Configure Kibana in DNA Center for Log Visualization

Contents

Introduction

This document describes how to use Kibana in order to search specific logs among different Cisco DNA Center services.

Prerequisites

Requirements

You must have access Cisco DNA Center through GUI with ADMIN ROLE also, you must be familiar with the names and use of Cisco DNA Center services.

Components Used

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Kibana is an open source data visualization plugin for Elasticsearch. It provides visualization capabilities on top of the content indexed on an Elasticsearch cluster that are available in Cisco DNA Center.

You can access Kibana in two ways:

- **https://<Cisco DNA Center ip>/kibana**
- **Main Menu > System > System 360 -> Cluster Tools -> Log Explorer**

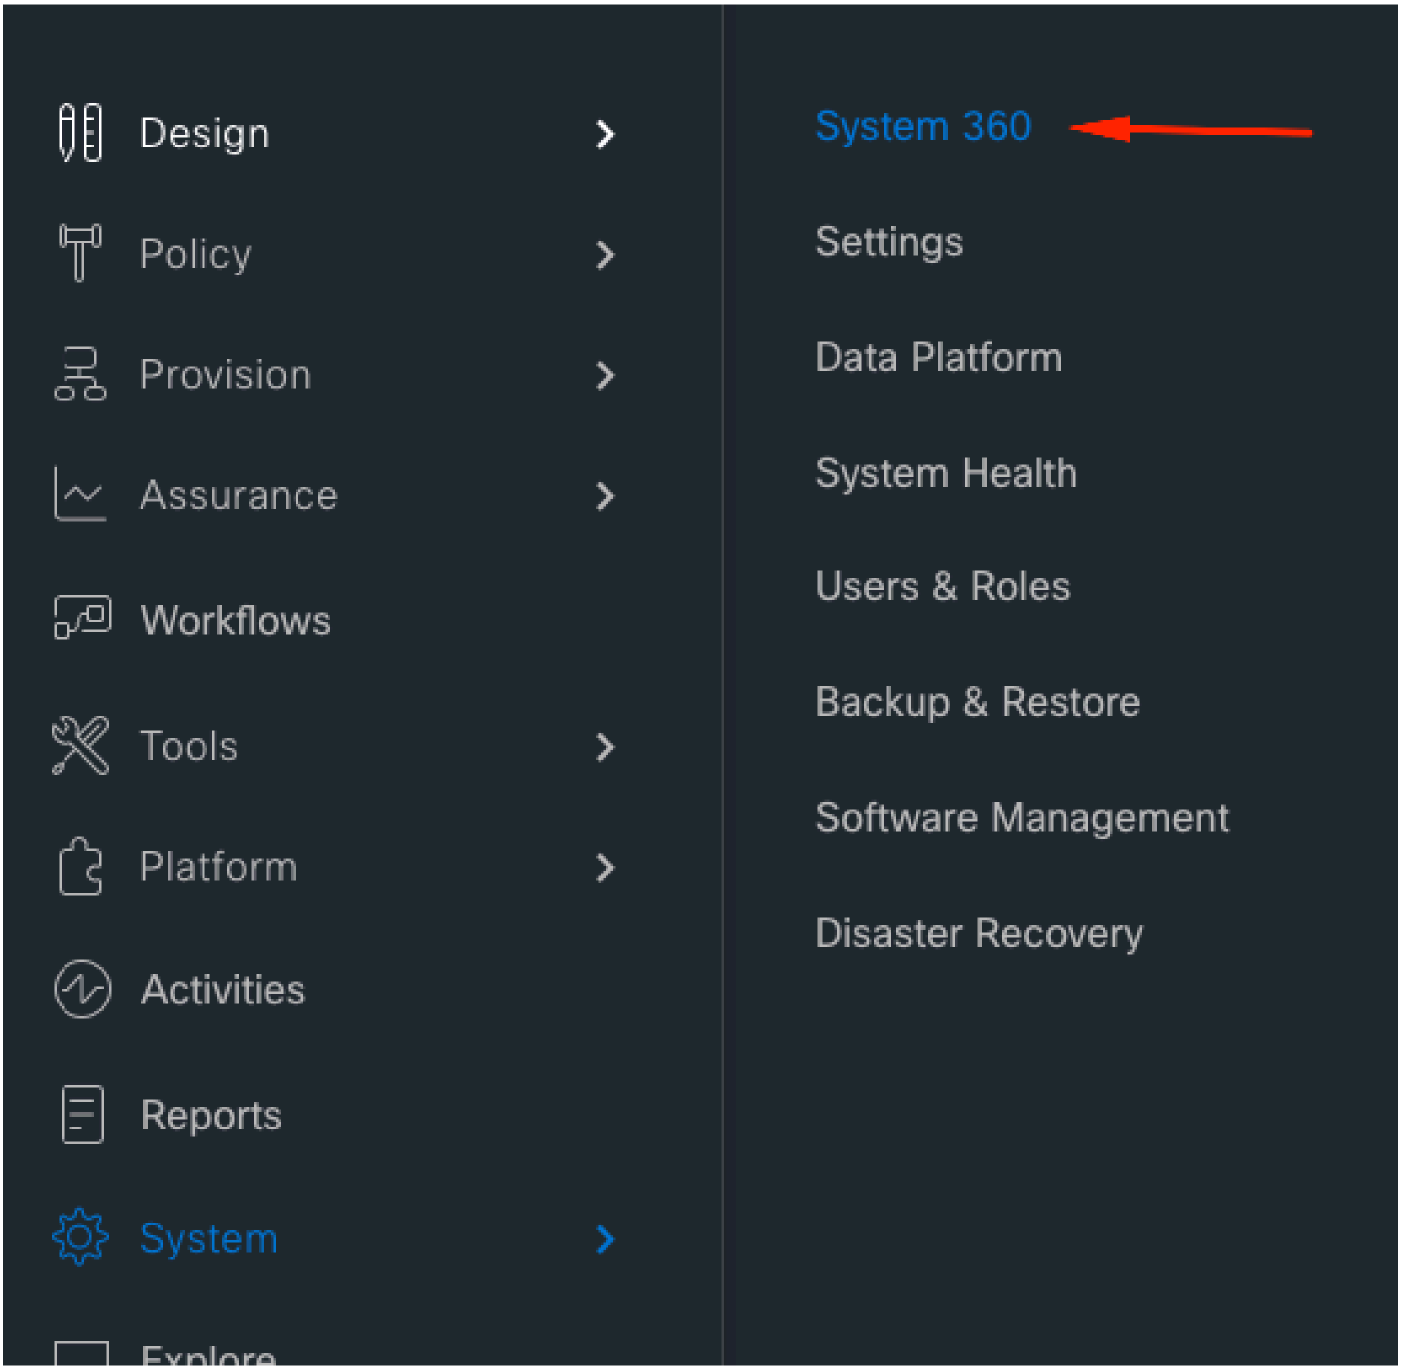

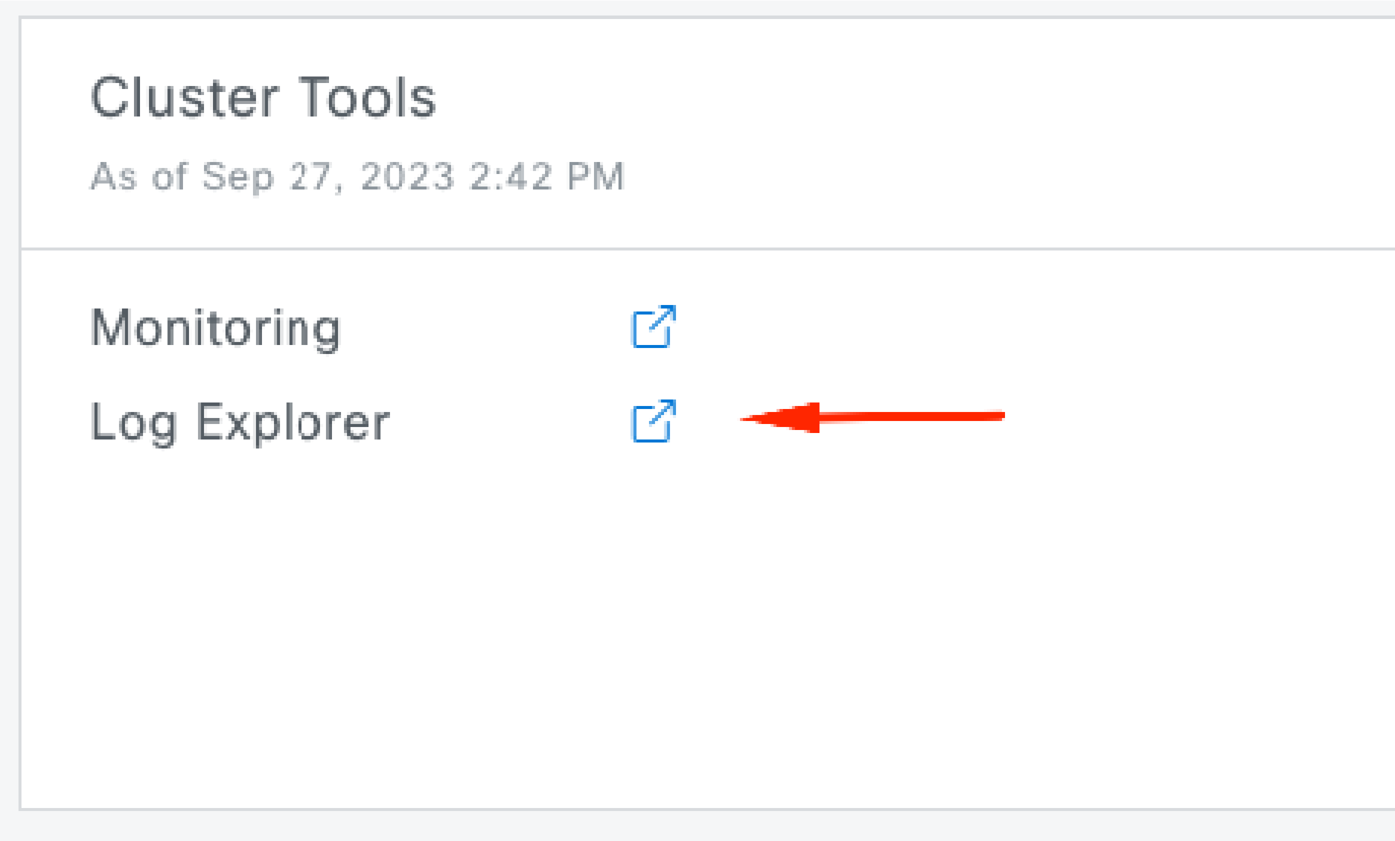
Cluster Tools

As of Sep 27, 2023 2:42 PM

Default Kibana web page

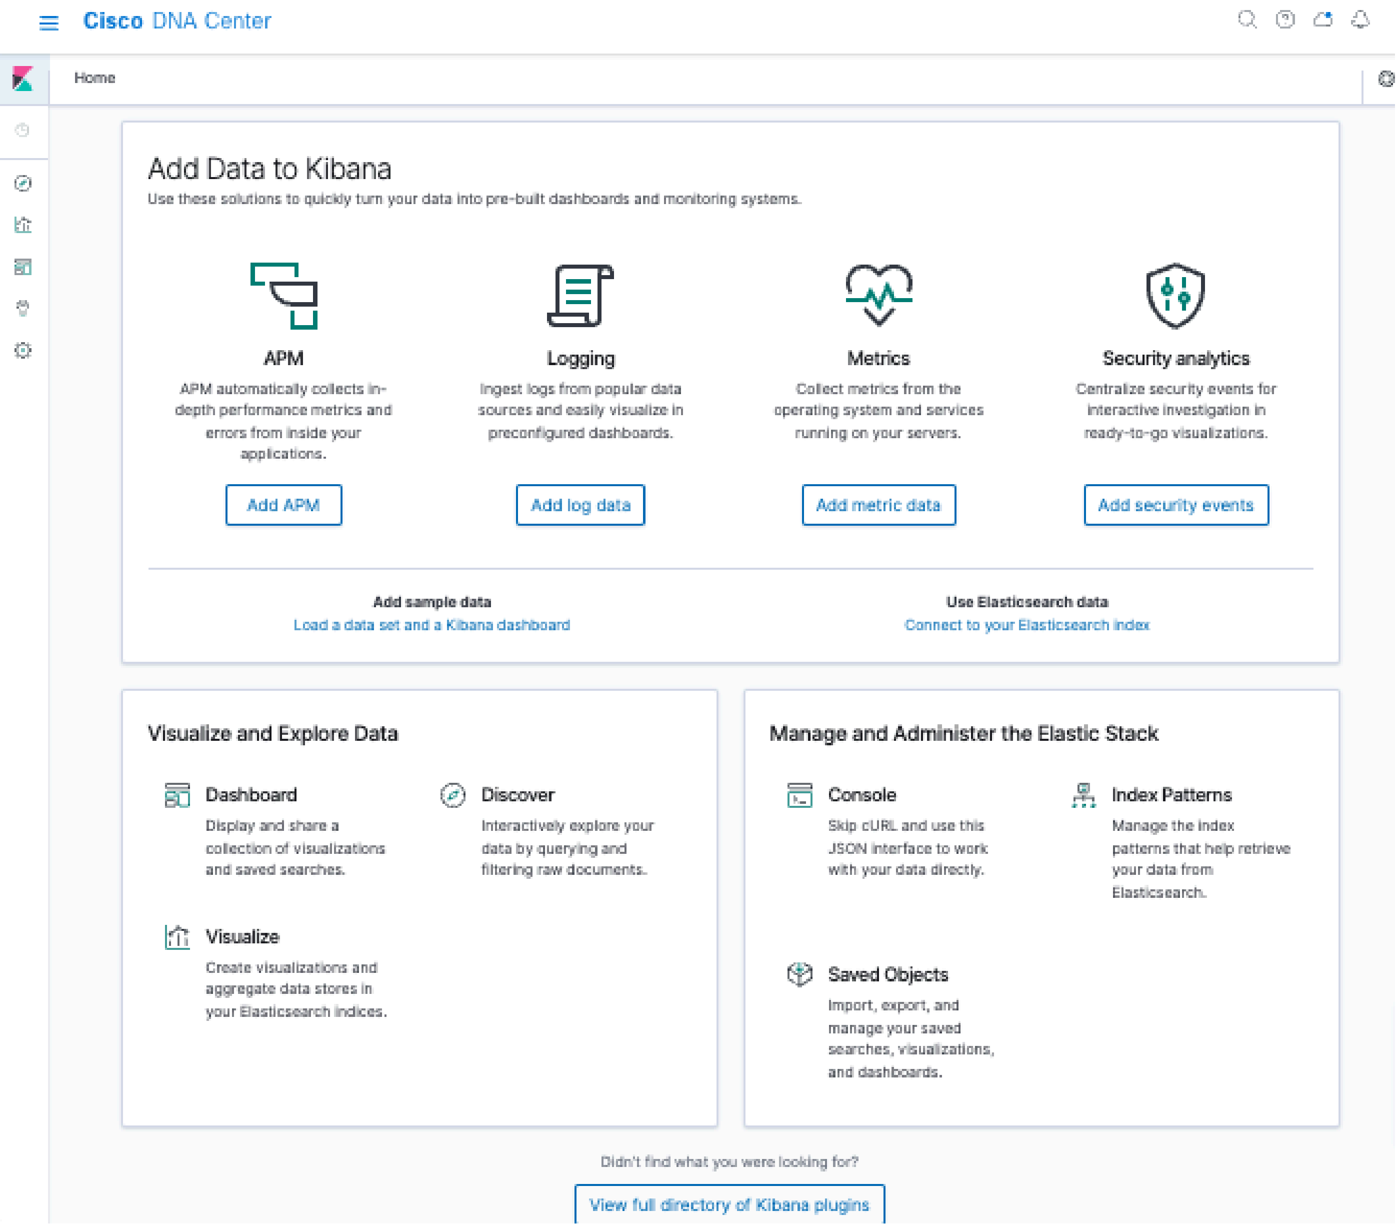

Add Data to Kibana

Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.

APM

APM automatically collects in-depth performance metrics and errors from inside your applications.

Add APM

Logging

Ingest logs from popular data sources and easily visualize in preconfigured dashboards.

Add log data

