video ok

Marcel

same here... all good

elasticguest8940

ok for me

elasticguest1970

working for me also

elasticguest2625

Reload

elasticguest4772

I'm back at the Grand Canyon ....<echo> 'Hello ... hello ... </echo>
Video good ... echo bad

ShivNarayan

till now video is not played

elasticguest4303

@shivnarayan try refreshing the page ?

elasticguest4772

Good now.

elasticguest2625

Disable Ghostery or AdBlock.

ShivNarayan

ok fine

elasticguest4772

It's an honor having this presentation being given by Late Night talk show host Craig Ferguson
ⓘ ChanServ set mode +v colinsurprenant

DanielLord

In which product is more easy to parse data! Splunk or Logstash?

mysery

some times i have multiple ip like 1.1.1.1,1.1.1.2 ... that can do to?

FelipeTavares

I have problems to generate a custom mmdb, that I can use for private IPs. This worked fine with easier version of mmdb (on logstash 2.X) but since 6.22 it won't go fine. do you have something to recommend? #logstash #geoip

earlier version*

andris

Question - you mentioned the high performance of geoip because it's in memory. Is there a good way to assess the performance impact / overhead of adding a given filter/enrichment to logstash? Like, if i add new enrichment filters, how can i tell how much extra work it's making my logstash instance do?

sitaBryn

andris: set the 'id' field for your filter, then use the API to query logstash, it'll report the number of events processed by each filter, as well as the time spent processing them
https://www.elastic.co/guide/en/logstash/current/monitoring.html

andris

nice @sitaBryn - I did not know about this

a1exus

sitaBryn, are you from Elastic?

sitaBryn

Nope, just a well versed customer...

a1exus

gotcha

elasticguest4772

Where in the log files can I find the F-bombs that were dropped earlier in the presentation?

Martin_Wismer

just little note: Port 22 is for ssh / sushi; Port 23 is telnet

andris

Question: is there a good way to tell how big is "too big" for a dictionary to put in a flat file? (vs a database lookup, etc)

elasticguest6613

on the translate filter, can i have multiple destinations? so for a given input, translate it to geoip but other items from the json object in a dictionary

andris

@elasticguest6613 - why not just have a geoip filter followed by a translate filter in your pipeline?

elasticguest5031

I can only understand 10% what they are talking about

elasticguest6613

am under the impression that it might add additional processing time?

a1exus

is that part of filter?
or input?

FelipeTavares

filter

andris

i think all changes to pipeline filtering has the potential to add processing time....but it doesn't sound unreasonable to me to have two filters in a row like that.

a1exus

ok

elasticguest5340

so you don't submit an elastic query?

FelipeTavares

yes, the "query" is the query you use

a1exus

it'll take results from elasticsearch and modify event

andris

every incoming event that goes through the filter will trigger an elasticsearch query

elasticguest5340

oh, he just answered what i was asking ,

Mathieu

@andris If you use X-Pack, you can also visualize your pipelines from Kibana https://www.elastic.co/guide/en/logstash/6.3/logstash-pipeline-viewer.html

a1exus

not every, but if you set special condition, it'd do it only then

elasticguestmv

any example of using wildcards in the search?

andris

thanks @mathieu - I currently run ES on AWS elasticsearch as a service so I can't get X-Pack, unfortunately.

FelipeTavares

You guys could talk a bit about Redis cache filter no? ,

elasticguest5031

Is this about configuring logstash?

andris

yes - all of these are logstash filters

elasticguest5031

thk

thank you

elasticguest6613

ok - althought would've been more convenient to have multiple "destinations" generated based on a single input

thanks andris

elasticguest1970

does the JDBC streaming filter makes sense for syncing ES with a database (postgre, mysql)

elasticguest8726

@andris I believe Elastic's stance is to not use AWS's ES as a service since they bypass a lot of the security enhancements.

andris

@elasticguest8726 - I *think* that would be more like a JDBC input (I do a lot of this)

elasticguest5966

how do we tail the parsing as similar tail file linux?

Vicente

greetings, some page where you can find examples

andris

I don't fully understand the difference between JDBC streaming for enrichment vs JDBC static for enrichment.

Mathieu

@elasticguestmv

@elasticguestmv query with wildcards can simply be `myfield:cust*` or such

andris

@elasticguest8726 yeah - I know there's some animosity there....but the pricing is hard to beat.

elasticguest5031

Andris, is the language GROK or something?

ⓘ elasticguest2625 is now known as Vlad

elasticguestmv

thx

ⓘ Vlad is now known as Guest8447

andris

@elasticguest5031 - they differentiate between parsing and enrichment. there is a logstash grok filter which I'd imagine they'd call "parsing" not "enrichment"

Mathieu

"grok" is kind of an enriched regular expressions engine

ⓘ Frank is now known as Guest15949

elasticguest8886

for elasticsearch filter, where credentials are stored?

Mathieu

@andris correct. One can parse a long text field (e.g. a raw log line) with grok or dissect, to extract bits of data to separate fields.

andris

yup

elasticguest5031

Ok, so we are talking about "enrichment" NOT "parsing"?

andris

yeah - I think that's the topic today

elasticguest8886

can password be encrypted in JDBC settings?

elasticguest8684

for elasticsearch filter, can you specify index?

