
CERN

CERN Site Report

Jérôme Belleman

CERN

European Organisation for Nuclear Research
(**C**onseil **E**uropéen pour la **R**echerche **N**ucléaire)

- Founded in 1954, 22 member states today
- World's largest particle physics laboratory
- Located at Franco-Swiss border near Geneva
- $\approx$ 2 500 staff members, $>$ 16 000 users

IT Department at CERN

Enabling the laboratory to fulfil its mission

- Main data centre in Meyrin site
- Wigner data centre in Budapest, 23 ms away
- Connected via 3 dedicated 100-Gb/s links
- Service resources in both sites where possible (+ disaster recovery)

IT Department at CERN

Enabling the laboratory to fulfil its mission

- Main data centre in Meyrin site
- Wigner data centre in Budapest, 23 ms away
- Connected via 3 dedicated 100-Gb/s links
- Service resources in both sites where possible (+ disaster recovery)

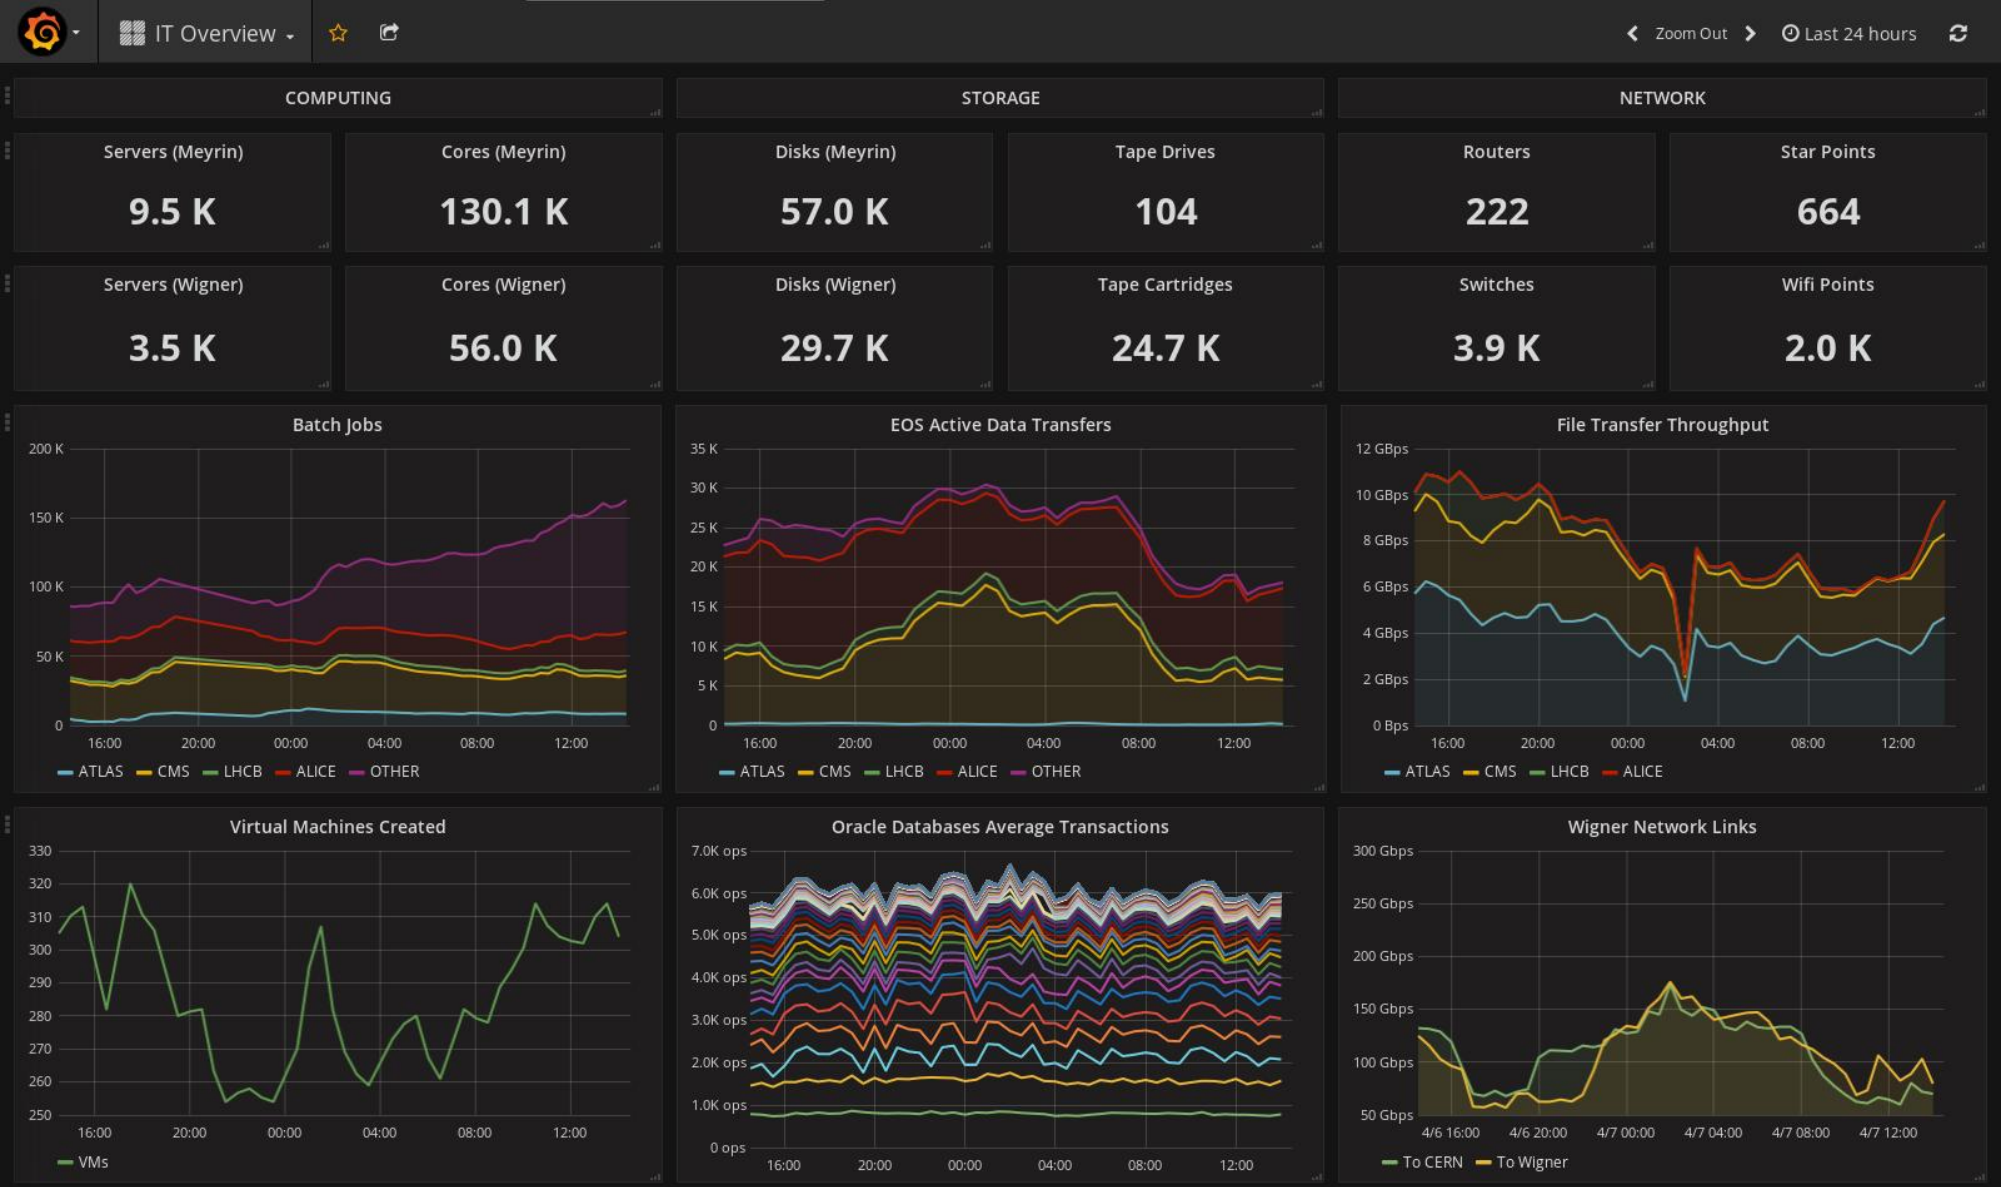
IT Overview
Zoom Out
Last 24 hours
COMPUTING
STORAGE
NETWORK
Servers (Meyrin)
9.5 K
Cores (Meyrin)
130.1 K
Disks (Meyrin)
57.0 K
Tape Drives
104
Routers
222
Star Points
664
Servers (Wigner)
3.5 K
Cores (Wigner)