Metrics

Collect metrics from the operating system and services running on your servers.

Add metric data

Security analytics

Centralize security events for interactive investigation in ready-to-go visualizations.

Add security events

Add sample data

Load a data set and a Kibana dashboard

Use Elasticsearch data

Connect to your Elasticsearch index

Visualize and Explore Data

Dashboard

Display and share a collection of visualizations and saved searches.

Discover

Interactively explore your data by querying and filtering raw documents.

Visualize

Create visualizations and aggregate data stores in your Elasticsearch indices.

Manage and Administer the Elastic Stack

Console

Skip cURL and use this JSON interface to work with your data directly.

Index Patterns

Manage the index patterns that help retrieve your data from Elasticsearch.

Saved Objects

Import, export, and manage your saved searches, visualizations, and dashboards.

Didn't find what you were looking for?

View full directory of Kibana plugins

Configure Kibana for log visualization

Navigate to the left bar menu and click on Discover:

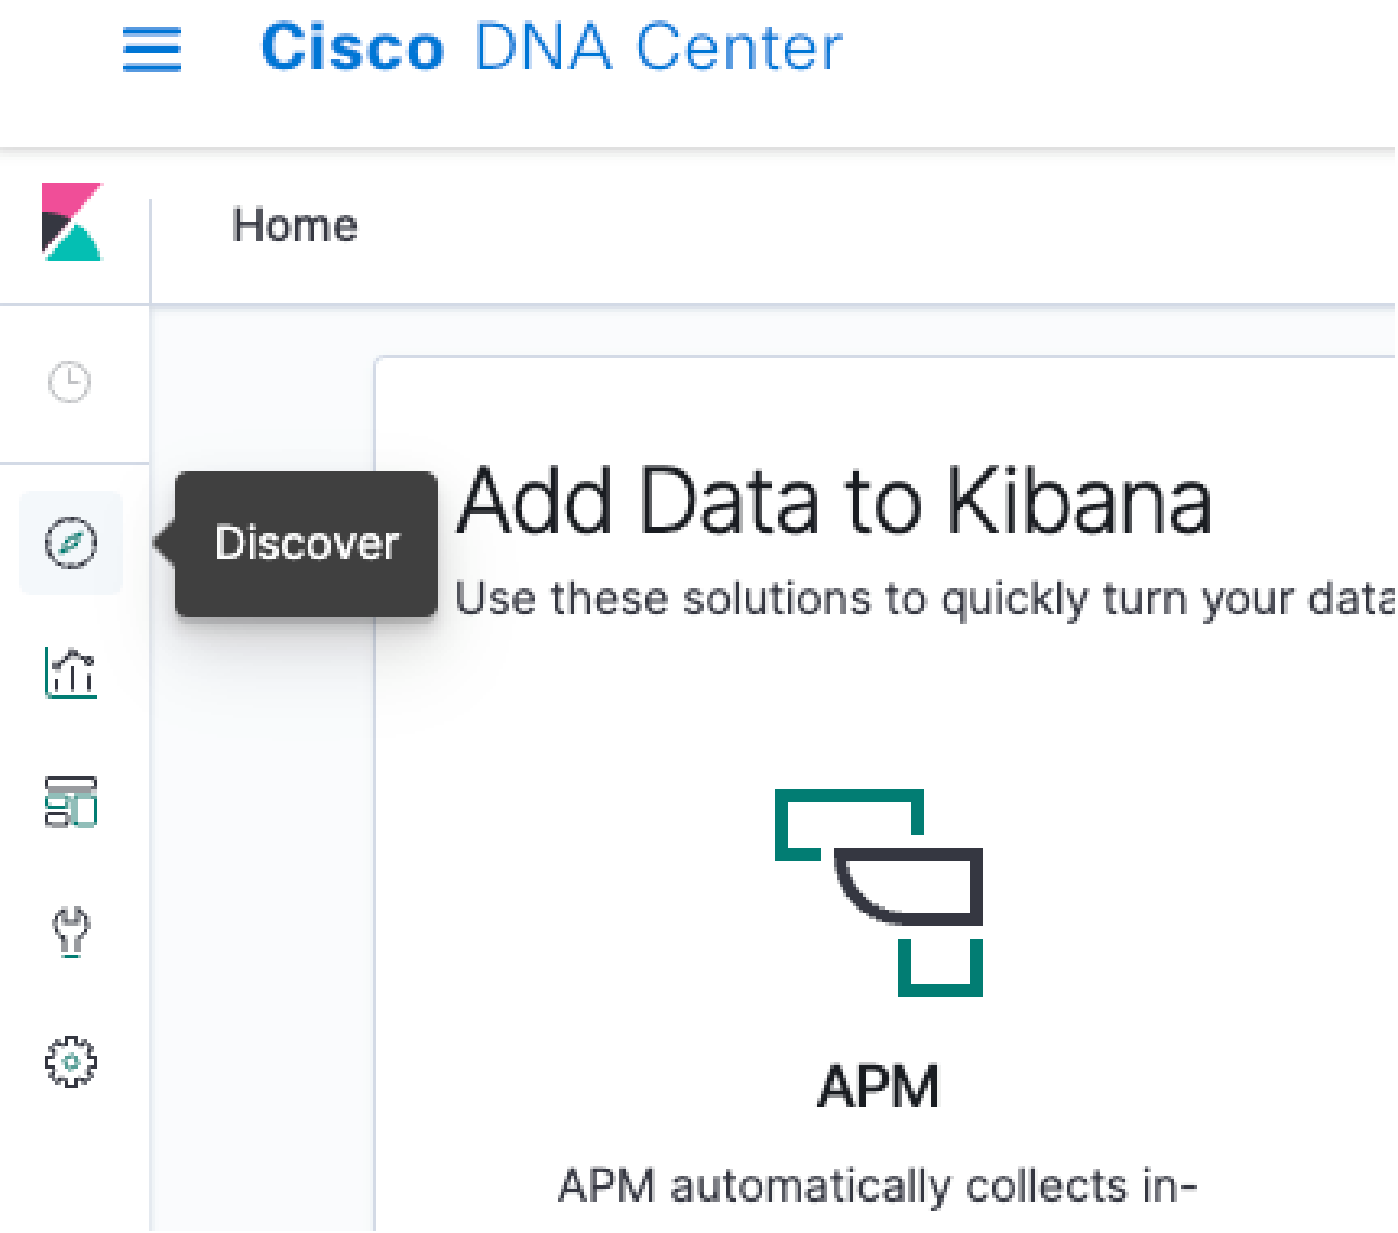

Cisco DNA Center
Home
Discover
Add Data to Kibana
Use these solutions to quickly turn your data
APM
APM automatically collects in-

