Elasticsearch Service@ CERN

WLCG SOC working group workshop

Pablo Saiz

28th June 2018

Outline

- Centralised Elasticsearch Service at CERN
  - Security
  - Kibana own home
- Tools provided:
  - Template management
  - Curator
  - Cross-cluster searches
- Setup of security clusters
- Lessons learned

Centralised Elasticsearch @CERN

- Provide Elasticsearch/Kibana clusters
  - 4 versions: 5.5.2, 5.6.4, 5.6.9, 6.2.4

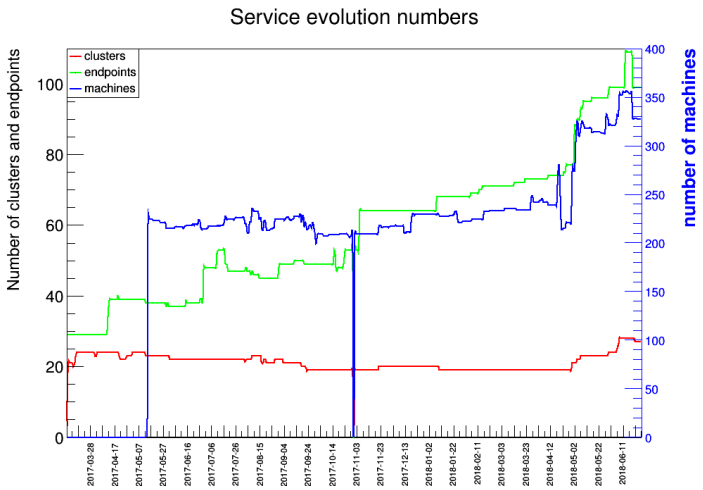

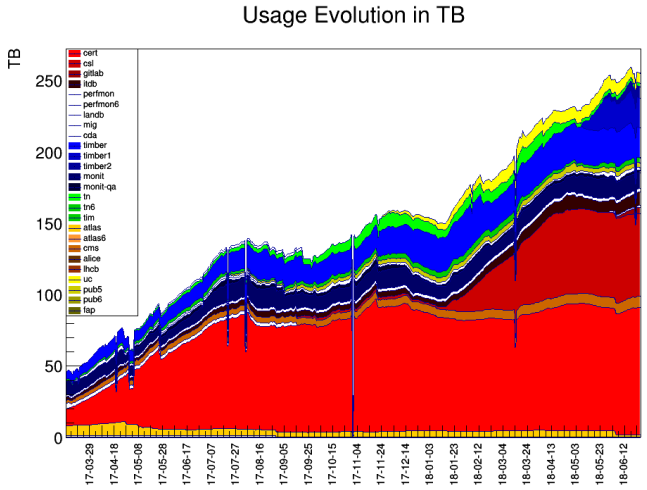

http://cern.ch/esdocs

Setup of a cluster

- Running on virtual machines with Openstack
- 4 dedicated tenants with distinct switches
  - Availability zones
- Multiple users on the same cluster:
  - Identified by alias
  - Restricted by ACL
- Apache proxy
  - Authentication
    - SSO, Kerberos, user/pass