56.0 K
Disks (Wigner)
29.7 K
Tape Cartridges
24.7 K
Switches
3.9 K
Wifi Points
2.0 K
Batch Jobs
200 K
150 K
100 K
50 K
0
16:00
20:00
00:00
04:00
08:00
12:00
ATLAS
CMS
LHCB
ALICE
OTHER
EOS Active Data Transfers
35 K
30 K
25 K
20 K
15 K
10 K
5 K
0
16:00
20:00
00:00
04:00
08:00
12:00
ATLAS
CMS
LHCB
ALICE
OTHER
File Transfer Throughput
12 GBps
10 GBps
8 GBps
6 GBps
4 GBps
2 GBps
0 Bps
16:00
20:00
00:00
04:00
08:00
12:00
ATLAS
CMS
LHCB
ALICE
Virtual Machines Created
330
320
310
300
290
280
270
260
250
16:00
20:00
00:00
04:00
08:00
12:00
VMs
Oracle Databases Average Transactions
7.0K ops
6.0K ops
5.0K ops
4.0K ops
3.0K ops
2.0K ops
1.0K ops
0 ops
16:00
20:00
00:00
04:00
08:00
12:00
Wigner Network Links
300 Gbps
250 Gbps
200 Gbps
150 Gbps
100 Gbps
50 Gbps
4/6 16:00
4/6 20:00
4/7 00:00
4/7 04:00
4/7 08:00
4/7 12:00
To CERN
To Wigner

External AFS Disconnection Test

- In the context of the AFS phaseout
- Disconnection test: 15 March, 9:00 CET
  https://twiki.cern.ch/twiki/bin/view/IT/AfsPhaseoutExtDisco
- Flush AFS dependencies
- Few sites affected, little feedback from experiments
- Will do again, for longer