Kibana has several fields, which are highlighted in the next image:

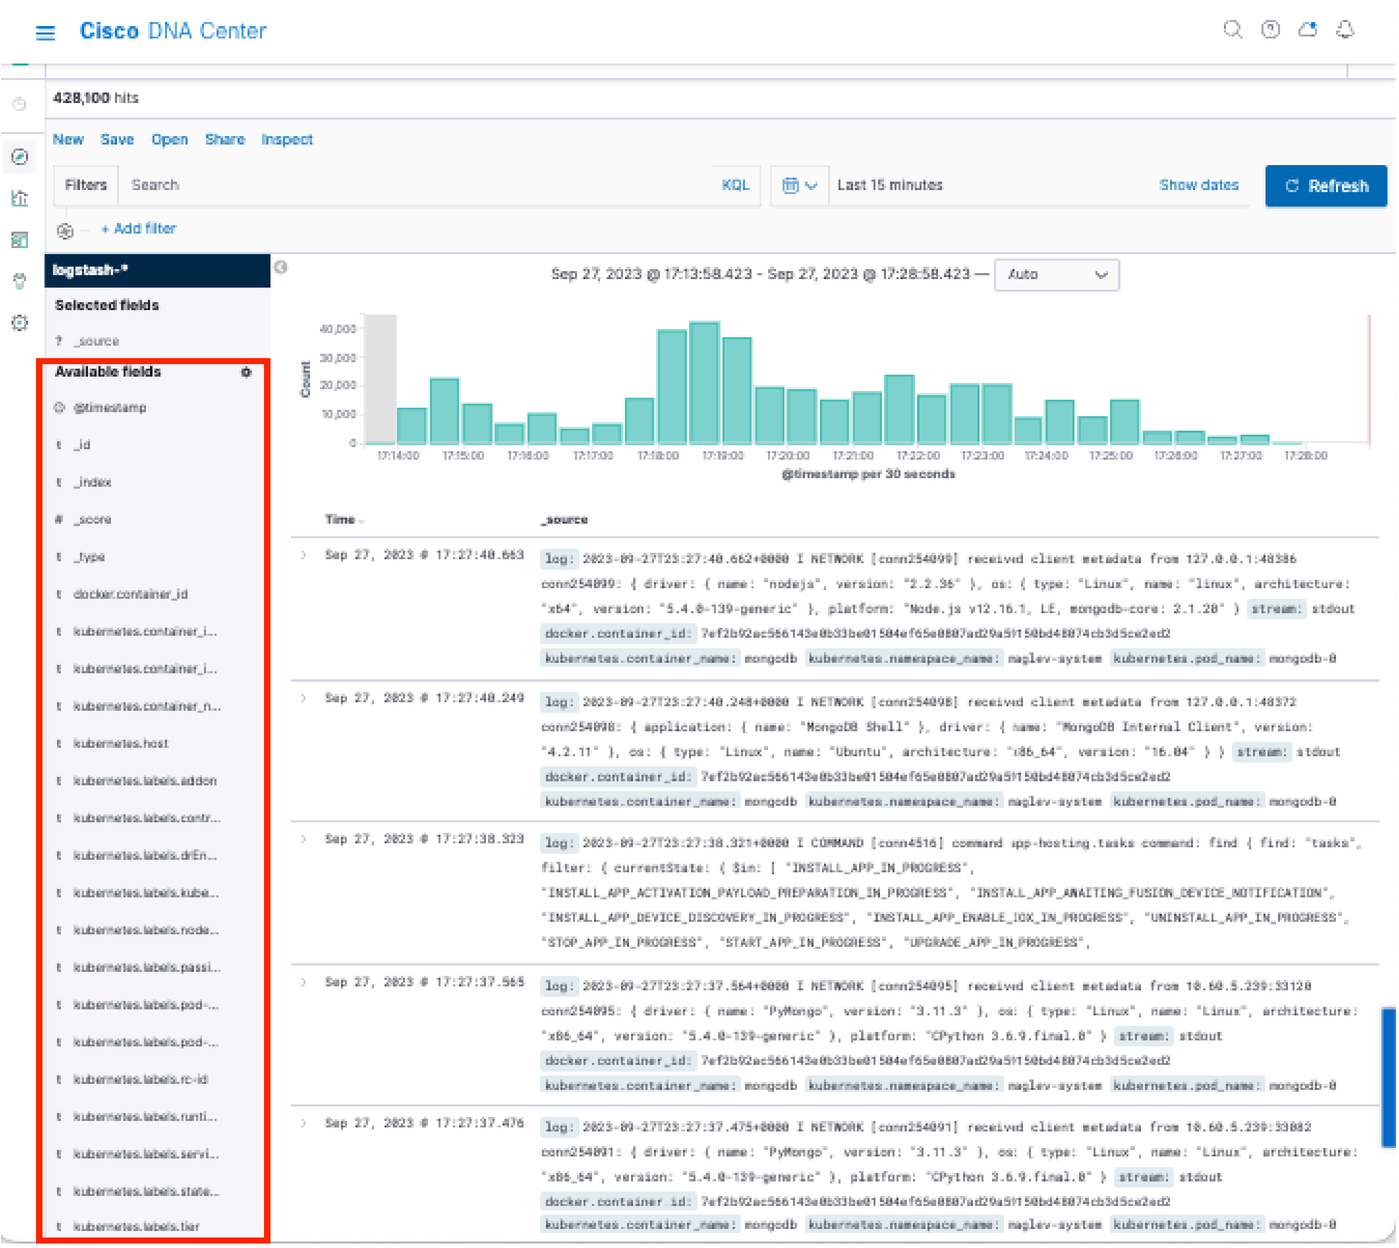
Cisco DNA Center
428,100 hits
New
Save
Open
Share
Inspect
Filters
Search
KQL
Last 15 minutes
Show dates
Refresh
+ Add filter
logstash-*
Selected fields
? _source
Available fields
@timestamp
t _id
t _index
# _score
t _type
t docker.container_id
t kubernetes.container_i...
t kubernetes.container_i...
t kubernetes.container_n...
t kubernetes.host
t kubernetes.labels.addon
t kubernetes.labels.contr...
t kubernetes.labels.drEn...
t kubernetes.labels.kube...
t kubernetes.labels.node...
t kubernetes.labels.passi...
t kubernetes.labels.pod-...
t kubernetes.labels.pod-...
t kubernetes.labels.rc-id
t kubernetes.labels.runti...
t kubernetes.labels.servi...
t kubernetes.labels.state...
t kubernetes.labels.tier
Sep 27, 2023 @ 17:13:58.423 - Sep 27, 2023 @ 17:28:58.423 —
Auto
Count
40,000
30,000
20,000
10,000
0
17:14:00
17:15:00
17:16:00
17:17:00
17:18:00
17:19:00
17:20:00
17:21:00
17:22:00
17:23:00
17:24:00
17:25:00
17:26:00
17:27:00
17:28:00
@timestamp per 30 seconds
Time
_source

```
Sep 27, 2023 @ 17:27:40.663   log: 2023-09-27T23:27:40.662+0000 I NETWORK [conn254099] received client metadata from 127.0.0.1:48386
                              conn254099: { driver: { name: "nodejs", version: "2.2.36" }, os: { type: "Linux", name: "linux", architecture:
                              "x64", version: "5.4.0-139-generic" }, platform: "Node.js v12.16.1, LE, mongodb-core: 2.1.20" } stream: stdout
                              docker.container_id: 7ef2b92ac566143e0b33be01504ef65e0807ad29a5f150bd48074cb3d5ce2ed2
                              kubernetes.container_name: mongodb kubernetes.namespace_name: maglev-system kubernetes.pod_name: mongodb-0

Sep 27, 2023 @ 17:27:40.249   log: 2023-09-27T23:27:40.248+0000 I NETWORK [conn254098] received client metadata from 127.0.0.1:48372
                              conn254098: { application: { name: "MongoDB Shell" }, driver: { name: "MongoDB Internal Client", version:
                              "4.2.11" }, os: { type: "Linux", name: "Ubuntu", architecture: "x86_64", version: "16.04" } } stream: stdout
                              docker.container_id: 7ef2b92ac566143e0b33be01504ef65e0807ad29a5f150bd48074cb3d5ce2ed2
                              kubernetes.container_name: mongodb kubernetes.namespace_name: maglev-system kubernetes.pod_name: mongodb-0

Sep 27, 2023 @ 17:27:38.323   log: 2023-09-27T23:27:38.321+0000 I COMMAND [conn4516] command app-hosting.tasks command: find { find: "tasks",
                              filter: { currentState: { $in: [ "INSTALL_APP_IN_PROGRESS",
                              "INSTALL_APP_ACTIVATION_PAYLOAD_PREPARATION_IN_PROGRESS", "INSTALL_APP_AWAITING_FUSION_DEVICE_NOTIFICATION",
                              "INSTALL_APP_DEVICE_DISCOVERY_IN_PROGRESS", "INSTALL_APP_ENABLE_IOX_IN_PROGRESS", "UNINSTALL_APP_IN_PROGRESS",
                              "STOP_APP_IN_PROGRESS", "START_APP_IN_PROGRESS", "UPGRADE_APP_IN_PROGRESS",

Sep 27, 2023 @ 17:27:37.565   log: 2023-09-27T23:27:37.564+0000 I NETWORK [conn254095] received client metadata from 10.60.5.239:33120
                              conn254095: { driver: { name: "PyMongo", version: "3.11.3" }, os: { type: "Linux", name: "Linux", architecture:
                              "x86_64", version: "5.4.0-139-generic" }, platform: "CPython 3.6.9.final.0" } stream: stdout
                              docker.container_id: 7ef2b92ac566143e0b33be01504ef65e0807ad29a5f150bd48074cb3d5ce2ed2
                              kubernetes.container_name: mongodb kubernetes.namespace_name: maglev-system kubernetes.pod_name: mongodb-0

Sep 27, 2023 @ 17:27:37.476   log: 2023-09-27T23:27:37.475+0000 I NETWORK [conn254091] received client metadata from 10.60.5.239:33082
                              conn254091: { driver: { name: "PyMongo", version: "3.11.3" }, os: { type: "Linux", name: "Linux", architecture:
                              "x86_64", version: "5.4.0-139-generic" }, platform: "CPython 3.6.9.final.0" } stream: stdout
                              docker.container_id: 7ef2b92ac566143e0b33be01504ef65e0807ad29a5f150bd48074cb3d5ce2ed2
                              kubernetes.container_name: mongodb kubernetes.namespace_name: maglev-system kubernetes.pod_name: mongodb-0
```