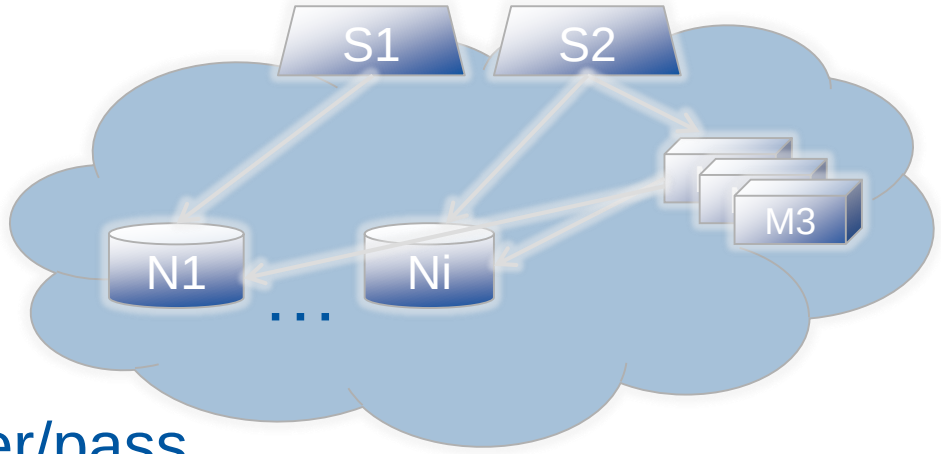

Elasticsearch security

- Open-source solution: Read-Only-Rest
- Define ACL based on:
  - Indices
  - Actions
  - User/groups
  - Host (alias)
- First match applies

```
Name: es-cds reader
   x_forwarded_for:   es-cds
   groups:            cds_ro
   indices:           cds_*
   actions:
     - cluster:monitor/*
     - indices:data/read/*
 type:        allow
```

http://readonlyrest.com

Kibana own-home

- Open source kibana plugin
- Provides different kibana indices
- https://github.com/wtakase/kibana-own-home

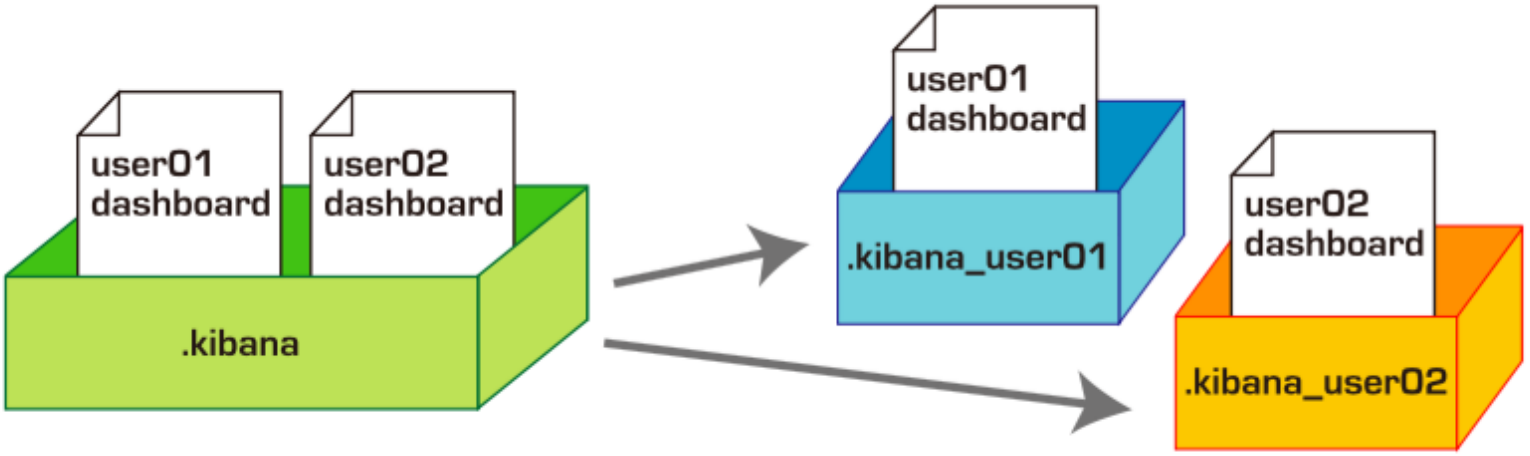

Management and Support

- Gitlab repo per endpoint
  - Settings
- Cluster and ACL configuration with Puppet
- Dedicated tools for upgrades, reboots…
- Dedicated cluster for monitoring
- Interventions scheduled with users
- SNOW as ticketing system
- Best effort support

More settings

- Managed by user:
  - Curator
  - Template management
  - Kibana backup
  - Default landing page
  - ES backup (experimental)
- Managed by admins:
  - Cross-cluster searches
  - Users/egroups access
  - Multiple kibana (ro/rw)