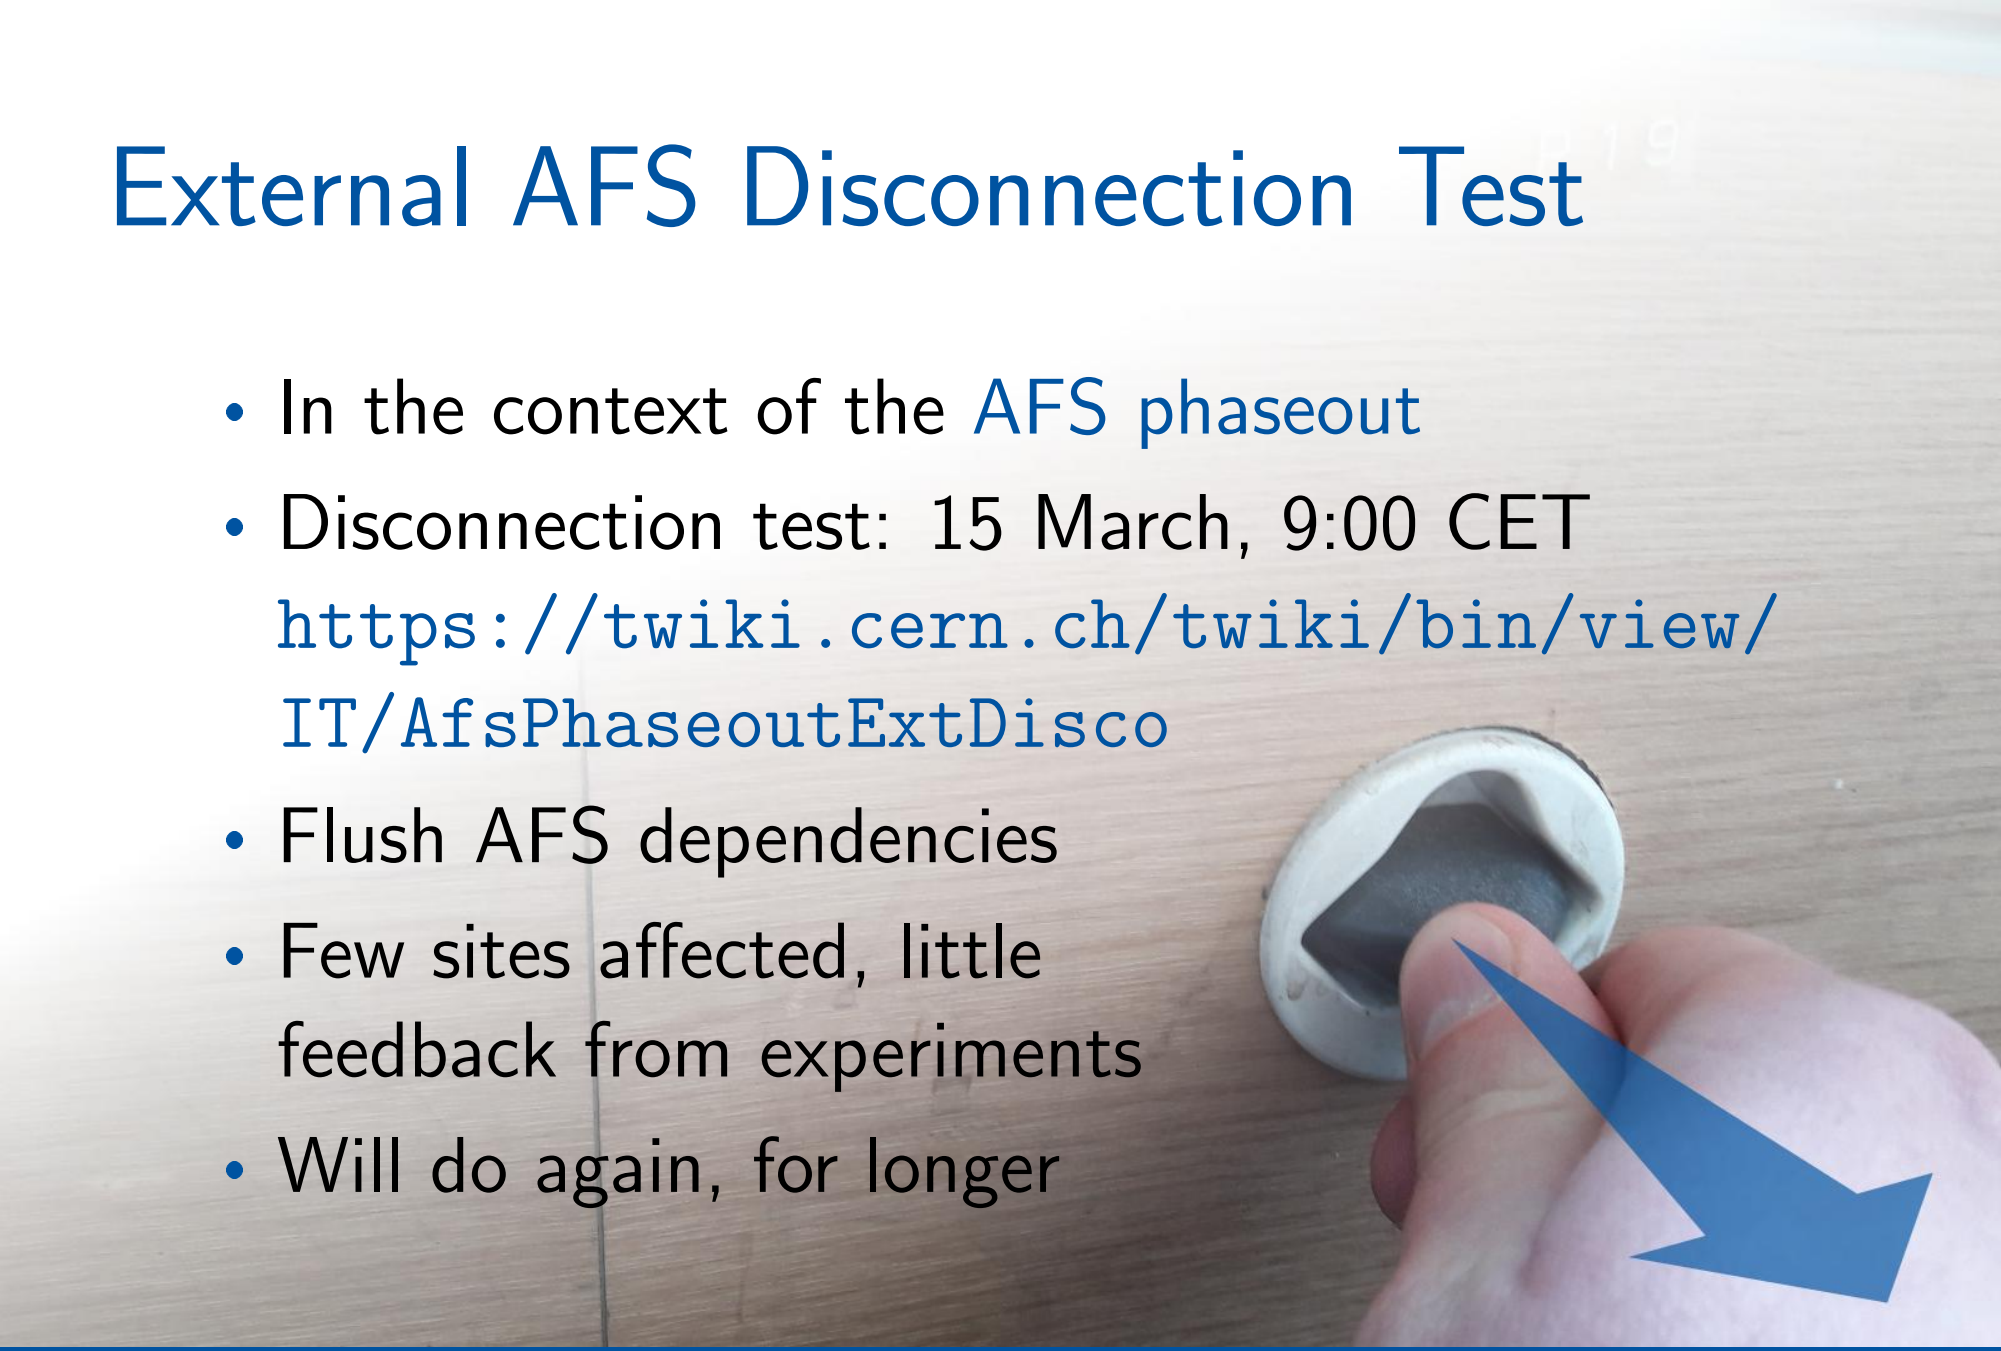

Storage: Tape Ops and Development

- CERN Tape Archive: 190 PB
  - 49 PB LHC data in 2016, peak of 11 PB in July
- Working on CERN Tape Archive (CTA)