Add fields in Kibana

Navigate to Filters > Available fields

The fields that you must need to add for logs visualization are:

- **Kubernetes.labels.serviceName** - Service that displays the specific log
- **Log** - Raw content of the log

Click on the add button

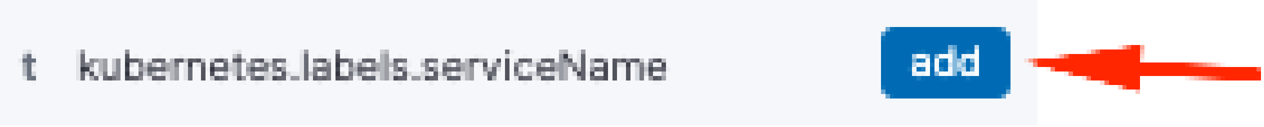

Ensure that you have the next configuration:

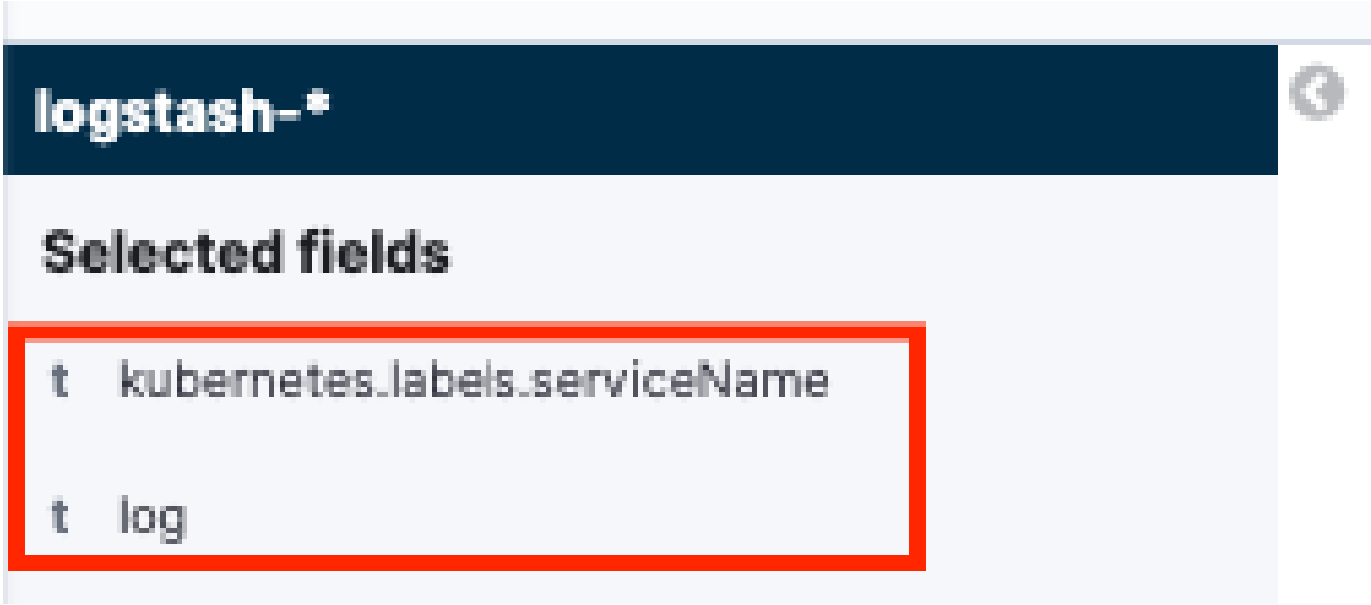

logstash-*

Selected fields

t kubernetes.labels.serviceName

t log

Note: Time field is added by default.

Add and edit filters in Kibana

In order to add a filter, perform the next activity:

- Click in **Add filter**
- **Field** select: **Kubernetes.labels.serviceName**
- **Operator** select: **is**
- Value: select the service of your interest
- Click in **Save** button

Take a look the next example where the service selected is apic-em-inventory-manager-service:

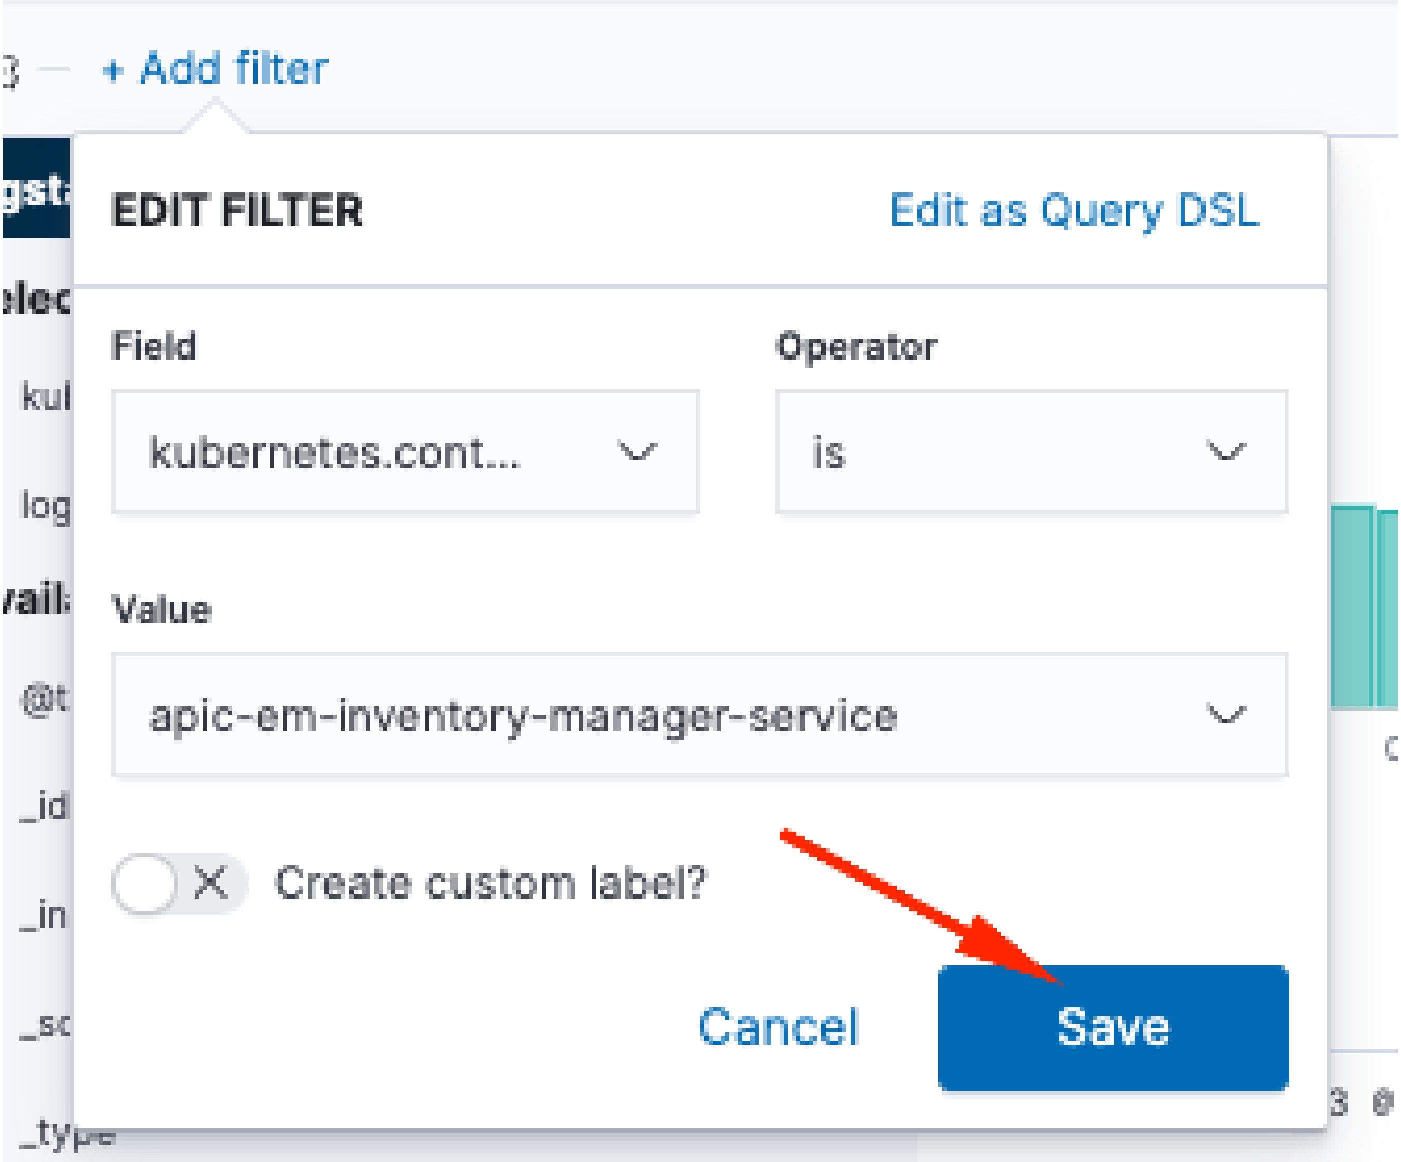

EDIT FILTER

Field: kubernetes.cont...

Operator: is

Value: apic-em-inventory-manager-service

You can add more filters as you need.

