




CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



CHIPP - CSCS F2F meeting

UNIBE, Berna

October 4th, 2017



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Tier 2 status and plans

CSCS



1. Statistics

- a. Availability/Reliability
- b. CPU usage
- c. Storage usage


2. Operations


- a. Updates
- b. Main Issues

3. Plans


- a. Resources Overview

Statistics

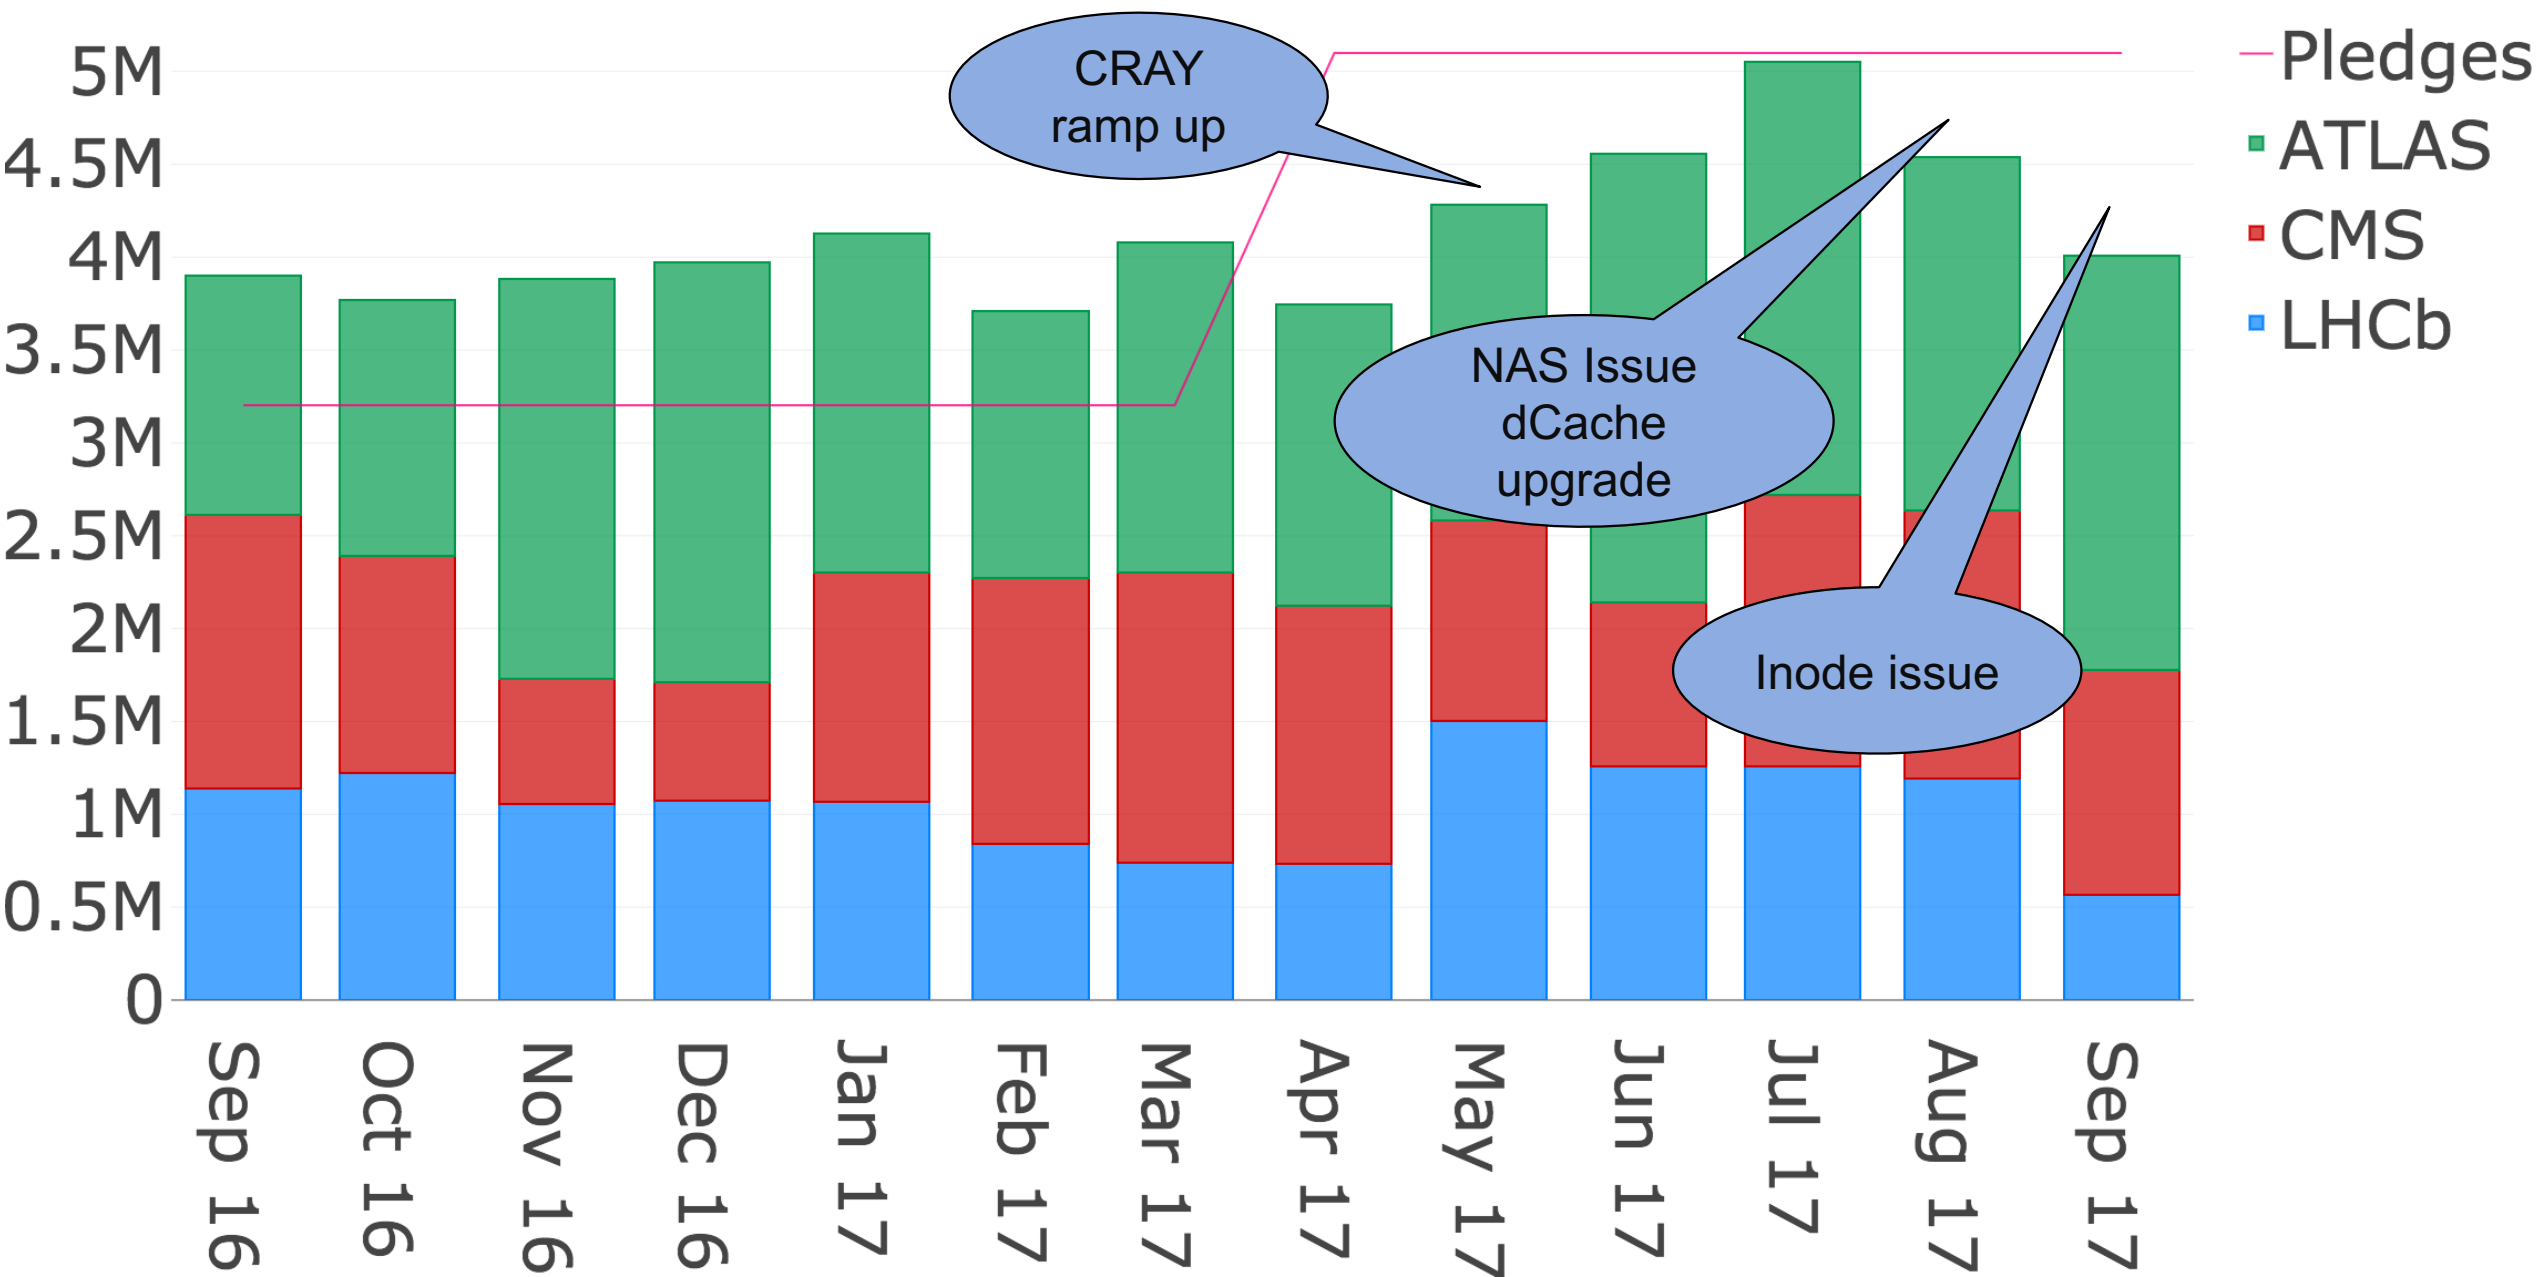





Reliability and Availability (%)