  - Tape back-end for EOS
  - Thin layer on top of existing CASTOR tape server software

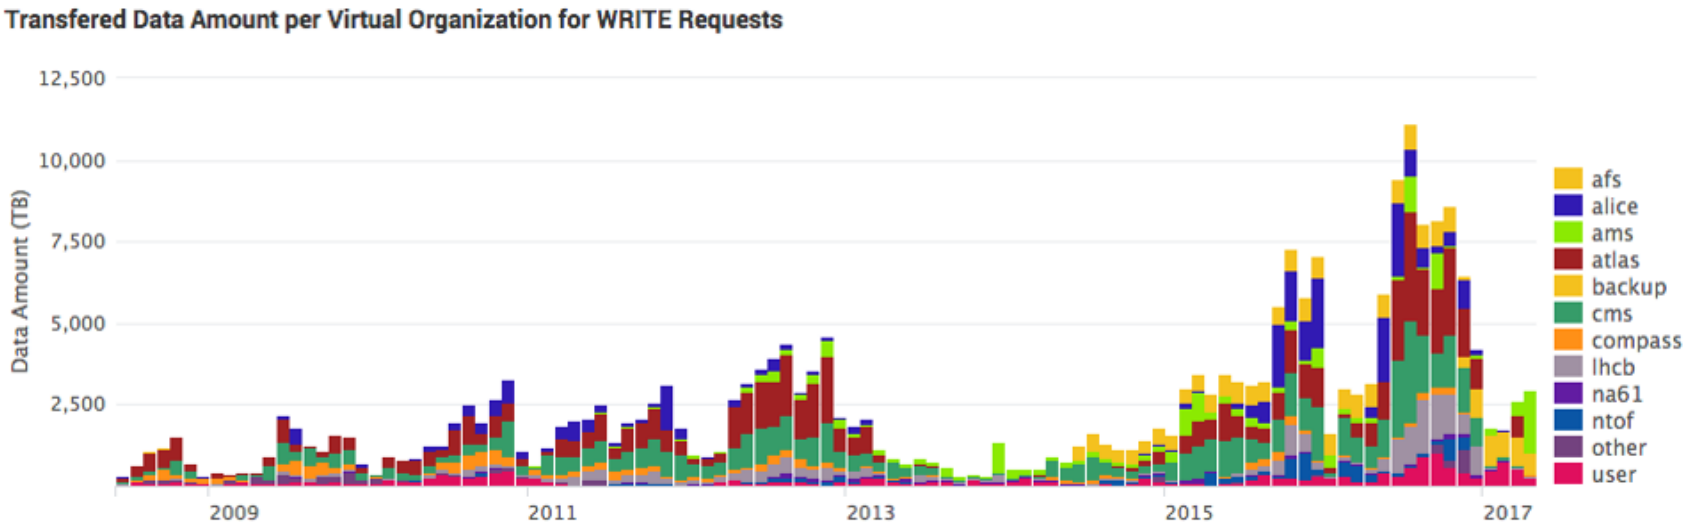

Storage: Disk Operations (I)