The next example, a new filter was added where the Field:log, operator:is and Value: error:

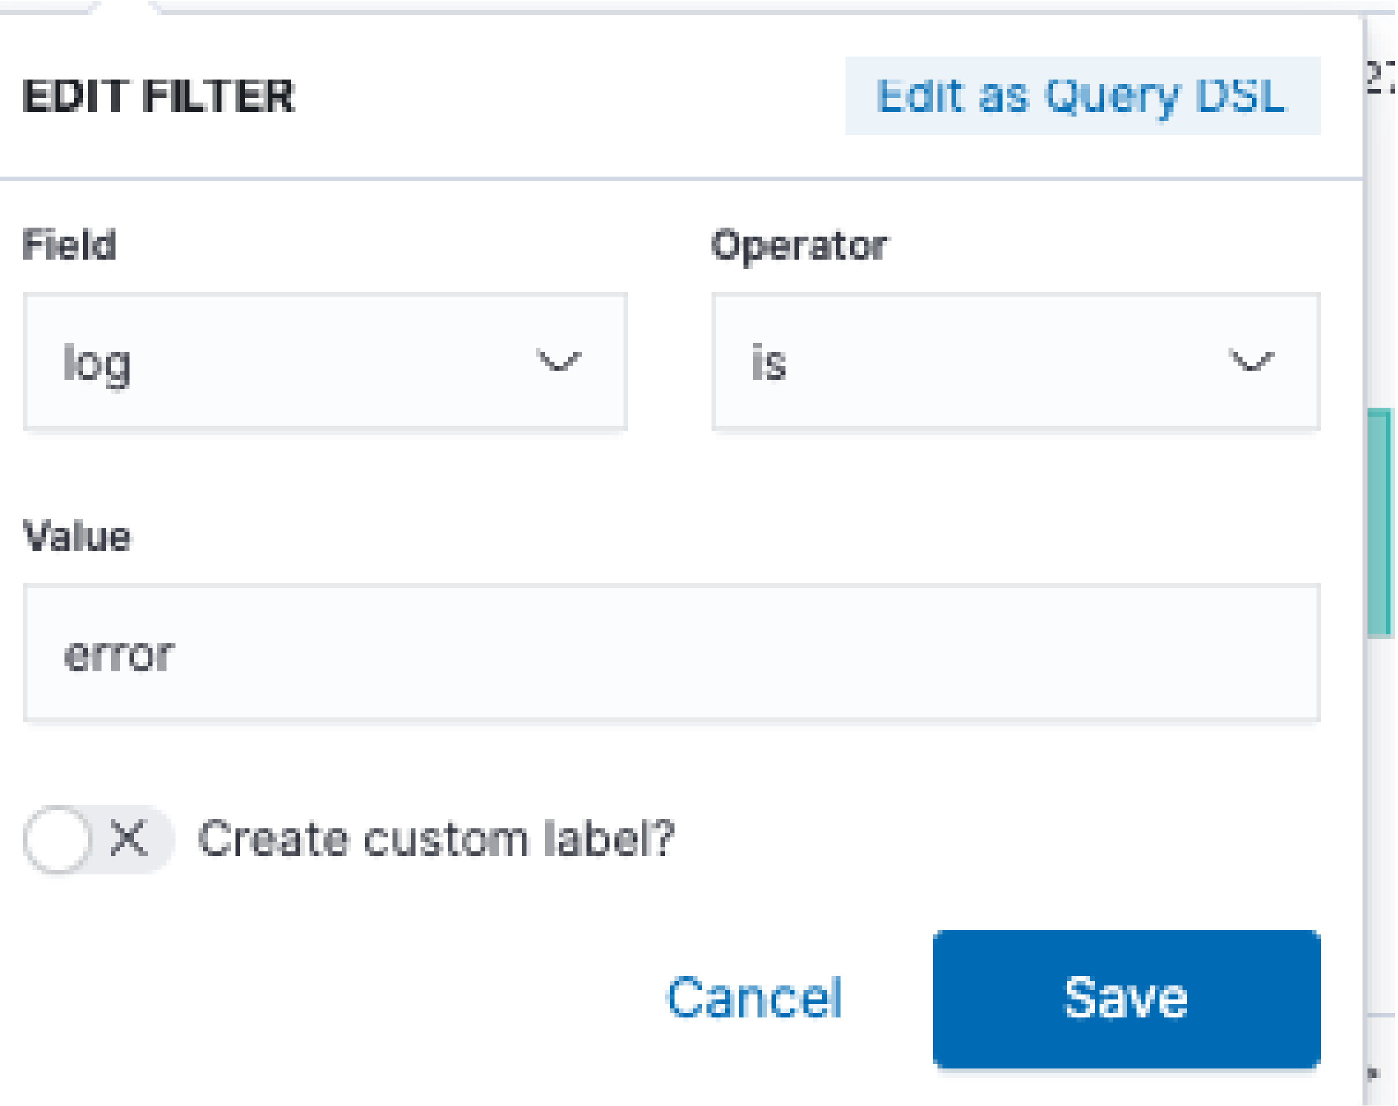

Field: log

Operator: is

Value: error

Create custom label?

Get logs from a specific date

You can add a time element to your search criteria.

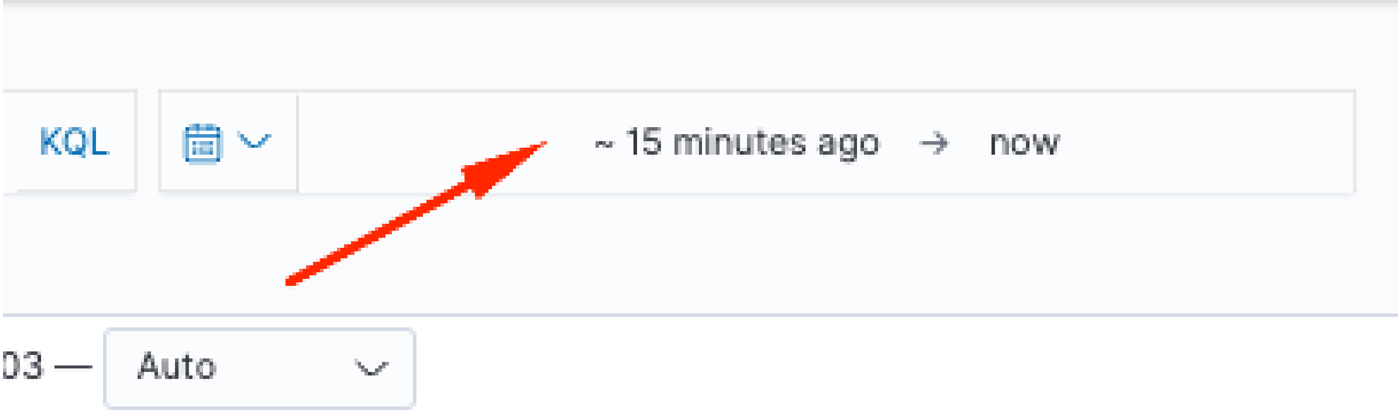

Use one of the next options from the Time Range field:

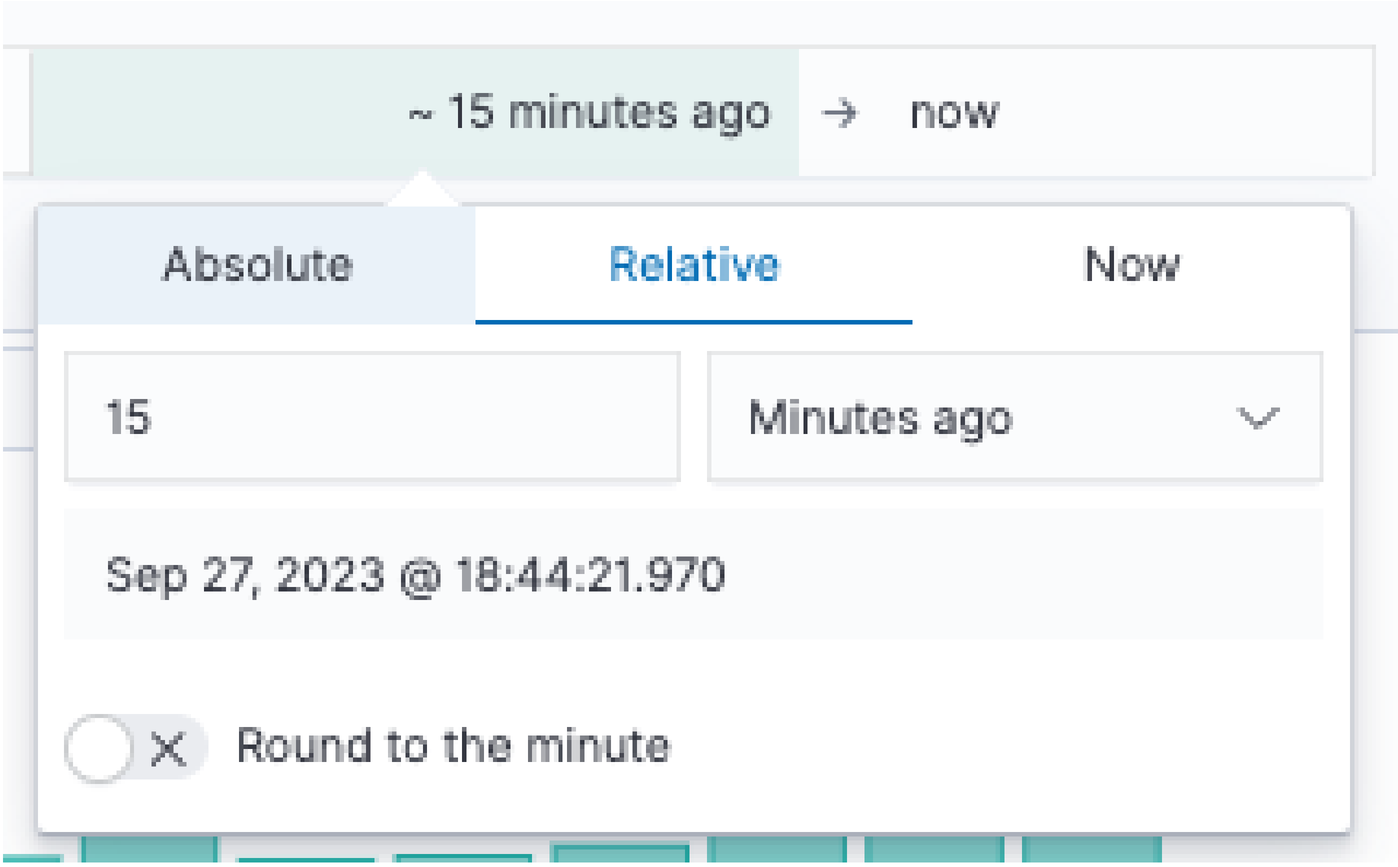

- **Absolute** - From a specific date to another specific date.
- **Relative** - From the last X minutes, hours, days, or weeks to a specific date.
- **Now** - Setting the time to "now" means that on every refresh this time is going to be set to the time of the refresh.

Use cases with Lucene

Lucene is a high-performance, full-featured text search engine library. It is a technology suitable for nearly any application that requires full-text search.

