CERN

Large elasticsearch cluster management

Pablo Saiz

on behalf of the Centralised Elasticsearch team

Summary

- Elasticsearch service at CERN
- User perspective
- Ongoing work
- Summary

Elastic ecosystem

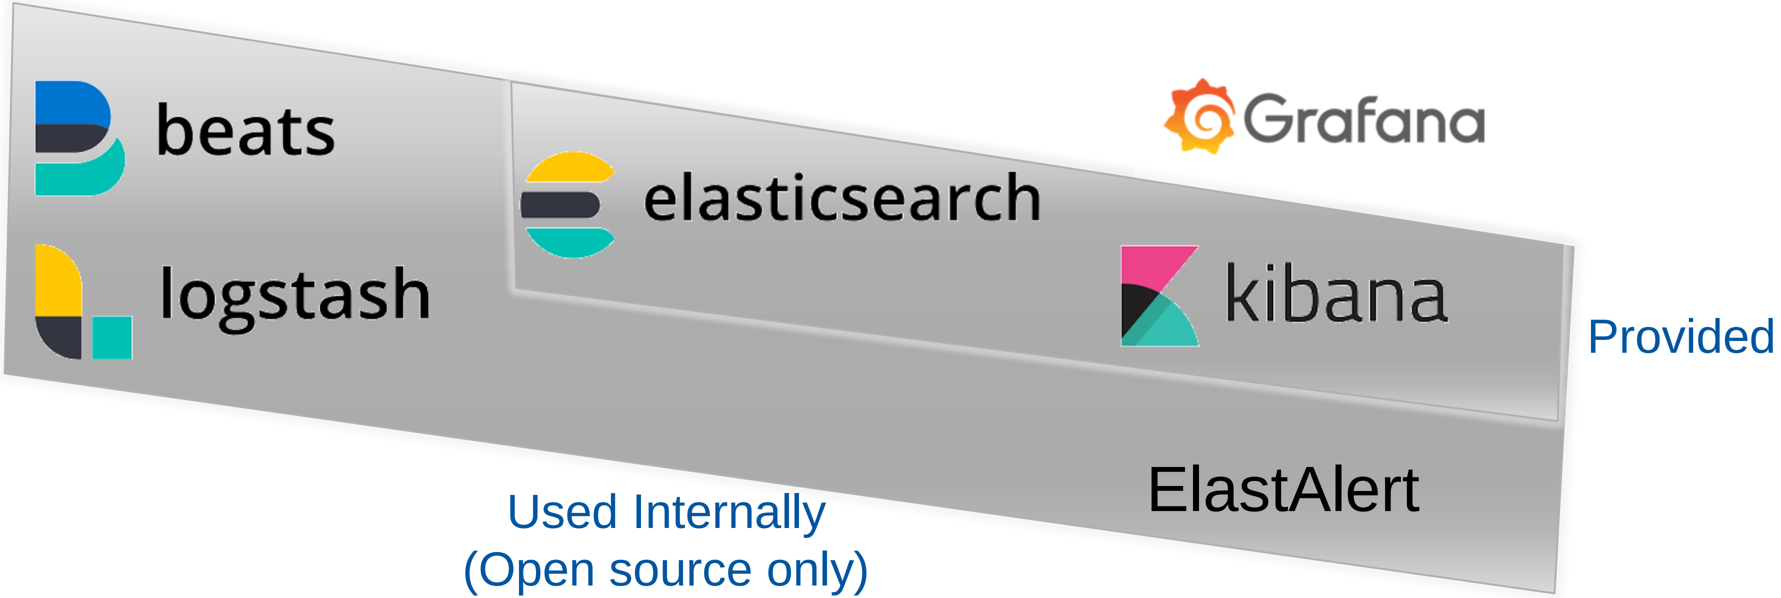

beats

logstash

elasticsearch

kibana

Grafana

ElastAlert

Provided

Used Internally
(Open source only)