Impressive growth of EOS, CERNBox

- 170 PB, 1.3 billion files
- Preparing next delivery of 100 PB
- 9500 CERNBox users

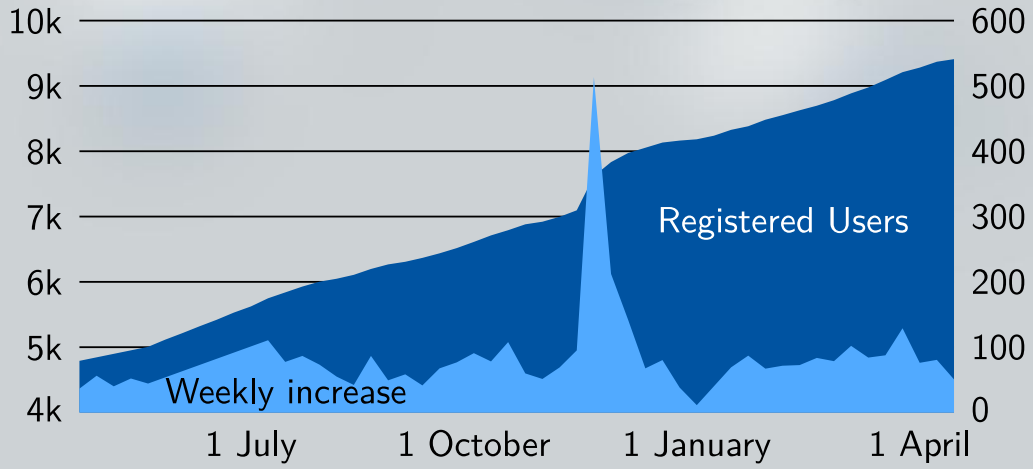

Storage: Disk Operations (II)

CASTOR:

- Preparation for next run
- Ceph-pool for ALICE

CVMFS improvements

→ Luca/Julien's talks

Storage: Developments and Analytics

- EOS developments (CITRINE)
  - High availability and name space scalability
  - Next generation FUSE
  - Workflow extensions
- DPM (now $\approx$ 77 PB/138 sites)
  - Minor DMLite update, working towards further consolidation
  - Dynafed (based on same DMLite) to get updated GeoIP plugin
- FTS3@CERN ($\approx$ 18 M files, $\approx$ 5.7 PB moved/week)
  - Upgraded to v3.6.8
  - Work on tape monitoring features
- HADOOP
  - HA setup and routine backup to CASTOR
  - Upgraded to CC7
  - Additional test and QA clusters

Linux and CentOS

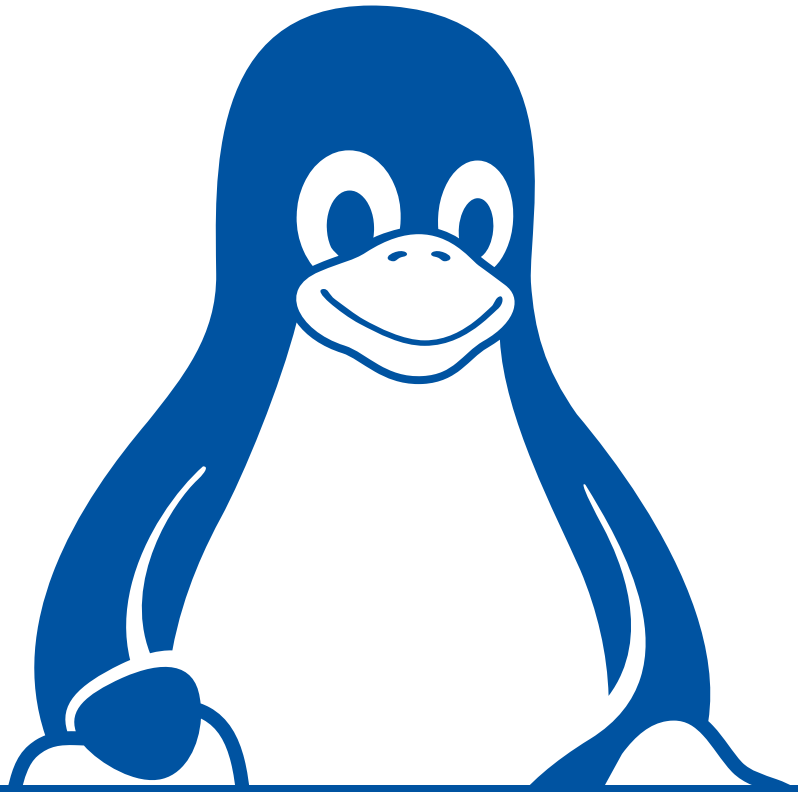

→ Ulrich's talk

Batch Services: High-Throughput

- 130k CPU cores, expecting 200k end of 2017
- 650k jobs/day
- 50% LSF, 50% Condor
- Moving major users and all LHC by Q3 2017
- Training, staff to help migrate
- Understanding CPU efficiency
- Better HW, multicore support, CC7, cgroups
- Investigating expanding the pool with lower SLA work, by using under-utilised capacity, e.g. disk servers

Batch Services: Fifemon

Batch Services: High-Performance

- Goal: MPI, shared memory across nodes, InfiniBand
- 5k-cores SLURM batch system being deployed
- Backfill via HTCondor/SLURM interface