Navigate to search bar and disable KQL in order to enable Lucene:

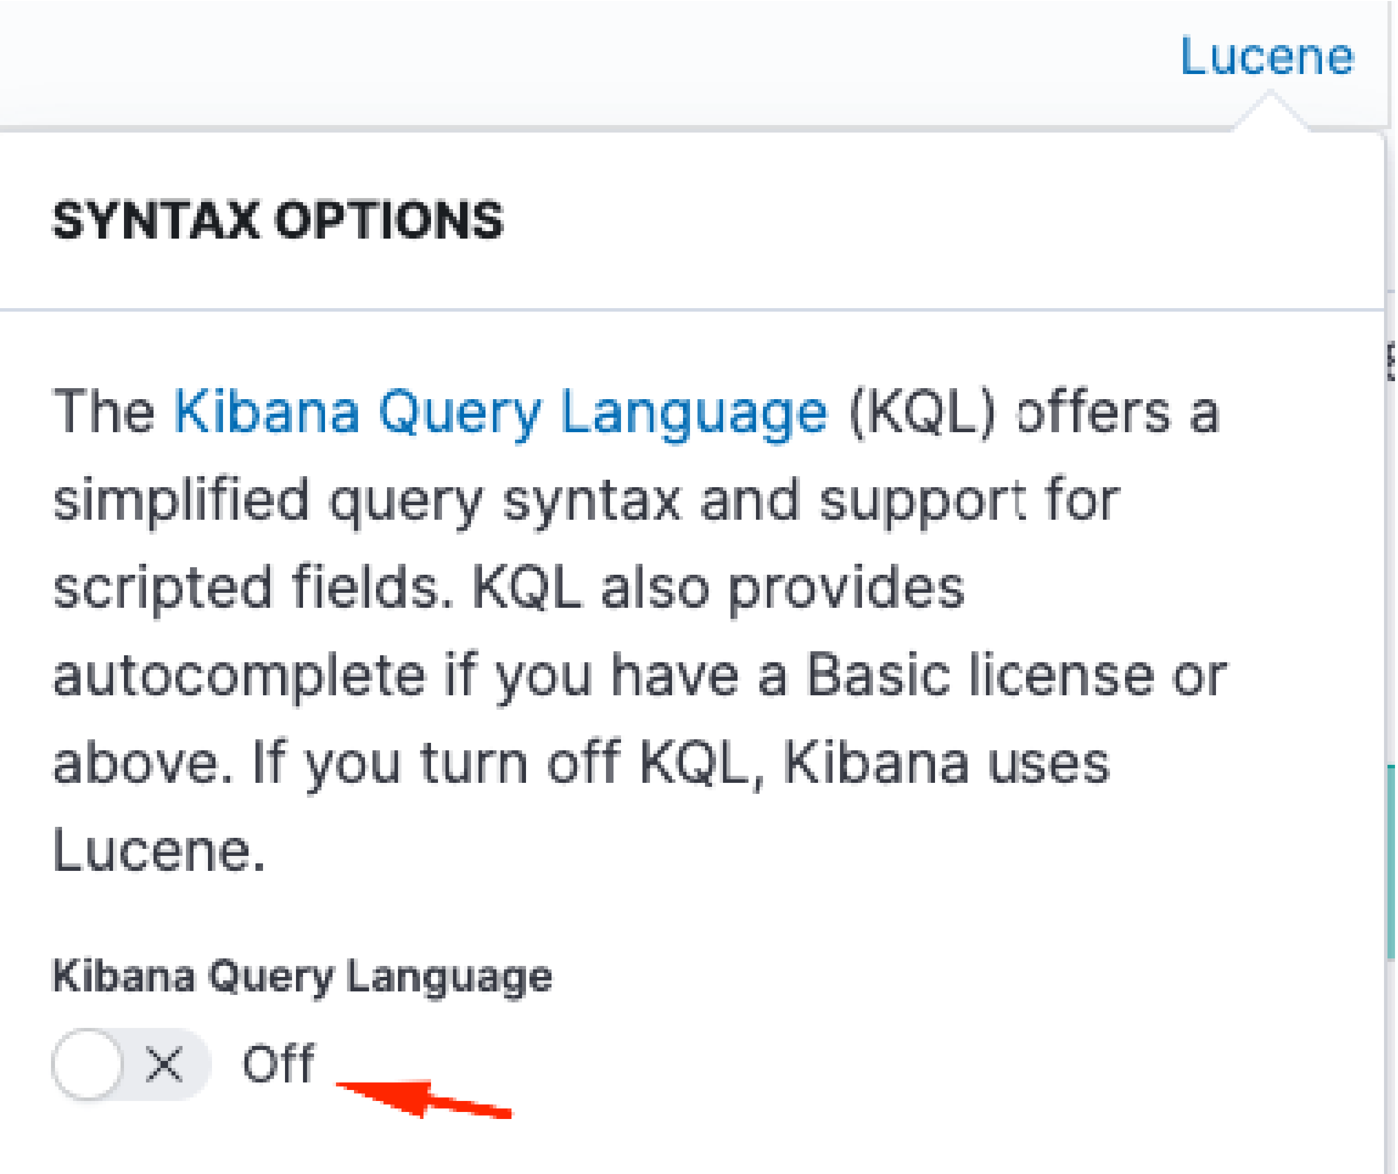

SYNTAX OPTIONS

The Kibana Query Language (KQL) offers a simplified query syntax and support for scripted fields. KQL also provides autocomplete if you have a Basic license or above. If you turn off KQL, Kibana uses Lucene.

Kibana Query Language

○ ✕ Off

Get logs for a specific service

Type the next query into the filter bar and press **Refresh** button

```
kubernetes.labels.serviceName:<service-name>
```

Take a look the next example with task-service:

```
kubernetes.labels.serviceName:task-service
```

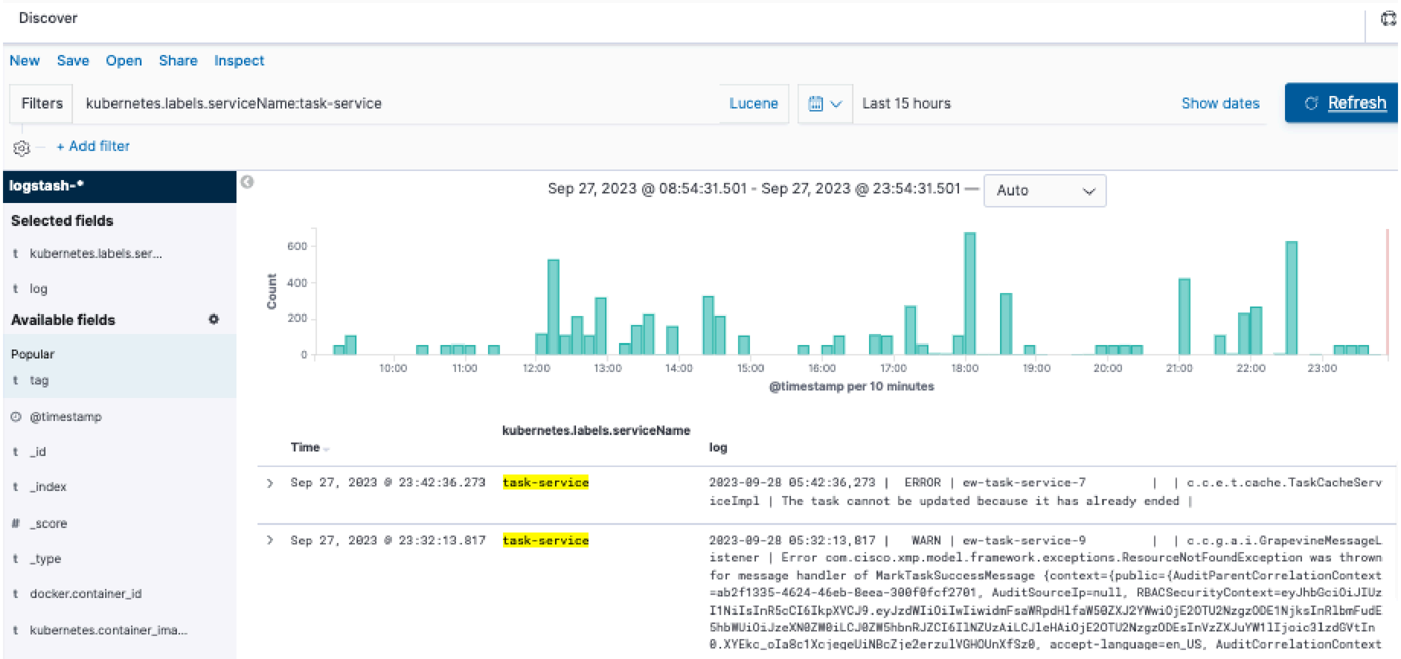

```
Time                        kubernetes.labels.serviceName  log

Sep 27, 2023 @ 23:42:36.273 task-service                   2023-09-28 05:42:36,273 |  ERROR | ew-task-service-7         |  | c.c.e.t.cache.TaskCacheServ
                                                           iceImpl | The task cannot be updated because it has already ended |

Sep 27, 2023 @ 23:32:13.817 task-service                   2023-09-28 05:32:13,817 |   WARN | ew-task-service-9         |  | c.c.g.a.i.GrapevineMessageL
                                                           istener | Error com.cisco.xmp.model.framework.exceptions.ResourceNotFoundException was thrown
                                                           for message handler of MarkTaskSuccessMessage {context={public={AuditParentCorrelationContext
                                                           =ab2f1335-4624-46eb-8eea-300f0fcf2701, AuditSourceIp=null, RBACSecurityContext=eyJhbGciOiJIUz
                                                           I1NiIsInR5cCI6IkpXVCJ9.eyJzdWIiOiIwIiwidmFsaWRpdHlfaW50ZXJ2YWwiOjE2OTU2NzgzODE1NjksInRlbmFudE
                                                           5hbWUiOiJzeXN0ZW0iLCJ0ZW5hbnRJZCI6I1NZUzAiLCJleHAiOjE2OTU2NzgzODEsInVzZXJuYW1lIjoic3lzdGVtIn
                                                           0.XYEkc_oIa8c1XcjeqeUiNBcZje2erzulVGHOUnXfSz0, accept-language=en_US, AuditCorrelationContext
```

Get logs that contain a specific word

Tip: The most common log entries that indicate problems contain " Error " , " Failed " and " Exception ", however, you can modify the string to be any other common string that can guide you in your troubleshooting.

Type the next query into the filter bar and press **Refresh** button

```
log:error
```

