Beats 动手实践

刘晓国
Elastic 社区布道师
2020年5月16日

elasticstack.blog.csdn.net

Beats 入门教程（一） https://elasticstack.blog.csdn.net/article/details/104432643
Beats 入门教程（二） https://elasticstack.blog.csdn.net/article/details/104473684

议程

- Elastic 简介
- Beats 是什么?
- 使用 Beats
  - Filebeat
  - Metricbeat

Elastic 概述

Elastic 产品生态

elastic

Beats 是什么？

轻量级的数据摄入器

elastic

- Beats是一个轻量级的数据摄入器或代理，这些代理收集并运送各种运营数据到 Elasticsearch
- Beats使将数据输入到 Elasticsearch 变得容易
- 在许多操作系统上可用，例如Debian，Redhat，Linux和Mac

https://www.elastic.co/products/beats

Beats 家族

Packetbeat
Network data

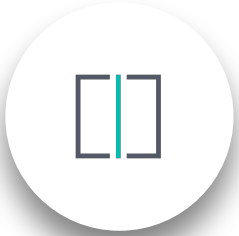

Auditbeat
Audit data

Metricbeat
Metrics

Filebeat
Log files

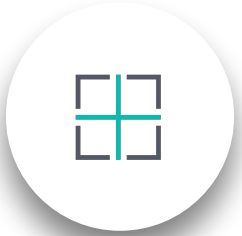

Winlogbeat
Windows Event Logs

Heartbeat
Uptime monitoring

Functionbeat
Serverless shipper for cloud

+90 community Beats

8

elastic

Beats 是如何接入到Elasticsearch中的?