Security clusters

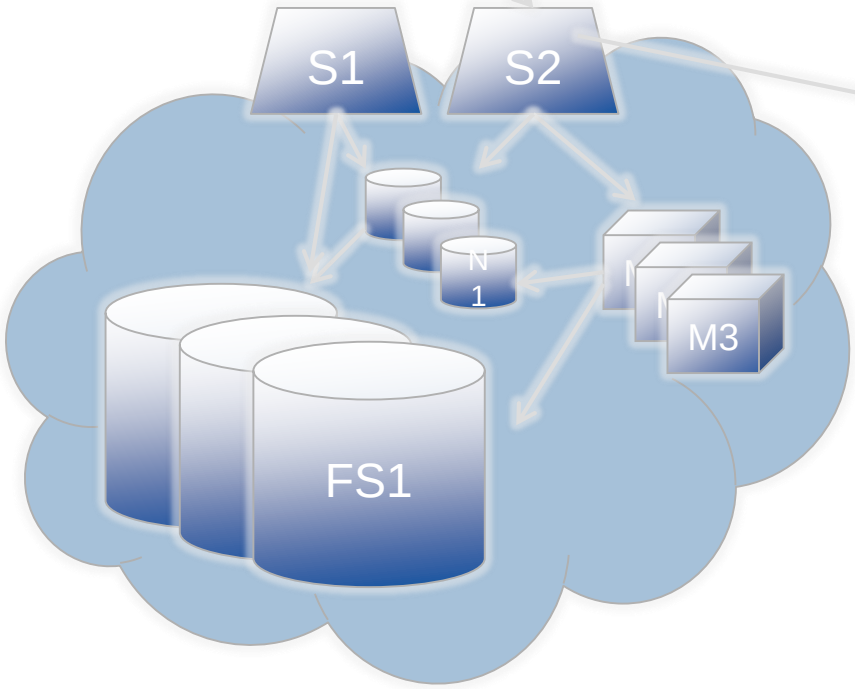

Cert

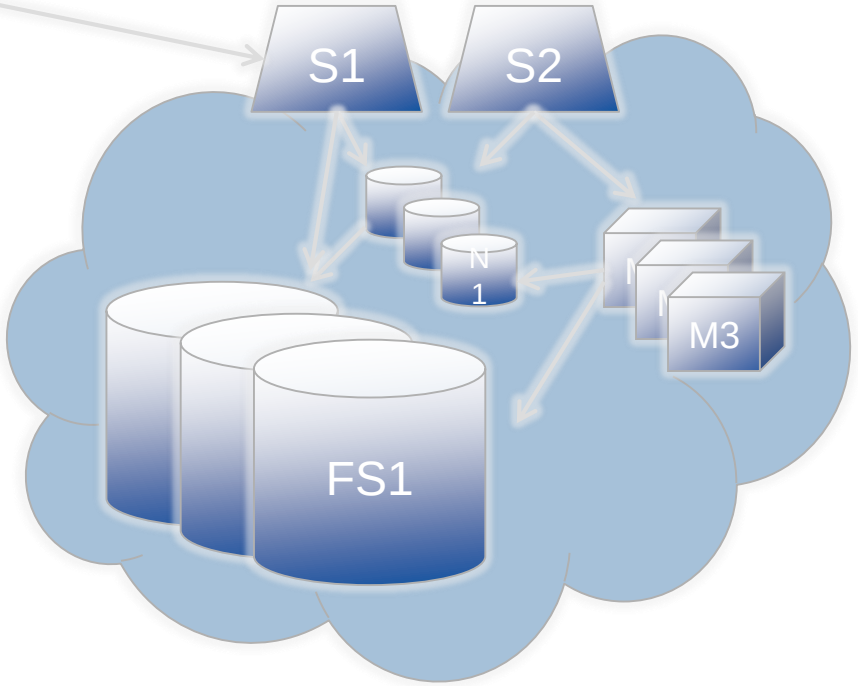

Csl

Security cluster configuration

- 2 types of data nodes:
  - SSD (1.6TB per node)
  - File servers (120 TB per node)
- Daily indices written to SSD
- Curator configured for:
  - Moving data from SSD to File Servers
  - Close indices after 30 days
  - Delete indices after 90 days
  - Keep aliases of ‘last week’
- Index templates with field types
- Cross cluster searches
  - Single Kibana querying both clusters
- Index templates available at http://cern.ch/go/k8F8

