

# MetricBeat Config

```
##### Metricbeat Configuration Example #####

# This file is an example configuration file highlighting only the most common
# options. The metricbeat.reference.yml file from the same directory contains all the
# supported options with more comments. You can use it as a reference.
#
# You can find the full configuration reference here:
# https://www.elastic.co/guide/en/beats/metricbeat/index.html

#==== Modules configuration =====

metricbeat.config.modules:
  # Glob pattern for configuration loading
  path: ${path.config}/modules.d/*.yml

  # Set to true to enable config reloading
  reload.enabled: false

  # Period on which files under path should be checked for changes
  #reload.period: 10s

#==== Elasticsearch template setting =====

setup.template.settings:
  index.number_of_shards: 1
  index.codec: best_compression
  #_source.enabled: false

#==== General =====

# The name of the shipper that publishes the network data. It can be used to group
# all the transactions sent by a single shipper in the web interface.
#name:

# The tags of the shipper are included in their own field with each
# transaction published.
```

```
#tags: ["service-X", "web-tier"]

# Optional fields that you can specify to add additional information to the
# output.
#fields:
#  env: staging

#===== Dashboards =====
# These settings control loading the sample dashboards to the Kibana index. Loading
# the dashboards is disabled by default and can be enabled either by setting the
# options here or by using the `setup` command.
#setup.dashboards.enabled: false

# The URL from where to download the dashboards archive. By default this URL
# has a value which is computed based on the Beat name and version. For released
# versions, this URL points to the dashboard archive on the artifacts.elastic.co
# website.
#setup.dashboards.url:

#===== Kibana =====

# Starting with Beats version 6.0.0, the dashboards are loaded via the Kibana API.
# This requires a Kibana endpoint configuration.
setup.kibana:
  host: "10.1.0.5:5601"

  # Kibana Host
  # Scheme and port can be left out and will be set to the default (http and 5601)
  # In case you specify an additional path, the scheme is required:
  http://localhost:5601/path
  # IPv6 addresses should always be defined as: https://[2001:db8::1]:5601
  #host: "localhost:5601"

  # Kibana Space ID
  # ID of the Kibana Space into which the dashboards should be loaded. By default,
  # the Default Space will be used.
  #space.id:

#===== Elastic Cloud =====
```

```
# These settings simplify using Metricbeat with the Elastic Cloud (https://cloud.elastic.co/).

# The cloud.id setting overwrites the `output.elasticsearch.hosts` and
# `setup.kibana.host` options.
# You can find the `cloud.id` in the Elastic Cloud web UI.
#cloud.id:

# The cloud.auth setting overwrites the `output.elasticsearch.username` and
# `output.elasticsearch.password` settings. The format is `:<pass>`.
#cloud.auth:

#===== Outputs =====

# Configure what output to use when sending the data collected by the beat.

#----- Elasticsearch output -----
output.elasticsearch:
  # TODO: Change the hosts IP address to the IP address of your ELK server
  # TODO: Change password from `changem` to the password you created
  hosts: ["10.1.0.5:9200"]
  username: "elastic"
  password: "Password"

#----- Logstash output -----
#output.logstash:
  # The Logstash hosts
  #hosts: ["localhost:5044"]

  # Optional SSL. By default is off.
  # List of root certificates for HTTPS server verifications
  #ssl.certificate_authorities: ["/etc/pki/root/ca.pem"]

  # Certificate for SSL client authentication
  #ssl.certificate: "/etc/pki/client/cert.pem"

  # Client Certificate Key
  #ssl.key: "/etc/pki/client/cert.key"

#===== Processors =====

# Configure processors to enhance or manipulate events generated by the beat.
```

processors:

- add\_host\_metadata: ~
- add\_cloud\_metadata: ~

#===== Logging =====

# Sets log level. The default log level is info.

# Available log levels are: error, warning, info, debug

#logging.level: debug

# At debug level, you can selectively enable logging only for some components.

# To enable all selectors use ["\*"]. Examples of other selectors are "beat",

# "publish", "service".

#logging.selectors: ["\*"]

#===== X-Pack Monitoring =====

# metricbeat can export internal metrics to a central Elasticsearch monitoring  
# cluster. This requires xpack monitoring to be enabled in Elasticsearch. The  
# reporting is disabled by default.

# Set to true to enable the monitoring reporter.

#monitoring.enabled: false

# Sets the UUID of the Elasticsearch cluster under which monitoring data for this  
# Metricbeat instance will appear in the Stack Monitoring UI. If output.elasticsearch  
# is enabled, the UUID is derived from the Elasticsearch cluster referenced by  
output.elasticsearch.

#monitoring.cluster\_uuid:

# Uncomment to send the metrics to Elasticsearch. Most settings from the  
# Elasticsearch output are accepted here as well.

# Note that the settings should point to your Elasticsearch \*monitoring\* cluster.

# Any setting that is not set is automatically inherited from the Elasticsearch  
# output configuration, so if you have the Elasticsearch output configured such  
# that it is pointing to your Elasticsearch monitoring cluster, you can simply  
# uncomment the following line.

#monitoring.elasticsearch:

#===== Migration =====

```
# This allows to enable 6.7 migration aliases
#migration.6_to_7.enabled: true
```

---

Revision #1

Created 2025-11-25 18:23:59 UTC by David Rizzo

Updated 2025-11-25 18:24:08 UTC by David Rizzo