lo / hola! :-) Is there gonna be a recording of this webinar available?
BrandonPresenter
Yes Marcinsz, it'll be available with slides too.
marcinsz
Fantastic, thank you :-)
Guest62376
where is the link to webinar??
a1exus
Is there any ETA for fixing rabbitmq module inside of metricbeat?
a1exus
https://discuss.elastic.co/t/0-rabbitmq-nodes-no-results-found/119312
elasticguest4455
YouTube link to webinar: https://www.youtube.com/watch?v=UQbhBoBvGQ4
ythalorossy
I am using this link:
https://www.elastic.co/webinars/elasticsearch-log-collection-with-kubernetes-docker-and-containers
BrandonPresenter
a1exus, I'll remind the team to have a look at your reply. Thanks for raising it!
cdieringer
question: i have filebeat streaming streaming docker enhanced logs as just discussed. however, now i have an issue of huge log files and no rotation. is that a solved problem? otherwise, i need to instead stream logs to systemd, which handles this for me, and stream from there
a1exus
BrandonPresenter, thanks)
nestrada
k8s rotates logs
cdieringer
ya, unlike his example, im single host docker-composing* ATM, with migration to swarm soon
BrandonPresenter
cdieringer let's take that question at the end and see if Carlos has suggestions. Filebeat doesn't have any capability to rotate or delete logs - and that's on purpose.
chris12345
Single agent per node per concern? ie, 1 log beat container per worker node
BrandonPresenter
chris12345 I'm not sure I understand the question. Are you asking about how many Beats to have deployed and where?
nestrada
daemonset = 1 pod / node
elasticguest8051
how many beat processes will run in the pod, when I start 100 nginx pods?
nestrada
they're not sidecars, they're daemonsets

nestrada
so if you have 5 nodes, you'll have 5 fb pods
BrandonPresenter
Correct. As a daemonset per Beat, you would have one Filebeat pod on each node for instance.
elasticguest1668
0 beats processes run in the nginx pods; 1 beats pod runs per each k8s-worker node
chris12345
^^, that answers my question
elasticguest8051
So one process in Beat pod will scrape metrics from 100 nginx pods?
a1exus
i'm little bit confused how to enable autodiscovery for beats with docker..
BrandonPresenter
a1exus, You'll be setting the filebeat.autodiscover: configuration object (or Metricbeat) and you can see the full docs here:
https://www.elastic.co/guide/en/beats/filebeat/6.2/configuration-autodiscover.html
BrandonPresenter
This is in the filebeat.yml or metricbeat.yml
Ignacio314
I can hardly read the text, even when I maximize the youtube video.
elasticguest7080
@Ignacio314 have you tried to change the resolution of the youtube video?
a1exus
BrandonPresenter, thanks, great! i'll do some reading)
BrandonPresenter
No problem!
elasticguest9566
Do you already have pre-configured dashboards that displays Kubernets details?
nestrada
how difficult is it to implement own modules for autodiscovery?
skearns64
Ignacio314 - sometimes the youtube video defaults to a lower resolution - try clicking the gear icon in the lower right, and selecting 720p under Quality
a1exus
Ignacio314, your provider is limiting traffic to youtube (net-neutrality), try adjusting it to HD
elasticguest8051
Is there a JVM / JMX beat e.g. for Cassandra monitoring? How does metric beat find JMX port of cassandra?
elasticguest4455
Auto discovery Feature in Beats 6.2 (10 min video)
https://www.youtube.com/watch?v=1-iUoWGfByE&t=300s
chris12345

For time series data does elastic solve this problem completely? for example, with influx the concept of continuous queries and different retention policies allows for the capture of fine-grained to more coarse grained metrics
BrandonPresenter
elasticguest9566 We do have an overview dashboard for Metricbeat's module for K8S
dano99
can you please organize a presentation on containers that log to files with dynamic names (ie appfilelog-a3c87f.log or appname-20180215.log)
agup
For JMX https://www.elastic.co/blog/brewing-in-beats-add-support-for-jolokia-lmx
a1exus
I'm running JBOSS/wildfly, can I use Beats* to monitor ?
a1exus
pick me! pick me!))
elasticguest7267
what is pack licensing thing you mentioned in the beginning?
elasticguest7267
which part is free to use if I like it
elasticguest9223
does filebeat support multiline from docker log
dadak
how to monitor the application specific logs in kubernetes which are not written in the files but are getting written in standard console output?
nestrada
x-pack, monitoring is free
elasticguest8140
do you plan to support consul or/and nomad?
agup
elasticguest7267 all the opensource and anything under X-Pack basic license
elasticguest7267
so when do I have to pay?
Jconlon
Excellent presentation. Thanks.
elasticguest9574
I am using the 6.1 version of the Filebeat daemonset which used the "file" prospector and picked up messages from /var/lib/docker/containers/*/*.log
vjsamuel
@dadak, by default stdout/stderr are what are monitored by default. the docker prospector collects logs from only stdout/stderr
elasticguest8094
Is it possible to manage docker container logs on jelastic too?
elasticguest9574
what would be the advantage to use the docker prospector ?