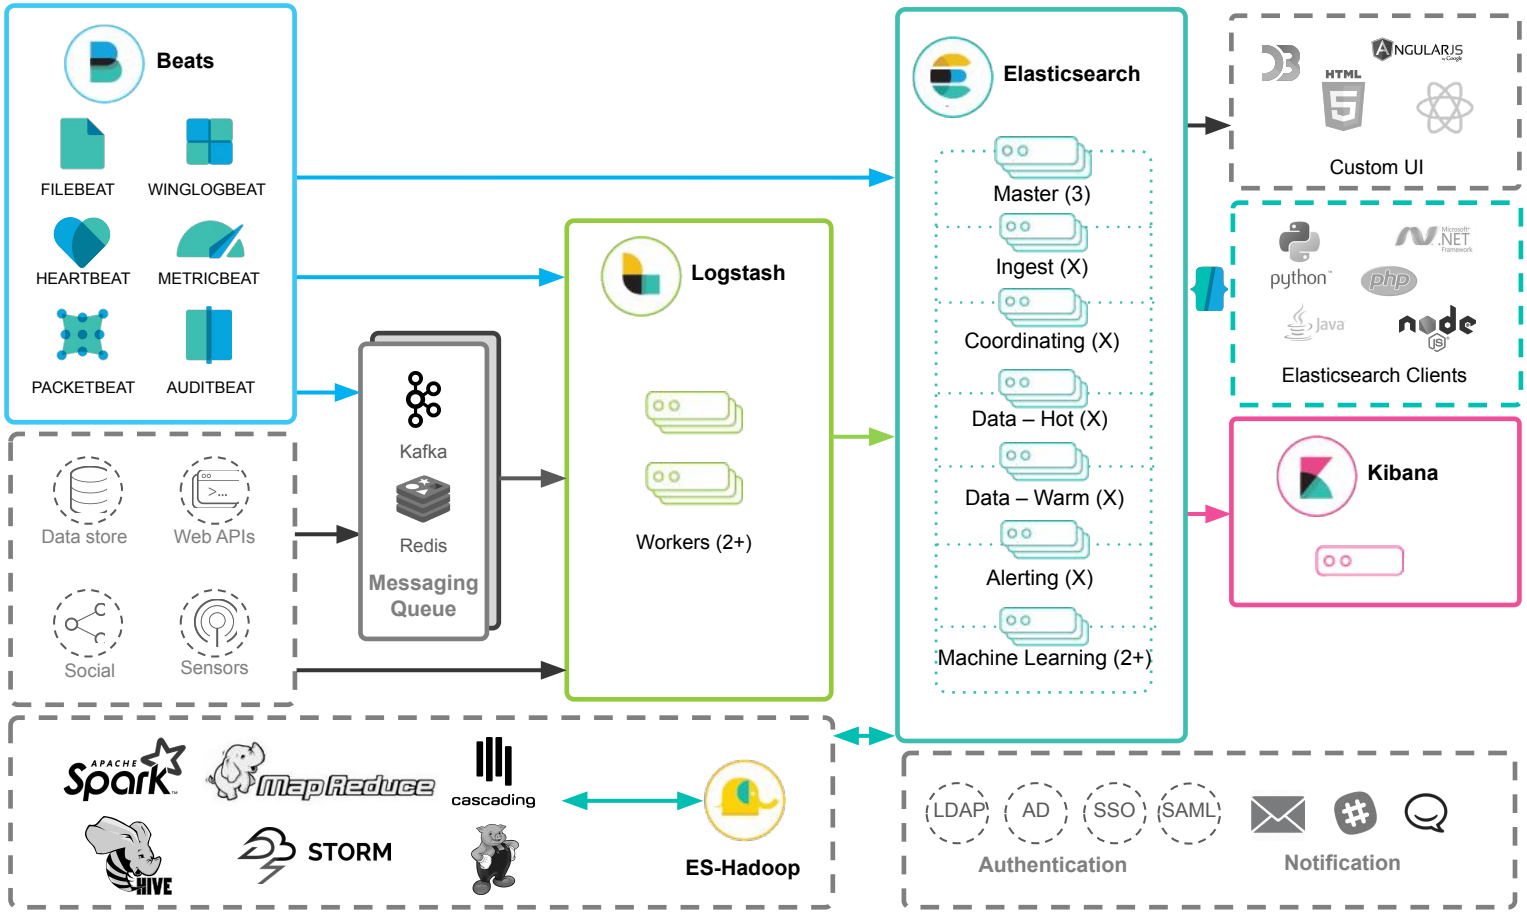

使用 Beats

Filebeat

- 正确处理日志轮转
- 背压机制
- 至少一次数据消费保证
- 结构化日志
- 多行事件处理
- 调节过滤

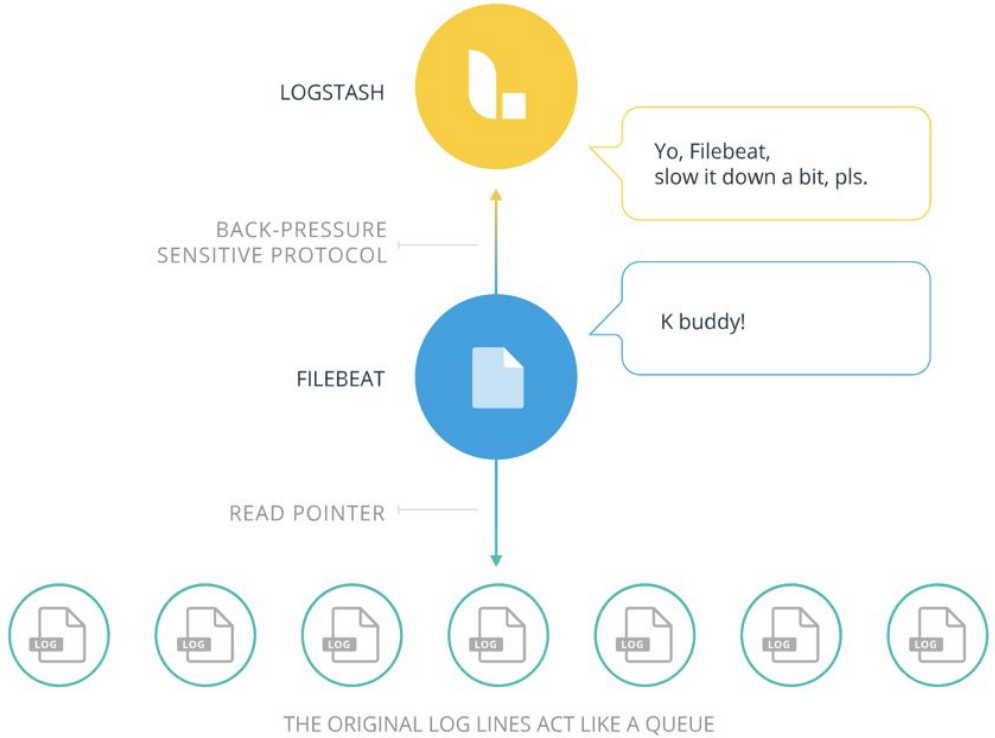

elastic

Filebeat 概述

- 开始一个或多个输入，查找你为日志数据指定的位置
- 对于Filebeat所找到的每个日志，Filebeat都会启动havester
- 每个havester都会读取一个日志以获取新内容，并将新日志数据发送到libbeat，libbeat会汇总事件并将汇总的数据发送到你为Filebeat配置的输出

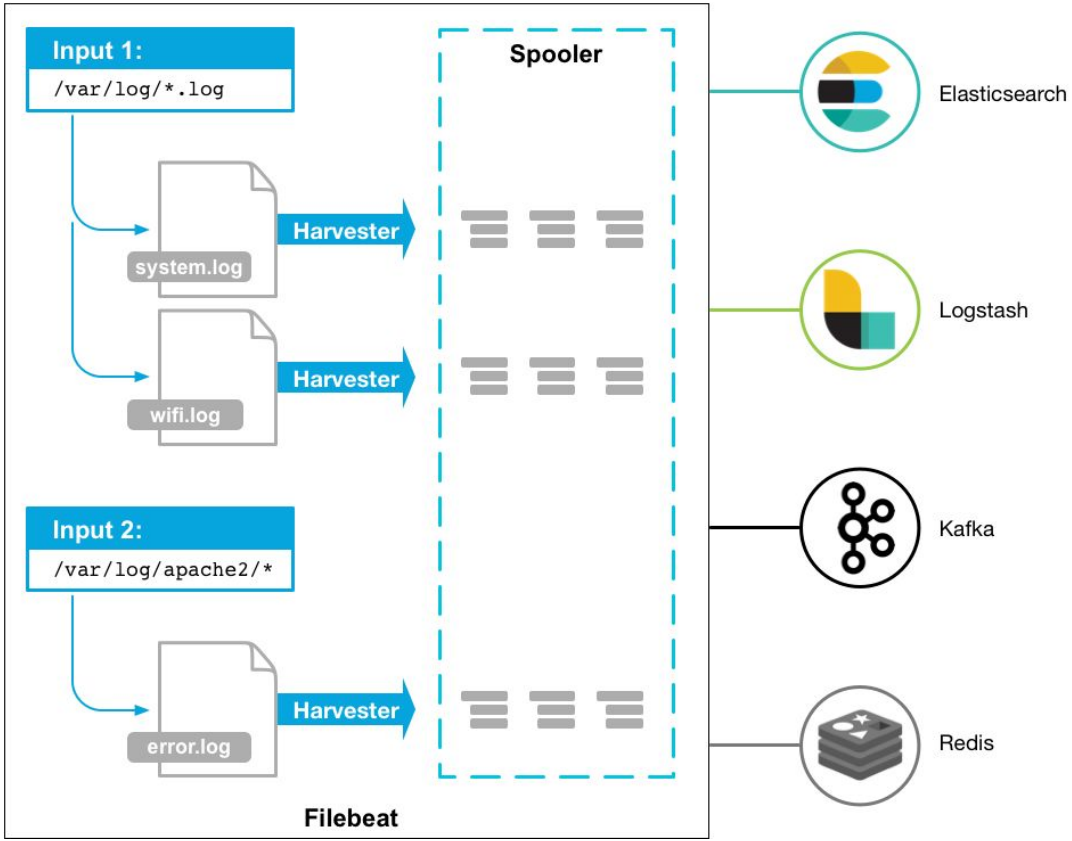

elastic

什么是 Filebeat 模块?