Structure of a cluster

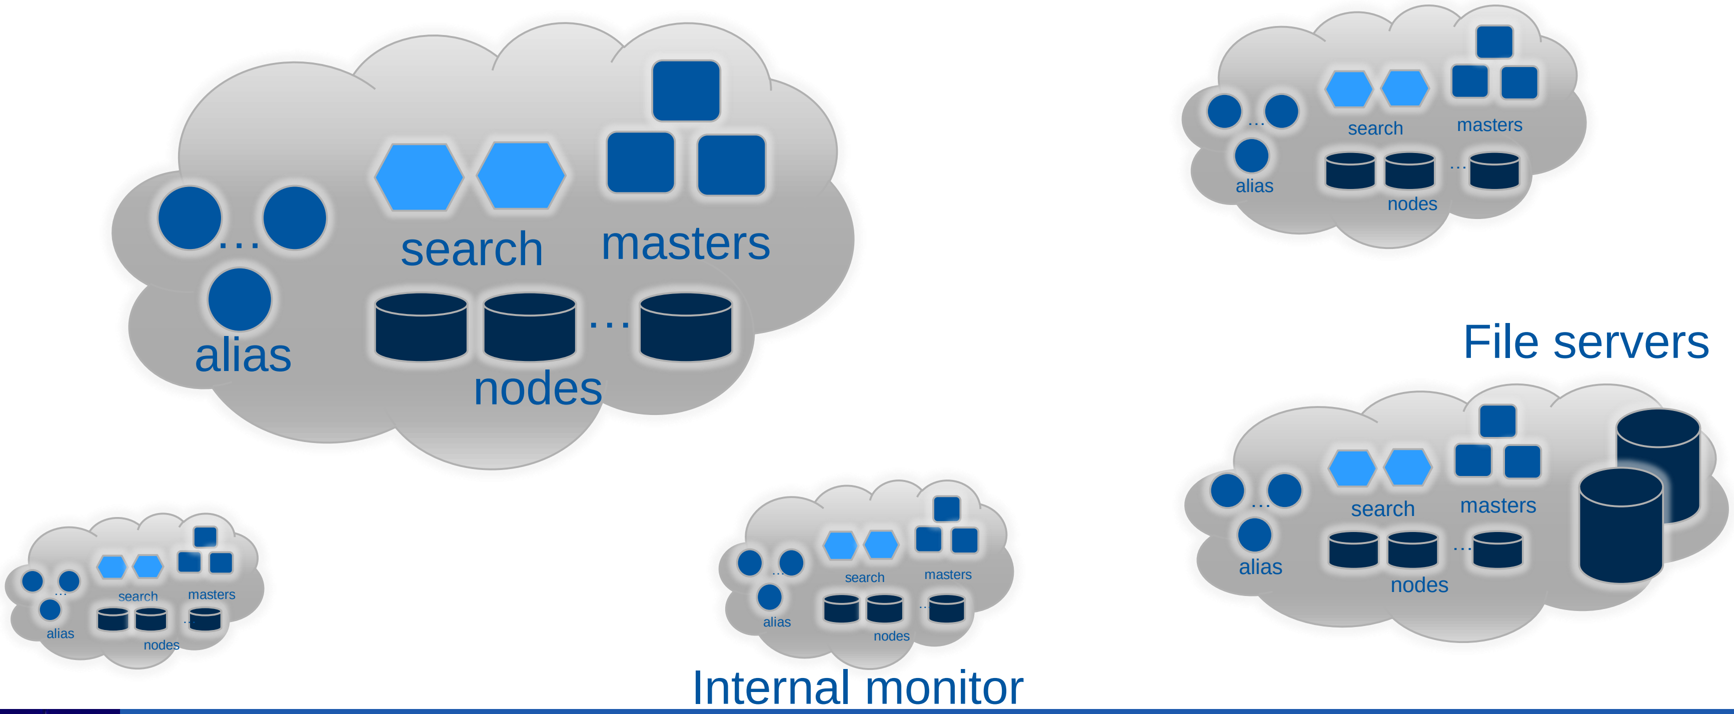

Centralised Elasticsearch in numbers

- 30 shared clusters
  - 160 dedicated aliases
- 250 data nodes
  - 5000 cores
  - 500 TB SSD
  - 600 TB disk servers
- 200 extra nodes
  - 800 cores
- 4.5 kHz access rate

