

Centre de Calcul de l'Institut National de Physique Nucleaire et de Physique des Particules

ELASTICSEARCH

UPDATE


LOGS AND METRICS

- we need a backend for storing and querying
- ES is awesome for logs
- ES is awesome for metrics
- ES is awesome for everything

SECURITY FIRST

- everybody wants to see their logs
- everybody wants to see their metrics
- free ES is insecure
- X-pack (*fka* Shield) is expensive

Architecture



LOGS

- *Avg: 1k/s*
- *Peak: 20k/s*
- *Retention: 1 year*

METRICS

- *Avg: 15k/s*
- *Retention: 5 years (aggregated)*

TRANSPORT

- end-to-end encryption
- ES data nodes
- `syslog_ng` (logs)
- `riemann` (metrics)

AUTHENTICATION

- Kerberos / **GSSAPI SPNEGO** (API/CLI)
- Client Certificate (API/CLI)
- Web SSO (Browser)
 - **CAS**

AUTHORIZATION

- 2016: flat file (managed by puppet)
- 2017: LDAP (!)
- User / Team

PAST WORK WITH HELP FROM KEK

- SearchGuard v1 deployed on ES v1.7
- Kibana patch for dynamic index
- NodeJS proxy (ldap *facsimile*)
- apache CAS & KRB5 proxy

PAST CONCERNS

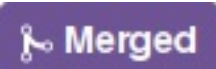

- not using SearchGuard-SSL
- floragunn (SearchGuard) commercial?
- elastic.co products moving too fast

minimize work for elastic stack version N+1

SEARCH GUARD

- floragunn is awesome
 - [free searchguard license](#)
 - [great support](#)
 - development catches up on ES faster than lightning

SYSLOG-NG

- syslog-ng ES support
 - [elastic-v2](#) destination (Balabit)
 - `client_mode(searchguard)` (CCIN2P3 )
 - `client_mode(https)` (CCIN2P3 )
 - two other PRs (bugfixes) to be merged ([1](#),[2](#))

What's new?



floragunn UG (haftungsbeschränkt)
Tempelhofer Ufer 16
109 63 Berlin
Germany
info@floragunn.com
+49 30 89379249

Invoice

INV-000153

Balance Due
€0,00

Bill To
CCIN2P3
21 av. Pierre de Coubertin, Campus de la Doua
Villeurbanne, 69 100
France

Invoice Date : 05/10/2016
Due Date : 05/10/2016

#	Item & Description	Qty	Rate	Amount
1	Search Guard® 2 Academic / Scientific Enterprise License SKU : SG2-EL-A Academic/Scientific license for Search Guard 2, compatible with Elasticsearch 2.x, for an unlimited amount of production clusters. Development, staging and test clusters included. Grants usage rights for all enterprise features for non-commercial projects. The license does not expire.	1,00	0,00	0,00
Sub Total				0,00
Total				€0,00

Notes
Your licence reference number: 05fd330e-8aee-11e6-a83b-e35f37c24a59. Thanks for using Search Guard.


Claudia Kressin, Managing Director


floragunn UG (haftungsbeschränkt), Tempelhofer Ufer 16, 109 63 Berlin, Germany
Managing Director: Claudia Kressin
Phone: +49 30 89379249, Email: info@floragunn.com
Registered at: Amtsgericht Charlottenburg, Company ID: HRB 147010 (VAT-ID: DE287373363, Tax-ID: 37/288/31435)
Deutsche Bank, IBAN: DE32 100 700 240 129 416 400, BIC: DEU103300000

LESSONS LEARNED

- Search Guard config is tricky but worth it
- No elaborate mappings: they'll be deprecated soonish
- Use REST in favor of Transport
- Monitor your cluster (heap usage, indexing rate, ...)
- Use bare-metal nodes only for data
- Use tiered cluster (*shard allocation awareness*)
- Use saner flush_interval (120s)
- Use store: false whenever possible (metrics)
- Throttle sources!

