which Enrichment topology to execute

Enrichment Topology: Split-Join

Split Join Topology

Enrichment Join Max Cache Size: 1000

Threat Intel Join Max Cache Size: 1000

Enrichment Spout Parallelism

   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.

1.8. Upgrading to Elasticsearch 5.6.2

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

- Updating Elasticsearch Templates
- Updating Existing Indexes

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 1.8.1, “Type Mapping Changes” [5].

For more information, see Upgrade Elasticsearch.

1.8.1. 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"`
		`}`
		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

1.9. Third-Party Licenses

Global: Apache 2.0

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.

1.10. Known Issues

The HCP 1.5.0 release has the following known issues:

- **BUG-104383** - When you add a comment to an alert in the Alerts UI, the comment might not display immediately. To work around this issue, close and reopen the status tab for the alert.
- 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 Importing the Apache Zeppelin Notebook Manually for more information.

1.10.1. Known Differences Between HCP 1.5.0 and HCP 1.4.2

The following bugs identify known differences between HCP 1.5.0 and HCP 1.4.2.

Table 1.6. Known Differences Between HCP 1.5.0 and HCP 1.4.2

Feature	Description
METRON-1419	Create a SolrDao
METRON-1421	Solr Metaalerts
METRON-1423	Ambari work to handle Solr configuration
METRON-1424	Kerberos: Solr
METRON-1436	Manually Install Solr Cloud in Full Dev
METRON-1441	Create complementary Solr schemas for the main sensors
METRON-1448	Update SolrWriter to conform to new collection strategy

Feature	Description
METRON-1464	Convert schemas to be compatible with Solr 5.5.2
METRON-1482	Update REST to work with Solr
METRON-1503	Alerts are not getting populated in alerts UI when search engine is Solr
METRON-1526	Location field types cause DocValuesField appear more than once error
METRON-1540	Solr Integration tests should use actual schemas
METRON-1548	Remove hardcoded source:type from Alerts UI

1.10.2. Known Differences Between HCP 1.5.0 and Apache Metron 0.4.2

The following bugs identify known differences between HCP 1.5.0 and Apache Metron 0.4.2.

Table 1.7. Known Differences Between HCP 1.5.0 and Apache Metron 0.4.2

Feature	Description
METRON-1577	Solr searches don't include the index of the result
METRON-1548	Remove hardcoded source:type from Alerts UI
METRON-1564	Full dev kafka has offsets.topic.replication.factor set to 3 instead of 1
METRON-1552	Add gzip file validation check to the geo loader
METRON-1551	Profiler Should Not Use Java Serialization
METRON-1549	Add empty object test to WriterBoltIntegrationTest implementation
METRON-1541	Mvn clean results in git status having deleted files
METRON-1461	MIN MAX stellar function should take a stats or list object and return min/max
METRON-1184	EC2 Deployment - Updating control_path to accommodate for Linux
METRON-1530	Default proxy config settings in metron-contrib need to be updated
METRON-1421	Solr Metaalerts
METRON-1545	Upgrade Spring and Spring Boot
METRON-1543	Unable to Set Parser Output Topic in Sensor Config
METRON-1540	Solr Integration tests should use actual schemas
METRON-1526	Location field types cause DocValuesField appear more than once error
METRON-1539	Specialized RENAME field transformer
METRON-1520	Add caching for stellar field transformations
METRON-1529	CONFIG_GET Fails to Retrieve Latest Config When Run in Zeppelin REPL
METRON-1511	Unable to Serialize Profiler Configuration
METRON-1528	Fix missing file in metron.spec
METRON-1445	Update performance tuning guide with more explicit parameter instructions
METRON-1502	Upgrade Doxia plugin to 1.8
METRON-1527	Remove dead test file sitting in source folder

Feature	Description
METRON-1499	Enable Configuration of Unified Enrichment Topology via Ambari
METRON-1515	Errors loading stellar functions currently bomb the entire topology, they should be recoverable
METRON-1522	Fix the typo errors at profile debugger readme
METRON-1519	Indexing Error Topic Property Not Displayed in MPack
METRON-1347	Indexing Topology should fail tuples without a source.type
METRON-1503	Alerts are not getting populated in alerts UI when search engine is Solr
METRON-1521	JSONMapParser is no longer serializable
METRON-1516	Support for Ansible 2.5.0
METRON-1494	Profiler Emits Messages to Kafka When Not Needed
METRON-1510	Update Metron website to include info about github update subscription
METRON-1518	Build Failure When Using Profile HDP-2.5.0.0
METRON-1465	Support for Elasticsearch X-pack
METRON-1504	Enriching missing values does not match the semantics between the new enrichment topology and old
METRON-1505	Intermittent Profiler Integration Test Failure
METRON-1449	Set ZooKeeper URL for Stellar Running in Zeppelin Notebook
METRON-1462	Separate ES and Kibana from Metron Mpack
METRON-1501	Parser messages that fail to validate are dropped silently
METRON-1497	Rest endpoint '/api/v1/search/search' needs to handle null when elastic search response return null for getAggregations
METRON-1500	Enhance 'prepare-commit' to Support Feature Branches
METRON-1424	Kerberos: Solr
METRON-1491	The indexing topology restart logic is wrong
METRON-590	Enable Use of Event Time in Profiler
METRON-1483	Create a tool to monitor performance of the topologies
METRON-1487	Define Performance Benchmarks for Enrichment Topology
METRON-1493	Unhelpful Error Message When Assignment Expressions Fail
METRON-1397	Support for JSON Path and complex documents in JSONMapParser
METRON-1299	In MetronError tests, don't test for HostName if getHostName wouldn't work
METRON-1485	Upgrade vagrant for dev environments
METRON-1488	user_settings hbase table does not have acls set up for kerberos
METRON-1490	Better error message when user specifies an enrichment type that doesn't exist
METRON-1468	Add support for apache/metron-bro-plugin-kafka to prepare-commit
METRON-1471	Migrate shuffle connections to local or shuffle
METRON-1482	Update REST to work with Solr
METRON-1464	Convert schemas to be compatible with Solr 5.5.2