Batch Services: LHC@Home

Volunteer computing: LHC@Home, a single BOINC project, multiple applications

- Accelerator physics simulations (SixTrack) as native BOINC app
- Simulations from ATLAS, CMS, LHCb and Theory running under CernVM and VirtualBox
- Job management back-end integrated with Condor – potentially wide range of (low-IO/high CPU) applications

Cloud Services (I)

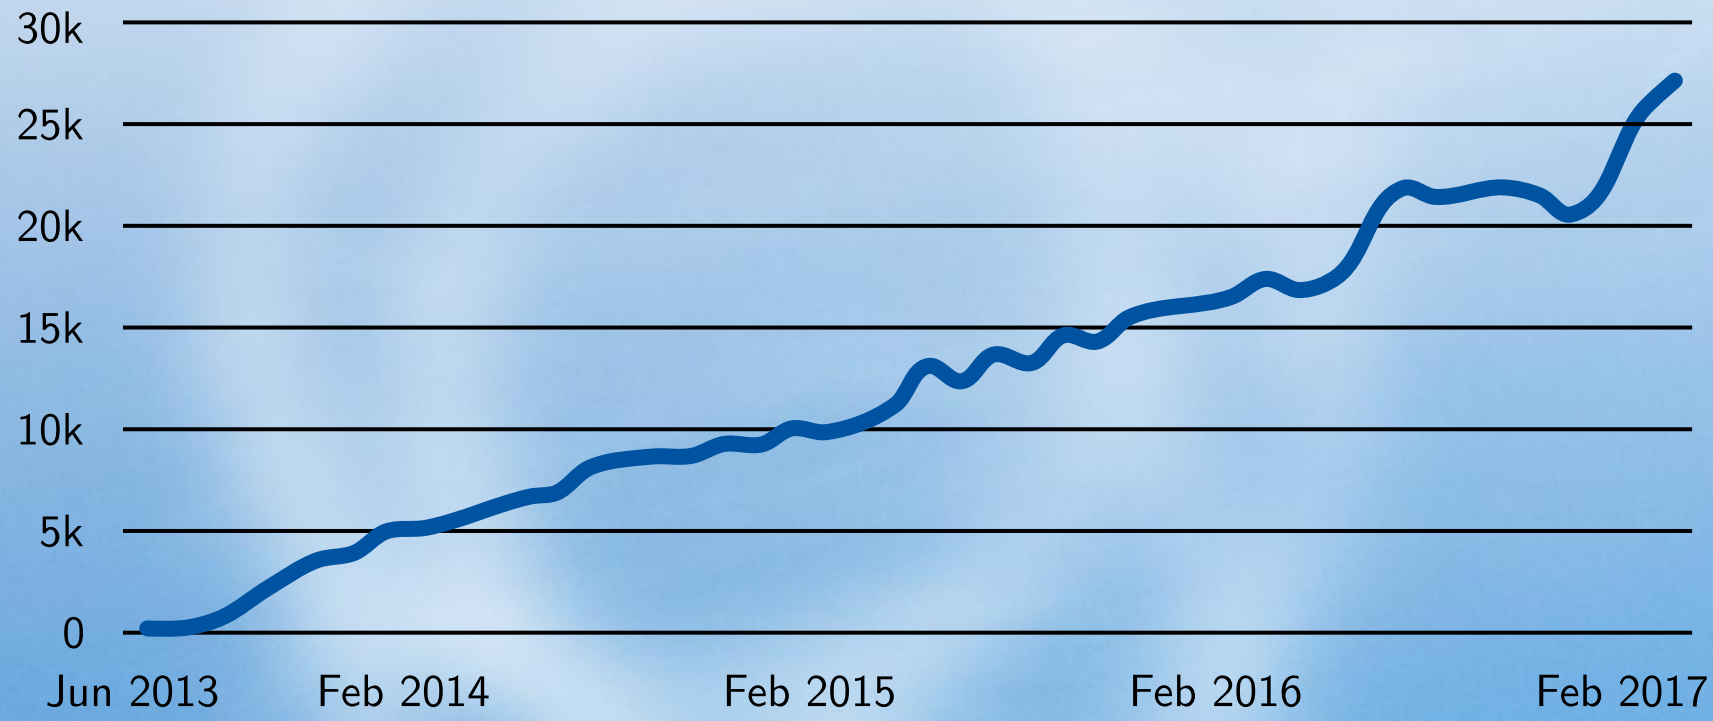

Number of running VMs

Cloud Services (II)

Operations:

- H/W retirement, live/cold migration of 5k VMs
- Default hypervisors will run SSDs
- nova-network → Neutron this year

→ Luis's talk

Configuration Management

Puppet 4 → Jérôme's talk

Databases (I)

SQL:

- Oracle Application Express 5: rapid web app dev
- InfluxDB now available on Database on Demand
- Large growth of PostgreSQL/MySQL/InfluxDB services

Hadoop service update:

- HDFS, YARN, Hive, Impala, HBase, ZooKeeper
- Off-line analytics, accelerator logging system
- Data mustn't be lost → High availability

Databases (II)

Databases applications:

- 560 Oracle WebLogic 12.1.3 servers
- High availability platform for critical services (AIS, EDMS, . . .)
- ORDS for APEX and PL/SQL applications
- SAML 2 native WebLogic implementation for SSO
- WebLogic 12.2.X on top of OpenStack Magnum (ongoing)

