




Introduction to Logging with the ELK Stack

Jamie Smith
Product Marketing, Observability



Housekeeping & Logistics

- Attendees are automatically muted when joining Zoom
- Q&A will be at the end of the webinar
 - Ask questions for us in the Zoom chat during the webinar
 - Set chat settings **To: All panelists and attendees**



- Ask more questions on our discuss forum: **discuss.elastic.co**
- **Recording** and **Slides** will be available after the webinar and emailed to all registrants



logs

+



metrics

=

+



apm



ObservaBLT
Observability

Elastic Approach to Observability

Dev & Ops Teams



Log Data

Web Logs
App Logs
Database Logs
Container Logs

Metrics Data

Container Metrics
Host Metrics
Database Metrics
Network Metrics
Storage Metrics

APM Data

Real User Monitoring
Txn Perf Monitoring
Distributed Tracing

Uptime Data

Uptime
Response Time



kibana



elasticsearch

Agenda

Things we're going to cover

- 1 Challenges with log analytics

- 2 Sending logs to Elasticsearch

- 3 Beyond logging: Observability

- 4 Leveraging Elastic security

Agenda

Challenges with log analytics

1 Challenges with log analytics

2 Sending logs to Elasticsearch

3 Beyond logging: Observability

4 Leveraging Elastic security

Logs for one host or app

This is fairly straightforward

Terminal — 100×19

```
$ > tail -f /var/log/messages
```

```
Dec 10 14:05:30 justa-build kernel: type=1326 audit(1575986730.517:383998660): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17069 comm="node" sig=0 arch=c000003e syscall=324 compat=0 ip=0x7efe9c254889 code=0x50000
Dec 10 14:05:30 justa-build kernel: type=1326 audit(1575986730.551:383998661): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17069 comm="node" sig=0 arch=c000003e syscall=332 compat=0 ip=0x7efe9c269171 code=0x50000
Dec 10 14:05:33 justa-build kernel: type=1326 audit(1575986733.110:383998662): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17179 comm="node" sig=0 arch=c000003e syscall=324 compat=0 ip=0x7fee1cf0f889 code=0x50000
Dec 10 14:05:33 justa-build kernel: type=1326 audit(1575986733.150:383998663): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17179 comm="node" sig=0 arch=c000003e syscall=332 compat=0 ip=0x7fee1cf24171 code=0x50000
Dec 10 14:05:35 justa-build kernel: type=1326 audit(1575986735.155:383998664): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17367 comm="node" sig=0 arch=c000003e syscall=324 compat=0 ip=0x7ffb3b7bf889 code=0x50000
Dec 10 14:05:35 justa-build kernel: type=1326 audit(1575986735.194:383998665): auid=4294967295 uid=0 gid=0 ses=4294967295 subj=system_u:system_r:container_runtime_t:s0 pid=17367 comm="node"
```

Interacting with logs


Built-in tools for log viewing

- grep
- tail
- cat / less / more / type
- sed / awk / perl
- vim / notepad / event viewer
- clever combinations of the above

Immediate needs for log analytics

What's missing from the previous desktop

- Easy setup for a variety of sources
- Correlating and cross referencing
- Searching, filtering, and highlighting
- Visualize
- Anomaly detection and alerting
- Flexible retention



The screenshot displays a log analytics interface. The top section shows a list of log entries with columns for time, source, type, and message. The bottom section shows a summary table with columns for source, count, and other metrics.

Source	Count	Other Metrics
20961 packer	20	0 2343260 457168 8832 S 0.2 1.5 200:11.16 near local
20414 packer	20	0 2204836 484268 8832 S 5.9 1.6 207:01.43 heartbeat
7457 packer	20	0 1084444 84636 9340 S 5.3 0.3 269:26.21 apm-server
9507 root	20	0 327592 51428 2880 S 4.6 0.2 49:24.77 gunicorn
19449 root	20	0 1389668 49652 8600 S 3.6 0.2 452:37.78 filebeat
7373 root	20	0 1002296 34340 10640 S 2.6 0.1 111:13.50 metricbeat
5190 root	20	0 1003968 22088 8328 S 2.3 0.1 81:30.29 heartbeat
10776 root	20	0 190336 95820 3356 S 2.3 0.3 148:04.35 bundle

Agenda

Things we're going to cover

1 Challenges with log analytics

2 Sending logs to Elasticsearch

3 Beyond logging: Observability

4 Leveraging Elastic security

We're running in Elastic Cloud

Works the same in the cloud or running the default distribution

  Home   



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](#)



SIEM

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

Upload data from log file
Import a CSV, NDJSON, or log file

Use Elasticsearch data
Connect to your Elasticsearch index

Visualize and Explore Data



APM

Automatically collect in-depth performance metrics and errors from inside your applications.



Canvas

Showcase your data in a pixel-perfect way.

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.

Click on the Logging Button

Works the same in the cloud or running the default distribution

 Home





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](#)



SIEM

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

Upload data from log file

Import a CSV, NDJSON, or log file

Use Elasticsearch data

Connect to your Elasticsearch index

Visualize and Explore Data



APM

Automatically collect in-depth performance metrics and errors from inside your applications.



Canvas

Showcase your data in a pixel-perfect way.

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.

We're going to ingest the **System logs**

We're going to ingest the **System logs**

All **Logging** Metrics SIEM Sample data

All **Logging** Metrics SIEM Sample data

Cloudwatch Logs

Collect Cloudwatch logs with Functionbeat

 **Elasticsearch logs**
Collect and parse logs created by Elasticsearch.

IIS logs

Collect and parse access and error logs created by the IIS HTTP server.

Logstash logs

Collect and parse debug and slow logs created by Logstash itself.

MySQL logs

Collect and parse error and slow logs created by MySQL.

Nats logs

Collect and parse logs created by Nats.

PostgreSQL logs

Collect and parse error and slow logs created by PostgreSQL.

Redis logs

Collect and parse error and slow logs created by Redis.

System logs

Collect and parse logs written by the local Syslog server.

Traefik logs

Collect and parse access logs created by the Traefik Proxy.

Detailed instructions

Context-aware instructions for cloud or on-prem installs

Home / Add data / System logs




System logs

The `system` Filebeat module collects and parses logs created by the system logging service of common Unix/Linux based distributions. This module is not available on Windows. [Learn more](#).

[View exported fields](#)

Self managed Elastic Cloud



Getting Started

macOS DEB RPM

1 Download and install Filebeat

First time using Filebeat? See the [Getting Started Guide](#).