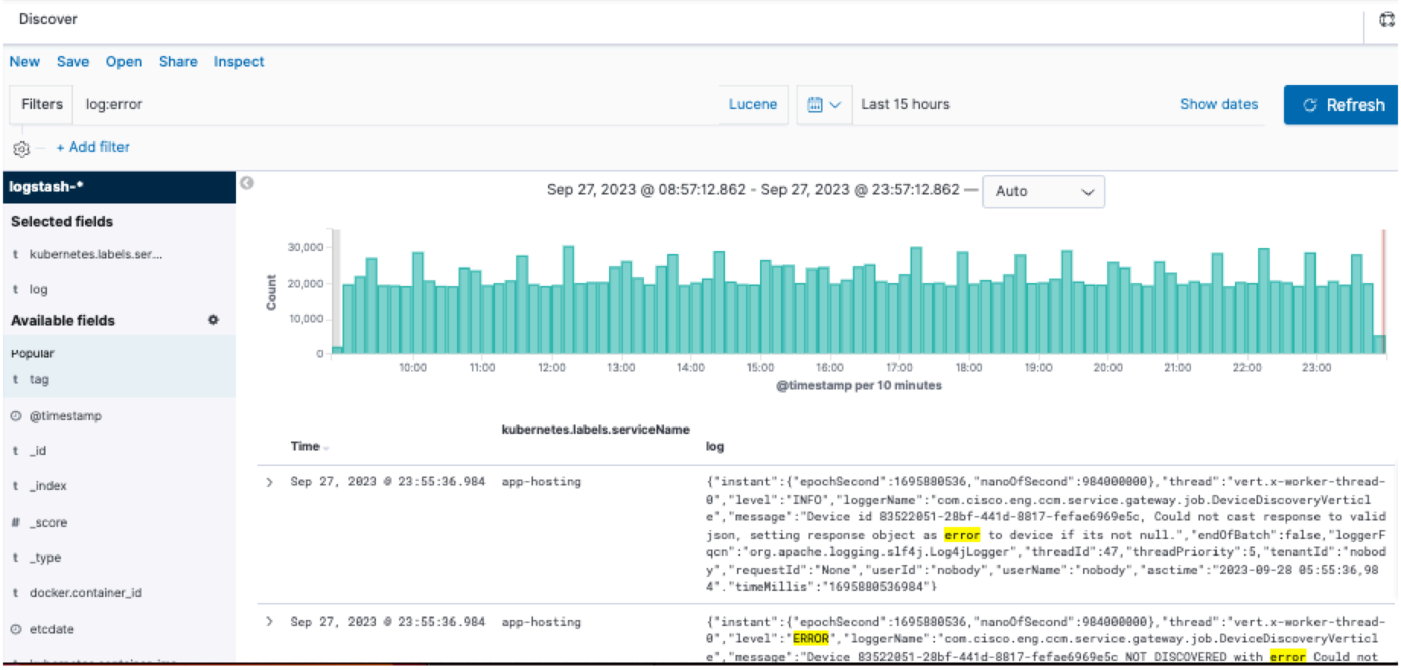

```
Time                        kubernetes.labels.serviceName  log

Sep 27, 2023 @ 23:55:36.984 app-hosting                    {"instant":{"epochSecond":1695880536,"nanoOfSecond":984000000},"thread":"vert.x-worker-thread-
                                                           0","level":"INFO","loggerName":"com.cisco.eng.ccm.service.gateway.job.DeviceDiscoveryverticl
                                                           e","message":"Device id 83522051-28bf-441d-8817-fefae6969e5c, Could not cast response to valid
                                                           json, setting response object as error to device if its not null.","endOfBatch":false,"loggerF
                                                           qcn":"org.apache.logging.slf4j.Log4jLogger","threadId":47,"threadPriority":5,"tenantId":"nobod
                                                           y","requestId":"None","userId":"nobody","userName":"nobody","asctime":"2023-09-28 05:55:36,98
                                                           4","timeMillis":"1695880536984"}

Sep 27, 2023 @ 23:55:36.984 app-hosting                    {"instant":{"epochSecond":1695880536,"nanoOfSecond":984000000},"thread":"vert.x-worker-thread-
                                                           0","level":"ERROR","loggerName":"com.cisco.eng.ccm.service.gateway.job.DeviceDiscoveryVerticl
                                                           e","message":"Device 83522051-28bf-441d-8817-fefae6969e5c NOT DISCOVERED with error Could not
```

Mix and match your search

You can search entries that match a combination of strings by using AND (or &&) between the strings.

<#root>

```
log:error
```

AND

```
 kubernetes.labels.serviceName:onboarding-service
```

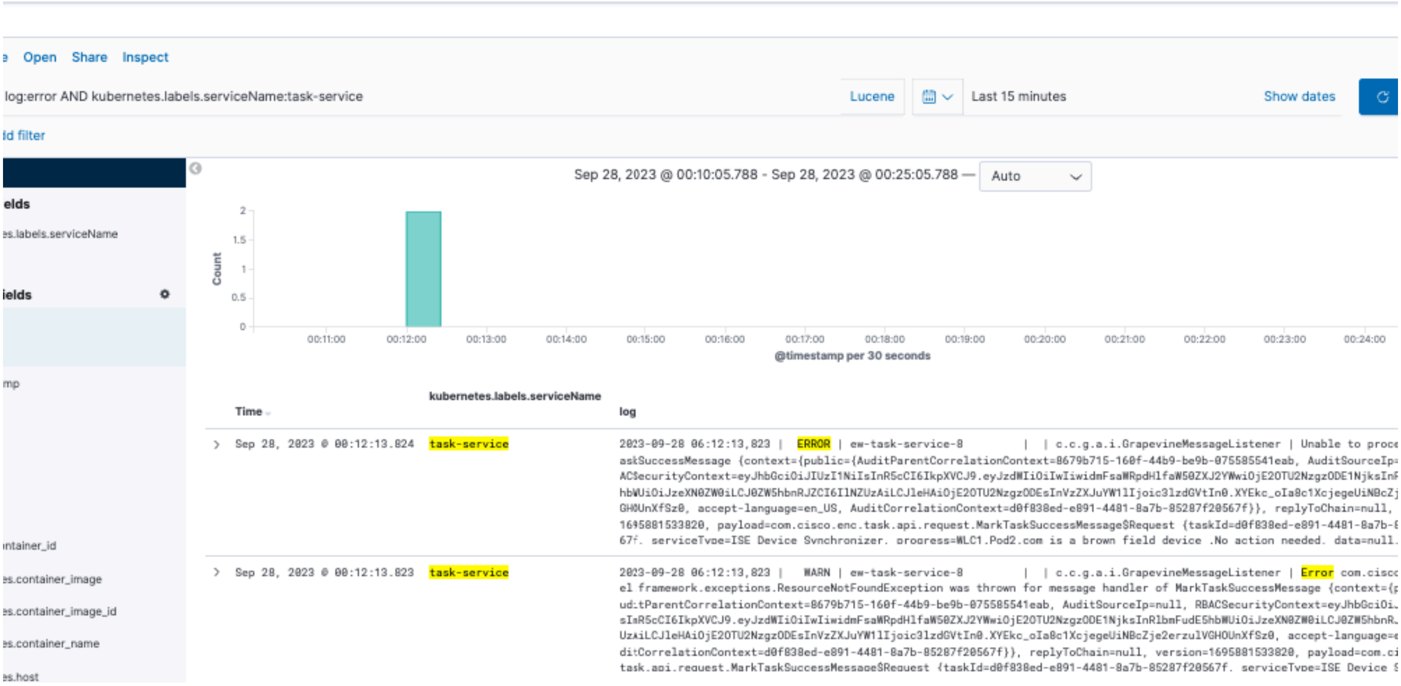

```
log:error AND kubernetes.labels.serviceName:task-service
```

```
Time                         kubernetes.labels.serviceName   log
> Sep 28, 2023 @ 00:12:13.824   task-service                 2023-09-28 06:12:13,823 |  ERROR | ew-task-service-8        |  | c.c.g.a.i.GrapevineMessageListener | Unable to proce
                                                             askSuccessMessage {context={public={AuditParentCorrelationContext=8679b715-160f-44b9-be9b-075585541eab, AuditSourceIp=
                                                             AC$ecurityContext=eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJzdWIiOiIwIiwidmFsaWRpdHlfaW50ZXJ2YWwiOjE2OTU2NzgzODE1NjksInF
                                                             hbWUiOiJzeXN0ZW0iLCJ0ZW5hbnRJZCI6IlNZUzAiLCJleHAiOjE2OTU2NzgzODEsInVzZXJuYW1lIjoic3lzdGVtIn0.XYEkc_oIa8c1XcjegeUiNBcZj
                                                             GH0UnXfSz0, accept-language=en_US, AuditCorrelationContext=d0f838ed-e891-4481-8a7b-85287f20567f}}, replyToChain=null,
                                                             1695881533820, payload=com.cisco.enc.task.api.request.MarkTaskSuccessMessage$Request {taskId=d0f838ed-e891-4481-8a7b-8
                                                             67f, serviceType=ISE Device Synchronizer, progress=WLC1.Pod2.com is a brown field device .No action needed, data=null.
> Sep 28, 2023 @ 00:12:13.823   task-service                 2023-09-28 06:12:13,823 |   WARN | ew-task-service-8        |  | c.c.g.a.i.GrapevineMessageListener | Error com.cisco
                                                             el framework.exceptions.ResourceNotFoundException was thrown for message handler of MarkTaskSuccessMessage {context={p
                                                             ud:tParentCorrelationContext=8679b715-160f-44b9-be9b-075585541eab, AuditSourceIp=null, RBACSecurityContext=eyJhbGciOi.
                                                             sInR5cCI6IkpXVCJ9.eyJzdWIiOiIwIiwidmFsaWRpdHlfaW50ZXJ2YWwiOjE2OTU2NzgzODE1NjksInRlbmFudE5hbWUiOiJzeXN0ZW0iLCJ0ZW5hbnR.
                                                             UzAiLCJleHAiOjE2OTU2NzgzODEsInVzZXJuYW1lIjoic3lzdGVtIn0.XYEkc_oIa8c1XcjegeUiNBcZje2erzulVGHOUnXfSz0, accept-language=e
                                                             ditCorrelationContext=d0f838ed-e891-4481-8a7b-85287f20567f}}, replyToChain=null, version=1695881533820, payload=com.ci
                                                             task.api.request.MarkTaskSuccessMessage$Request {taskId=d0f838ed-e891-4481-8a7b-85287f20567f, serviceType=ISE Device S
```

Note: Not all fields are searchable.

If you want to see only searchable fields in the **Available Fields** pane, select the cogwheel and customize the view. You can also define the type of search that you want to use, for example, string, Boolean, number, and so on.