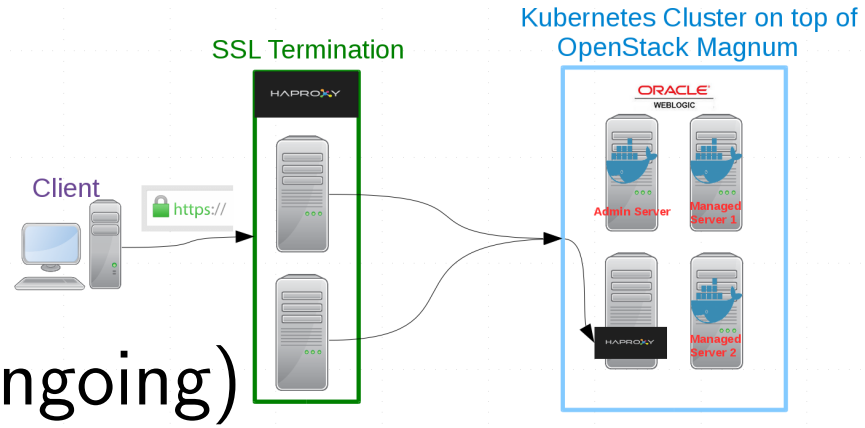

Databases (III)

Java web hosting service:

- Platform as a Service
- 205 Apache Tomcat servers
- Phase-out and replace with OpenShift (ongoing)

Databases (IV)

Apache Kafka pilot on-demand service:

- Distributed streaming
- Scalable, fault-tolerant, low latency
- Processes streams of records as they occur
- Dedicated/shared clusters
- Kerberos, ACLs, SSL
- Accelerator Logging, SCADA Systems, Security Team, Accelerator Post-Mortem, IT Monitoring

Unified Monitoring

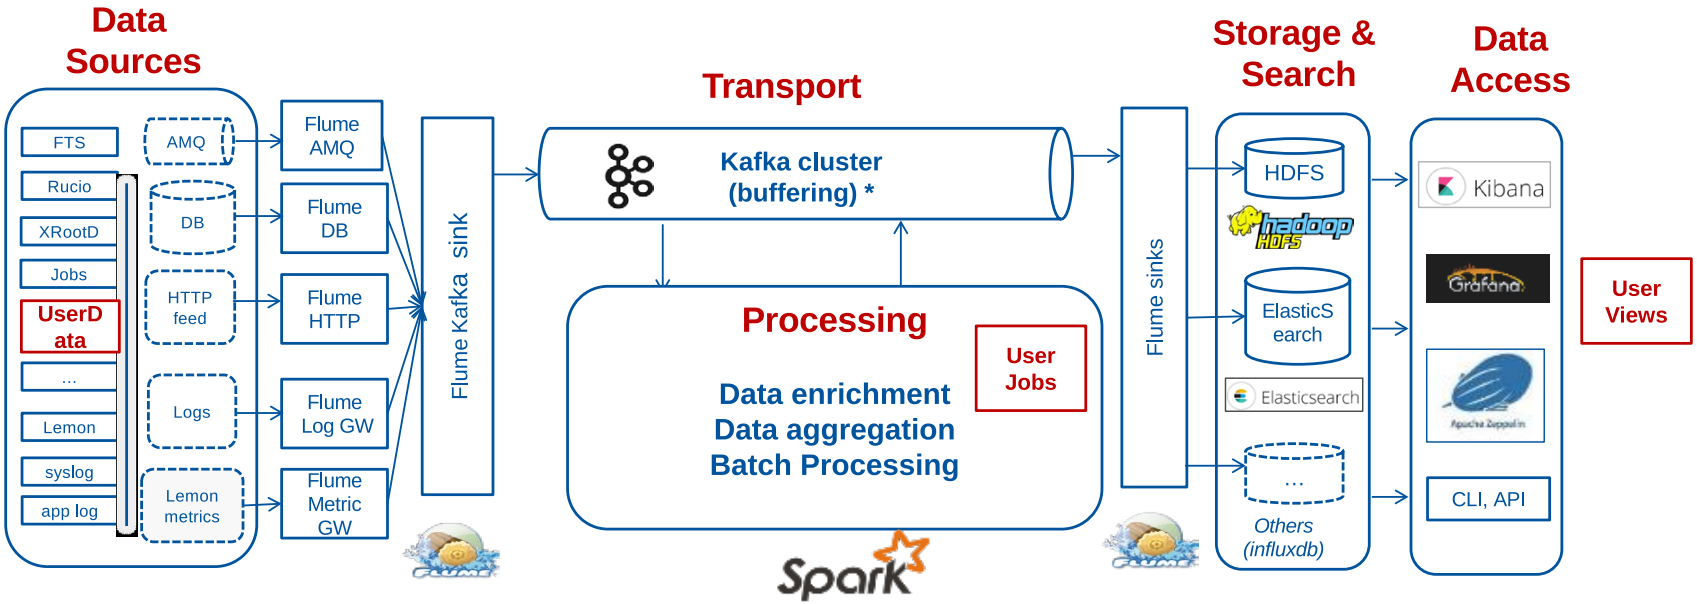

500 GB/day, 48h Kafka

http://monit.cern.ch

→ Jarka's talk

Centralised Elasticsearch Service

→ Ulrich's talk

Security

SSO for trusted organisations:

- eduGAIN supports 1000s of organisations
- CERN supports Sirtfi-compliant organisations

Future authorisation service:

- To manage authorisation for apps
- App-specific roles

WhiteHat training:

- Pen-testing for improving security and resilience
- Trainings given regularly

TWiki (I)

