```
######################### Filebeat Configuration #############################

# This file is a full configuration example documenting all non-deprecated
# options in comments. For a shorter configuration example, that contains only
# the most common options, please see filebeat.yml in the same directory.
#
# You can find the full configuration reference here:
# https://www.elastic.co/guide/en/beats/filebeat/index.html

filebeat.inputs:
#------------------------------ Log input --------------------------------
- type: log

  # Change to true to enable this input configuration.
  enabled: true

  # Paths that should be crawled and fetched. Glob based paths.
  # To fetch all ".log" files from a specific level of subdirectories
  # /var/log/*/*.log can be used.
  # For each file found under this path, a harvester is started.
  # Make sure not file is defined twice as this can lead to unexpected behaviour.
  paths:
    #- /var/log/*.log
    - <absolute\path\to\log\*>

  # Configure the file encoding for reading files with international characters
  # following the W3C recommendation for HTML5 (http://www.w3.org/TR/encoding).
  # Some sample encodings:
  #   plain, utf-8, utf-16be-bom, utf-16be, utf-16le, big5, gb18030, gbk,
  #    hz-gb-2312, euc-kr, euc-jp, iso-2022-jp, shift-jis, ...
  #encoding: plain


  # Exclude lines. A list of regular expressions to match. It drops the lines that are
  # matching any regular expression from the list. The include_lines is called before
  # exclude_lines. By default, no lines are dropped.
  #exclude_lines: ['^DBG']

  # Include lines. A list of regular expressions to match. It exports the lines that are
  # matching any regular expression from the list. The include_lines is called before
  # exclude_lines. By default, all the lines are exported.
  #include_lines: ['^ERR', '^WARN']

  # Exclude files. A list of regular expressions to match. Filebeat drops the files that
  # are matching any regular expression from the list. By default, no files are dropped.
  #exclude_files: ['.gz$']

  # Optional additional fields. These fields can be freely picked
  # to add additional information to the crawled log files for filtering
  #fields:
  #  level: debug
  #  review: 1

  # Set to true to store the additional fields as top level fields instead
  # of under the "fields" sub-dictionary. In case of name conflicts with the
  # fields added by Filebeat itself, the custom fields overwrite the default
  # fields.
  #fields_under_root: false

  # Ignore files which were modified more then the defined timespan in the past.
  # ignore_older is disabled by default, so no files are ignored by setting it to 0.
  # Time strings like 2h (2 hours), 5m (5 minutes) can be used.
  #ignore_older: 0

  # How often the input checks for new files in the paths that are specified
  # for harvesting. Specify 1s to scan the directory as frequently as possible
```

```
  # without causing Filebeat to scan too frequently. Default: 10s.
  #scan_frequency: 10s

  # Defines the buffer size every harvester uses when fetching the file
  #harvester_buffer_size: 16384

  # Maximum number of bytes a single log event can have
  # All bytes after max_bytes are discarded and not sent. The default is 10MB.
  # This is especially useful for multiline log messages which can get large.
  #max_bytes: 10485760

  ### Recursive glob configuration

  # Expand "**" patterns into regular glob patterns.
  #recursive_glob.enabled: true

  ### JSON configuration

  # Decode JSON options. Enable this if your logs are structured in JSON.
  # JSON key on which to apply the line filtering and multiline settings. This key
  # must be top level and its value must be string, otherwise it is ignored. If
  # no text key is defined, the line filtering and multiline features cannot be used.
  json.message_key: message

  # By default, the decoded JSON is placed under a "json" key in the output document.
  # If you enable this setting, the keys are copied top level in the output document.
  json.keys_under_root: true

   ### Harvester closing options

  # Close inactive closes the file handler after the predefined period.
  # The period starts when the last line of the file was, not the file ModTime.
  # Time strings like 2h (2 hours), 5m (5 minutes) can be used.
  close_inactive: 15m

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

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

#-------------------------- Elasticsearch output ------------------------------
output.elasticsearch:
  # Boolean flag to enable or disable the output module.
  enabled: true

  # Array of hosts to connect to.
  # Scheme and port can be left out and will be set to the default (http and 9200)
  # In case you specify and additional path, the scheme is required: http://localhost:9200/path
  # IPv6 addresses should always be defined as: https://[2001:db8::1]:9200
  hosts: ["ElasticsearchHost:9200"]

  # Set gzip compression level.
  #compression_level: 0

  # Configure escaping html symbols in strings.
  #escape_html: true

  # Optional protocol and basic auth credentials.
  protocol: "https"
  username: "filebeat"
  password: "filebeat"

  # Dictionary of HTTP parameters to pass within the url with index operations.
  #parameters:
    #param1: value1
    #param2: value2
```

```
  # Number of workers per Elasticsearch host.
  #worker: 1

  # Optional index name. The default is "filebeat" plus date
  # and generates [filebeat-]YYYY.MM.DD keys.
  # In case you modify this pattern you must update setup.template.name and
  setup.template.pattern accordingly.
  index: "rpalog-%{+yyyy.MM.dd}"

  # Optional ingest node pipeline. By default no pipeline will be used.
  #pipeline: ""

  # Optional HTTP Path
  #path: "/elasticsearch"

  # Custom HTTP headers to add to each request
  #headers:
  #  X-My-Header: Contents of the header

  # Proxy server url
  #proxy_url: http://proxy:3128

  # The number of times a particular Elasticsearch index operation is attempted. If
  # the indexing operation doesn't succeed after this many retries, the events are
  # dropped. The default is 3.
  #max_retries: 3

  # The maximum number of events to bulk in a single Elasticsearch bulk API index request.
  # The default is 50.
  #bulk_max_size: 50

  # The number of seconds to wait before trying to reconnect to Elasticsearch
  # after a network error. After waiting backoff.init seconds, the Beat
  # tries to reconnect. If the attempt fails, the backoff timer is increased
  # exponentially up to backoff.max. After a successful connection, the backoff
  # timer is reset. The default is 1s.
  #backoff.init: 1s

  # The maximum number of seconds to wait before attempting to connect to
  # Elasticsearch after a network error. The default is 60s.
  #backoff.max: 60s

  # Configure http request timeout before failing a request to Elasticsearch.
  #timeout: 90

  # Use SSL settings for HTTPS.
  ssl.enabled: true

  # Configure SSL verification mode. If `none` is configured, all server hosts
  # and certificates will be accepted. In this mode, SSL based connections are
  # susceptible to man-in-the-middle attacks. Use only for testing. Default is
  # `full`.
  ssl.verification_mode: full

  # List of supported/valid TLS versions. By default all TLS versions 1.0 up to
  # 1.2 are enabled.
  #ssl.supported_protocols: [TLSv1.0, TLSv1.1, TLSv1.2]

  # SSL configuration. By default is off.
  # List of root certificates for HTTPS server verifications

  ssl.certificate_authorities: "absolute\\path\\to\\certificate_file"
  #e.g. ssl.certificate_authorities: "C:\\Programs\\ES\\filebeat\\cert.crt"

  # Certificate for SSL client authentication
  #ssl.certificate: "cert.p12"
```