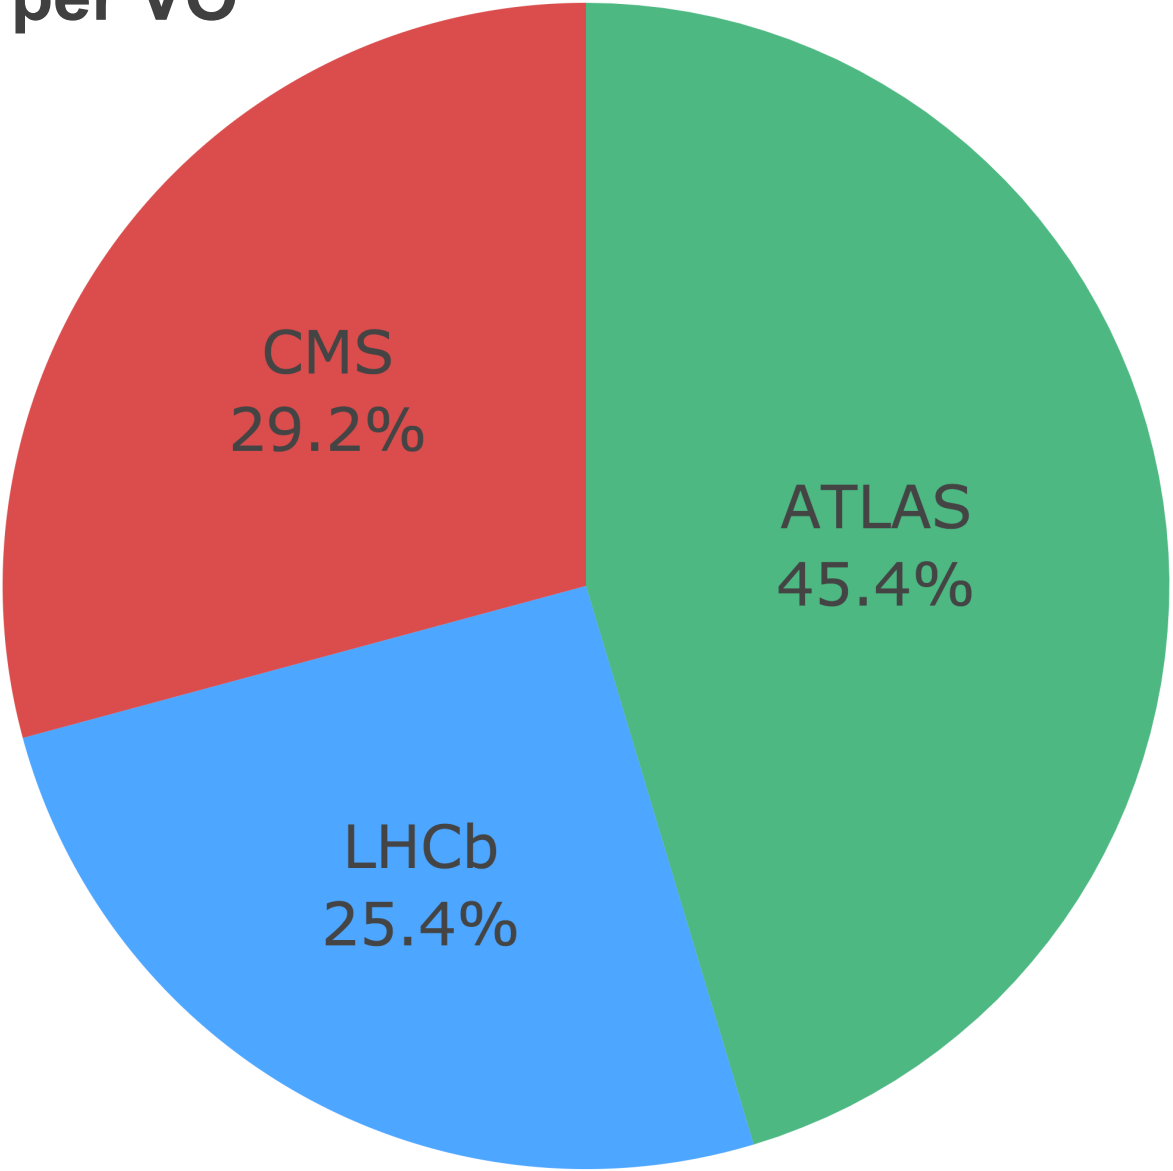
Accounting in CPU hours




Accounting in CPU hours