Amit

msg NickServ identify 1234Amit

Mathieu

yes. enrichment = take a field (e.g. IP) and enrich it (e.g. with geolocation) :-)

a1exus

Amit, change your password NOW

elasticguest5031

Thanks guys

Mathieu

@elasticguest8684 yes you can

elasticrob

@elasticguest8886 you can use the logstash secret store
https://www.elastic.co/guide/en/logstash/current/keystore.html

andris

ahh i see JDBC streaming vs static. JDBC static is like the flat file translate filter, but with a DB instead of a file to pull from.

Mathieu

Attributes supported by filter-elasticsearch are documented here: https://www.elastic.co/guide/en/logstash-versioned-plugins/current/v3.3.1-plugins-filters-elasticsea...

andris

Question: How can i tell the break point between where it makes more sense to use translate with a flat file dict vs JDBC static?

Mathieu

interesting ones: query or query_template, index mostly.

andris

Question: How can i tell the break point between where it makes more sense to use translate with a flat file dict vs JDBC static? (from a performance standpoint)

Mathieu

and 'fields'

elasticguest5031

Anybody knows a good link for "parsing"?

Mathieu

@andris if your data is in a database and likely to change (e.g. new customers every week), you may consider JDBC. If your data is pretty static (e.g. US zip codes refreshed a few times a year), then in memory makes sense

p4ulpc

i'm assuming the throughput of the JDBC might be a limiter (when let's say you're trying to enrich at around 8k EPS)

Mathieu

also if your data set is huge and you'll likely be querying for only a subset, JDBC makes sense

andris

@mathieu - yeah, that's what I'm trying to figure out. How big is "huge" I wonder.
totally get the use of JDBC streaming for stuff that changes a lot
since JDBC static and translate both support periodic updating - they seem really similar
(just different local data storage)

mysery

some know good integretion whit IDE like intelliJ for build good logstash json

a1exus

I possibly discovered a bug in Logstash, I posted it to discuss.elastic.co, haven't got any response yet.. what should I do next?

Mathieu

At this point, JDBC static vs translate will likely end up being a toss-up based on internal company decisions. Can you access that database? Is it easier to get a dump file?

elasticguest6613

@a1exus wait ,

@ycombinator

a1exus: do you have a link to your post?

a1exus

ycombinator, I do..
https://discuss.elastic.co/t/error-nomethoderror-undefined-method-truncate-for-iowriter-0x3e9b038-io...

Mathieu

Logstash is not JSON. It looks somewhat ruby-ish. I personally use https://github.com/robbles/logstash.vim :-)

mysery

@Mathieu thks for that.

a1exus

elasticguest6613, wait for what? I feel like nobody even knows about it, therefor I should tell...

andris

good point - we're a small enough startup where access isn't yet a problem. I think I gather translate with flat file is probably simple and not too bad for what i need. Sounds fast because it's in memory. If my number of dictionary elements is in the few hundred to few thousand range....seems like yaml dictionary for translate is the way to go.

Mathieu

@mysery it's not perfect. For example, no automated indent when opening a new block. But at least gives the proper visual cues
@andris Yes, a few thousand items sounds file for in memory

andris

thanks!

re syntax highlighting - this looks slick
https://github.com/nir0s/sublime-logstash-syntax-highlighter I use sublimetext so I'll have to check this out

elasticguest6613

a1exus - ok, apologies

Mathieu

security recommendation: don't use "logstash" and "logstash!!" for actual credentials on your database ;-)

elasticguest6613

still not figuring out difference between the two JDBC filters

andris

@elasticguest6613 i didn't get this at first but i do now.

jeffoule

is the geoip filter available with the base install or is it a plugin that i need to install before I can use it?

andris

streaming -> go look up in the external db each time

a1exus

no need for apologizes) i'd like to contribute back to open source community and make product better for everyone)

andris

static -> check the external db every once in a while, and cache to local. Use local cache for lookup

Mathieu

@elasticguest6613 JDBC static caches in memory at load time (and refreshes periodically) but the translation is made based on what's in memory
streaming makes an SQL query for every event

elasticguest6613

@andris - thanks! so in case i'm trying to synchronize data - how to ensure that it doesn't repopulate previous data in the streaming filter?

Mathieu

@jeffoule Included on install
(geoip)

andris

do you just want the db data to get into your ES database as documents? Or are you using data in the external DB to add fields/data to data coming in from another source? (elasticguest6613
)

jeffoule

ok thanks @Matthieu

ExcelliumSA

Thanks

Guest8447

So Logstash syntax is not JSON but Ruby?

Vicente

everybody, how do I solve this error: ERROR Status logger no log4j2 configuration file found. using default configuration...

achaussier

Thanks

@ycombinator

a1exus: just responded: https://discuss.elastic.co/t/error-nomethoderror-undefined-method-truncate-for-iowriter-0x3e9b038-io...