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

  # Optional passphrase for decrypting the Certificate Key.
  #ssl.key_passphrase: ''

  # Configure cipher suites to be used for SSL connections
  #ssl.cipher_suites: []

  # Configure curve types for ECDHE based cipher suites
  #ssl.curve_types: []

  # Configure what types of renegotiation are supported. Valid options are
  # never, once, and freely. Default is never.
  #ssl.renegotiation: never



#================================ Paths ======================================

# The home path for the filebeat installation. This is the default base path
# for all other path settings and for miscellaneous files that come with the
# distribution (for example, the sample dashboards).
# If not set by a CLI flag or in the configuration file, the default for the
# home path is the location of the binary.
#path.home:

# The configuration path for the filebeat installation. This is the default
# base path for configuration files, including the main YAML configuration file
# and the Elasticsearch template file. If not set by a CLI flag or in the
# configuration file, the default for the configuration path is the home path.
#path.config: ${path.home}

# The data path for the filebeat installation. This is the default base path
# for all the files in which filebeat needs to store its data. If not set by a
# CLI flag or in the configuration file, the default for the data path is a data
# subdirectory inside the home path.
#path.data: ${path.home}/data

# The logs path for a filebeat installation. This is the default location for
# the Beat's log files. If not set by a CLI flag or in the configuration file,
# the default for the logs path is a logs subdirectory inside the home path.
#path.logs: ${path.home}/logs

#================================ Keystore ==========================================
# Location of the Keystore containing the keys and their sensitive values.
#keystore.path: "${path.config}/beats.keystore"

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

# The directory from where to read the dashboards. The default is the `kibana`
# folder in the home path.
#setup.dashboards.directory: ${path.home}/kibana

# The URL from where to download the dashboards archive. It is used instead of
# the directory if it has a value.
#setup.dashboards.url:

# The file archive (zip file) from where to read the dashboards. It is used instead
# of the directory when it has a value.
#setup.dashboards.file:
```

```
# In case the archive contains the dashboards from multiple Beats, this lets you
# select which one to load. You can load all the dashboards in the archive by
# setting this to the empty string.
#setup.dashboards.beat: filebeat

# The name of the Kibana index to use for setting the configuration. Default is ".kibana"
#setup.dashboards.kibana_index: .kibana