Monitoring the clusters

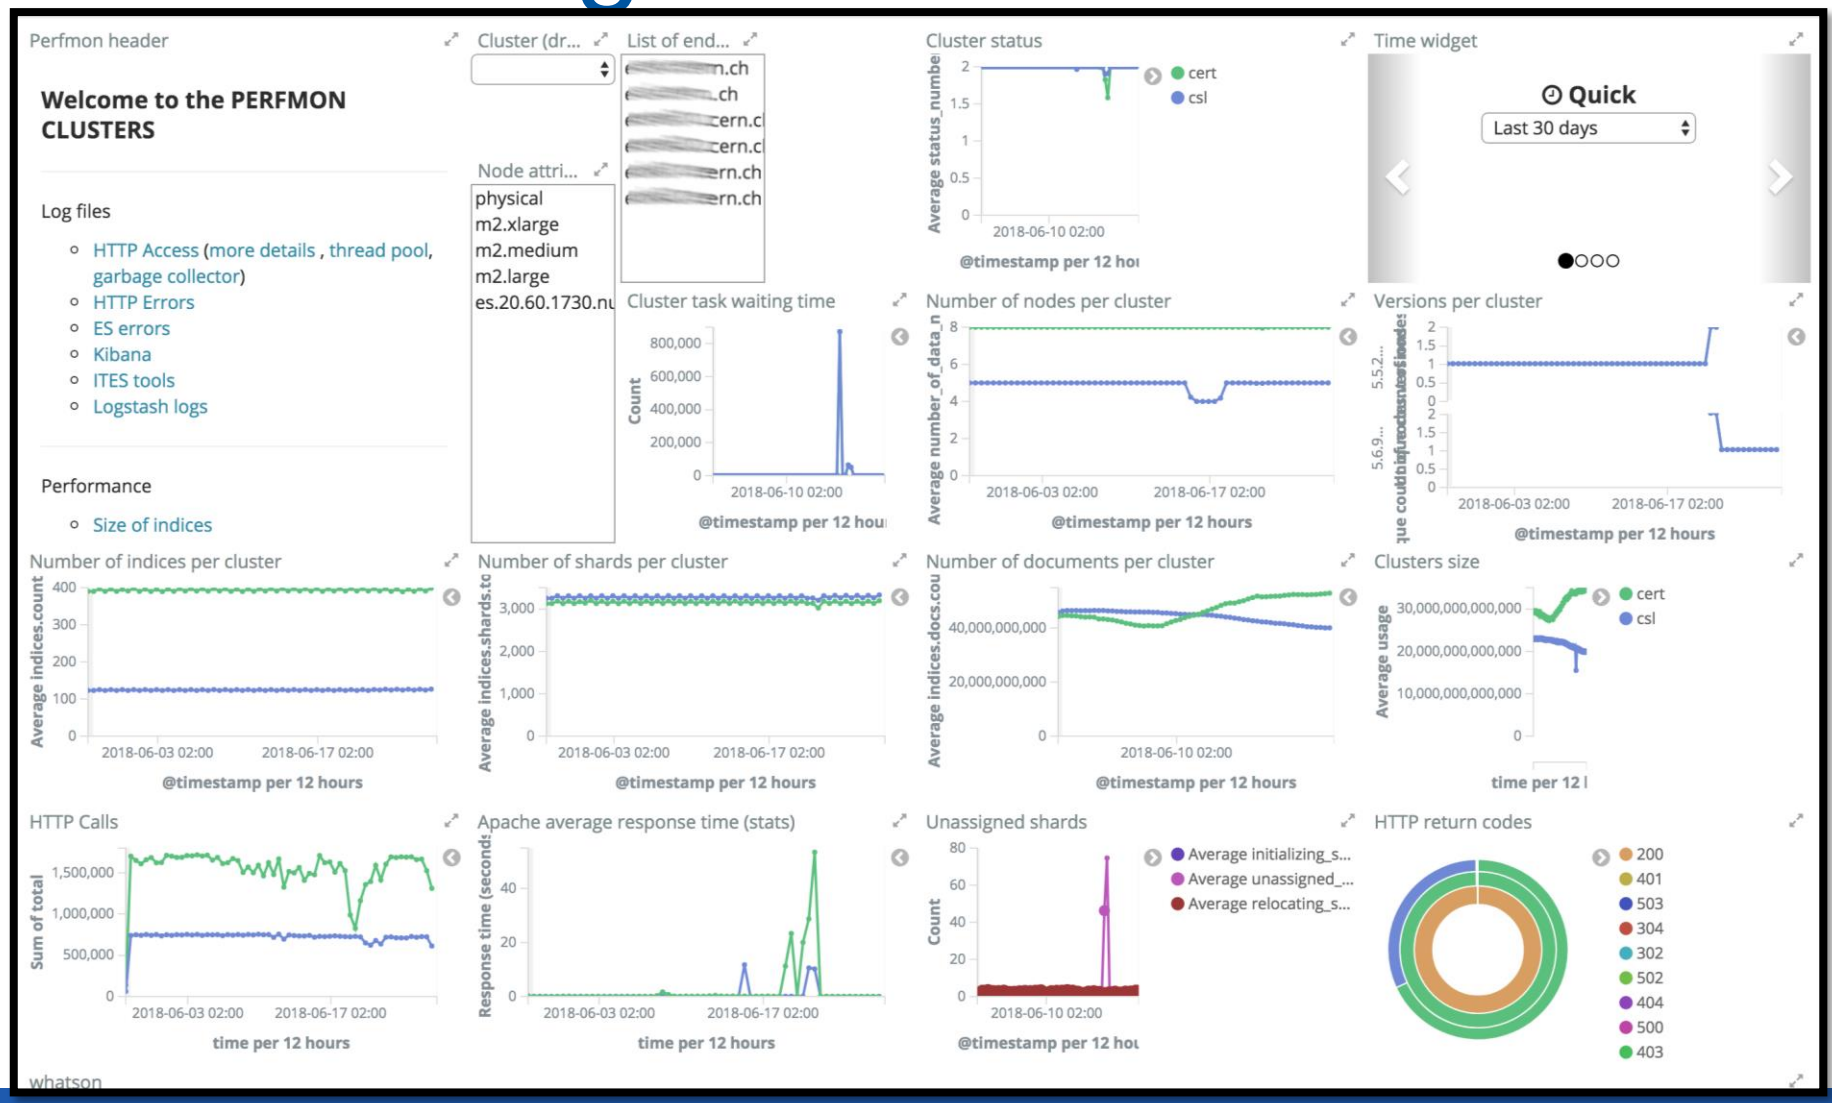

Lessons learned (I)

- Better machines → better performance
  - SSD whenever possible
- Number of indices/number of shards
  - O(10GB) data per shard
    - Indices < 1GB → single shard
    - Indices =~ 100GB → ~10shards
  - ~ 20 shards per 1GB memory of node
    - 3x 32 GB data nodes→~2000 shards
- Field types
  - Define mappings in advance
  - Use index templates

Lessons learned (II)

- Amount of data
  - Too much data → slow cluster
  - Keep summaries in different indices

Apache access statistics:

- 50 GB per day
- Fields: date, cluster, alias, client, time, size, request, response, verb…
- Kept for seven days

Apache summary statistics

- Fields: date (trunc to hour), cluster, alias, response, avg_time, avg_size, count
- Kept forever
- 200MB per year

Summary

- Goal:
  - Provide Elasticsearch/Kibana clusters
  - Most instances shared by multiple users
  - With ACL to isolate users
- How:
  - Open-source tools:
    - Apache (authentication)
    - Read-Only-Rest (authorization)
    - Kibana-Own-Home (multiple tenants)
  - Currently: ~20 clusters, 120 endpoints