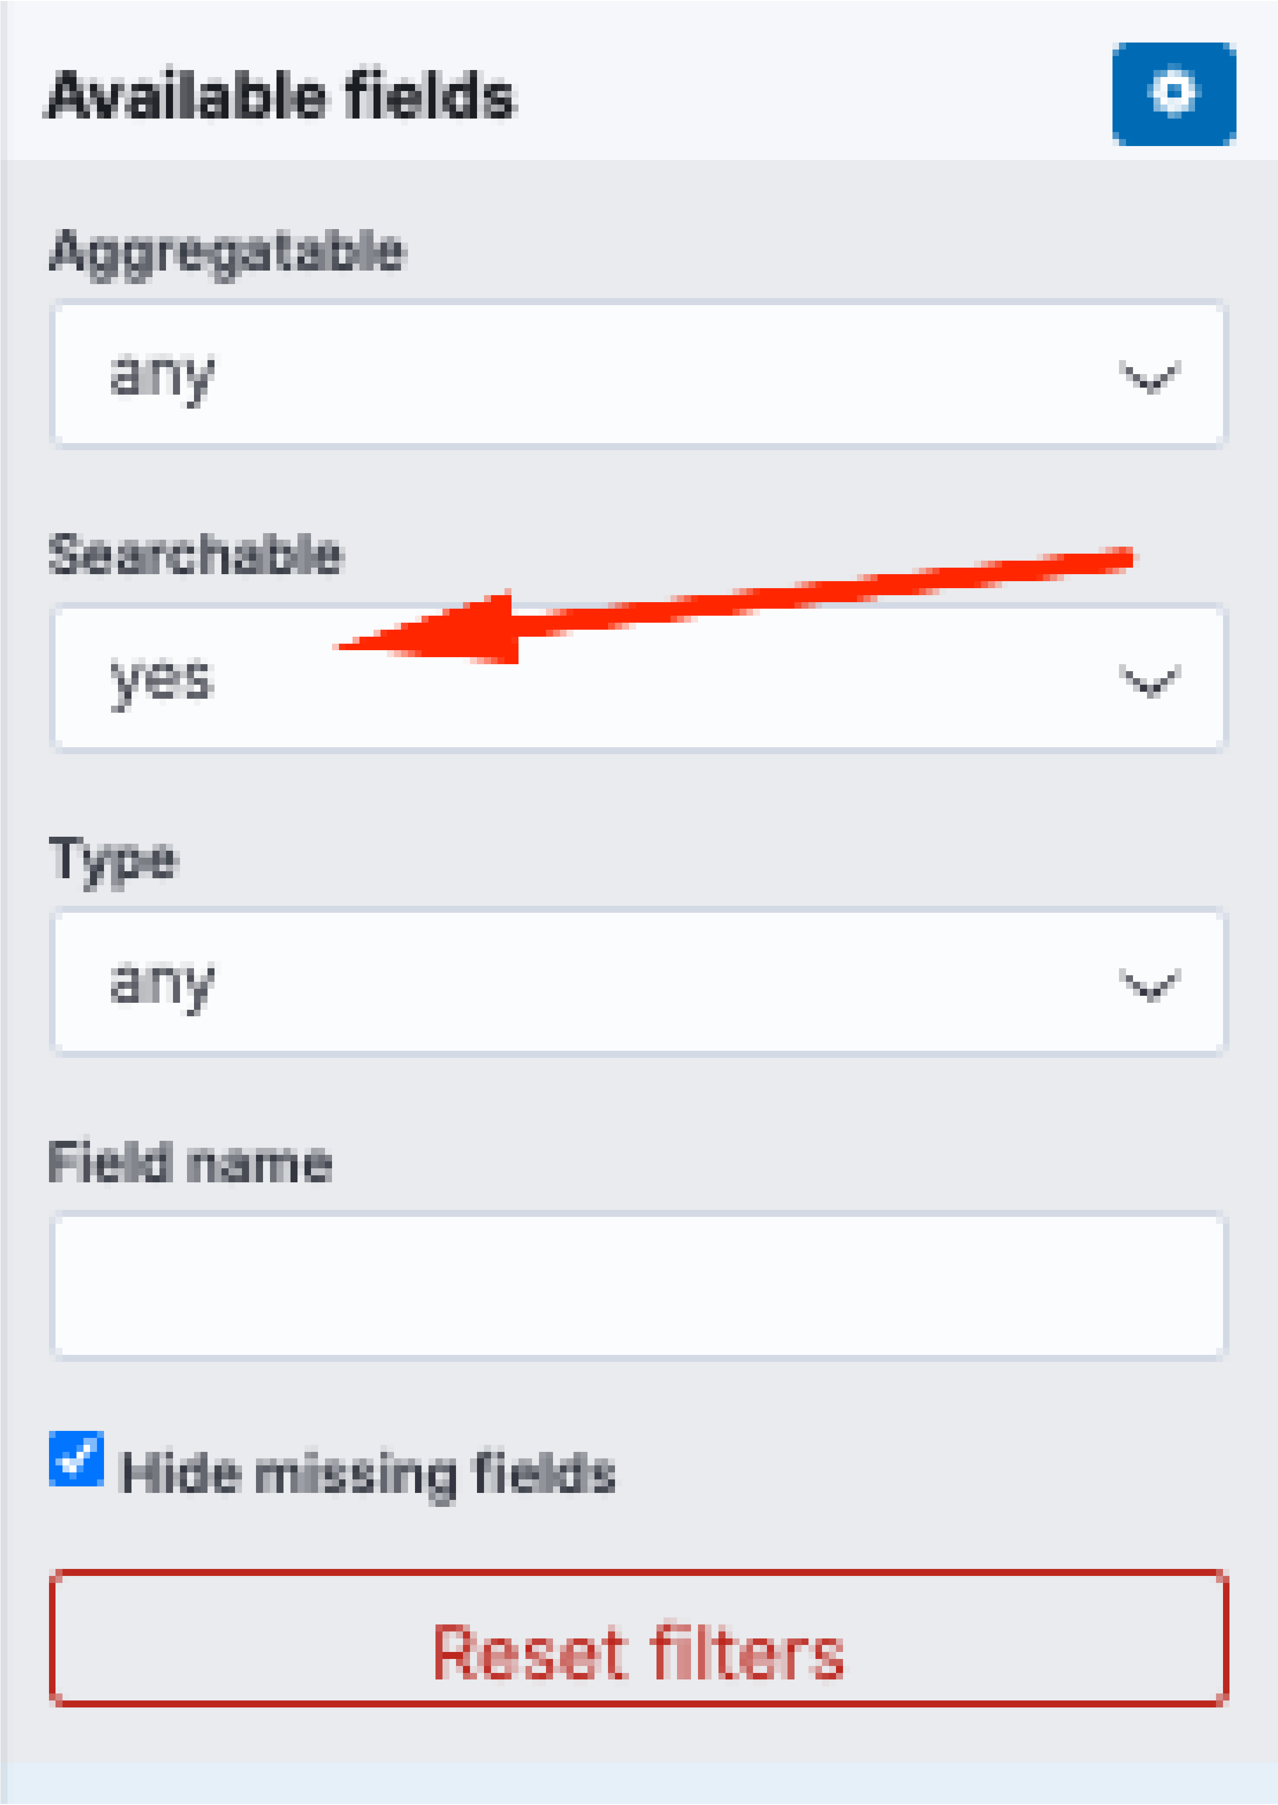

Available fields

Aggregatable: any

Searchable: yes

Type: any

Field name:

☑ Hide missing fields

Search two different services at the same time for an error

Include two or more services in your search criteria. Ensure that the services names are entered in parenthesis and separate them with **OR**.

```
log:error && (kubernetes.labels.serviceName:onboarding-service OR kubernetes.labels.serviceName:telemet
```

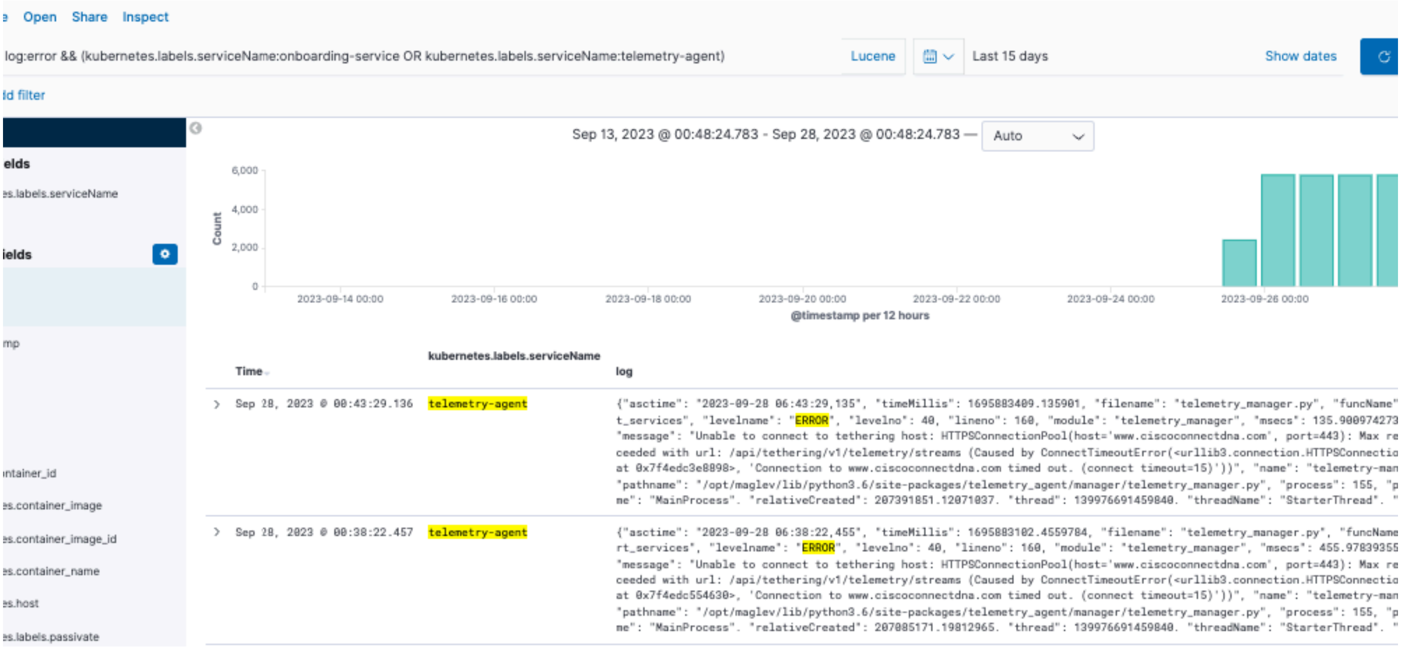

log:error && (kubernetes.labels.serviceName:onboarding-service OR kubernetes.labels.serviceName:telemetry-agent)

Sep 13, 2023 @ 00:48:24.783 - Sep 28, 2023 @ 00:48:24.783 —

```
Time          kubernetes.labels.serviceName   log
> Sep 28, 2023 @ 00:43:29.136   telemetry-agent   {"asctime": "2023-09-28 06:43:29,135", "timeMillis": 1695883409.135901, "filename": "telemetry_manager.py", "funcName"
                                                  t_services", "levelname": "ERROR", "levelno": 40, "lineno": 160, "module": "telemetry_manager", "msecs": 135.900974273
                                                  "message": "Unable to connect to tethering host: HTTPSConnectionPool(host='www.ciscoconnectdna.com', port=443): Max re
                                                  ceeded with url: /api/tethering/v1/telemetry/streams (Caused by ConnectTimeoutError(<urllib3.connection.HTTPSConnectio
                                                  at 0x7f4edc3e8898>, 'Connection to www.ciscoconnectdna.com timed out. (connect timeout=15)'))", "name": "telemetry-man
                                                  "pathname": "/opt/maglev/lib/python3.6/site-packages/telemetry_agent/manager/telemetry_manager.py", "process": 155, "p
                                                  me": "MainProcess". "relativeCreated": 207391851.12071037. "thread": 139976691459840. "threadName": "StarterThread". "
> Sep 28, 2023 @ 00:38:22.457   telemetry-agent   {"asctime": "2023-09-28 06:38:22,455", "timeMillis": 1695883102.4559784, "filename": "telemetry_manager.py", "funcName
                                                  rt_services", "levelname": "ERROR", "levelno": 40, "lineno": 160, "module": "telemetry_manager", "msecs": 455.97839355
                                                  "message": "Unable to connect to tethering host: HTTPSConnectionPool(host='www.ciscoconnectdna.com', port=443): Max re
                                                  ceeded with url: /api/tethering/v1/telemetry/streams (Caused by ConnectTimeoutError(<urllib3.connection.HTTPSConnectio
                                                  at 0x7f4edc554630>, 'Connection to www.ciscoconnectdna.com timed out. (connect timeout=15)'))", "name": "telemetry-man
                                                  "pathname": "/opt/maglev/lib/python3.6/site-packages/telemetry_agent/manager/telemetry_manager.py", "process": 155, "p
                                                  me": "MainProcess". "relativeCreated": 207085171.19812965. "thread": 139976691459840. "threadName": "StarterThread". "
```

Reference

- [Elastic searh common options]
- [Apache Lucene - Query Parser Syntax]