



The Centralised Elasticsearch Service at CERN

- Project goals, challenges
- Service status and limitations
- Access control (ACLs)
- Summary



Project goals/mandate

Setup a centralised Elasticsearch service

- New, centrally managed service
- Consolidate existing clusters



Challenges

Consolidation:

- Centralised management
- Resource sharing
- Use of standard hardware
- Use of virtualisation



Expectations:

- Special requirements
- Privacy and security
- Performance
- Scalability

Challenges

- **Elasticsearch advantages:**
 - Build-in fail-safeness via (user-defined) replicas
 - Many knobs for tuning
- **Elasticsearch intrinsic limitations:**
 - No intrinsic concept of quota
 - Neither on space nor on search sizes
 - Individual users can bring the system down
 - Outages can cause data loss
 - I/O intensive
 - Requires careful tuning, depending on the use case
 - Hardware must be good enough to support the individual use case



Solution

- Share resources where possible
 - Consolidate smaller use cases
 - Put users with similar needs on the same cluster
- Use dedicated clusters where needed
 - Special networking requirements (eg. Technical network (TN) trusted)
 - High demand use cases (eg. CERN IT monitoring)


Service status: ES clusters

- ~20 Clusters up and running:
 - ~40 use cases (entry-points) supported
 - Currently up to 6 entry points on a single cluster
- Elasticsearch 2.4.1 or 5.2.1, 5.4.0 in preparation
- Kibana 4.6.1 or 5.2.1 (5.4.0)
- Planning upgrades to 5.X with our users

Service status

Description	Details	Link	Status
Service status monitoring	SLS	https://cern.service-now.com/service-portal?cis2showall=true	OK
Internal service monitoring	Accesses, errors, disk usage, ...	https://es-perfmon.cern.ch/	OK
Integration into AI infrastructure	Puppet, lemon monitoring, CI, ...		OK
Support structures	Functional Elements, SNOW integration, Mattermost channels		OK
Documentation	Enduser documentation Service rota person documentation Service manager documentation	https://cern.ch/esdocs https://cern.ch/esops	OK ACLs to be documented
Work flow automation	Removing/adding machines, ES restarts, upgrades, reboots	https://itesrundeck.cern.ch , https://gitlab.cern.ch/it-elasticsearch-project/itestools	OK
Index level security	ACL settings on index level		OK for ES 5.X
Accounting	Aggregation by Accounting group		Per cluster OK, per accounting group ongoing
Disaster recovery	Replication of the service at eg. Wigner		Missing
Backup			On the road map, for part of the data only

Access control (ACL) implementation (1)

- Why ?
 - Privacy and security requirements
 - Needed for efficient consolidation of resources
- Commercial plugins: 
 - Offer for XPACK from Elastic (shield) too expensive
 - SearchGuard: concerns about performance and integration
- Desired model
 - Pure OpenSource solution (Apache2 license or similar)
 - Index-level security is enough



ACL Implementation (2)



1) Apache proxies and virtual hosts

2) Readonlyrest Elasticsearch plugin (from Simone Scarduzio)