Statistics – Usage per VO




Statistics – Storage usage (2/2) dCache data distribution



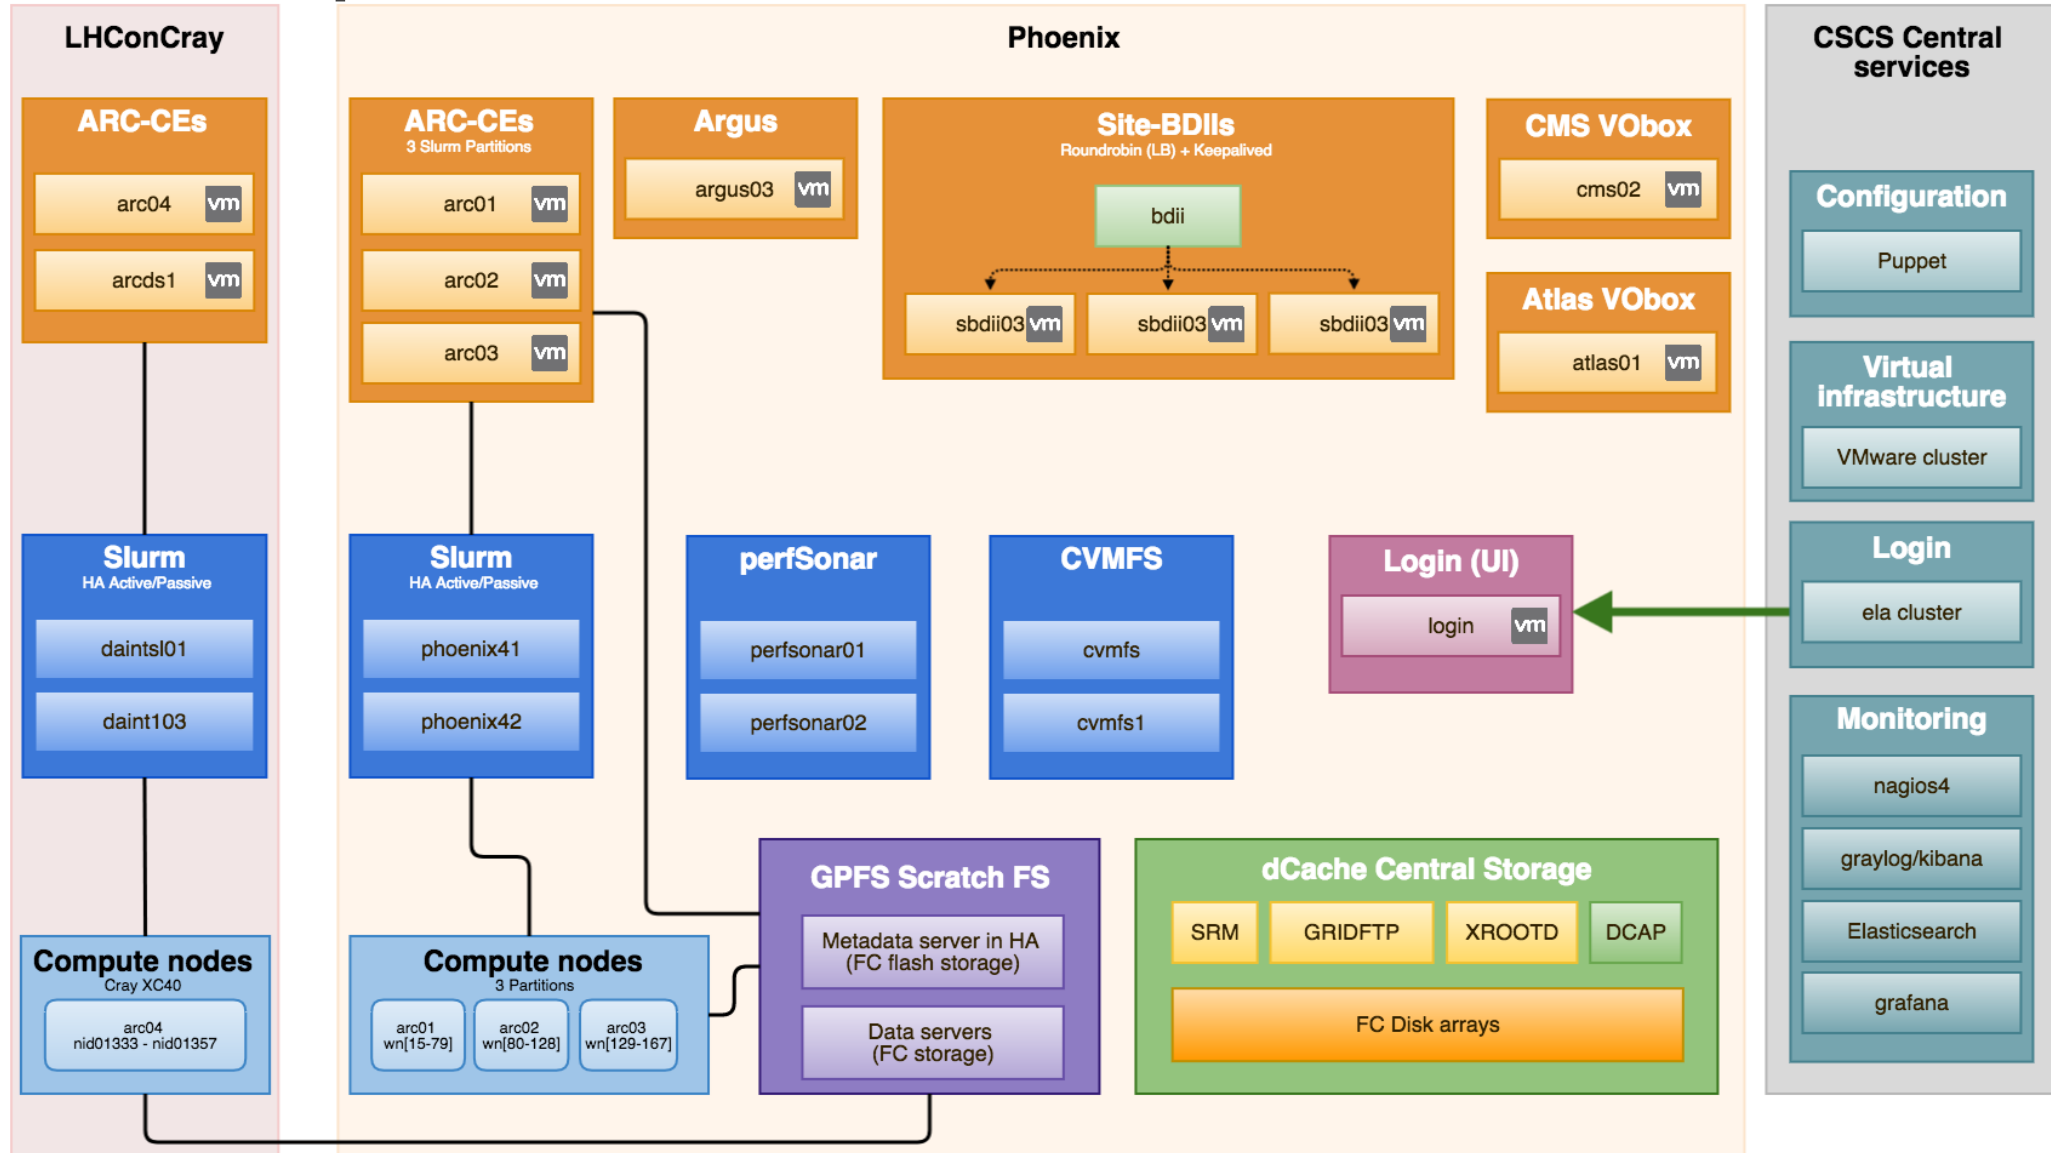
VO	Available space
CMS	1'686'705 GB
ATLAS	1'734'730 GB
LhCB	698'589 GB