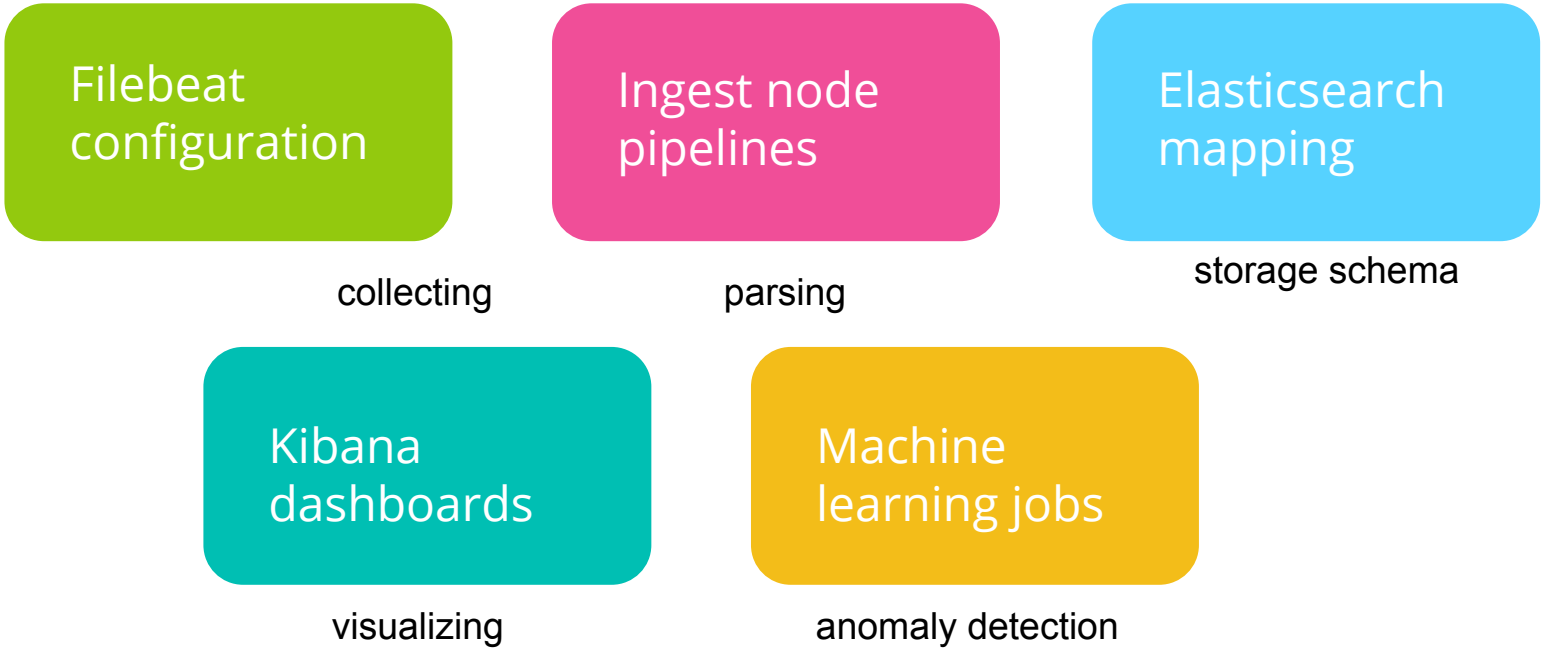

elastic

模块 - 收集,解析, 可视化为一体的预制方案

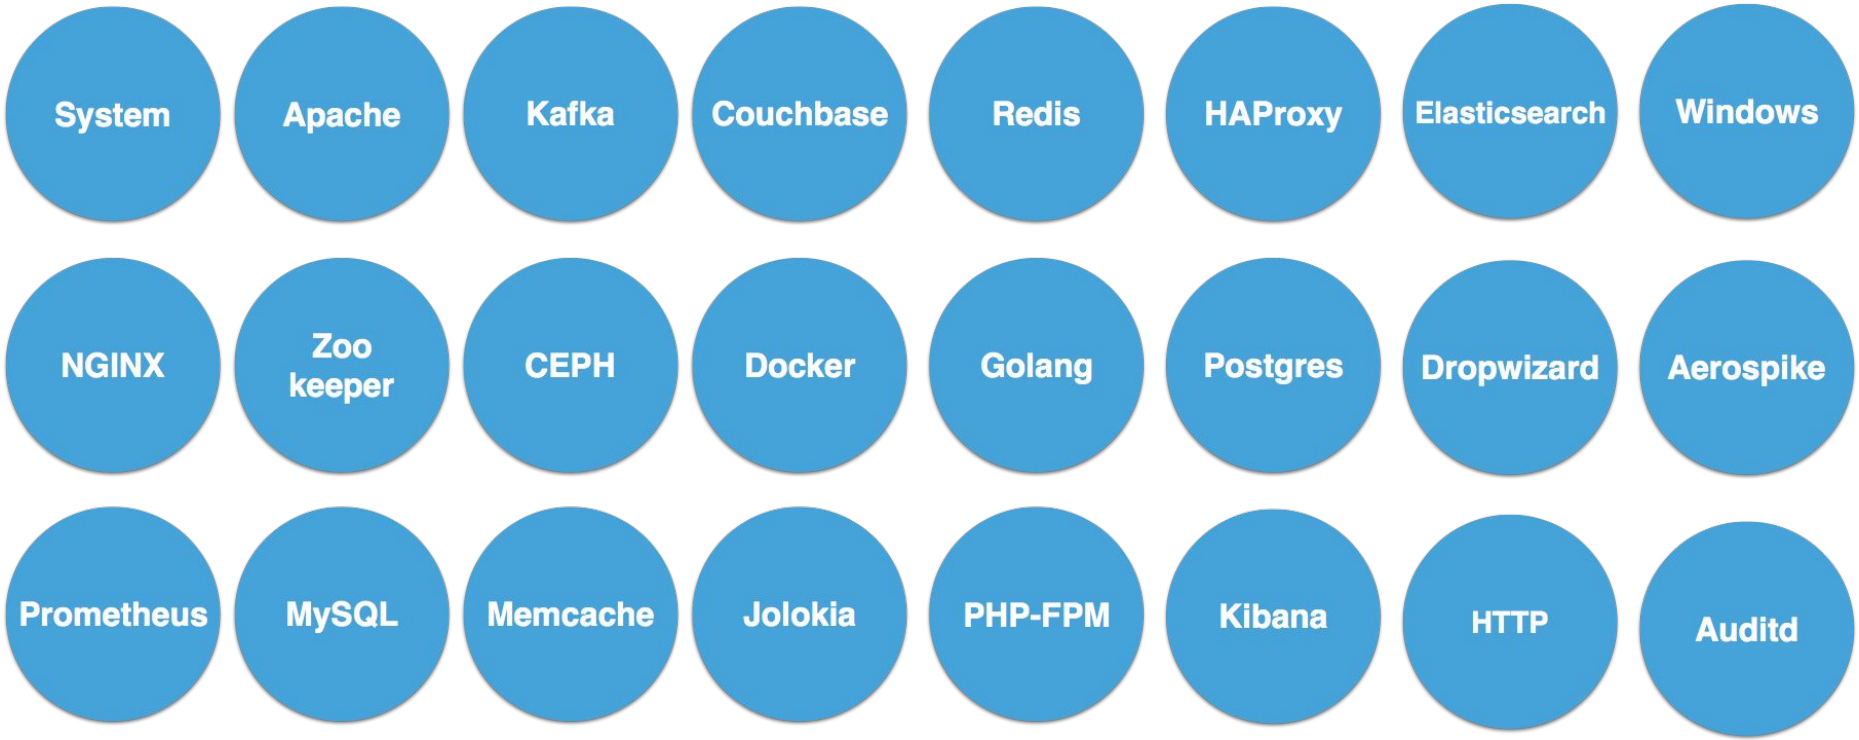

elastic

让我们来展示 Filebeat 吧

Installation for Filebeat