[Copy snippet](#)

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```

2 Edit the configuration

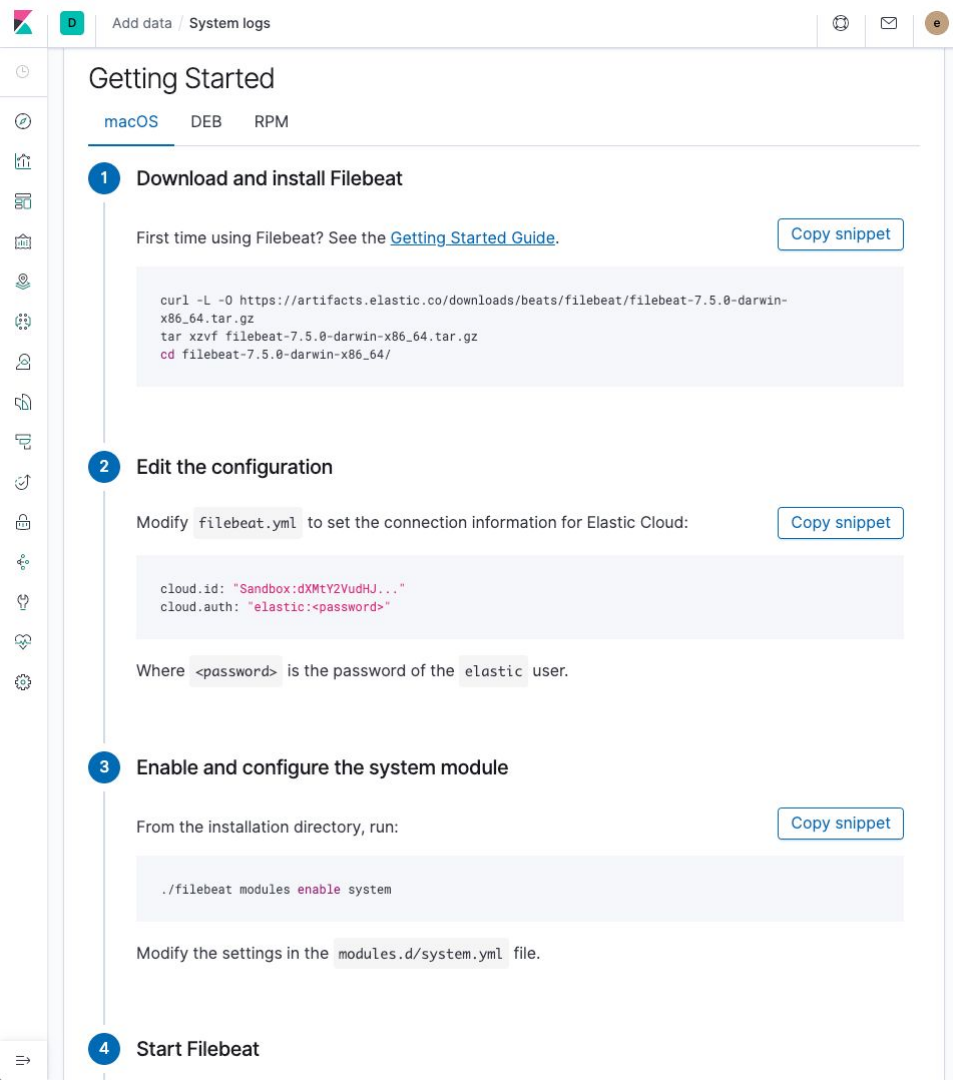
Modify `filebeat.yml` to set the connection information:

[Copy snippet](#)

Getting Started

Cloud or on-prem installs

- **Download and install Filebeat**
- Edit the configuration
- Enable and configure the system module
- Start Filebeat
- Check out the dashboard!



The screenshot shows the 'Getting Started' page for Filebeat on macOS. It features a sidebar with icons for various tools and a main content area with four numbered steps. Each step includes a title, a brief instruction, and a code snippet with a 'Copy snippet' button.

Add data / System logs

Getting Started

macOS DEB RPM

- 1 Download and install Filebeat**

First time using Filebeat? See the [Getting Started Guide](#). [Copy snippet](#)

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```
- 2 Edit the configuration**

Modify `filebeat.yml` to set the connection information for Elastic Cloud: [Copy snippet](#)

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.
- 3 Enable and configure the system module**

From the installation directory, run: [Copy snippet](#)

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.
- 4 Start Filebeat**

Steps

Download and install Filebeat

```
$ >curl -LO --silent \
https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz

$ >tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
$ >cd filebeat-7.5.0-darwin-x86_64
$ >ls -l
LICENSE.txt
NOTICE.txt
README.md
fields.yml
filebeat*
filebeat.reference.yml
filebeat.yml
kibana/
module/
modules.d/
```

macOS DEB RPM

1 Download and install Filebeat

First time using Filebeat? See the [Getting Started Guide](#).

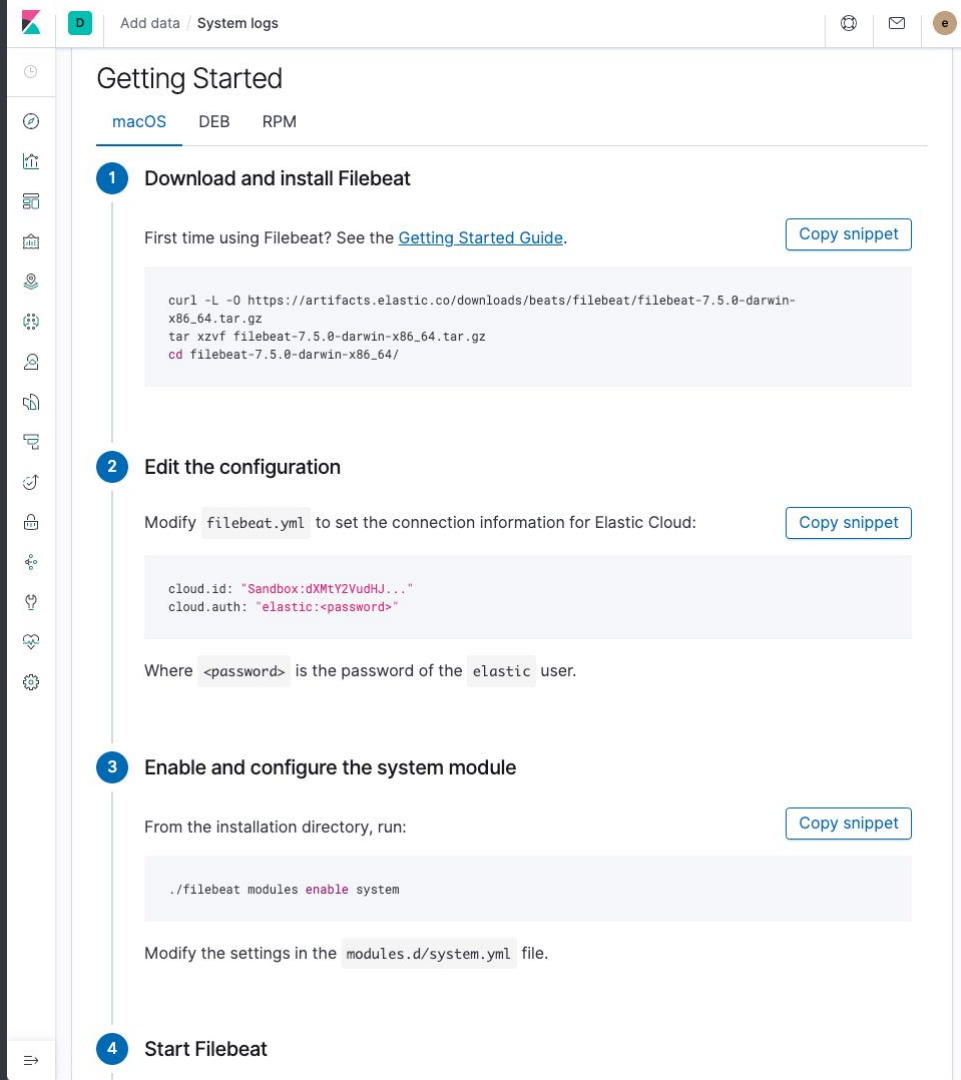
Copy snippet

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-
x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```

Steps

Edit the configuration

- Download and install Filebeat
- **Edit the configuration**
- Enable and configure the system module
- Start Filebeat
- Check out the dashboard!