BL2

Would the Elasticsearch filter be a good solution to combining Cisco ISE authentication logs with say, netflow?

elasticguest1632

Any plans to incorporate the HTML plugin into future Kibana updates? It was there in 5.x but no longer in 6.x ... much better than Markdown widgets.

FelipeTavares

Question: I have problems to generate a custom mmdb, that I can use for private IPs. This worked fine with easier version of mmdb (on logstash 2.X) but since 6.22 it won't go fine. do you have something to recommend? #logstash #geoip

elasticguest6613

no its the former @andris

Mathieu

@Guest8447 No, it looks like ruby, but it's a Logstash-specific language

ricknc

how would you support encrichment from a mongodb

andris

you probably want the JDBC input then
https://www.elastic.co/guide/en/logstash/current/plugins-inputs-jdbc.html

FelipeTavares

Question: The Redis cache filter is also supported by you guys or it is only community made?

HokieKev

I need to do the equivalent of a query->sub query->sub query. In a database I'd do this with a stored procedure - could I implement that functionality by using log stash to store intermediate results of queries to a JDBC connected database and write the result back to Elasticsearch?

Vini

Is there an OOTB filter to remove http tags from the incoming stream?

elasticguest8752

I am experiencing very frequent interruption in the stream (voice & video)

elasticguest6613

@andris thanks will check it out

elasticguest4271

What about enrichment from an API?

andris

np - its pretty easy to get going. Basically you define a query to run regularly on the external db - and you define how the query can tell whats new data since the last time it was run
Do ES guys have comments on:

FelipeTavares

Ok!

andris

--when to use JDBC static vs flat file dictionary with translate?

Mathieu

"Discuss" being https://discuss.elastic.co :-)

elasticguest6613

@andris - ok, that's the challenge...need to think this through
ⓘ raja is now known as Guest88733

andris

How would i know what is "too big' for a dict in a flat file?

TiagoPinto

Thank you guys, my doubts are about performance on multiple jdbc streaming after a jdbc input of like 2M reccords. I guess i'll hit the discuss forum.

jeffoule

Question : is there a filter that can translate a postal code into latitude/longitude?

andris

oh that makes sense: re scaling JDBC static to multiple LS

elasticguest8386

How different is CIDR filter if I use IP data type mapping in Elasticearch?

sitaBryn

Any tips on using filters like the JDBC filter in a high-availability scenario? Multiple logstash instances would mean multiple identical queries, no?

Mathieu

@jeffoule If you get access to a postal code database (e.g. likely a CSV), you could use translate for that
If you have access to an API, you could use logstash-filter-http instead

andris

Thanks! That's a really useful discussion on JDBC static vs translate, appreciate it!

Kofi

I run JDBC_static filter in for my LS as an application run off a scheduled cron instead of as a service. So everytime it runs, it re creates the DB to cache each time it runs right? Is there any way I can get around this?

FaqESK

hi guys, I need to parse a message field , into this message field, the data is separate with | Character.... how is the best way to do it? thanks

Mathieu

To figure out which plugins are included by default with Logstash, you can start from your Logstash version here https://www.elastic.co/guide/en/logstash/current/index.html then look at the "input", "filter", "codec" and "output" filters.
If instead you want the full list from the community, you should start here https://www.elastic.co/guide/en/logstash-versioned-plugins/current/index.html

HokieKev

Can you use log stash to read from Kafka topics?

Mathieu

I meant "input", "filter", "codec" and "output" plugins. :-)

p4ulpc

re: kafka - yes https://www.elastic.co/guide/en/logstash/current/plugins-inputs-kafka.html

jamiguet

Can you perform a rest call as part of any of the filters?

mysery

if i have multiples IPs in same source, ejem: "127.0.1.1, 127.0.1.2", (if i don't have local i will take the IPS), i can filter whit one IP filter or i will need other filter first.?

p4ulpc

unless you mean as a filter plugin with kafka tables, in which case, that would be interesting

Mathieu

@jamiguet Nothing prevents that from happening per se. We don't have a logstash-filter-http, but this is something we have been discussing. Performance discussion about that is happening in audio ;-)

a1exus

https://www.elastic.co/blog/just_enough_redis_for_logstash

andris

super useful! Thanks to the presenters!

a1exus

thanks Christian and Boertje

jamiguet

Thanks a lot.

mysery

thx

TiagoPinto

thanks

gln

thnx

Sudheer

Thank you

elasticguest1947

thanks

jeffoule

thanks

Vini

thank you

mickmill54

thanks!

cantipop

I think you should make a webinar for big scale logging. Like we have multiple accounts in AWS, Azure, , private datacenter. And how can we have all the logs in a central location. Logstash-> kafka>elastic.. something like that. Best practices.. ok. thanks

jamiguet

@Mathieu I do agree that the caching is an issue and that the marshalling of the data to make the request may be non trivial. But the built in cache for free could be very appealing

Mathieu

@cantipop For that kind of discussion, you should get in touch with an Elastic solutions architect. There will be a lot of "it depends" in that discussion :-)

cantipop

@Mathieu you are right, there are many details to considers for this subject and many "if"s.

jamiguet

Many thanks to the presenters and many thanks @Mathieu for your answers.

Mathieu

My pleasure :-)