- Click on Kibana logo in the top left of Kibana, and select "**Add log data**"

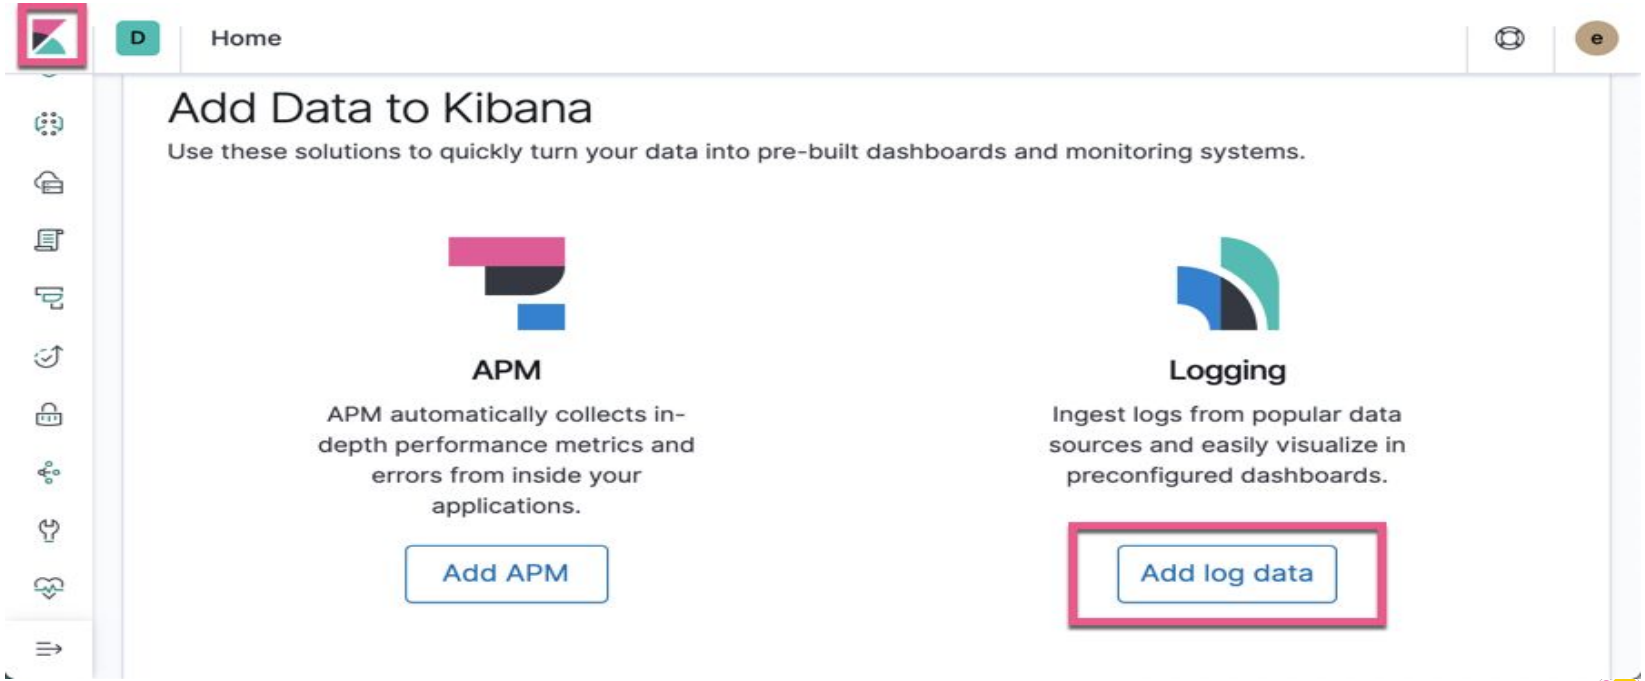

elastic

Select “System logs”