Feature	Description
METRON-1467	Replace guava caches in places where the keyspace might be large
METRON-1463	Adjust the groupings and shuffles in enrichment to be more efficient
METRON-1460	Create a complementary non-split-join enrichment topology
METRON-1470	Update jquery to version 3+
METRON-1450	Add REST endpoint docs for index topology split
METRON-1337	List of facets should not be hardcoded
METRON-1452	Rebase Dev Environment on Latest CentOS 6
METRON-1423	Ambari work to handle Solr configuration
METRON-1457	Move ASF links to main page in the Metron website
METRON-1394	Create Rest endpoint to add the ACL for current user to kafka topics
METRON-941	native PaloAlto parser corrupts message when having a comma in the payload
METRON-1386	Fix Metron Website Required Links
METRON-1455	Patch and Replace methods in the REST UpdateController return 400
METRON-1318	Update MacOS Instructions for AWS
METRON-1448	Update SolrWriter to conform to new collection strategy
METRON-1451	On Centos full dev, Metron Indexing shows up as stopped
METRON-1444	Add Ubuntu Repositories for Elasticsearch to the Mpack
METRON-1273	Website documentation link should point to the current site-book
METRON-1447	Heap Size Not Set Correctly by MPack for ES 5.x
METRON-1441	Create complementary Solr schemas for the main sensors
METRON-1446	Fix openjdk issue with Ubuntu
METRON-1442	Split rest end points for indexing topology into random access indexing and batch indexing
METRON-1443	Missing Critical MPack Install Instruction for Ubuntu
METRON-1436	Manually Install Solr Cloud in Full Dev
METRON-1438	STELLAR: Move shell functions to common from metron-management
METRON-1435	Management UI cannot save json objects in advanced config
METRON-1419	Create a SolrDao
METRON-1439	Turn off git pager in platform-info script
METRON-1091	STELLAR Shell: Stand Alone installation
METRON-1427	Add support for storm 1.1 and hdp 2.6
METRON-1391	Typos in Documentation/Examples within metron-management/README.md
METRON-1389	Zeppelin notebook import does not work with Ambari 2.6
METRON-1432	JDK Install Fails on Ubuntu Development Environment
METRON-1431	Add REGEXP_REPLACE function to Stellar
METRON-1410	Some more upgrade fallout... Can't restart Metron Indexing
METRON-1370	Create Full Dev Equivalent for Ubuntu

Feature	Description
METRON-1430	Isolate jackson from being used as arguments or returns from JSONUtils
METRON-1398	Exclude the basic-error-controller from being added to the swagger description
METRON-1392	Fix a test case to expect an Exception when replication factor more than number of brokers while creating topic
METRON-1413	Add Metron Commit Tool
METRON-1429	SearchIntegrationTest refactor
METRON-1426	SensorIndexingConfigControllerIntegrationTest fails intermittently
METRON-1417	Disable pcap-service by default in Monit
METRON-1400	Elasticsearch service check fails in Ambari
METRON-1428	Travis build failing from metron-config
METRON-1302	Split up Indexing Topology into batch and random access sections
METRON-1395	Documentation missing for Produce a message to a Kafka topic Rest API endpoint
METRON-1411	Fix sed command in Upgrading.md
METRON-1326	Metron deploy with Kerberos fails on Ambari 2.5 during ES service stop
METRON-1380	Create a typosquatting use-case
METRON-1230	As a stopgap prior to METRON-777, add more simplistic sideloading of custom Parsers
METRON-1378	Create a summarizer
METRON-1231	Separate Sensor name and topic in the Management UI
METRON-1382	Run Stellar in a Zeppelin Notebook
METRON-1396	Fix .gitignore files to not ignore themselves
METRON-1366	Add an entropy stellar function
METRON-1390	Swagger UI for "Web Security Config" Controller needs request method
METRON-1393	Fix bro Elasticsearch template
METRON-1379	Add an OBJECT_GET stellar function
METRON-939	Upgrade ElasticSearch and Kibana
METRON-1377	Stellar function to generate typosquatted domains
METRON-1385	Missing "properties" in index template causes ElasticsearchColumnMetadataDao.getColumnMetadata to fail
METRON-1388	update public web site to point at 0.4.2 new release
METRON-1362	Improve Metron Deployment README
METRON-1384	Increment master version number to 0.4.3 for on-going development
METRON-1381	Add Apache license to MD files and remove the Rat exclusion
METRON-1071	Create CONTRIBUTING.md
METRON-1373	RAT failure for metron-interface/metron-alerts
METRON-1351	Create Installable Packages for Ubuntu Trusty
METRON-1376	RC Check Script should have named parameters
METRON-1365	Allow PROFILE_GET to return a default value for a profile and entity that does not have a value written

Feature	Description
METRON-1348	Metron Service Checks Use Wrong Hostname
METRON-1350	Add reservoir sampling functions to Stellar
METRON-1374	Script the RC checking process
METRON-1372	Validate JIRA for Releases
METRON-1345	Update EC2 README for custom Ansible
METRON-1349	Full Dev Builds Metron Twice
METRON-1343	Swagger UI for User Controller needs request method
METRON-1306	When index template install fails, we should fail the install
METRON-1341	Projection FieldTransformation (simonellistonball