



The Elasticsearch Service at CERN

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



Project goals/mandate

Setup a centralised Elasticsearch service

- Setup a **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)
 - Dedicated clusters for ALICE, ATLAS, CMS, LHCb

Service status: ES clusters

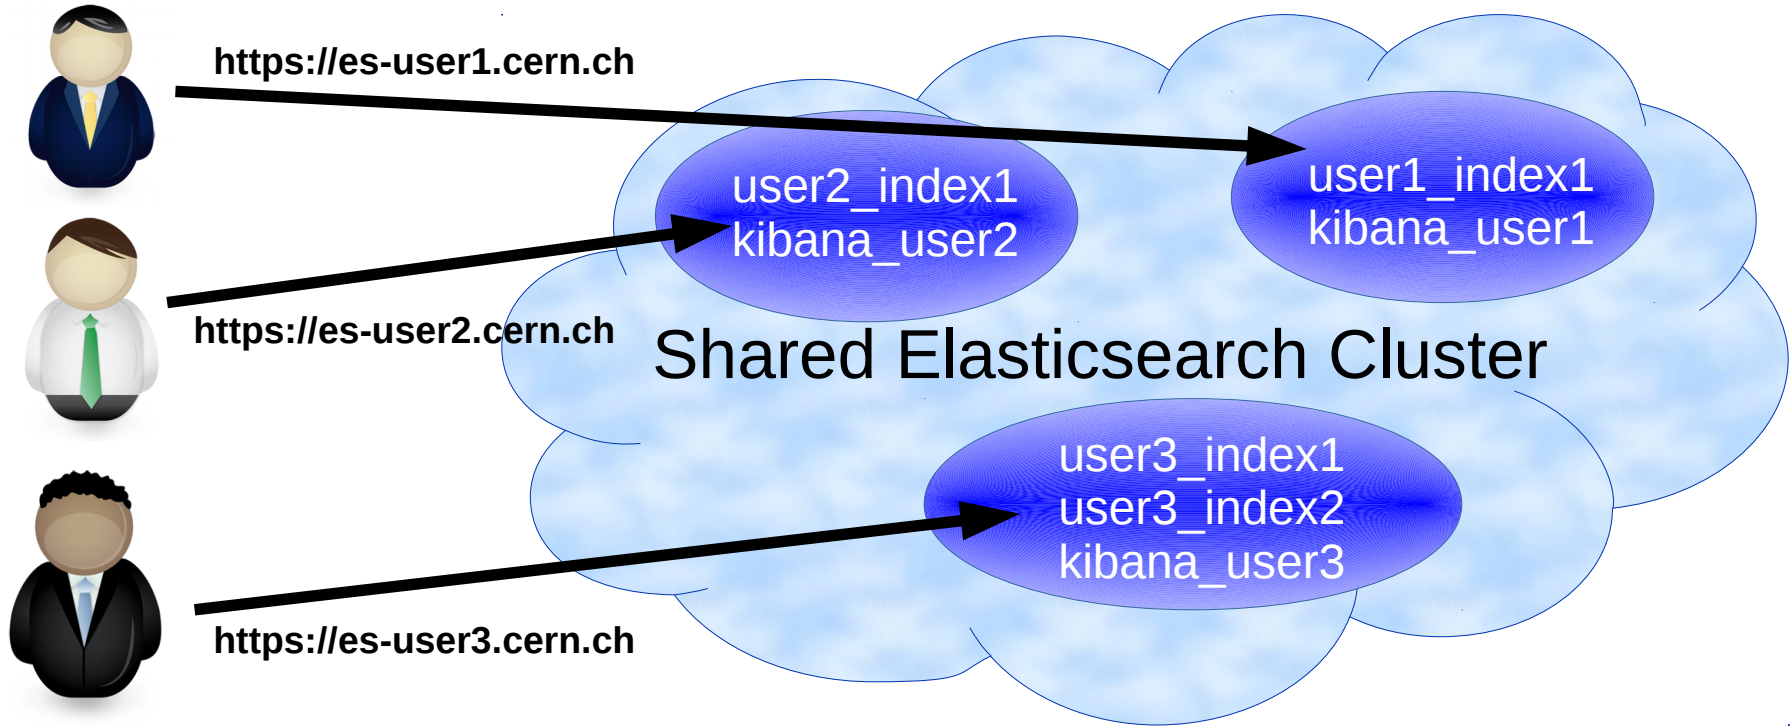
- ~20 Clusters up and running:
 - O(40) use cases (entry-points)
 - Currently up to 6 entry points on a single cluster
 - Including test access to ES 5.x
- Elasticsearch 2.4.1 or 5.2.1 (moving to 5.4.0)
- Kibana 4.6.1 or 5.2.1 (5.4.0)
- Planning upgrades to 5.X with our users

Access control (ACL) implementation (1)

- Why ?
 - Privacy and security requirements
 - Needed for **efficient consolidation** of resources
- Commercial plugins: 
 - Tested XPACK and SearchGuard
 - Concerns about costs and performance
- Implemented model (ES 5.x only)
 - Pure OpenSource solution (Apache 2, GPL V3)
 - Index-level security ensured by Readonlyrest and kibana-ownhome plugins



ACL implementation (2)



Summary

- Running a centralised Elasticsearch service at CERN
- Support 2.X and 5.X versions
 - Moving to 5.X
 - Index level security for 5.X Elasticsearch
- Lessons learned
 - Very different use cases and requirements
 - Careful tunings are needed on **both** client and service side
- Contact: elasticsearch-support@cern.ch



www.cern.ch