The screenshot shows the 'Getting Started' page for Filebeat on macOS in the Elastic Cloud interface. The page is divided into four numbered steps:

- 1 Download and install Filebeat**

First time using Filebeat? See the [Getting Started Guide](#). Copy snippet

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```
- 2 Edit the configuration**

Modify `filebeat.yml` to set the connection information for Elastic Cloud: Copy snippet

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.
- 3 Enable and configure the system module**

From the installation directory, run: Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.
- 4 Start Filebeat**

Configuration

Cloud aware - using superuser

2 Edit the configuration

Modify `filebeat.yml` to set the connection information for Elastic Cloud:

[Copy snippet](#)

```
output.elasticsearch:
  hosts: ["<es_url>"]
  username: "elastic"
  password: "<password>"
setup.kibana:
  host: "<kibana_url>"
```

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.

Where `<password>` is the password of the `elastic` user, `<es_url>` is the URL of Elasticsearch, and `<kibana_url>` is the URL of Kibana.

Edit the configuration

Copy the snippet, paste in the password

2 Edit the configuration

Modify `filebeat.yml` to set the connection information for Elastic Cloud:

Copy snippet

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.

Terminal — 100x19

```
#===== Elastic Cloud =====
# These settings simplify using Filebeat 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: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:long-random-password" # because we are using Elastic Cloud
```

```
output.elasticsearch:
  # Array of hosts to connect to.
  hosts: ["localhost:9200"]      ← If we were not using Elastic Cloud
  #username: "elastic"          ←
  #password: "long-random-password" ←
```

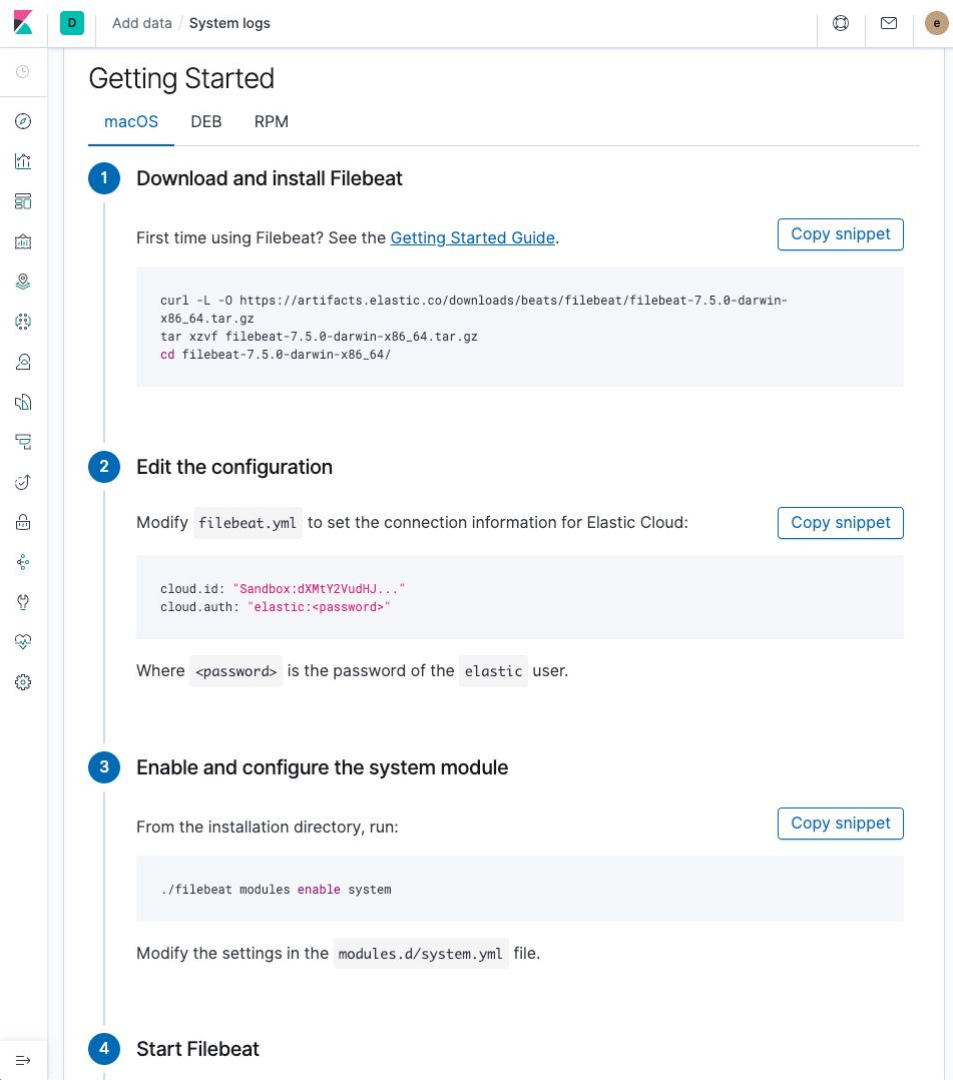
-UU-:----F1 filebeat.yml

(YAML)

Steps

Set up the system module

- Download and install Filebeat
- Edit the configuration
- **Enable and configure the system module**
- Start Filebeat
- Check out the dashboard!



The screenshot shows the 'Getting Started' page for Filebeat on macOS. It contains four numbered steps: 1. Download and install Filebeat, 2. Edit the configuration, 3. Enable and configure the system module, and 4. Start Filebeat. Each step includes a code snippet and a 'Copy snippet' button. The page is part of a larger application with a sidebar on the left and a top navigation bar.

Add data / System logs

Getting Started

macOS DEB RPM

- 1 Download and install Filebeat**

First time using Filebeat? See the [Getting Started Guide](#). [Copy snippet](#)

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```
- 2 Edit the configuration**

Modify `filebeat.yml` to set the connection information for Elastic Cloud: [Copy snippet](#)

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.
- 3 Enable and configure the system module**

From the installation directory, run: [Copy snippet](#)

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.
- 4 Start Filebeat**

Enable the system module

Again, just copy and paste the snippet

Terminal — 100×19

```
$ > ./filebeat modules enable system
```

3 Enable and configure the system module

From the installation directory, run:

Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.