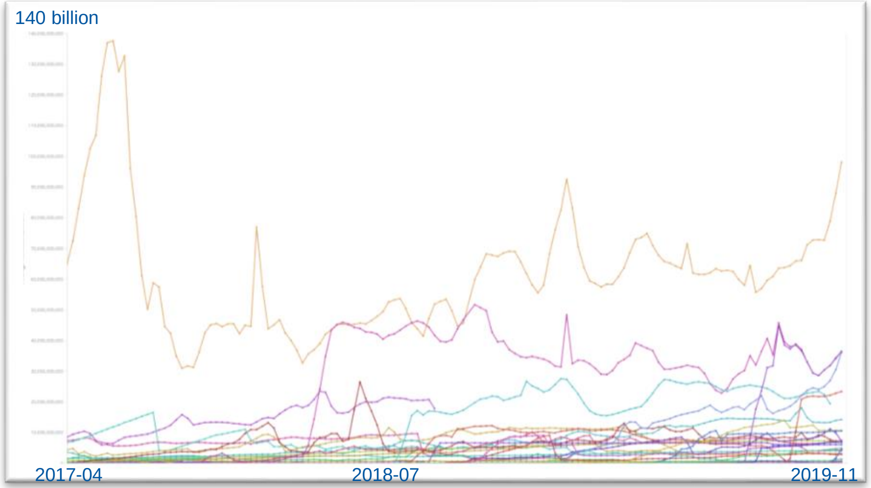

Number of documents per cluster

Other plugins offered

- Elasticsearch:
  - ReadOnlyRest: index level security
  - SQL
- Kibana:
  - Own home: multi tenancy
  - 3 in-house visualizations: relational filter, list of indices, logout. Available on GitHub
- Other applications:
  - Curator: cleanup old data
  - Template management: git repo for index templates
  - Kibana backup: copy documents to git

Service deployment

- Virtual machines on openstack
  - Multiple tenants for high availability
- Puppet managed
- Cluster definition and settings on yaml
- Ruby code
- Service operated by 1 FTE, spread over multiple people

Client requests a cluster:

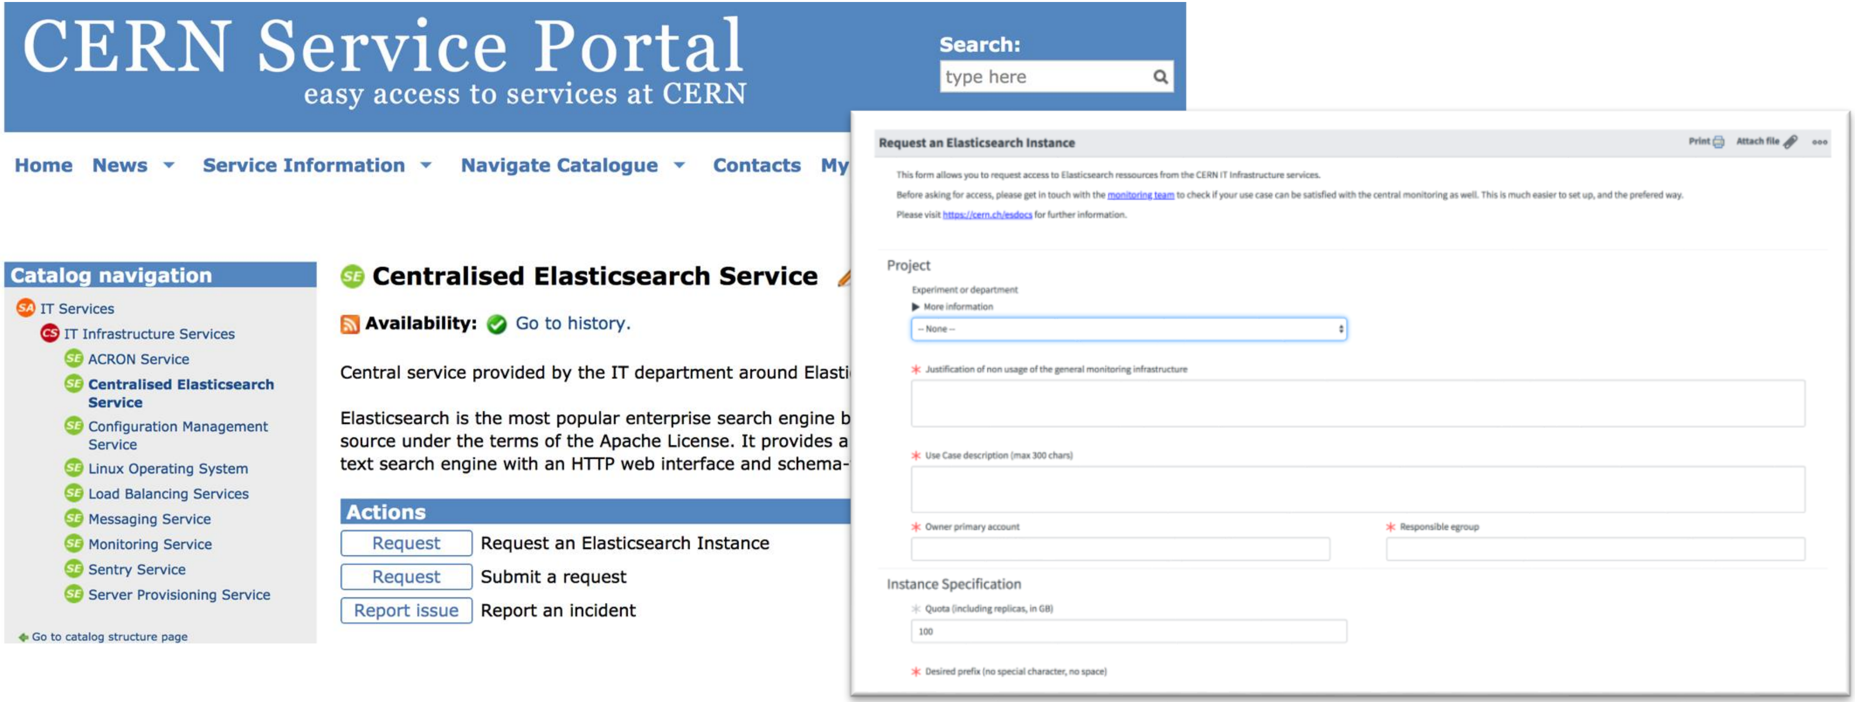

Centralised Elasticsearch Service

Availability: ✓ Go to history.

Central service provided by the IT department around Elasti

Elasticsearch is the most popular enterprise search engine b
source under the terms of the Apache License. It provides a
text search engine with an HTTP web interface and schema-

Actions

Request	Request an Elasticsearch Instance
Request	Submit a request
Report issue	Report an incident

Request an Elasticsearch Instance

This form allows you to request access to Elasticsearch ressources from the CERN IT Infrastructure services.

Before asking for access, please get in touch with the monitoring team to check if your use case can be satisfied with the central monitoring as well. This is much easier to set up, and the prefered way.

Please visit https://cern.ch/esdocs for further information.

Project

Experiment or department

▶ More information

-- None --

* Justification of non usage of the general monitoring infrastructure

* Use Case description (max 300 chars)

* Owner primary account

* Responsible egroup

Instance Specification

* Quota (including replicas, in GB)

100

* Desired prefix (no special character, no space)