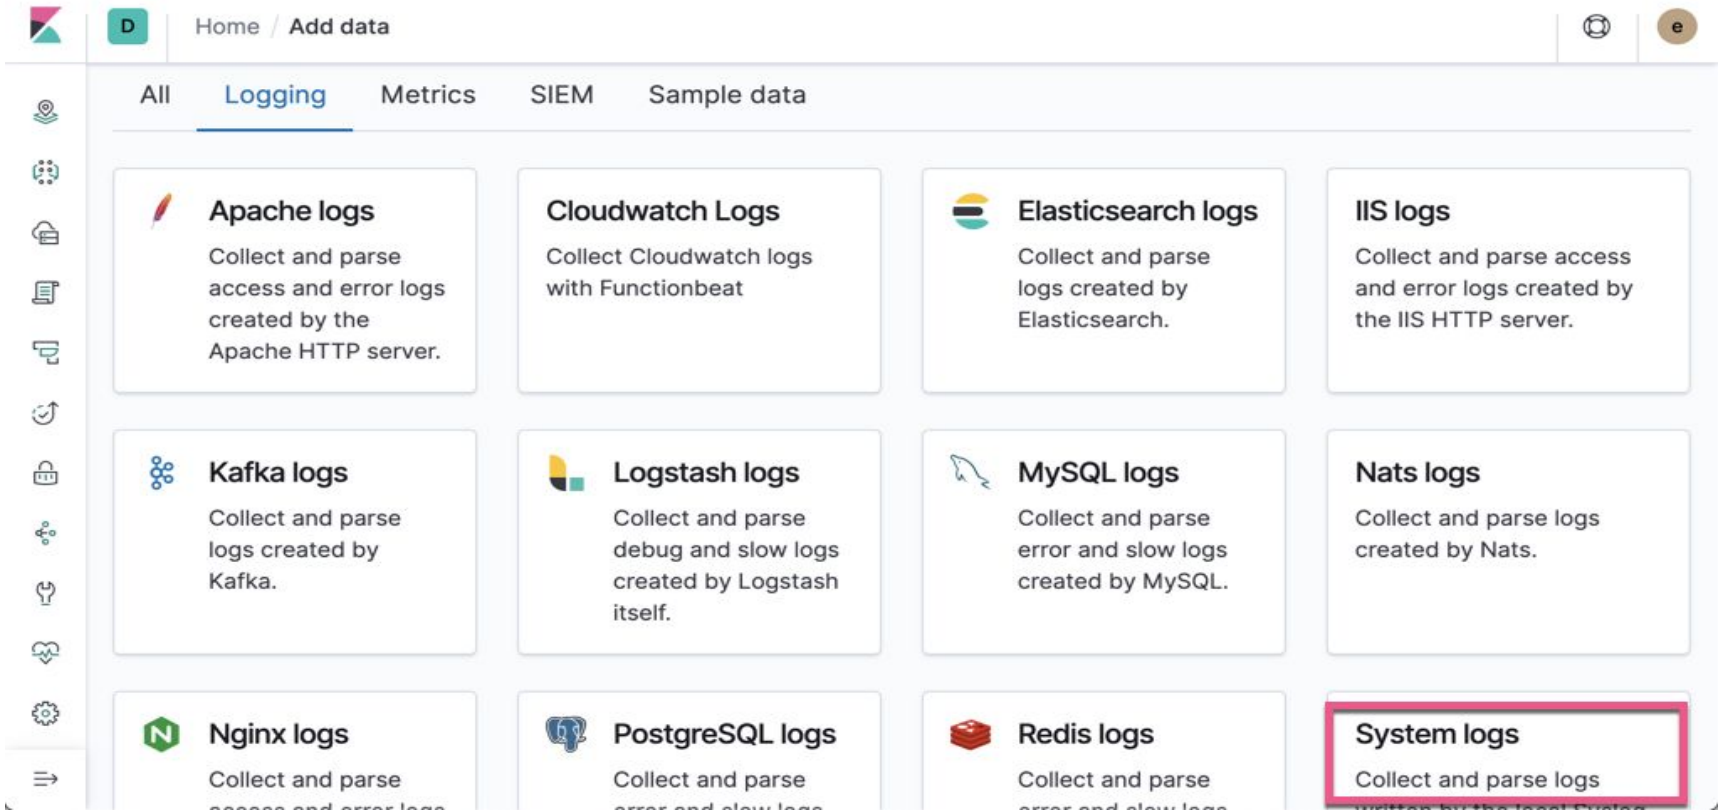

elastic

Select your platform

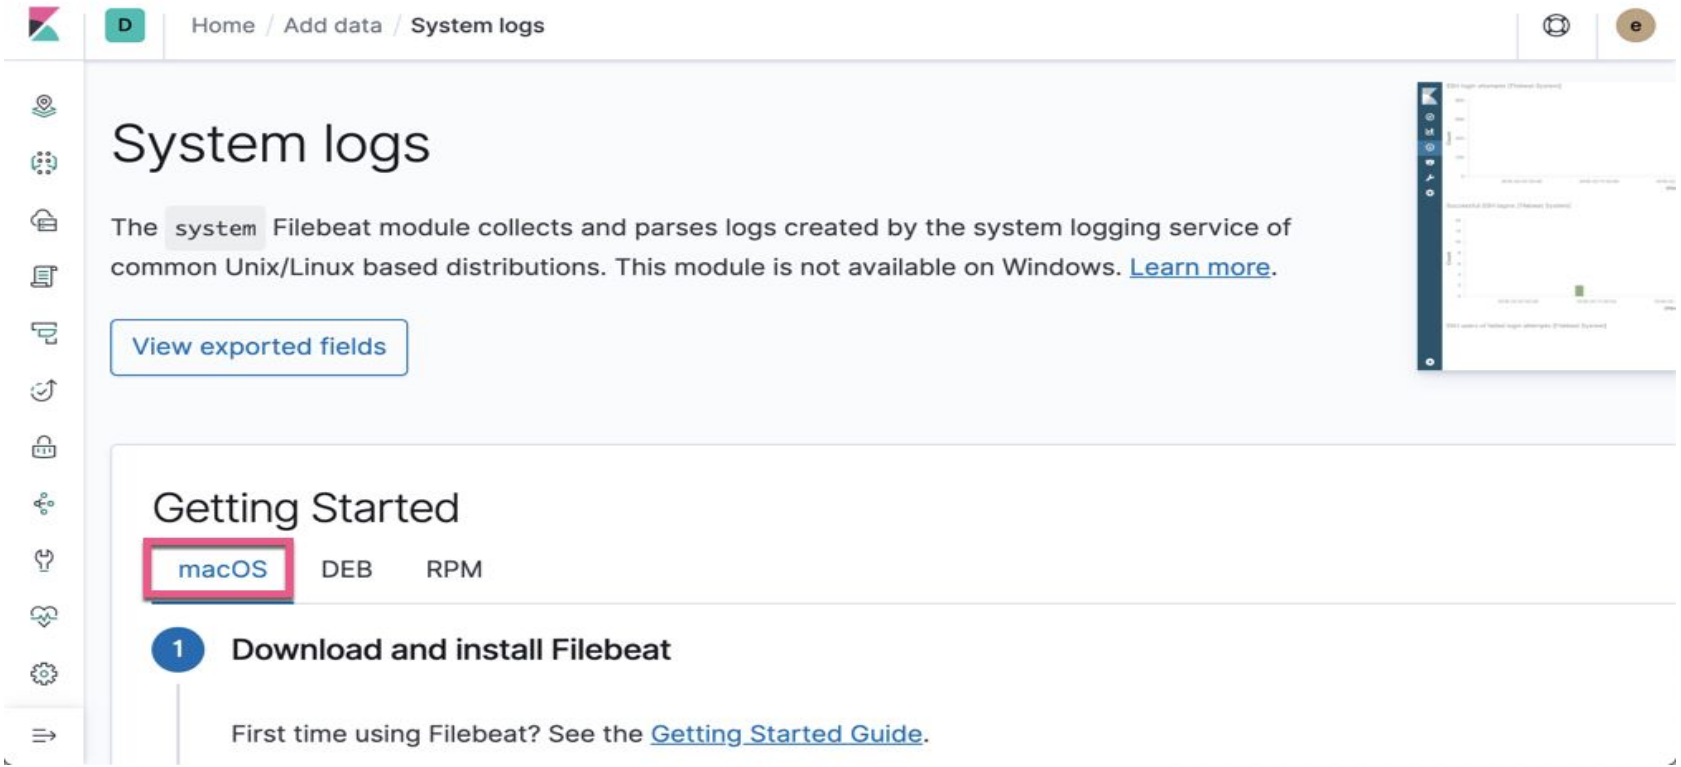

elastic

Filebeat commands - list modules

- List all of the available modules (enabled/disabled)
  - `./filebeat modules list`

```
liuxg-2:filebeat-7.4.2-darwin-x86_64 liuxg$ ./filebeat modules list
Enabled:
system
Disabled:
apache
auditd
aws
cef
cisco
coredns
elasticsearch
envoyproxy
googlecloud
haproxy
ibmmq
icinga
iis
iptables
kafka
kibana
logstash
```

Enabled nodules

Filebeat commands - enable/disable modules

- Enable modules
  - `./filebeat modules enable nginx apache`
- Disable modules
  - `./filebeat modules disable nginx apache`

```
liuxg-2:filebeat-7.4.2-darwin-x86_64 liuxg$ ./filebeat modules enable nginx a
pache
Module nginx is already enabled
Enabled apache
liuxg-2:filebeat-7.4.2-darwin-x86_64 liuxg$ ./filebeat modules disable nginx
apache
Disabled nginx
Disabled apache
```

Filebeat commands - setup and run

- Set up the Kibana dashboards
  - `./filebeat setup`
- Run Filebeat
  - `./filebeat -e`
- Run Filebeat for a customized filebeat configuration file
  - `./filebeat -e -c myfilebeatconfig.yml`
- To view the published transactions, you can start Filebeat with the publish selector like this
  - `./filebeat -e -d "publish"`
- If you want all the debugging output (fair warning, it's quite a lot), you can use *
  - `filebeat -e -d "*"`
- Test config and output
  - `./filebeat config`
  - `./filebeat output`

How to configure filebeat modules?