elasticguest7367
q: Do you have any plans to add the "metadata processors" into the Logstash pipeline if e.g. a customer want to continue to use docker's fluentd log export -> logstash -> elasticsearch?
elasticguest9097
Signing off... Thank you both for this very good presentation. Will definately play with it on my DEV environment today.
elasticguest2515
How did he build a histogram of response codes?
elasticguest7267
how would you parse the "log line" with log stash in this workflow?
a1exus
jboss/wildfly monitoring
a1exus
plz
nestrada
I also use file prospector and autodiscovery works fine
elasticguest845
Support for Docker swarm is available?
nestrada
- type: log
nestrada
  paths:
nestrada
  - /var/lib/docker/containers/*/*.log
nestrada
  json:
nestrada
    message_key: log
nestrada
    keys_under_root: true
nestrada
    add_error_key: true
nestrada
  processors:
nestrada
  - add_kubernetes_metadata:
nestrada
      in_cluster: true
nestrada
      namespace: ${POD_NAMESPACE}
nestrada
  - decode_json_fields:
nestrada

fields: ['log']

nestrada

target: ''

elasticguest845

using gelf to logstash is the right way or should we use gelf to filebeat is correct?

nestrada

k8s doesn't support gelf

dadak

how to separate a log of one application module from another if all the logs are getting written in the standard output?

cdieringer

thx

cdieringer

:)

elasticguest6383

great demo, thanks

elasticguest7267

@dadak its done via metadata

elasticguest7267

then you can query

dadak

great demo

dadak

:)

nestrada

thanks!

Exekias_

Thank you everyone for coming, will try to answer some questions now, but will have to leave soon, If I don't get to answer yours, please go to our discuss and post it there *:)*

BrandonPresenter

Thanks folks! Ok, now that we're off the air, Carlos and I will try to get back through the history here and pull out questions. For anyone that isn't sticking around right now, please feel free to reach out in the usual ways and we'll help with those questions.

elasticguest6033

really nice

elasticguest845

how to get the application specific logs running in docker into elk?

Jconlon

If I am using Docker Swarm instead of Kubernetics.   Will these tools still be helpful?

ythalorossy

Thanks a lot for this explanation.
elasticguest7614
Is there a best practice to follow for promoting between enviroments: dev/test/prod.   Especially for rollback configuration
exekias_
For Docker Swarm: we support it at the Docker level, but we don't have a module for Swarm at the moment, it may be interesting to add it though
exekias_
the Docker module gives you a lot already, and Autodiscover should work there too
BrandonPresenter
a1exus I'm not very familiar with JBOSS/wildfly but my impression is that these can generate log files which you would soak up using Filebeat - in the most generic situation. I would think then that you can simply configure your docker/k8s deployment of Filebeat to pull in the relevant log files?
elasticguest7614
Most specifically for Kibana dashboard
NateS
Thanks so much, awesome presentation! Learned a lot!
elasticguest1668
Can the Beats send to kafka rather than Elastic (so I can consume from kafka using e.g. Logstash later to get to Elastic)?
a1exus
BrandonPresenter, I'd like to monitor JVM status/performance not so much as filebeat logs
nestrada
is it possible to use filebeat 6.2 with elasticsearch 5.6.x? filebeat 6.1 works
NicolasG
thank you for presentation
exekias_
FB 6.2 should work with Elasticsearch 5.6.X, yes
exekias_
about Java apps, we have a jolokia module to get metrics from JMX
a1exus
i'll look into that, thanksw
nestrada
thanks *:)* great presentation thanks!
elasticguest3930
Do you think the auto discovery will work just the same way in DC/OS as in Kubernetes?
elasticguest845
the application logs resides inside the ephermal container and are not exposed to the host. in this configuration, how to the filebeat to load the logs to elasticsearch
BrandonPresenter
elasticguest8051
BrandonPresenter

elasticguest8051
BrandonPresenter
elasticguest7267 There are paid features based on the license level you need of Xpack. This is clearer in a table than a sentence, so have a look at this page:
https://www.elastic.co/subscriptions
BrandonPresenter
(Sorry, meant to copy just 7267)
BrandonPresenter
a1exus, this might be what you're looking for then:
https://www.elastic.co/guide/en/beats/metricbeat/6.2/metricbeat-module-jolokia.html
elasticguest7267
thanks
BrandonPresenter
elasticguest1668 Yes, Beats can send to Kafka, here's a doc page that should help:
https://www.elastic.co/guide/en/beats/filebeat/6.2/kafka-output.html
BrandonPresenter
There are similar pages on each Beats docs
elasticguest1668
thank you!
elasticguest8051
Do beats support InfluxDB?
BrandonPresenter
elasticguest8051 Beats don't send to InfluxDB themselves. However, you could send to Logstash and use the InfluxDB output plugin to achieve this. We do plan to add features to Elasticsearch over time which might make it better for you to use ES instead of Influx, but we definitely want you to have options either way: https://www.elastic.co/guide/en
BrandonPresenter
/logstash/current/plugins-outputs-influxdb.html
elasticguest8051
Thx
BrandonPresenter
NP
elasticguest3930
Thank you!
BrandonPresenter
elasticguest7614 We don't have one particular practice for dev/qa/prod rollout procedures for how people manage all the details of the stack. Instead, we have many ways you could achieve this based on your needs. We have seen many customers have success using source control for configurations and orchestration/containerization to handle deployment.
exekias_
Yep, one cool pattern I've seen there is using different indexes per namespace
exekias_
you leverage the metadata to inject the namespace name in the index you use to store logs

BrandonPresenter
elasticguest845 It sounds like maybe it would be best if the logs were exposed so they could be collected centrally. Perhaps Carlos can comment more if he's seen instances of other workarounds users have used.
elasticguest8051
Well, influx/chronograph has authentication, will you add authentication to X-PAck basic?