Enable the system module

Again, just copy and paste the snippet

Terminal — 100×19

```
$ >./filebeat modules enable system
Enabled system
```

3

Enable and configure the system module

From the installation directory, run:

Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.

Enable the system module

Check your work

Terminal — 100×19

```
$ >./filebeat modules enable system  
Enabled system
```

```
# Can also verify
```

3

Enable and configure the system module

From the installation directory, run:

Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.

Enable the system module

Check your work

Terminal — 100×19

```
$ >./filebeat modules enable system
```

```
Enabled system
```

```
# Can also verify
```

```
$ >./filebeat modules list
```

3 Enable and configure the system module

From the installation directory, run:

Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.

Enable the system module

All good

Terminal — 100×19

```
$ > ./filebeat modules enable system
```

```
Enabled system
```

```
# Can also verify
```

```
$ > ./filebeat modules list
```

```
Enabled:
```

```
system
```

```
Disabled:
```

```
apache
```

```
auditd
```

```
aws
```

```
azure
```

```
(...)
```

3 Enable and configure the system module

From the installation directory, run:

Copy snippet

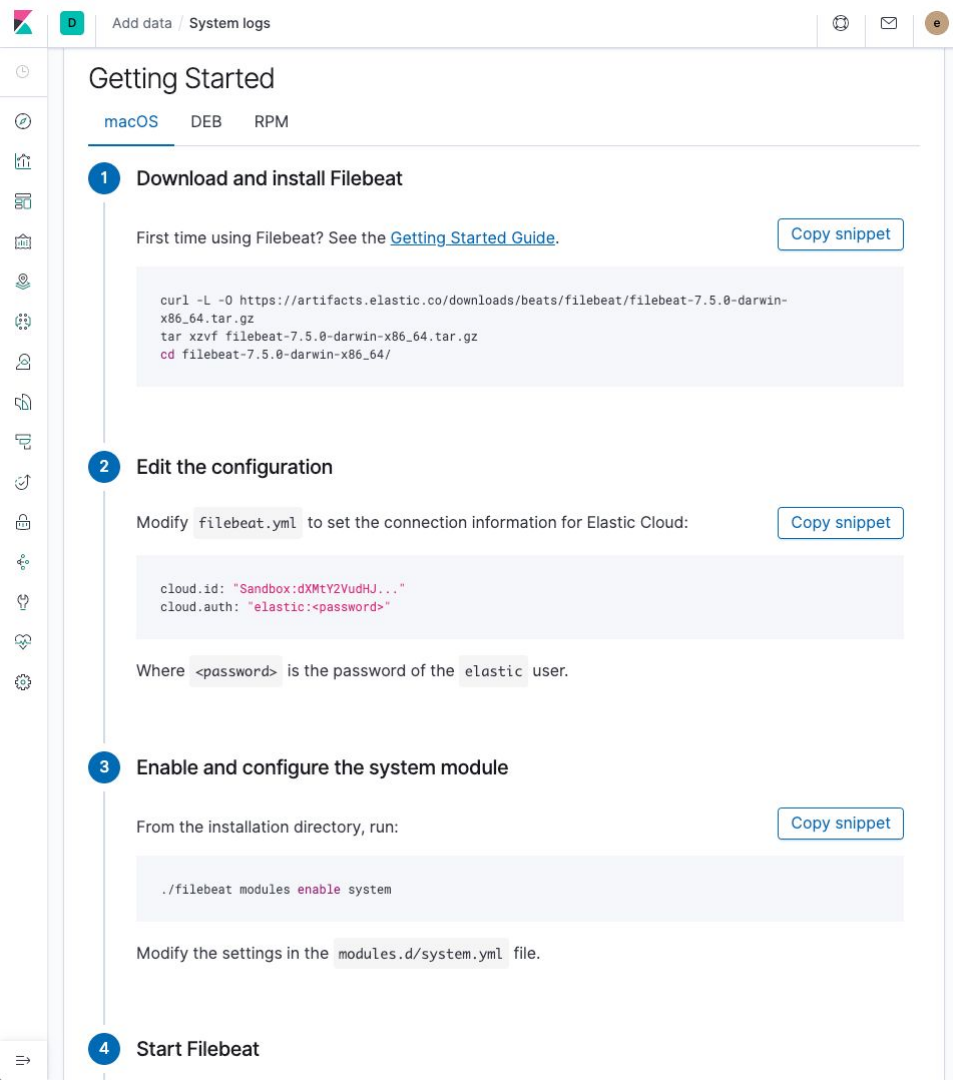
```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.

Steps

Start Filebeat

- Download and install Filebeat
- Edit the configuration
- Enable and configure the system module
- **Start Filebeat**
- Check out the dashboard!



The screenshot shows the 'Getting Started' page for Filebeat on macOS in the Elastic Cloud interface. The page is divided into four numbered steps:

- 1 Download and install Filebeat**

First time using Filebeat? See the [Getting Started Guide](#). Copy snippet

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```
- 2 Edit the configuration**

Modify `filebeat.yml` to set the connection information for Elastic Cloud: Copy snippet

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.
- 3 Enable and configure the system module**

From the installation directory, run: Copy snippet

```
./filebeat modules enable system
```

Modify the settings in the `modules.d/system.yml` file.
- 4 Start Filebeat**

And start it up!

Startup steps

Terminal — 100×19

\$ >

4 Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup
./filebeat -e
```

First run the setup process

Setup preps dashboards and indices



Terminal — 100×19

```
$ >./filebeat setup
```

4

Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup  
./filebeat -e
```

First run the setup process

Setup preps dashboards and indices



Terminal — 100×19

```
$ >./filebeat setup  
Index setup finished.
```

4

Start Filebeat

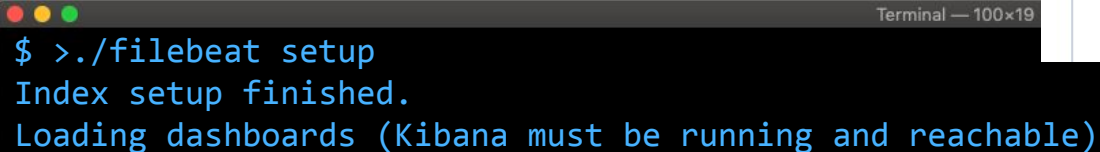
The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup  
./filebeat -e
```

First run the setup process

Setup preps dashboards and indices



A terminal window with a dark background and light blue text. The title bar shows three colored window control buttons (red, yellow, green) on the left and the text "Terminal — 100x19" on the right. The terminal content shows the command `$ >./filebeat setup` being executed, followed by the output `Index setup finished.` and `Loading dashboards (Kibana must be running and reachable)`.

```
$ >./filebeat setup  
Index setup finished.  
Loading dashboards (Kibana must be running and reachable)
```