- Find the .yml file under the sub-dir **modules.d** and edit it

```
liuxg-2:filebeat-7.4.2-darwin-x86_64 liuxg$ ls
LICENSE.txt             fields.yml              kibana
NOTICE.txt              filebeat                logs
README.md               filebeat.reference.yml  module
data                    filebeat.yml            modules.d
liuxg-2:filebeat-7.4.2-darwin-x86_64 liuxg$ ls modules.d
apache.yml.disabled         mongodb.yml.disabled
auditd.yml.disabled         mssql.yml.disabled
aws.yml.disabled            mysql.yml.disabled
cef.yml.disabled            nats.yml.disabled
cisco.yml.disabled          netflow.yml.disabled
coredns.yml.disabled        nginx.yml.disabled
elasticsearch.yml.disabled  osquery.yml.disabled
envoyproxy.yml.disabled     panw.yml.disabled
googlecloud.yml.disabled    postgresql.yml.disabled
haproxy.yml.disabled        rabbitmq.yml.disabled
ibmmq.yml.disabled          redis.yml.disabled
icinga.yml.disabled         santa.yml.disabled
iis.yml.disabled            suricata.yml.disabled
iptables.yml.disabled       system.yml
kafka.yml.disabled          traefik.yml.disabled
kibana.yml.disabled         zeek.yml.disabled
```

Metricbeat

- 通过 API 获取服务的各项指标
- 高效地把数据存于到 Elasticsearch 之中
- 应用指标获取协议支持 JMX/Jolokia, Prometheus, Dropwizard, Graphite
- 智能标签（AWS, Docker）

Metricbeat 概述

- 由 modules 及 metricset 组成. 一个 Metricbeat module 定义了一个基本的逻辑来收集从特定服务，比如 Redis, MySQL 等 如何收集数据
- 每个 module 含有一个或更多的 metricset。一个 metricset 是 module 的一部分。它用于获取数据，并结构化数据。metricset 不是将每个指标采集作为一个单独的事件，而是获取一组多个相关的指标在一个请求中。

