



CERN's Centralised Elasticsearch Service

The service has been stable since April 2018, number changed a bit though due to an increase of clients

- Basic facts
- Service design and ACLs
- Other features
- References
- Plans

Basic facts

- About 30 clusters and 160 endpoints/use cases
 - On average 5 user cases per cluster
 - In reality few clusters with many use cases
- Using open source tools only
- ACLs: index level security with ReadonlyREST
- Elasticsearch 6.x, 5.x being phased out
- Initial 7.x support (experimental)
- In the process of switching to OSS version (for > 6.3 versions)
- Plans to evaluate “Open Distro for Elasticsearch”

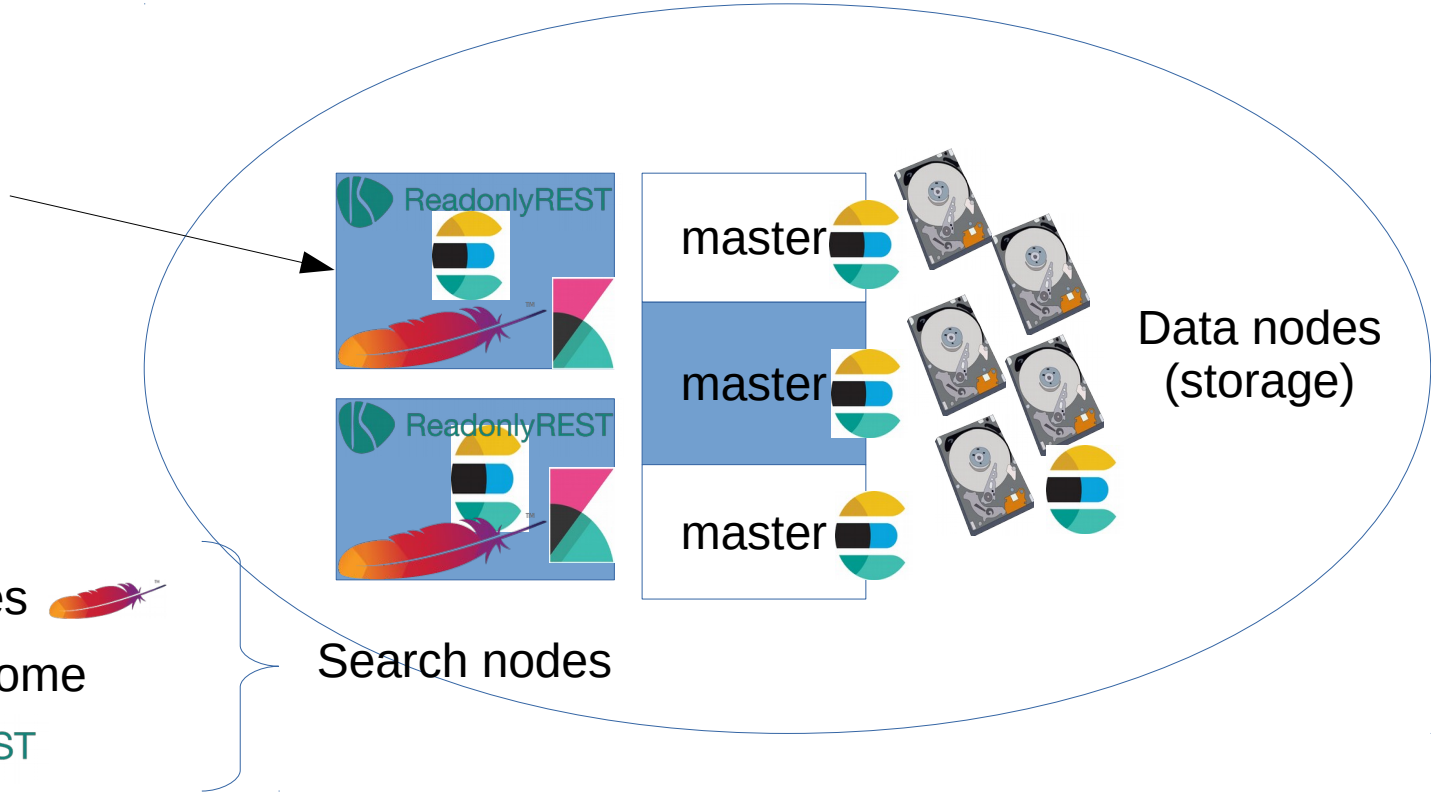
Basic facts

- Apache access rate (total): ~1.2 kHz
- Nodes counts: 250 (data), 75(search), 108 (master)
- Data node sizes: 1.8TB or 900TB each
- Puppetized VMs running on our OpenStack internal cloud
- Total capacity
 - ~500 TB on SSD
 - an additional ~600 TB (on disk servers) on 2 special clusters
- Usage: up to 80%

Use cases and customers

- Sample use cases:
 - IT monitoring infrastructure
 - Security infrastructure
 - Searches/search engines
 - ...
- Customers:
 - IT
 - Experiments (ALICE, ATLAS, CMS, LHCb, and others)
 - Beams department
 - ...