4 Start Filebeat

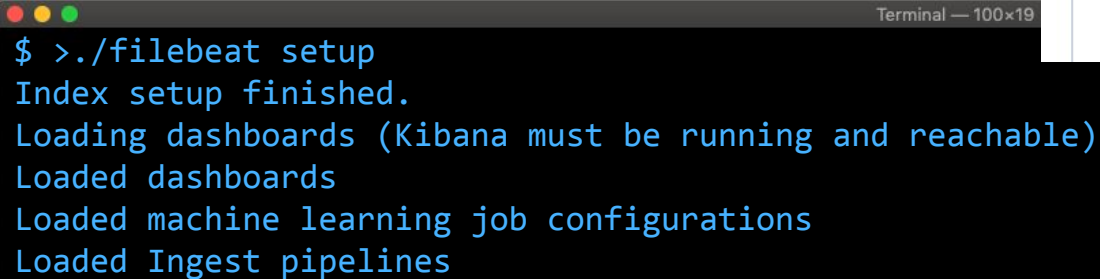
The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

[Copy snippet](#)

```
./filebeat setup  
./filebeat -e
```

First run the setup process

Setup preps dashboards and indices



Terminal — 100×19

```
$ >./filebeat setup
Index setup finished.
Loading dashboards (Kibana must be running and reachable)
Loaded dashboards
Loaded machine learning job configurations
Loaded Ingest pipelines
```

4 Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

[Copy snippet](#)

```
./filebeat setup
./filebeat -e
```

Finally, start it!

-e tells it to send messages to console



Terminal — 100×19

```
$ >./filebeat -e
```

4

Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup  
./filebeat -e
```

Finally, start it!

-e tells it to send messages to console

```
$ >./filebeat -e
```

```
2019-12-09T18:02:42.500Z INFO instance/beat.go:610Home path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64] Config path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64] Data path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64/data] Logs path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64/logs]
2019-12-09T18:02:42.501Z INFO instance/beat.go:618Beat ID: 04e276d0-79bd-40e3-9c83-3cdc4a64f791
2019-12-09T18:02:42.513Z INFO add_cloud_metadata/add_cloud_metadata.go:93 add_cloud_metadata:
hosting provider type detected as gcp,
metadata={"availability_zone":"us-east1-b","instance":{"id":"8271592631829869565","name":"user-smi
th-build"},"machine":{"type":"n1-standard-8"},"project":{"id":"elastic-product-marketing"},"provid
er":"gcp"}
2019-12-09T18:02:42.564Z INFO [seccomp] seccomp/seccomp.go:124 Syscall filter successfully
installed
(...)
```

4 Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

[Copy snippet](#)