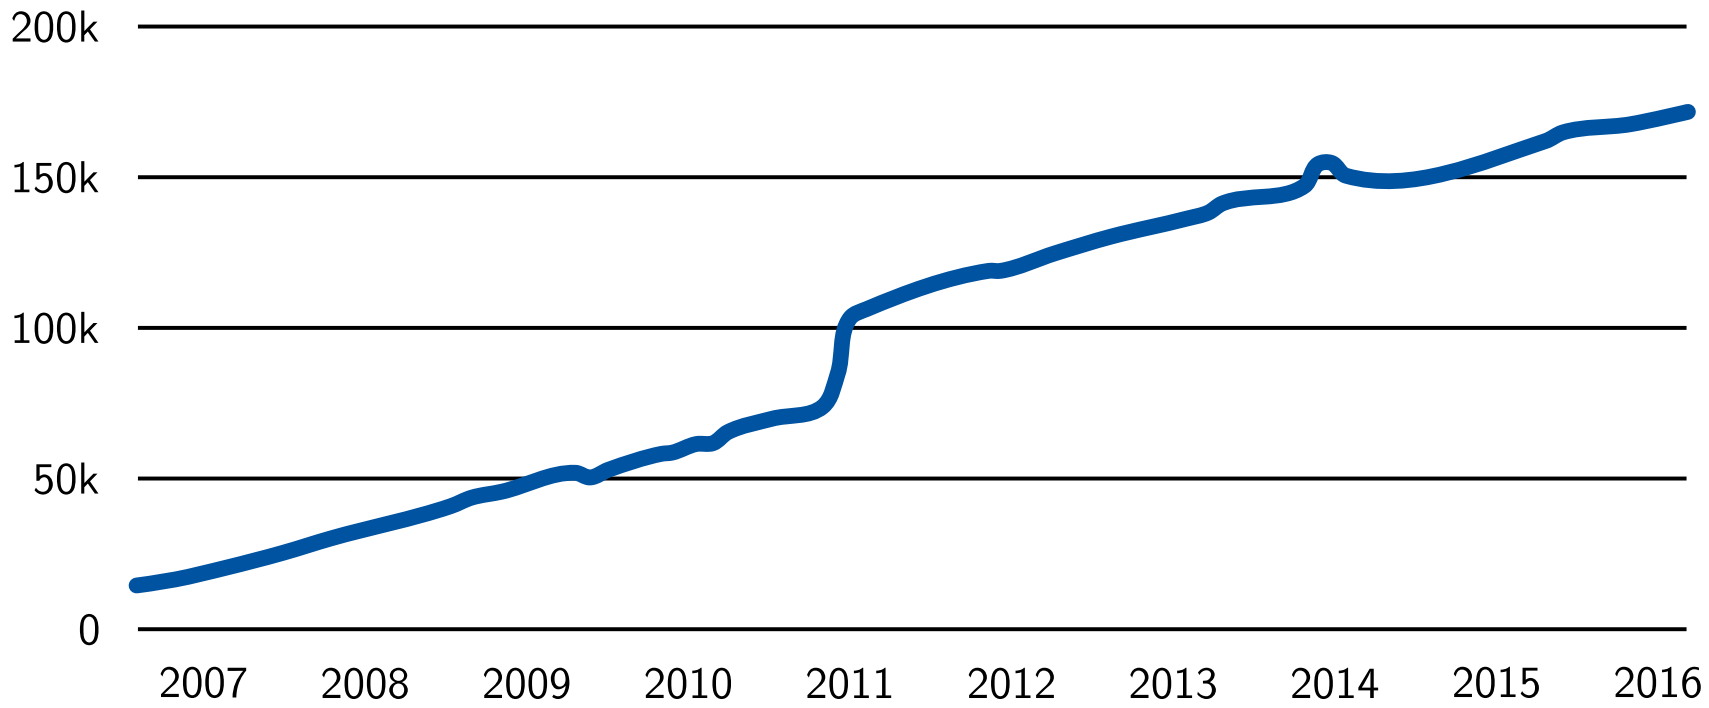

Number of topics (pages) → Used more than ever

TWiki (II)

- TWiki 6: better editor, dashboards, columns
- Main users: LHC Experiments
- People tend not to remove old topics → Archive

Next year is the 15^{th} year of TWiki at CERN

Talks from CERN This Week (I)

- *The Computing Resource Information Catalog*, Alessandro Di Girolamo
- *Computer Security Update*, Liviu Valsan
- *Building and operating a large scale Security Operations Center*, Liviu Valsan
- *A Hard Puppexit from 3 to 4*, Jérôme Belleman
- *System testing service developments using Docker and Kubernetes: EOS + CTA use case*, Julien Leduc
- *CERN IT-Storage Strategy Outlook*, Luca Mascetti, Julien Leduc

Talks from CERN This Week (II)

- *Understanding performance: optimisation activities in WLCG*, Andrey Kiryanov and Andrea Sciaba
- *CERN Computing Facilities Update*, Wayne Salter
- *Centralising Elasticsearch*, Ulrich Schwickerath
- *Understanding the performance of benchmark applications*, Luca Atzori
- *CERN Linux services status update*, Ulrich Schwickerath
- *CERN Cloud service update: Containers, migrations, upgrades, etc.*, Luis Pigueiras

Talks from CERN This Week (III)

- *Unified Monitoring Architecture for CERN IT and Grid Services*, Jaroslava Schovancova
- *HammerCloud extension for Data Centre commissioning*, Jaroslava Schovancova
- *EOS and CERNBox Update*, Luca Mascetti
- *Federated data storage system prototype for LHC experiments and data intensive science*, Andrey Kirianov
- *Deployment of IPv6-only CPU on WLCG – an update from the HEPiX IPv6 Working Group*, Andrea Sciaba

www.cern.ch