Nominal activity


error rate




syslog-ng queue




load




net rx




net tx




heap




syslog-ng




index rate




query rate




disk used




inodes



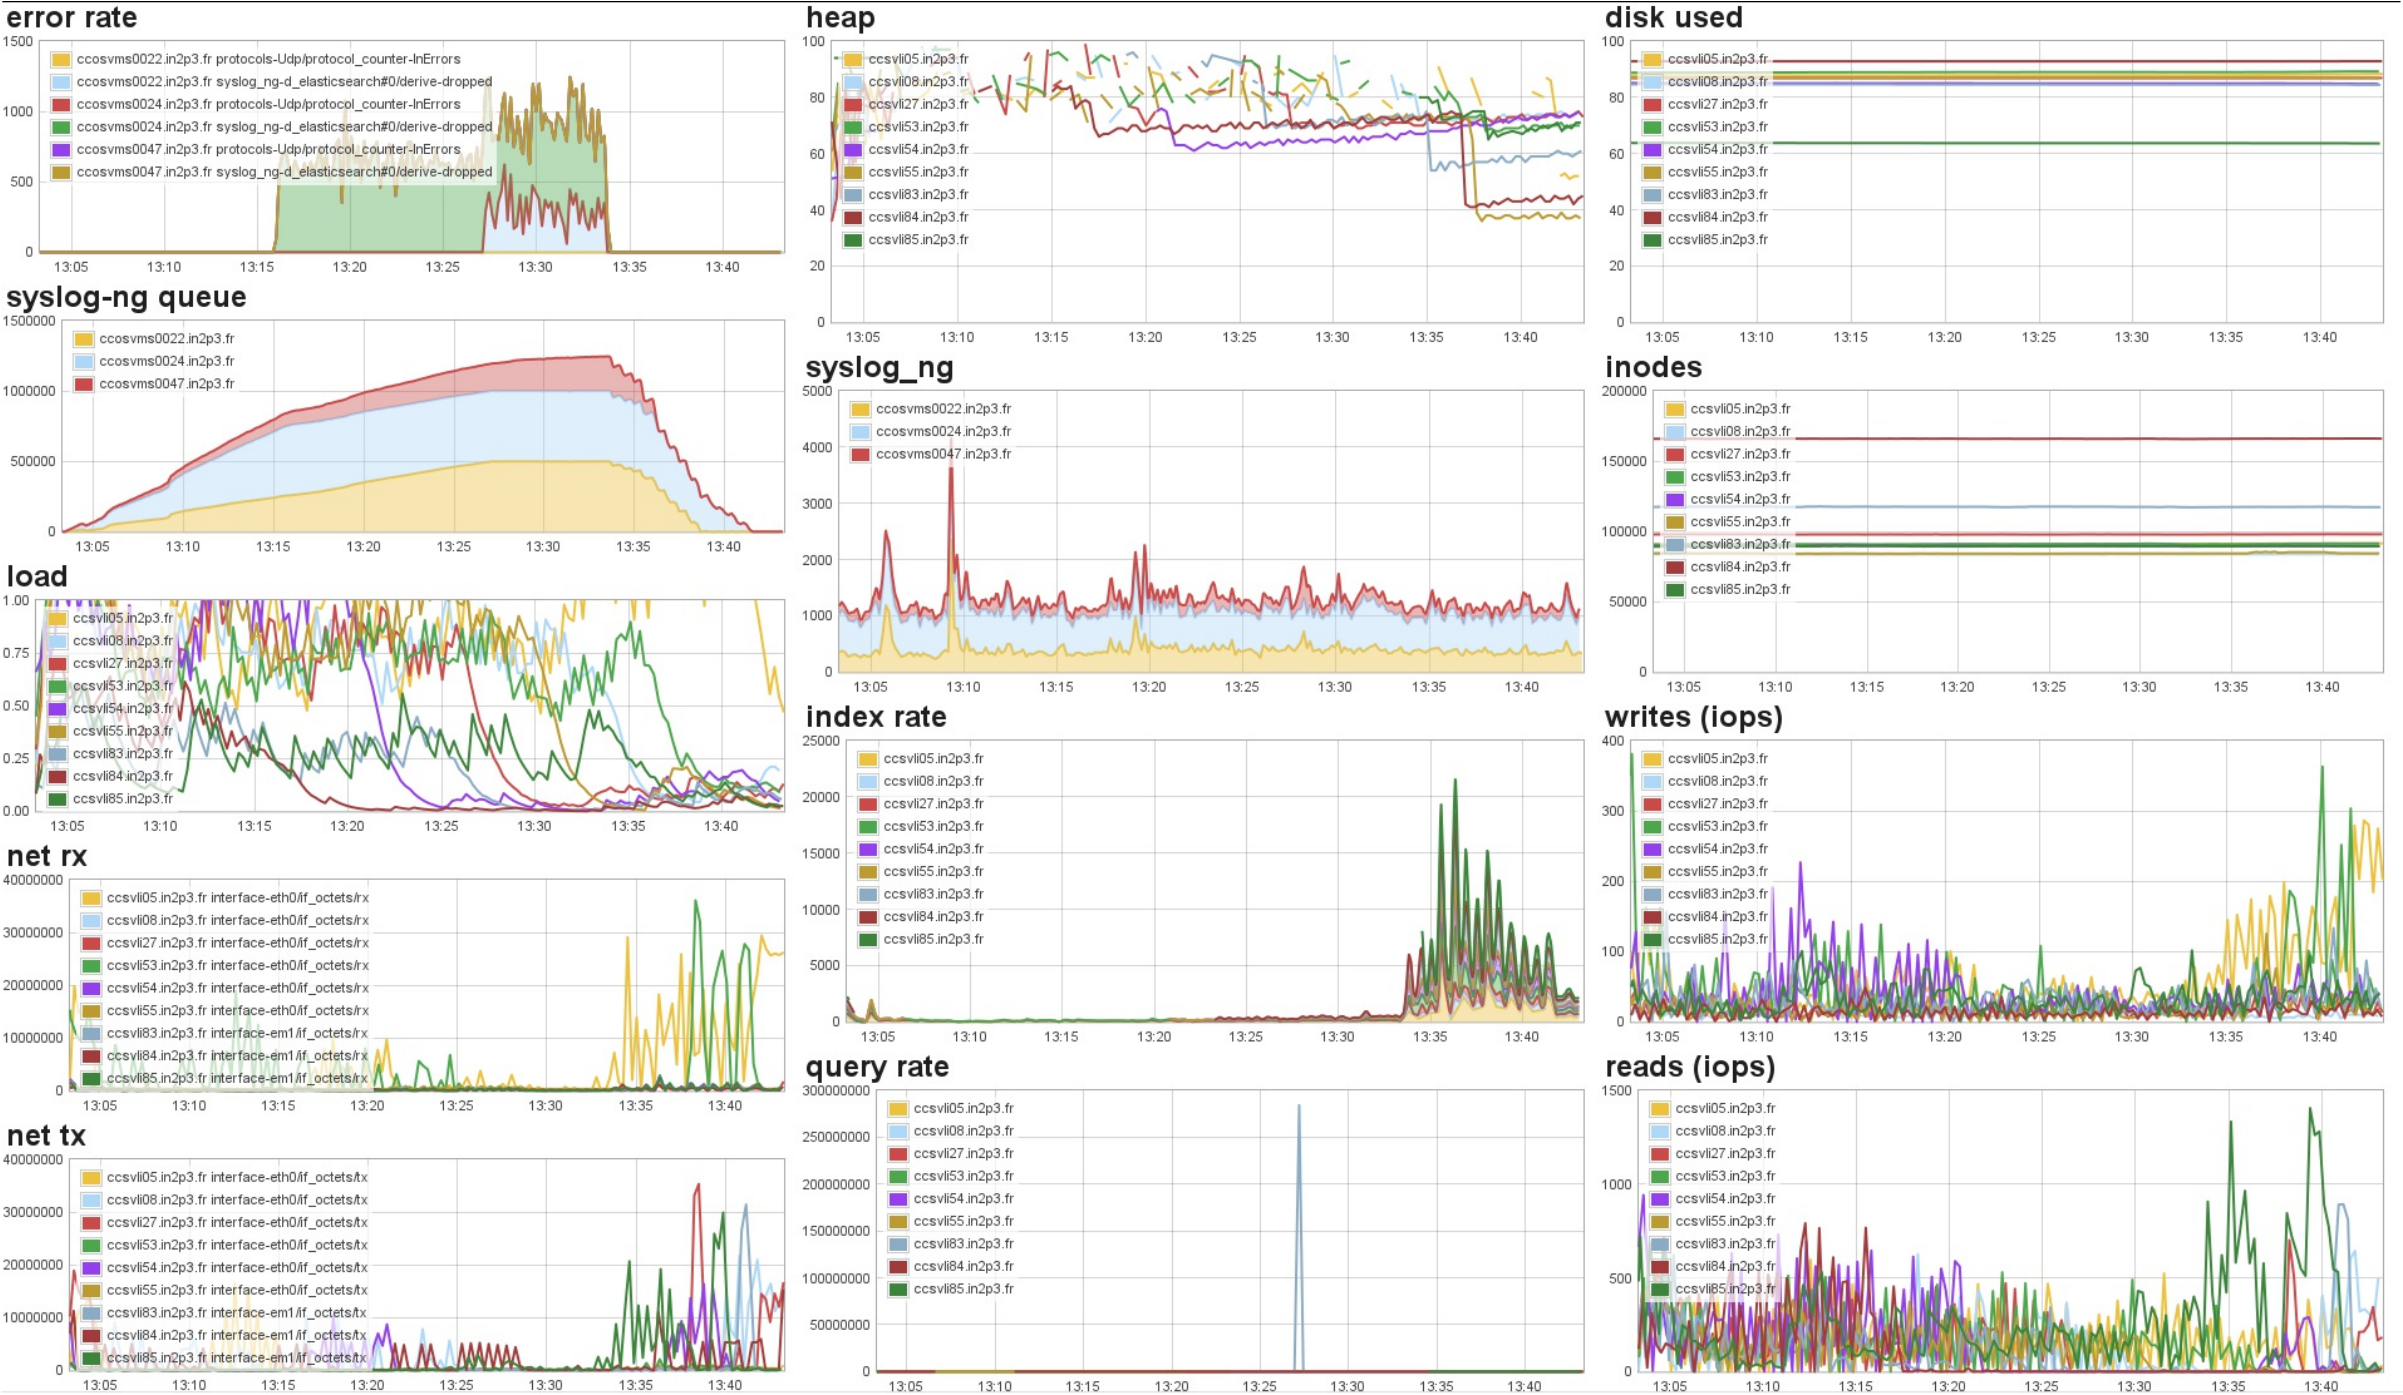
writes (iops)