Service design and ACLs



Apache proxies 

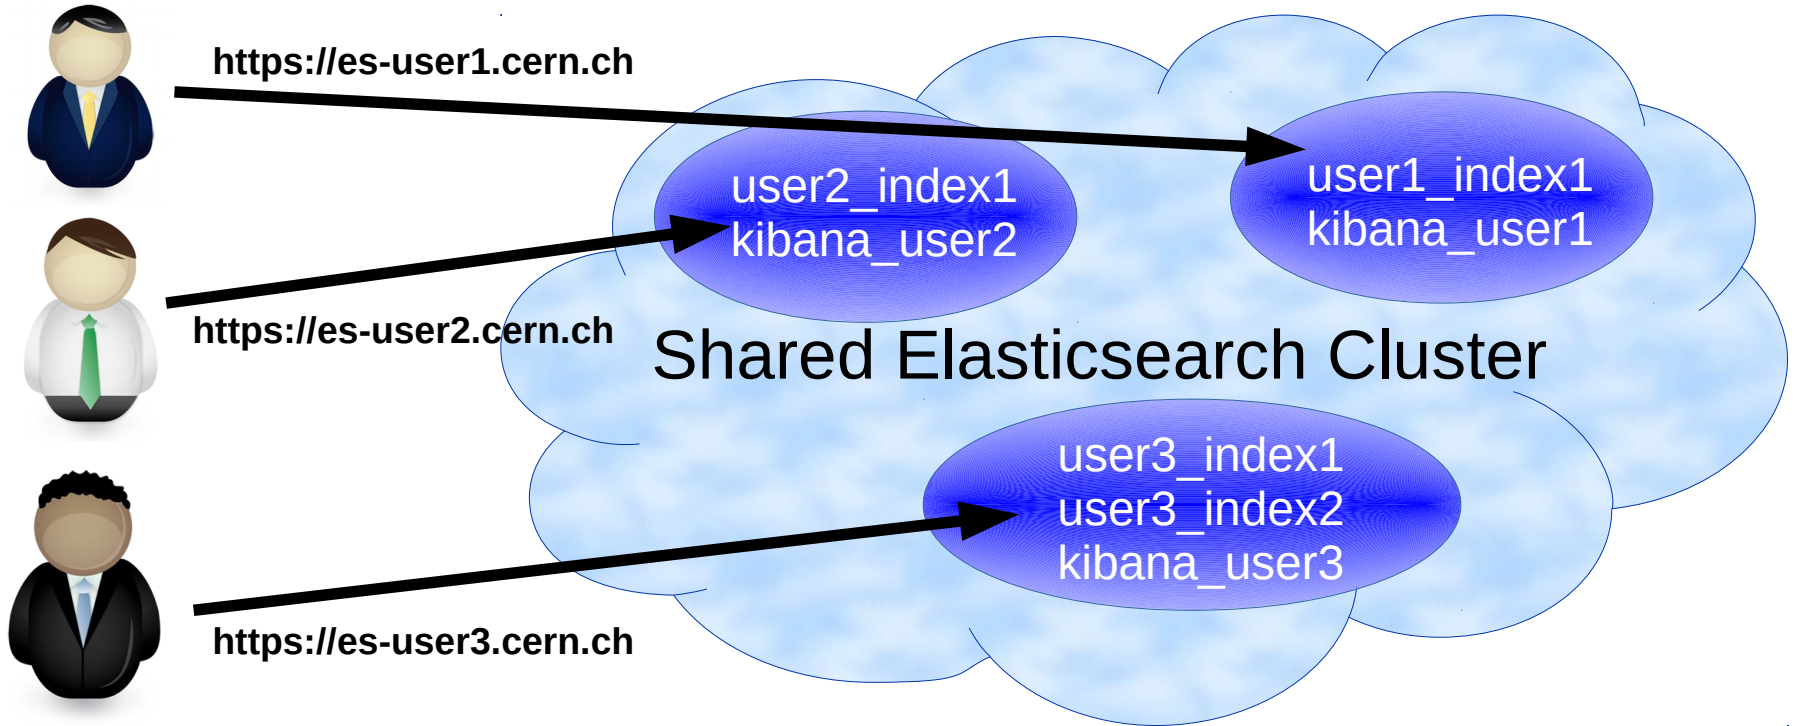
Kibana OwnHome



Search nodes

Data nodes
(storage)

Sharing resources



Other features

- Full integration into the CERN IT infrastructure
- A gitlab repository per endpoint allows for user management of
 - Curator configurations
 - Templates
- Integration with accounting based on disk usage
- Detailed internal service monitoring
 - Independent from centralized monitoring
 - Partly accessible for users as well
 - Recently added experimental anomaly detection for specific clusters

References

- User documentation: <https://cern.ch/esdocs>
- Readonlyrest Elasticsearch plugin (from Simone Scarduzio)
<https://readonlyrest.com/>
<https://github.com/sscarduzio/elasticsearch-readonlyrest-plugin>
- Kibana ownhome plugin (from Wataru Takase)
<https://github.com/wtakase/kibana-own-home>

Plans

- Phase out of ES 5.X
- Add support for ES 7.X (started)
- Evaluate Open Distro for Elasticsearch
 - Hoping for a significant simplification of the service as it contains build-in security
- Vega

We are very interested in any possibly exiting experiences you might have on these items!

Discussion/ Questions



www.cern.ch