# The Elasticsearch index name. This overwrites the index name defined in the
# dashboards and index pattern. Example: testbeat-*
#setup.dashboards.index:

# Always use the Kibana API for loading the dashboards instead of autodetecting
# how to install the dashboards by first querying Elasticsearch.
#setup.dashboards.always_kibana: false

# If true and Kibana is not reachable at the time when dashboards are loaded,
# it will retry to reconnect to Kibana instead of exiting with an error.
#setup.dashboards.retry.enabled: false

# Duration interval between Kibana connection retries.
#setup.dashboards.retry.interval: 1s

# Maximum number of retries before exiting with an error, 0 for unlimited retrying.
#setup.dashboards.retry.maximum: 0


#================================ Template =====================================

# A template is used to set the mapping in Elasticsearch
# By default template loading is enabled and the template is loaded.
# These settings can be adjusted to load your own template or overwrite existing ones.

# Set to false to disable template loading.
#setup.template.enabled: true

# Template name. By default the template name is "filebeat-%{[beat.version]}"
# The template name and pattern has to be set in case the elasticsearch index pattern is
modified.
setup.template.name: "rpalog"

# Template pattern. By default the template pattern is "-%{[beat.version]}-*" to apply to the
default index settings.
# The first part is the version of the beat and then -* is used to match all daily indices.
# The template name and pattern has to be set in case the elasticsearch index pattern is
modified.
setup.template.pattern: "rpalog-*"

# Path to fields.yml file to generate the template
#setup.template.fields: "${path.config}/fields.yml"

# A list of fields to be added to the template and Kibana index pattern. Also
# specify setup.template.overwrite: true to overwrite the existing template.
# This setting is experimental.
#setup.template.append_fields:
#- name: field_name
#  type: field_type

# Enable json template loading. If this is enabled, the fields.yml is ignored.
#setup.template.json.enabled: false

# Path to the json template file
#setup.template.json.path: "${path.config}/template.json"

# Name under which the template is stored in Elasticsearch
```

```
#setup.template.json.name: ""

# Overwrite existing template
#setup.template.overwrite: false

# Elasticsearch template settings
setup.template.settings:

  # A dictionary of settings to place into the settings.index dictionary
  # of the Elasticsearch template. For more details, please check
  # https://www.elastic.co/guide/en/elasticsearch/reference/current/mapping.html
  #index:
    #number_of_shards: 1
    #codec: best_compression
    #number_of_routing_shards: 30

  # A dictionary of settings for the _source field. For more details, please check
  # https://www.elastic.co/guide/en/elasticsearch/reference/current/mapping-source-field.html
  #_source:
    #enabled: false

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

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

  # Kibana Host
  # Scheme and port can be left out and will be set to the default (http and 5601)
  # In case you specify and 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"

  # Optional protocol and basic auth credentials.
  #protocol: "https"
  #username: "elastic"
  #password: "changeme"

  # Optional HTTP Path
  #path: ""

  # Use SSL settings for HTTPS. Default is true.
  #ssl.enabled: true

  # Configure SSL verification mode. If `none` is configured, all server hosts
  # and certificates will be accepted. In this mode, SSL based connections are
  # susceptible to man-in-the-middle attacks. Use only for testing. Default is
  # `full`.
  #ssl.verification_mode: full

  # List of supported/valid TLS versions. By default all TLS versions 1.0 up to
  # 1.2 are enabled.
  #ssl.supported_protocols: [TLSv1.0, TLSv1.1, TLSv1.2]

  # SSL configuration. 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"

  # Optional passphrase for decrypting the Certificate Key.
  #ssl.key_passphrase: ''
```

```
  # Configure cipher suites to be used for SSL connections
  #ssl.cipher_suites: []

  # Configure curve types for ECDHE based cipher suites
  #ssl.curve_types: []



#================================ Logging =====================================
# There are four options for the log output: file, stderr, syslog, eventlog
# The file output is the default.

# Sets log level. The default log level is info.
# Available log levels are: error, warning, info, debug
#logging.level: info

# Enable debug output for selected components. To enable all selectors use ["*"]
# Other available selectors are "beat", "publish", "service"
# Multiple selectors can be chained.
#logging.selectors: [ ]

# Send all logging output to syslog. The default is false.
#logging.to_syslog: false

# Send all logging output to Windows Event Logs. The default is false.
#logging.to_eventlog: false

# If enabled, filebeat periodically logs its internal metrics that have changed
# in the last period. For each metric that changed, the delta from the value at
# the beginning of the period is logged. Also, the total values for
# all non-zero internal metrics are logged on shutdown. The default is true.
#logging.metrics.enabled: true