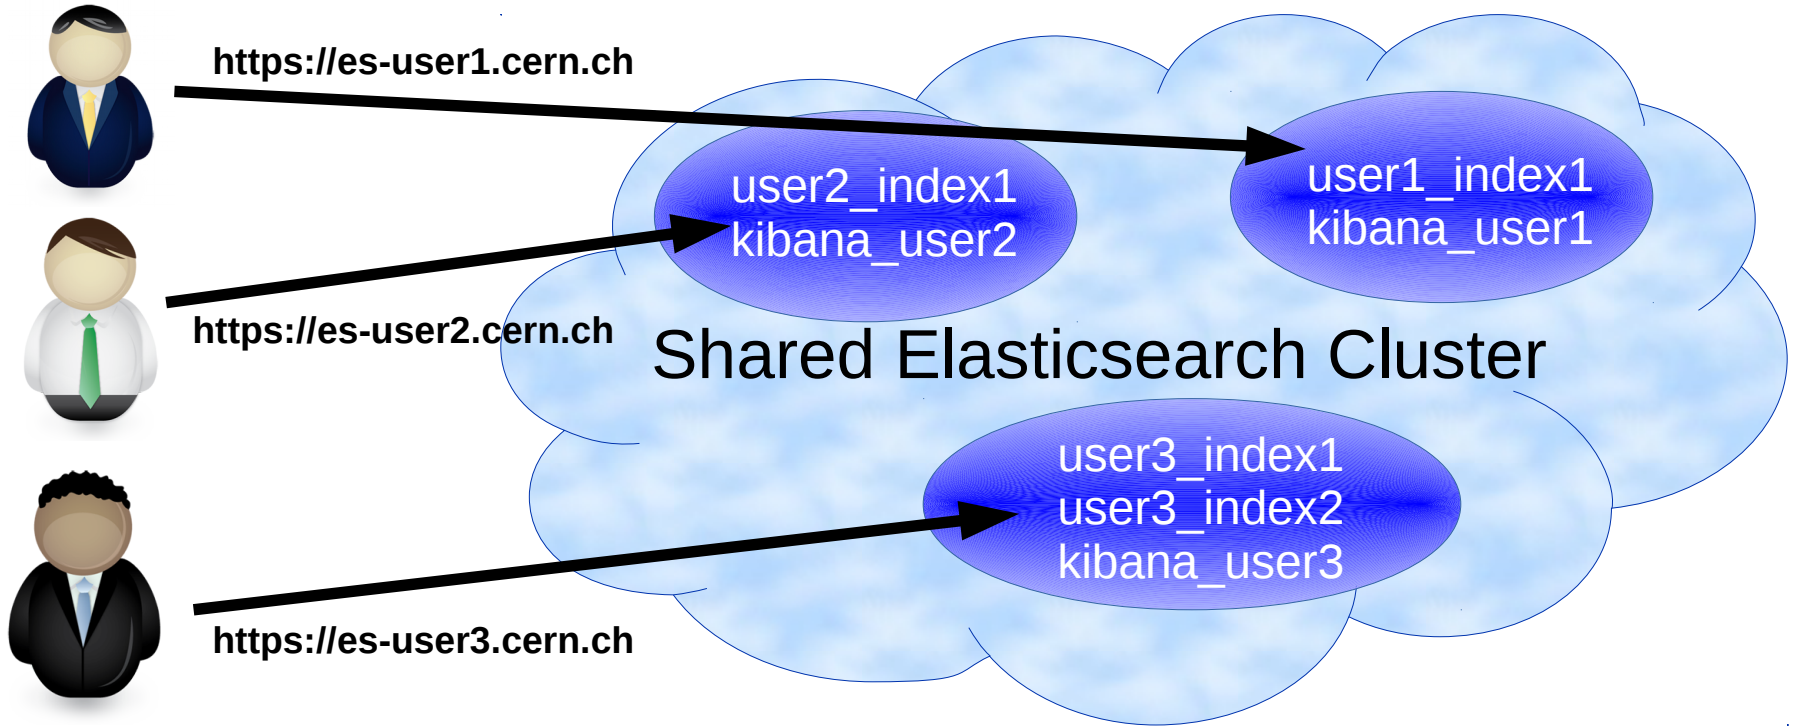
<https://readonlyrest.com/>

<https://github.com/sscarduzio/elasticsearch-readonlyrest-plugin>

3) Kibana ownhome plugin (from Wataru Takase)

<https://github.com/wtakase/kibana-own-home>

ACL implementation (3)



Summary

- Running a centralised Elasticsearch service at CERN
- Support 2.X and 5.X versions
 - Index level security only for 5.X Elasticsearch
- Lessons learned
 - Very different use cases and requirements
 - Careful tunings are needed on **both** client and service side

See also:

- “Centralising Elasticsearch”, HEPiX 2017, <https://indico.cern.ch/event/595396/contributions/2532588/>
- “Elasticsearch status and lessons learned”, ASDF Meeting, CERN, <https://indico.cern.ch/event/639585/contributions/2593332/attachments/1461957/2258542/asdf.pdf>
- Readonlyrest Elasticsearch plugin, Talk given by Simone Scarduzio at CERN, <https://cds.cern.ch/record/2261999>



www.cern.ch