```
./filebeat setup
./filebeat -e
```

Recall the earlier list

- [illegible]

Easy setup for variety of log sources

Add Data to Kibana

All [Logging](#) Metrics SIEM Sample data

Apache logs

Collect and parse access and error logs created by the Apache HTTP server.

Cloudwatch Logs

Collect Cloudwatch logs with Functionbeat



Elasticsearch logs

Collect and parse logs
created by Elasticsearch.

IIS logs

Collect and parse access and error logs created by the IIS HTTP server.



Kafka logs

Collect and parse logs created by Kafka.



Logstash logs

- Collect and parse debug and slow logs created by Logstash itself.



MySQL logs

Collect and parse error and slow logs created by MySQL.

Nats logs

Collect and parse logs created by Nats.



Nginx logs

Collect and parse access and error logs created by the Nginx HTTP server.



PostgreSQL logs

Collect and parse error and slow logs created by PostgreSQL.



Redis logs

Collect and parse error and slow logs created by Redis.

System logs

Collect and parse logs written by the local Syslog server.

Traefik logs

Collect and parse access logs created by the Traefik Proxy.

Needs for log analytics

Correlating and cross referencing

The screenshot displays a log analytics interface. On the left, a sidebar contains various icons for navigation. The main panel shows a list of log entries with columns for 'Timestamp' and 'Message'. The 'Log Rate' is indicated as 'BETA'. A search bar is present with the placeholder text 'Search for log entries... (e.g. host.name:host-1)'. A detailed view of a log entry is shown on the right, listing fields and their corresponding values.

Timestamp	Message
Jan 14, 2020 @ 08:37:08.790	[INFO] received ad request (context_words=[Cook
Jan 14, 2020 @ 08:37:08.790	[INFO] Cache miss for category: Cookware
Jan 14, 2020 @ 08:37:08.792	[redis.log][verbose] Accepted 10.48.4.11:36760
Jan 14, 2020 @ 08:37:08.800	[redis.log][verbose] Client closed connection
Jan 14, 2020 @ 08:37:08.811	[INFO] Adding 2 items to cache
Jan 14, 2020 @ 08:37:08.811	[INFO] Items 9081 now in cache
Jan 14, 2020 @ 08:37:08.811	[INFO] Returning 2 ads
Jan 14, 2020 @ 08:37:08.820	[INFO] received conversion request
Jan 14, 2020 @ 08:37:08.823	[INFO] conversion request successful
Jan 14, 2020 @ 08:37:08.829	[INFO] Getting supported currencies...
Jan 14, 2020 @ 08:37:08.836	[DEBUG] request complete
Jan 14, 2020 @ 08:37:08.838	[INFO] Adding 1 items to cache
Jan 14, 2020 @ 08:37:08.838	[INFO] Items 9082 now in cache
Jan 14, 2020 @ 08:37:08.838	[INFO] Returning 1 ads
Jan 14, 2020 @ 08:37:08.844	[DEBUG] request complete
Jan 14, 2020 @ 08:37:08.938	[DEBUG] request started
Jan 14, 2020 @ 08:37:08.946	[DEBUG] view user cart
Jan 14, 2020 @ 08:37:08.948	[INFO] GetCartAsync called with userId=\"59aee2
Jan 14, 2020 @ 08:37:09.280	[INFO] received conversion request
Jan 14, 2020 @ 08:37:09.304	[INFO] conversion request successful
Jan 14, 2020 @ 08:37:09.340	[INFO] Getting supported currencies...
Jan 14, 2020 @ 08:37:09.347	[INFO] listing products
Jan 14, 2020 @ 08:37:09.452	[INFO] [Recv ListRecommendations] product_ids=[
Jan 14, 2020 @ 08:37:09.520	[INFO] Getting product with ID 6E92ZMYFFZ

Field	Value
@timestamp	2020-01-14T13:37:08.838Z
_id	1xVFpG8BvYbYXIs0t_E1
_index	filebeat-7.5.1-2020.01.14-000048
agent.ephemeral_id	4b0464e0-cf97-498c-a2bb-61cdbaba36b6
agent.hostname	filebeat-6xktz
agent.id	520ddfa4-f182-44c9-b919-2f344cd00ca5
agent.type	filebeat
agent.version	7.5.1
cloud.availability_zone	us-central1-a
cloud.instance.id	2567286355104662140
cloud.instance.name	gke-edn-prod-default-pool-ef9bba0b-5bqx
cloud.machine.type	n1-standard-8
cloud.project.id	elastic-product
cloud.provider	gcp

Needs for log analytics

Searching, filtering, and highlighting

Stream

Log Rate

BETA

Settings

Search for log entries... (e.g. host.name:host-1)

Customize

Highlights


01/14/2020 8:37:08 AM

Stream live

Timestamp	Message	kubernetes.container.name
Jan 14, 2020 @ 08:37:08.790	[INFO] received ad request (context_words=[Cookware])	adservice
Jan 14, 2020 @ 08:37:08.790	[INFO] Cache miss for category: Cookware	adservice
Jan 14, 2020 @ 08:37:08.792	[redis.log][verbose] Accepted 10.48.4.11:36760	redis-master
Jan 14, 2020 @ 08:37:08.800	[redis.log][verbose] Client closed connection	redis-master
Jan 14, 2020 @ 08:37:08.811	[INFO] Adding 2 items to cache	adservice
Jan 14, 2020 @ 08:37:08.811	[INFO] Items 9081 now in cache	adservice
Jan 14, 2020 @ 08:37:08.811	[INFO] Returning 2 ads	adservice
Jan 14, 2020 @ 08:37:08.820	[INFO] received conversion request	currencyservice
Jan 14, 2020 @ 08:37:08.823	[INFO] conversion request successful	currencyservice
Jan 14, 2020 @ 08:37:08.829	[INFO] Getting supported currencies...	currencyservice
Jan 14, 2020 @ 08:37:08.836	[DEBUG] request complete	frontend
Jan 14, 2020 @ 08:37:08.838	[INFO] Adding 1 items to cache	adservice
Jan 14, 2020 @ 08:37:08.838	[INFO] Items 9082 now in cache	adservice
Jan 14, 2020 @ 08:37:08.838	[INFO] Returning 1 ads	adservice
Jan 14, 2020 @ 08:37:08.844	[DEBUG] request complete	frontend
Jan 14, 2020 @ 08:37:08.938	[DEBUG] request started	frontend
Jan 14, 2020 @ 08:37:08.946	[DEBUG] view user cart	frontend
Jan 14, 2020 @ 08:37:08.948	[INFO] GetCartAsync called with userId=\"59aee2be-5279-449f-86d0-a40733b41bcd\"	cartservice
Jan 14, 2020 @ 08:37:09.280	[INFO] received conversion request	currencyservice
Jan 14, 2020 @ 08:37:09.304	[INFO] conversion request successful	currencyservice
Jan 14, 2020 @ 08:37:09.340	[INFO] Getting supported currencies...	currencyservice
Jan 14, 2020 @ 08:37:09.347	[INFO] listing products	productcatalogservice
Jan 14, 2020 @ 08:37:09.452	[INFO] [Recv ListRecommendations] product_ids=[u'6E92ZMYFFZ', u'0PUK6V6EV0', u'2ZYFJ3GM2N', u'9SIQT8T0JO', u'0LJCSPC7Z']	recommendationservice
Jan 14, 2020 @ 08:37:09.520	[INFO] Getting product with ID 6E92ZMYFFZ	productcatalogservice

Needs for log analytics

Visualize



Needs for log analytics

Visualize




D

Dashboard / [Filebeat System] SSH login attempts ECS




SSH login attempts [Filebeat System] ECS



Successful SSH logins [Filebeat System] ECS



SSH users of failed login attempts [Filebeat System] ECS




SSH failed login attempts source locations [Filebeat System] ECS



Needs for log analytics

Visualize



Needs for log analytics

Flexible retention

 Kibana

 Elasticsearch

- Index Management
 - [Index Lifecycle Policies](#)
 - Rollup Jobs
 - Transforms
 - Watcher
 - Snapshot and Restore
 - 8.0 Upgrade Assistant
-  Kibana
 - Index Patterns
 - Saved Objects
 - Spaces
 - Reporting
 - Advanced Settings
-  Logstash
 - Pipelines
-  Beats
 - Central Management
-  Machine Learning
 - Jobs list
-  Security
 - Users
 - Roles

Edit index lifecycle policy filebeat-7.5.1

Use an index policy to automate the four phases of the index lifecycle, from actively writing to the index to deleting it. [Learn about the index lifecycle.](#)

You are editing an existing policy. Any changes you make will affect the indices that are attached to this policy. Alternatively, you can save these changes in a new policy.

☐ Save as new policy

Hot phase Active

This phase is required. You are actively querying and writing to your index. For faster updates, you can roll over the index when it gets too big or too old.

☒ Enable rollover

The new index created by rollover is added to the index alias and designated as the write index.
[Learn about rollover](#)

Maximum index size

Maximum documents

Maximum age