# The period after which to log the internal metrics. The default is 30s.
#logging.metrics.period: 30s

# Logging to rotating files. Set logging.to_files to false to disable logging to
# files.
logging.to_files: true
logging.files:
  # Configure the path where the logs are written. The default is the logs directory
  # under the home path (the binary location) or the directory specified as argument to
  Filebeat service
#path:

  # The name of the files where the logs are written to.
  #name: filebeat

  # Configure log file size limit. If limit is reached, log file will be
  # automatically rotated
  #rotateeverybytes: 10485760 # = 10MB

  # Number of rotated log files to keep. Oldest files will be deleted first.
  #keepfiles: 7

  # The permissions mask to apply when rotating log files. The default value is 0600.
  # Must be a valid Unix-style file permissions mask expressed in octal notation.
  #permissions: 0600

# Set to true to log messages in json format.
#logging.json: false


#============================== Xpack Monitoring ===============================
# filebeat 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.
#xpack.monitoring.enabled: false

# Uncomment to send the metrics to Elasticsearch. Most settings from the
# Elasticsearch output are accepted here as well. Any setting that is not set is
# automatically inherited from the Elasticsearch output configuration, so if you
# have the Elasticsearch output configured, you can simply uncomment the
# following line, and leave the rest commented out.
#xpack.monitoring.elasticsearch:

  # Array of hosts to connect to.
  # Scheme and port can be left out and will be set to the default (http and 9200)
  # In case you specify and additional path, the scheme is required: http://localhost:9200/path
  # IPv6 addresses should always be defined as: https://[2001:db8::1]:9200
  #hosts: ["localhost:9200"]

  # Set gzip compression level.
  #compression_level: 0

  # Optional protocol and basic auth credentials.
  #protocol: "https"
  #username: "beats_system"
  #password: "changeme"

  # Dictionary of HTTP parameters to pass within the url with index operations.
  #parameters:
    #param1: value1
    #param2: value2

  # Custom HTTP headers to add to each request
  #headers:
  #  X-My-Header: Contents of the header

  # Proxy server url
  #proxy_url: http://proxy:3128

  # The number of times a particular Elasticsearch index operation is attempted. If
  # the indexing operation doesn't succeed after this many retries, the events are
  # dropped. The default is 3.
  #max_retries: 3

  # The maximum number of events to bulk in a single Elasticsearch bulk API index request.
  # The default is 50.
  #bulk_max_size: 50

  # The number of seconds to wait before trying to reconnect to Elasticsearch
  # after a network error. After waiting backoff.init seconds, the Beat
  # tries to reconnect. If the attempt fails, the backoff timer is increased
  # exponentially up to backoff.max. After a successful connection, the backoff
  # timer is reset. The default is 1s.
  #backoff.init: 1s

  # The maximum number of seconds to wait before attempting to connect to
  # Elasticsearch after a network error. The default is 60s.
  #backoff.max: 60s

  # Configure http request timeout before failing an request to Elasticsearch.
  #timeout: 90

  # Use SSL settings for HTTPS.
  #ssl.enabled: true

  # Configure SSL verification mode. If `none` is configured, all server hosts
```

```
  # and certificates will be accepted. In this mode, SSL based connections are
  # susceptible to man-in-the-middle attacks. Use only for testing. Default is
  # `full`.
  #ssl.verification_mode: full

  # List of supported/valid TLS versions. By default all TLS versions 1.0 up to
  # 1.2 are enabled.
  #ssl.supported_protocols: [TLSv1.0, TLSv1.1, TLSv1.2]

  # SSL configuration. 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"

  # Optional passphrase for decrypting the Certificate Key.
  #ssl.key_passphrase: ''

  # Configure cipher suites to be used for SSL connections
  #ssl.cipher_suites: []

  # Configure curve types for ECDHE based cipher suites
  #ssl.curve_types: []

  # Configure what types of renegotiation are supported. Valid options are
  # never, once, and freely. Default is never.
  #ssl.renegotiation: never

  #metrics.period: 10s
  #state.period: 1m

#================================ HTTP Endpoint ======================================
# Each beat can expose internal metrics through a HTTP endpoint. For security
# reasons the endpoint is disabled by default. This feature is currently experimental.
# Stats can be access through http://localhost:5066/stats . For pretty JSON output
# append ?pretty to the URL.

# Defines if the HTTP endpoint is enabled.
#http.enabled: false

# The HTTP endpoint will bind to this hostname or IP address. It is recommended to use only
localhost.
#http.host: localhost

# Port on which the HTTP endpoint will bind. Default is 5066.
#http.port: 5066

#============================= Process Security ================================

# Enable or disable seccomp system call filtering on Linux. Default is enabled.
#seccomp.enabled: true
```