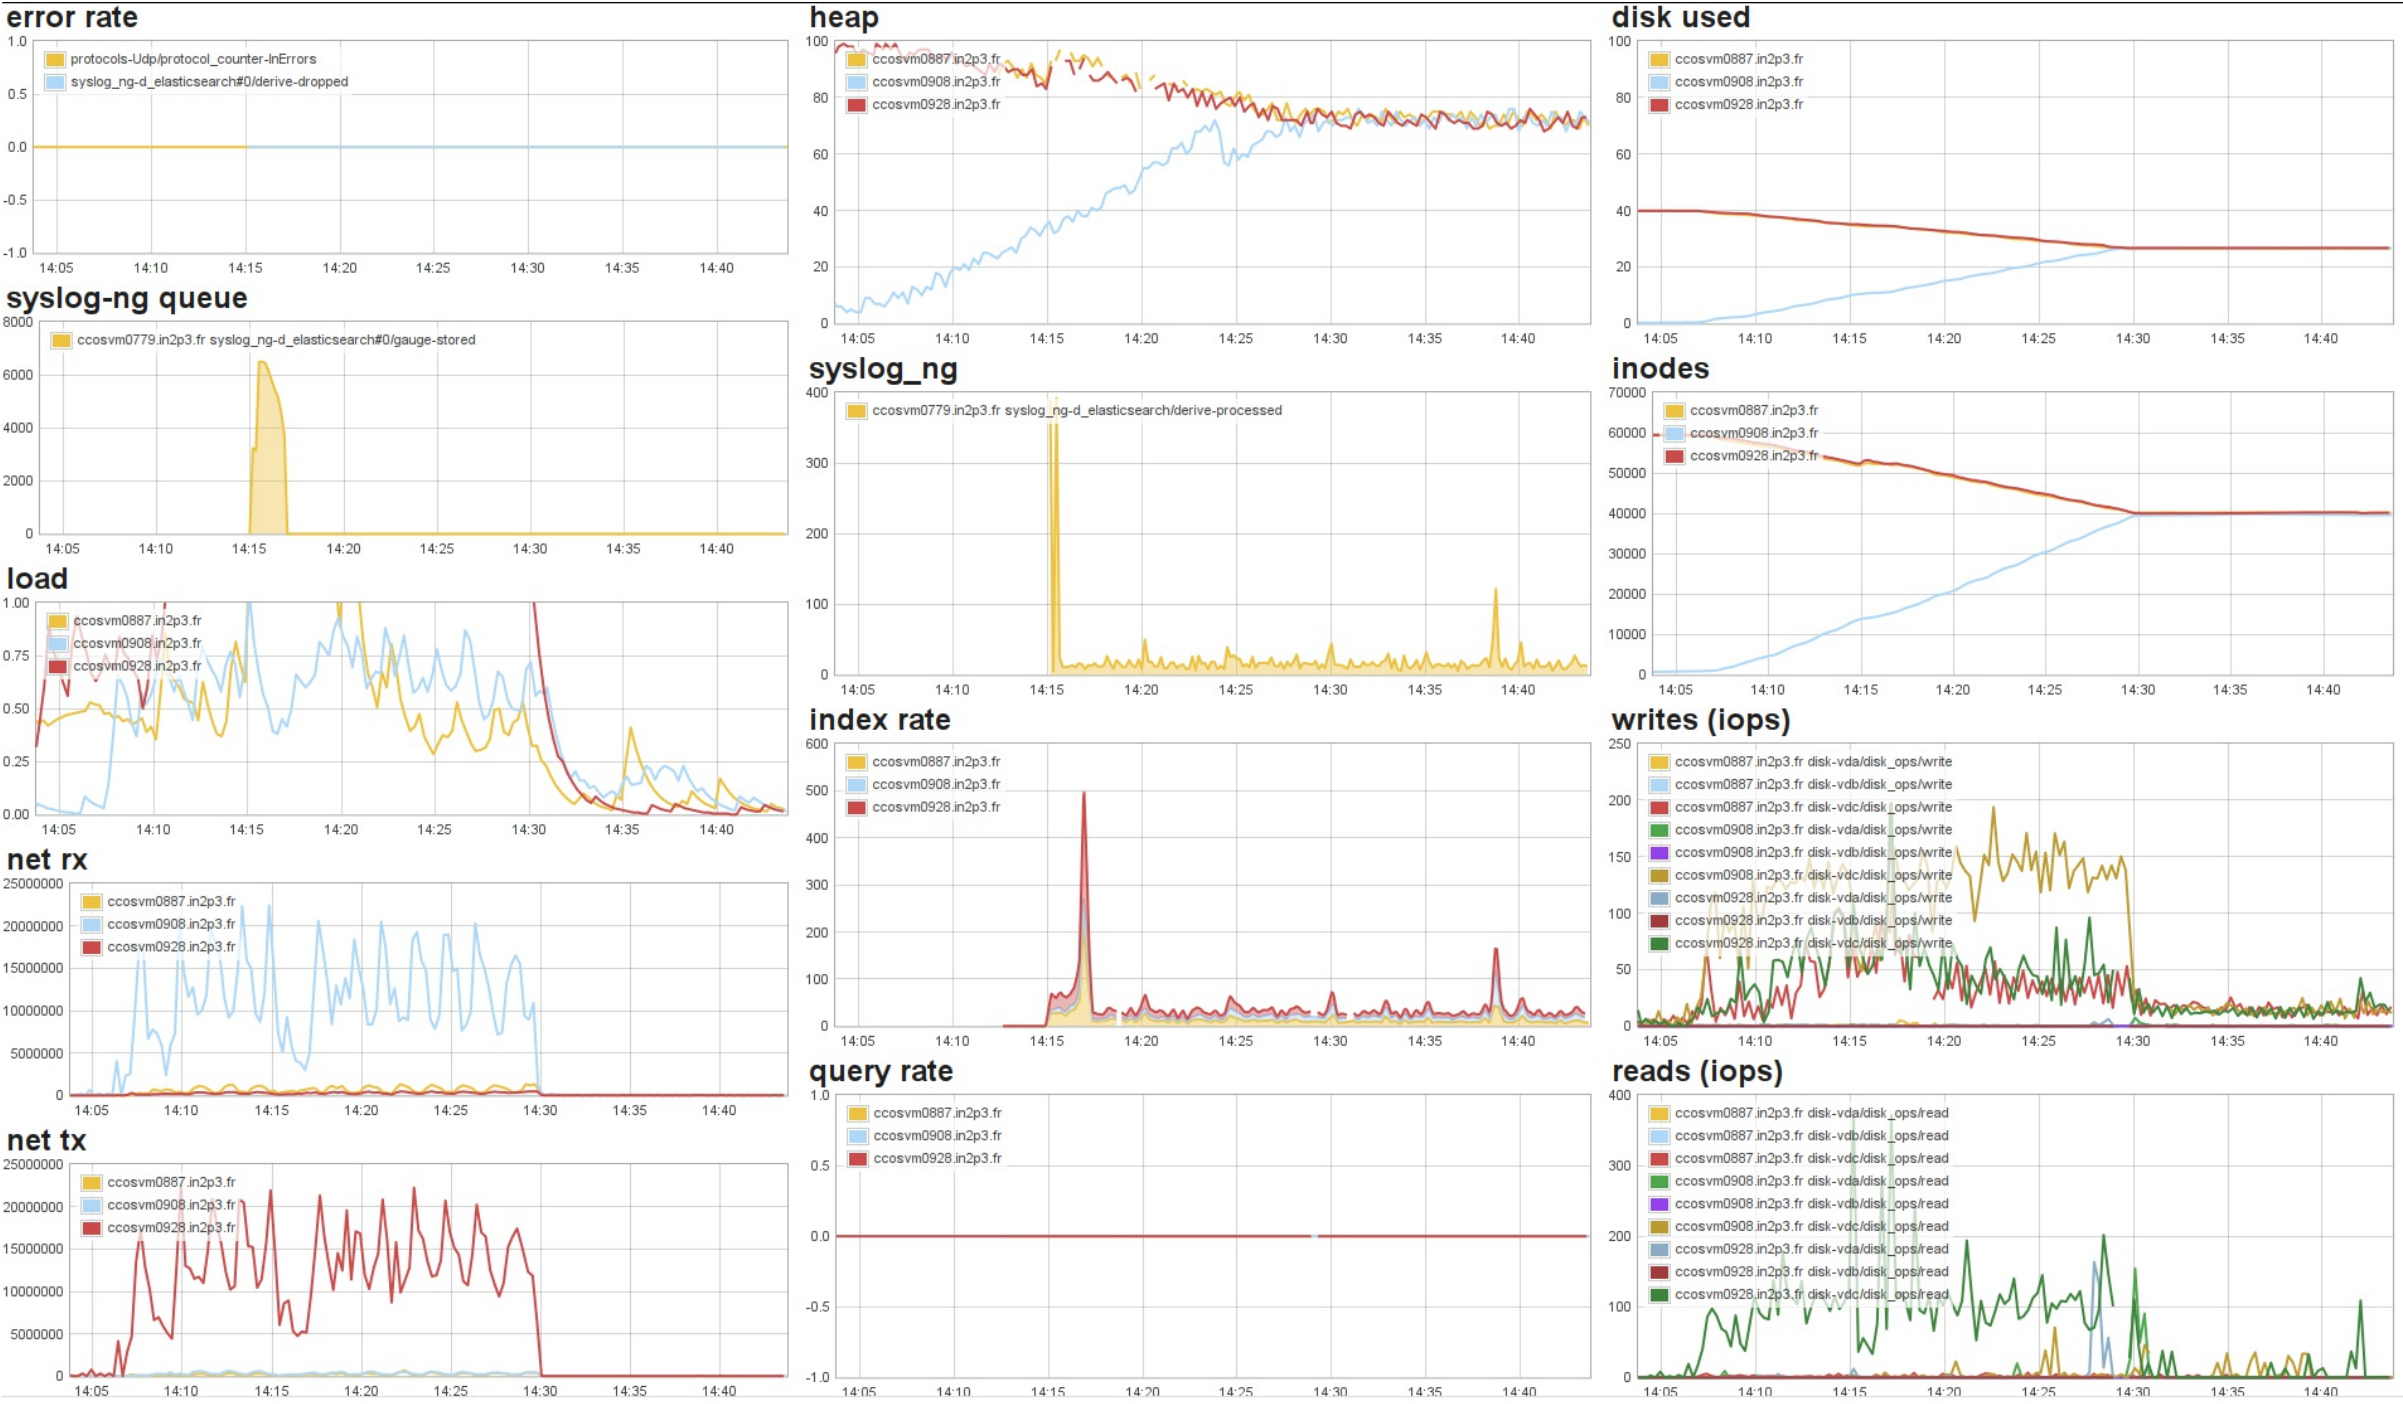
reads (iops)



Killer query



Add new node



TODO

- Kibana dynamic index upstream
 - needs [drawing attention / convincing](#)
 - larger audience than [kibana_own_home](#)
 - HTTP header (\$user) + query param (override)
- migrate existing data from ES v1.7 to v5.x
 1. split cluster
 2. reindex
 3. reunite cluster

REFERENCE

- SearchGuard
- floragunn
- syslog-ng
- riemann
- samplerr
- Securing your ESK stack for free using Search Guard

WHY USE ES FOR METRICS?

- we already have experience with ES
- it's very flexible (custom metadata)
- online aggregation to Elasticsearch: [samplerr](#)

Questions?