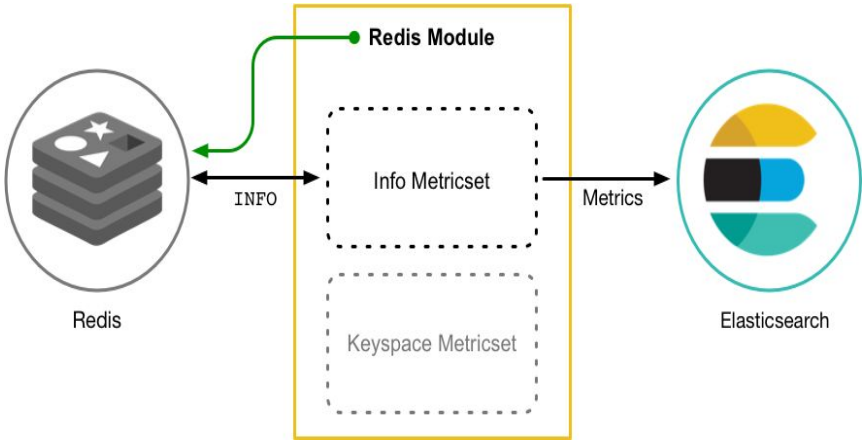

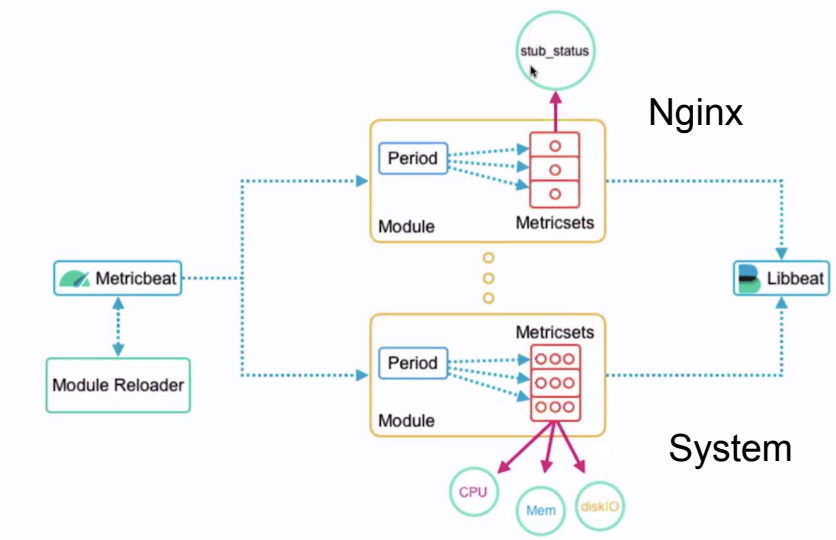

让我们来展示 Metricbeat 吧

Select “System metrics”