Index priority

Set the priority for recovering your indices after a node restart. Indices with higher priorities are recovered before indices with lower priorities. [Learn more](#)

Index priority (optional)

Needs for log analytics

Anomaly detection and alerting



From the earlier list

- [illegible]

Agenda

Beyond logging: Observability

- 1 Challenges with log analytics
- 2 Sending logs to Elasticsearch
- 3 Beyond logging: Observability
- 4 Leveraging Elastic security

You can add metrics in the same manner

Select your integration



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



SIEM

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

Upload data from log file

Import a CSV, NDJSON, or log file

Use Elasticsearch data

Connect to your Elasticsearch index

Visualize and Explore Data



APM

Automatically collect in-depth performance metrics and errors from inside your applications.



Canvas

Showcase your data in a pixel-perfect way.

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.

Many integrations

For example, system metrics



Add Data to Kibana

All Logging **Metrics** SIEM Sample data

The screenshot displays the Prometheus Metrics Explorer interface. At the top, there's a search bar with the text "network, and e host." and a "Filter" button. Below the search bar, a grid of 16 metrics is shown, each with an icon, a title, and a description. The metrics are arranged in a 4x4 grid. The first three metrics in the first row are "Aerospike metrics", "Apache metrics", and "AWS metrics". The first three metrics in the second row are "Ceph metrics", "CoreDNS metrics", and "Couchbase metrics". The first three metrics in the third row are "CouchDB metrics", "Docker metrics", and "Dropwizard metrics". The first three metrics in the fourth row are "Elasticsearch metrics", "Etcd metrics", and "Golang metrics". The first three metrics in the fifth row are "Kafka metrics", "Kibana metrics", and "Kubernetes metrics". The first three metrics in the sixth row are "Memcached metrics", "Microsoft SQL Server Metrics", and "MongoDB metrics".


Icon	Metric Name	Description
	Aerospike metrics	Fetch internal metrics from the Aerospike server.
	Apache metrics	Fetch internal metrics from the Apache 2 HTTP server.
	AWS metrics	Fetch monitoring metrics for EC2 instances from the AWS APIs and Cloudwatch.
	Ceph metrics	Fetch internal metrics from the Ceph server.
	CoreDNS metrics	Fetch monitoring metrics from the CoreDNS server.
	Couchbase metrics	Fetch internal metrics from Couchbase.
	CouchDB metrics	Fetch monitoring metrics from the CouchDB server.
	Docker metrics	Fetch metrics about your Docker containers.
	Dropwizard metrics	Fetch internal metrics from Dropwizard Java application.
	Elasticsearch metrics	Fetch internal metrics from Elasticsearch.
	Etcd metrics	Fetch internal metrics from the Etcd server.
	Golang metrics	Fetch internal metrics from a Golang app.
	Kafka metrics	Fetch internal metrics from the Kafka server.
	Kibana metrics	Fetch internal metrics from Kibana.
	Kubernetes metrics	Fetch metrics from your Kubernetes installation.
	Memcached metrics	Fetch internal metrics from the Memcached server.
	Microsoft SQL Server Metrics	Fetch monitoring metrics from a Microsoft SQL Server instance
	MongoDB metrics	Fetch internal metrics from MongoDB.

System metrics

Collect CPU, memory, network, and disk statistics from the host.

Metrics


Visualizing metrics




Metrics

Visualizing metrics


Dashboard / [Metricbeat System] Host overview ECS



● nice 0%
● irq 0%
● softirq 0%
● iowait 0%



Memory Usage [Metricbeat System] ECS




● Used 10GB
● Cache 4.9GB
● Free 1.1GB

Disk IO (Bytes) [Metricbeat System] ECS




Network Traffic (Packets) [Metricbeat System] ECS




● Inbound 39.3/s
● Outbound -39.37/s

Network Traffic (Bytes) [Metricbeat System] ECS



● Inbound 9.2KB/s
● Outbo... -21.2KB/s

Processes By Memory [Metricbeat System] ECS




Top Processes By CPU [Metricbeat System] ECS




Metrics

Exploring metrics




Metrics

Inventory view with multiple perspectives




Integrated Experience

Observability with one datastore



Application Performance Monitoring

Distributed Tracing



Setting up APM

Instructions in Kibana



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](#)



SIEM

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

Upload data from log file

Import a CSV, NDJSON, or log file

Use Elasticsearch data

Connect to your Elasticsearch index

Visualize and Explore Data



APM

Automatically collect in-depth performance metrics and errors from inside your applications.



Canvas

Showcase your data in a pixel-perfect way.



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.

Uptime Monitoring

Service availability

D

Uptime



Overview



Last 15 minutes

Show dates



Refresh

Up

Down

Location 0

Port 13

Scheme 3

1/22 monitors are down




● Down
● Up

1

21

Pings over time



Monitor status

Status	Name	URL	Downtime history	Integrations
● Up 8 minutes ago	adservice-7bcd956677-7ftmb	tcp://10.48.0.48:10000		...
● Down a few seconds ago	apm-server-969d845bc-pjc2d	http://10.48.3.115:8200/		...
● Up a few seconds ago	Unnamed - auto-http-OX4FA94B0313EE0BC4-78ebd16eaca0565d	https://www.bbc.com		...
● Up a few seconds ago	Unnamed - auto-http-OX4FA94B0313EE0BC4-c7eca2f6ea089820	https://github.com		...
● Up a few seconds ago	Unnamed - auto-http-OX4FA94B0313EE0BC4-debdfec4c62e9997	https://demo.elastic.co/status		...
● Up a few seconds ago	Unnamed - auto-http-OX4FA94B0313EE0BC4-f2bd00f715add916	https://www.elastic.co		...

Uptime Monitoring

Service availability

D

Uptime



a


7/33 monitors are down



● Down
● Up

7
25

Pings over time



Monitor status

Status	Name	URL	Downtime history	Integrations
● Up a few seconds ago	Unnamed - auto-http-0X14D5C52E77FA69FF	https://www.elastic.co/		... ▼
● Up a few seconds ago	Unnamed - auto-http-0X1BEDFC9AB574F394	http://192.168.64.11:3000		... ▼
● Down a few seconds ago	Website Monitor - Infra Error	https://www.elastic.co/products/infrastructure-monitoring		... ▼
● Up a few seconds ago	NodeJS	http://opbeans-node:3000/api/customers		... ▼
● Up a few seconds ago	NodeJS	http://opbeans-node:3000/api/stats		... ▼
● Up a few seconds ago	NodeJS	http://opbeans-node:3000/api/orders		... ▼
● Down a few seconds ago	SecurityContents	https://www.elastic.co/products/siem		... ▼
● Up	Unnamed - auto-http-0X418788B30A2375E3D	http://192.168.64.12:3000		... ▼


Uptime Monitoring

Integrated experience



Integrated Experience

Observability with one datastore



Integrated Experience

Observability with one datastore




Deployment Observability


Network

Traffic In

Traffic Out



CPU/Memory



Disk IO

Read

Write



Services

10

Containers

141

Errors

19413


Pods

70

- nginx
- redis
- mysql

Deployment
Status

63%



14%

CPU

68%

Memory

5%

Disk IO

Agenda

Securing your Beats

- 1 Challenges with log analytics
- 2 Sending logs to Elasticsearch
- 3 Beyond logging: Observability
- 4 Leveraging Elastic security

Recall the Filebeat steps

Use parameterized credentials

- Download and install Filebeat
- Edit the configuration
- Enable and configure the system module
- Start Filebeat

The screenshot shows the 'Getting Started' page for Filebeat on macOS in the Elastic Cloud interface. The page is divided into four numbered steps:

- 1 Download and install Filebeat**

First time using Filebeat? See the [Getting Started Guide](#). Copy snippet