The client gets

Dedicated alias …

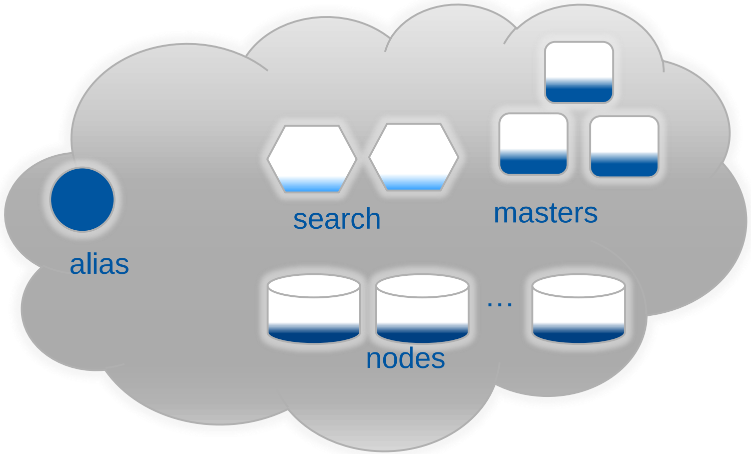

… on a shared cluster

Documentation

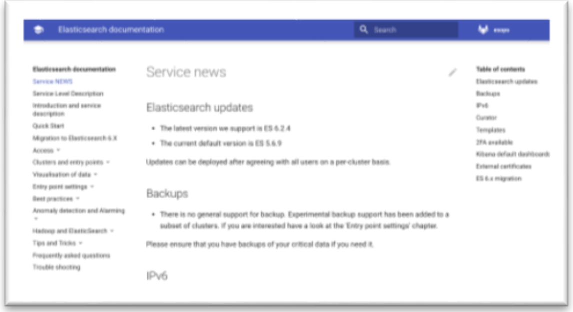

Elasticsearch documentation
- Service NEWS
- Service Level Description
- Introduction and service description
- Quick Start
- Migration to Elasticsearch 6.X
- Access
- Clusters and entry points
- Visualisation of data
- Entry point settings
- Best practices
- Anomaly detection and Alarming
- Hadoop and ElasticSearch
- Tips and Tricks
- Frequently asked questions
- Trouble shooting

Service news

Elasticsearch updates

- The latest version we support is ES 6.2.4
- The current default version is ES 5.6.9

Updates can be deployed after agreeing with all users on a per-cluster basis.

Backups

- There is no general support for backup. Experimental backup support has been added to a subset of clusters. If you are interested have a look at the 'Entry point settings' chapter.

Please ensure that you have backups of your critical data if you need it.

IPv6

Table of contents
- Elasticsearch updates
- Backups
- IPv6
- Curator
- Templates
- 2FA available
- Kibana default dashboards
- External certificates
- ES 6.x migration

Settings

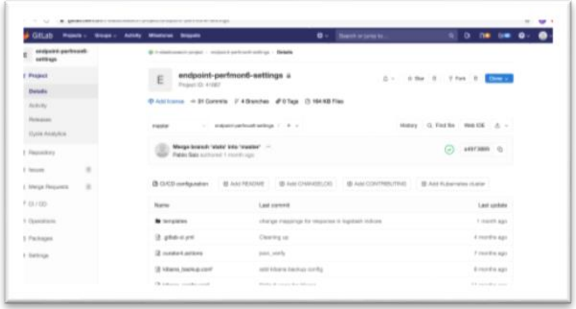

endpoint-perfmon6-settings

Add license · 31 Commits · 4 Branches · 0 Tags · 164 KB Files

Merge branch 'stats' into 'master'
Pablo Saiz authored 1 month ago

CI/CD configuration · Add README · Add CHANGELOG · Add CONTRIBUTING · Add Kubernetes cluster

Name	Last commit	Last update
templates	change mappings for response in logstash indices	1 month ago
.gitlab-ci.yml	Cleaning up	4 months ago
curator4.actions	pass_verify	7 months ago
kibana_backup.conf	add kibana backup config	8 months ago

Lessons learned

I. Use SSD only whenever possible
   - Cluster speed defined by slowest node

II. Find sweet spot of # nodes
   - Issues with large clusters (>20 data nodes)

III. Keep # indices/shards under control
   - Aim for ~10 GB shard size

IV. Reindexing is expensive

V. Beware of closed indices
   - Not replicated

Lessons learned (II)

VI. Index level security on shared clusters easy
   - Difficult to ensure isolation of clients

VII. Elasticsearch can guess data types
   - Better if already defined: index templates

VIII. Beware of large queries
   - Aggregations easily add up

IX. Need close communication with clients
   - Service Now, Mattermost

X. Plenty of parameters to tune/monitor
   - Need for advanced monitor

Ongoing work

- Service anomaly detection
- Transition to ES 7.X
- Accounting
- Evaluation of Open Distro Elasticsearch
  - Including change of security model
  - And container based solution

Centralised Elasticsearch service

- Providing dedicated aliases on shared clusters
  - 30 cluster, 160 endpoints, 500 TB SSD
- Using exclusively Open Source components
- Full isolation is challenging
- Automated management with puppet
  - Operated by 1 FTE
- Monitoring and recovery actions are crucial

CERN