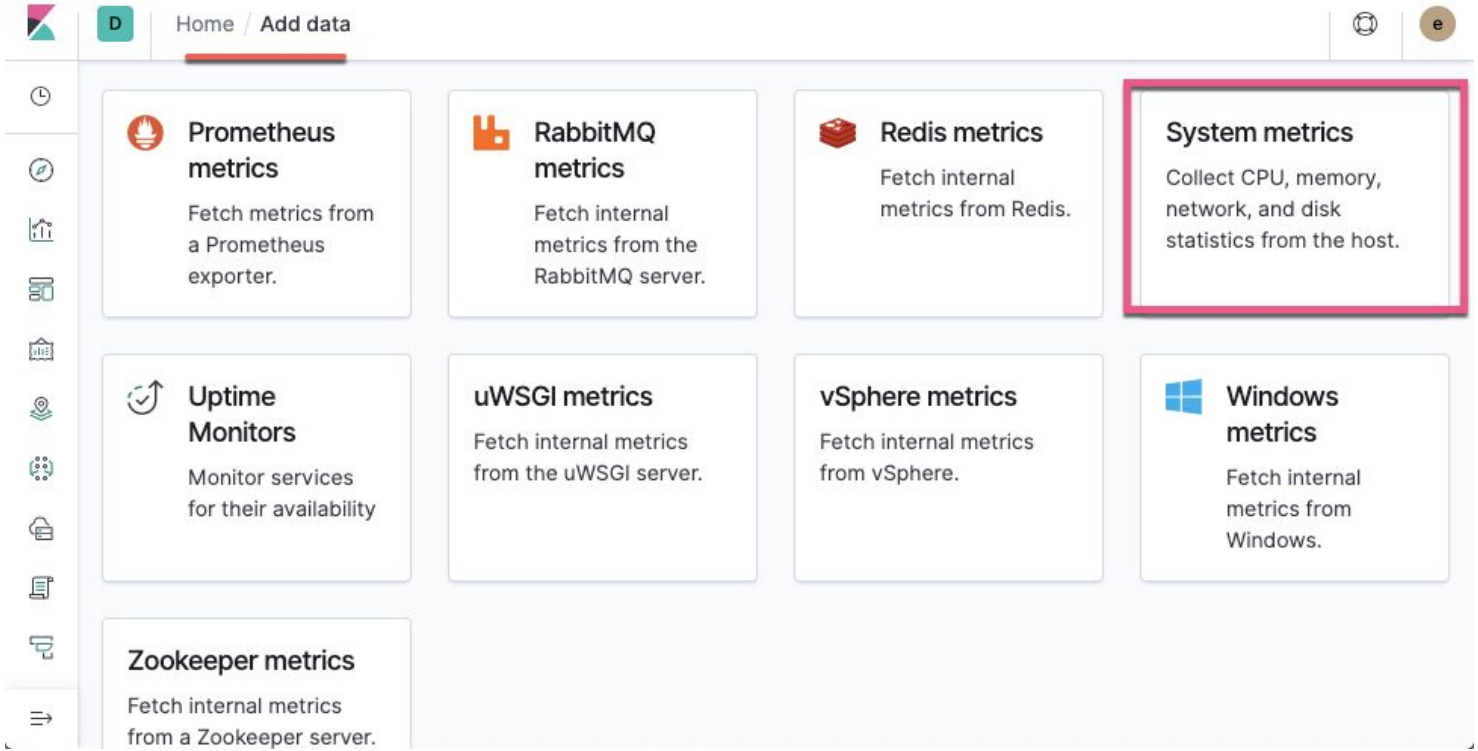

elastic

Select your platform

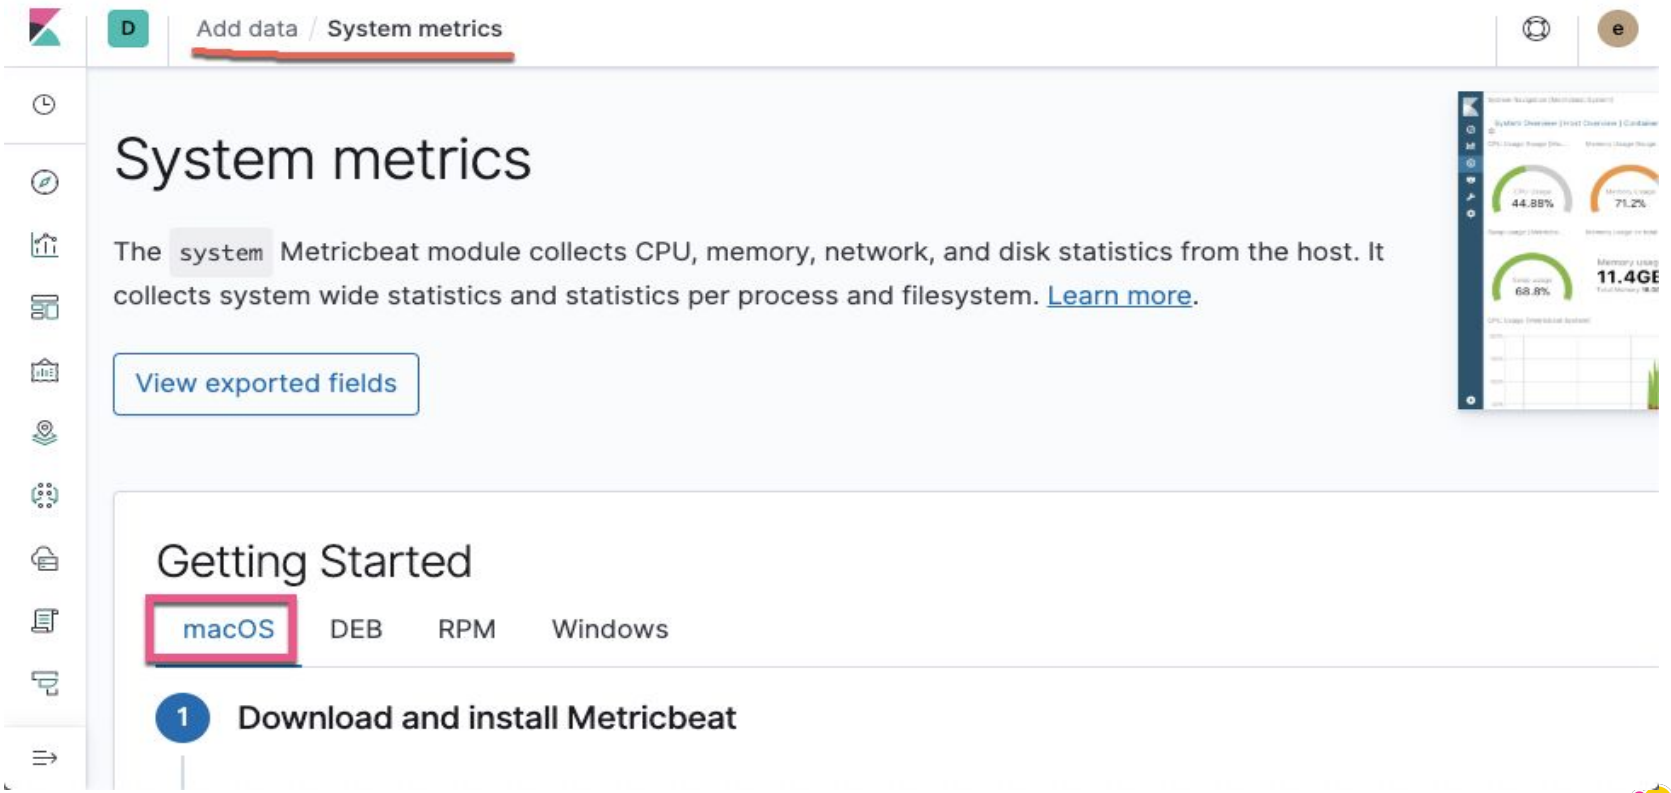

elastic

From Logs to Dashboards in One Command

elastic

Avoid password definition in metricbeat.yml

- Leaving your password in metricbeat.yml is risky sometimes since everybody can see it

```
#xpack.monitoring.elasticsearch:

cloud.id: "logs_dev:ZXVyb3BlLXdlc3QWI1YmNmNmFmZiRmZDdlYjc1MjQwMzI0N2M4OTk4NGY2MDI1NTgzODY5Mg=="
cloud.auth: "elastic:p6tD7eREdwi1azasTFRsJgqh"
```

- Key in the following command in your terminal
  `./metricbeat keystore create`
- Issue the following command, and paste your password
  `./metricbeat keystore add CLOUD_PWD`
- Change your metricbeat.yml to be like following using CLOUD_PWD

```
cloud.id: "logs_dev:YXAtbm9ydGhlYXN0LTEuYXdzLmZvdW5kLmlvJGE5MGNjYTgwODdiNTRlNTM5
ZmFkZDg5MjM2OTNiZmZkJDg0ZmJjYTI1NjM4ZDQwYjk5OWVjNDRjMzhlMDE1OGU3"
cloud.auth: "elastic:${CLOUD_PWD}"
```

Metricbeat commands - list modules

- List all of the modules
  - `./metricbeat modules list`

```
localhost:metricbeat-7.4.2-darwin-x86_64 liuxg$ ./metricbeat modules list
Enabled:
system

Disabled:
aerospike
apache
aws
beat
beat-xpack
ceph
cockroachdb
consul
coredns
couchbase
couchdb
docker
dropwizard
elasticsearch
elasticsearch-xpack
```

Metricbeat commands - enable/disable modules

- Enable modules
  - `./metricbeat modules enable nginx apache`
- Disable modules
  - `./metricbeat modules disable nginx apache`

```
localhost:metricbeat-7.4.2-darwin-x86_64 liuxg$ ./metricbeat modules enable
 nginx apache
Enabled nginx
Enabled apache
localhost:metricbeat-7.4.2-darwin-x86_64 liuxg$ ./metricbeat modules disabl
e nginx apache
Disabled nginx
Disabled apache
```

Metricbeat commands - setup and run

- Set up the Kibana dashboards
  - `./metricbeat setup`
- Run Metricbeat
  - `./metricbeat -e`
- Test modules
  - `./metricbeat test modules system`
  - `./metricbeat test config`
  - `./metricbeat test output`

How to configure metricbeat modules?

- Find the yml file under the sub-dir **modules.d** and edit it

```
localhost:metricbeat-7.4.2-darwin-x86_64 liuxg$ ls
LICENSE.txt              fields.yml               metricbeat.reference.yml
NOTICE.txt               kibana                   metricbeat.yml
README.md                logs                     module
data                     metricbeat               modules.d
localhost:metricbeat-7.4.2-darwin-x86_64 liuxg$ ls modules.d
aerospike.yml.disabled           kubernetes.yml.disabled
apache.yml.disabled              kvm.yml.disabled
aws.yml.disabled                 logstash-xpack.yml.disabled
beat-xpack.yml.disabled          logstash.yml.disabled
beat.yml.disabled                memcached.yml.disabled
ceph.yml.disabled                mongodb.yml.disabled
cockroachdb.yml.disabled         mssql.yml.disabled
consul.yml.disabled              munin.yml.disabled
coredns.yml.disabled             mysql.yml.disabled
couchbase.yml.disabled           nats.yml.disabled
couchdb.yml.disabled             nginx.yml.disabled
docker.yml.disabled              oracle.yml.disabled
dropwizard.yml.disabled          php_fpm.yml.disabled
elasticsearch-xpack.yml.disabled postgresql.yml.disabled
```

Thank you!