```
curl -L -O https://artifacts.elastic.co/downloads/beats/filebeat/filebeat-7.5.0-darwin-x86_64.tar.gz
tar xzvf filebeat-7.5.0-darwin-x86_64.tar.gz
cd filebeat-7.5.0-darwin-x86_64/
```
- 2 Edit the configuration**

Modify `filebeat.yml` to set the connection information for Elastic Cloud: Copy snippet

```
cloud.id: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "elastic:<password>"
```

Where `<password>` is the password of the `elastic` user.
- 3 Enable and configure the system module**

From the installation directory, run: Copy snippet

```
./filebeat modules enable system
```

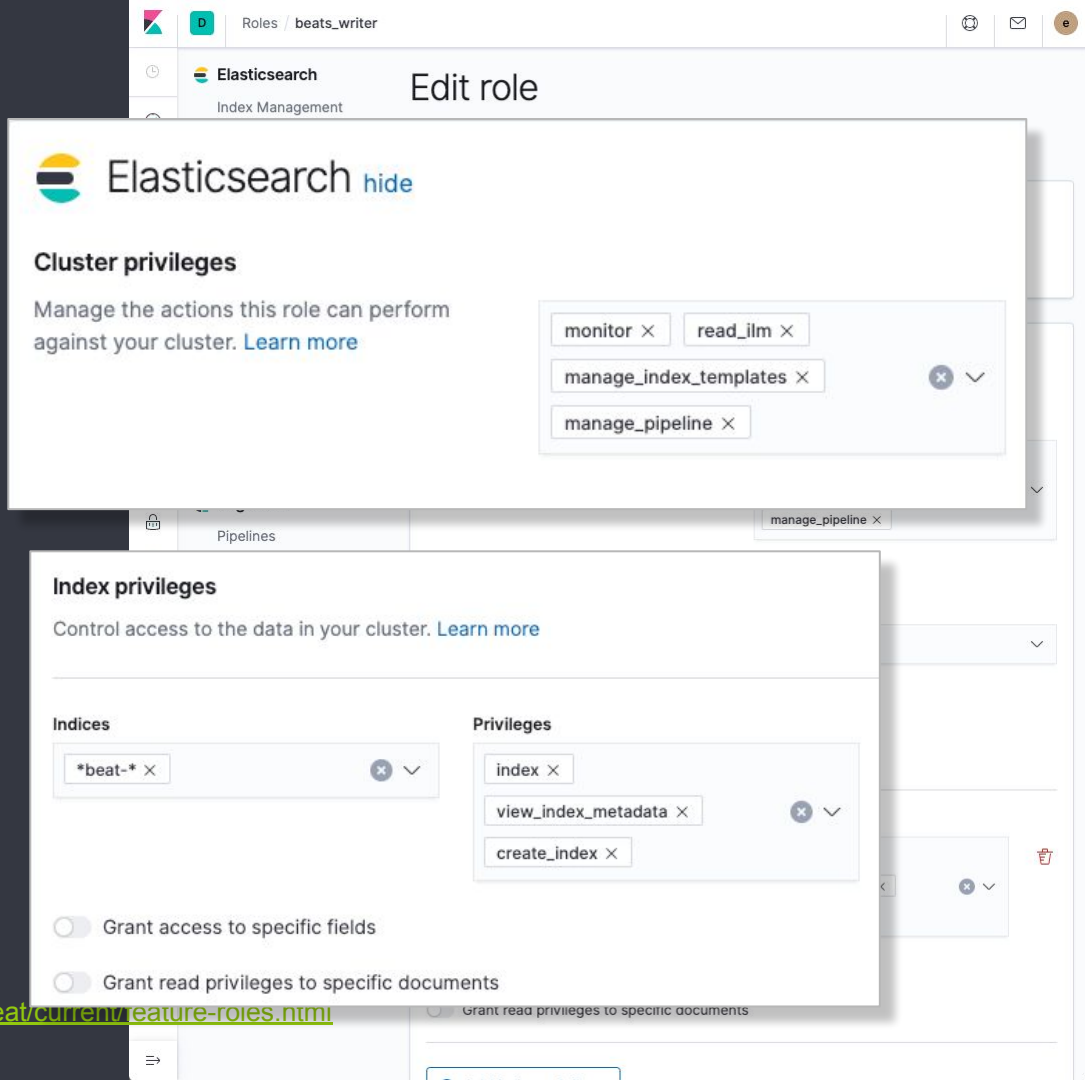
Modify the settings in the `modules.d/system.yml` file.
- 4 Start Filebeat**

beats_writer Role

Required permissions

- Cluster Permissions:
 - monitor
 - read_ilm
 - manage_index_templates
 - manage_pipeline
- Index Privileges (*beat-*)
 - create_index
 - index
 - view_index_metadata

<https://www.elastic.co/guide/en/beats/filebeat/current/feature-roles.html>



Corresponding User

Tying roles to users

- Give the user the corresponding roles
- Create a secure password
- beats-writer gets the writer role we created, plus the shipped beats_system role

Management / Users

Elasticsearch

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Transforms
- Watcher
- Snapshot and Restore
- 8.0 Upgrade Assistant

Kibana

- Index Patterns
- Saved Objects

Users

Create user

beats

☐	Full Name ↑	User Name	Email Address	Roles	Reserved
☐	Beats Writer	beats-writer	na@na.com	beats_system, beats_writer	

Rows per page: 20

Users

Create user

beats

☐	Full Name ↑	User Name	Email Address	Roles	Reserved
☐	Beats Writer	beats-writer	na@na.com	beats_system, beats_writer	

Security

[Users](#)

Roles

API Keys



Set up the keystore

Hiding credentials for beats-writer



Terminal — 100x19

```
$ >./filebeat keystore
```

```
Manage secrets keystore
```

```
Usage:
```

```
filebeat keystore [command]
```


```
Available Commands:
```

add	Add secret
create	Create keystore
list	List keystore
remove	Remove secret

- Command: filebeat keystore
- Create the keystore
- filebeat keystore add:
 - BEATS_WRITER_USER
 - BEATS_WRITER_PASSWORD
- Access keys via `${KEY_NAME}`

Previous Configuration

Had the user & password hardcoded



A terminal window titled "Terminal — 100x19" with a menu bar "File Edit Options Buffers Tools Help". The terminal displays the configuration for filebeat.yml, showing comments and settings for Elastic Cloud. The settings include cloud.id and cloud.auth.

```
#===== Elastic Cloud =====  
# These settings simplify using Filebeat 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: "Sandbox:dXMtY2VudHJ..."  
cloud.auth: "elastic:long-random-password" # because we are using Elastic Cloud
```

-UU-:----F1 filebeat.yml

(YAML)

Parameterize the user

Had the user & password hardcoded

```
Terminal — 100x19
File Edit Options Buffers Tools Help

#===== Elastic Cloud =====
# These settings simplify using Filebeat 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: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "${BEATS_WRITER_USER}:long-random-password" # because we are using Elastic Cloud
```



And the password

Had the user & password hardcoded

```
Terminal — 100x19
File Edit Options Buffers Tools Help

#===== Elastic Cloud =====
# These settings simplify using Filebeat 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: "Sandbox:dXMtY2VudHJ..."
cloud.auth: "${BEATS_WRITER_USER}:${BEATS_WRITER_PASSWORD}" # because we are using Elastic Cloud
```



Starts the same way

Picks up the keystore



Terminal — 100×19

```
$ >./filebeat -e
```

4

Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup  
./filebeat -e
```

Finally, start it!

assumes that you've run setup

```
$ >./filebeat -e
```

Terminal — 100x19

```
2019-12-09T18:02:42.500Z INFO instance/beat.go:610Home path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64] Config path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64] Data path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64/data] Logs path:
[/home/user/logs-demo/filebeat-7.5.0-linux-x86_64/logs]
2019-12-09T18:02:42.501Z INFO instance/beat.go:618Beat ID: 04e276d0-79bd-40e3-9c83-3cdc4a64f791
2019-12-09T18:02:42.513Z INFO add_cloud_metadata/add_cloud_metadata.go:93 add_cloud_metadata:
hosting provider type detected as gcp,
metadata={"availability_zone":"us-east1-b","instance":{"id":"8271592631829869565","name":"user-smi
th-build"},"machine":{"type":"n1-standard-8"},"project":{"id":"elastic-product-marketing"},"provid
er":"gcp"}
2019-12-09T18:02:42.564Z INFO [seccomp] seccomp/seccomp.go:124 Syscall filter successfully
installed
(...)
```

4 Start Filebeat

The `setup` command loads the Kibana dashboards. If the dashboards are already set up, omit this command.

Copy snippet

```
./filebeat setup
./filebeat -e
```

Continuing your Journey

Where to find more information

- Spin up a cluster
 - Hosted: cloud.elastic.co
 - Self managed - elastic.co/downloads
- Explore live examples @ elastic.co/demos
- Watch webinars @ elastic.co/videos
- Chat with us @ Forums : <https://discuss.elastic.co/>
- Go deeper with documentation @ elastic.co/guide
- Sign up for training @ elastic.co/training
- Attend a local meetup or Elastic{ON}



Q & A

—
Thank you!