Statistics – Storage usage (1/2)



Operations

Operations – Updates: CHiPP Services Overview



dCache Upgrades in 2017

May 2017 – update from 2.13.50 to 2.13.58

- Symptoms: slow performance / unusable service
- Action(s) to fix: rollback to 2.13.50
- dCache support: very unresponsive and not helpful

August 2017 – upgrade from 2.13.50 to 2.16.46

- Symptoms: slow performance / unusable service
- Action(s) to fix: disabled 'sequence scan' in PostgreSQL
- dCache support: absent (even worst than before)

NFS (New CES services over Spectrum Scale)

In 2017 CSCS moved from an appliance (EMC Isilon) to a GPFS cluster (with Cluster Export Services) in order to provide SMB and NFS in the center.

This has brought several benefits

- SAN integration, for more flexible storage allocation
- Easier scalability
- Knowledge consolidation (GPFS is widely used at CSCS)
- Simplified network model

NFS is used for

`/nas/shared/nfs/lcg/backup`

`/nas/shared/nfs/lcg/cm`

`/nas/shared/nfs/lcg/experiment_software`

`/nas/shared/nfs/lcg/home`

`/nas/shared/nfs/lcg/shared`

`/nas/shared/nfs/lcg/shared/phoenix4/users /nas/shared/nfs/lcg/shared/phoenix4/arc_rte /nas/shared/nfs/lcg/shared/phoenix4/slurm_shared`

`/nas/shared/nfs/lcg/shared/phoenix4/gridmapdir`

`/nas/shared/nfs/lcg/backup/apel_archiving /nas/shared/nfs/lcg/shared/cvmfs_cache_shared /nas/shared/nfs/lcg/shared/preprod/vo_tags`

`/nas/shared/nfs/sysbackup/lcg.cscs.ch`

Fetch CRL Issue

The symptom

After The dCache upgrade we noticed a lot of failed transfers due to certificate error.

Debug

We had a long debug session along with Derek, Gianfranco and Thomas

Solution

It ended up that the fetch-crl script (only on Puppet managed Storage Elements) was not updating the certificates correctly.

Outcome

It has been a great example of teamwork. Many thanks to Derek, Gianfranco and Thomas for the help!


VO Monitoring Dashboard

- First result: VO Monitoring Dashboard (work in progress)
 - <http://wiki.chipp.ch/twiki/bin/view/LCGTier2/VoMonitoringDashboard>
- Outcome of the 1 to 1 discussion of the VO with CSCS:
 - We need more interaction with VO in order to implement a proper Dashboard
 - At the moment we are checking individual links for each VO, using the chat to discuss/sync and then tickets (ggus/webirt) to track issues. The process works pretty well.
- Is it working?

Other Operations

Network Operation:

- All old QDR InfiniBand switches replaced by FDR switches (cable replacement on-going)
- 2 EDR switches deployed
- 4 Active/Active uplinks from Phoenix gateways to CSCS core switches (4x40Gb)




Other Operations

- This year operations:
 - ARC updated to 15
 - IB switch replacement to FDR/EDR
 - Twiki migration
 - LHConCray
- Future Operations:
 - CMS lost files
 - dCache upgrade (support for version 2.16 ends in May 2018)
 - IPv6
 - dCache alternatives investigation or we invest more on dCache and Tape

Plans

Plans – Pledges



Phase	Compute power pledged [HS06]	Storage pledged [TB]
Phase J – April 2015	35000	2300
Phase K – April 2016	49000	3070
Phase L – April 2017	78000	4000
Phase M – April 2018	96000	4000
Phase N – April 2019	110000	4300

Storage:

- dCache to keep the pledges
- Renew DDN storage Maintenance

Network:

- Standard network maintenance (ex.ethernet or infiniband cables)

Compute:


- Pledge capacity is 10% lower than install capacity
- Decision is needed for LHConCRAY



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you for your attention.