




**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich




# CHIPP - CSCS F2F meeting


CSCS, Lugano

September 13<sup>th</sup>, 2019


# Accounting per VOs in CPU hours (Aug 19 – Aug 19)




# Accounting in CPU hours (Aug 19 – Aug 19)




# CPU Efficiency EGI(%)




# Reliability and Availability (%)



## Statistics – Usage per VO 1 year (Aug 18 – Aug 19)



## Statistics – Usage per VO 3 months (Jun – Aug 19)



# Statistics – Storage usage (As of September '19)

## ATLAS

**Total:** 1'872 TiB (pledge: 1'840TiB)

**Free:** 207 TiB

## CMS


**Total:** 1'872 TiB (pledge: 1840TiB)

**Free:** 126 TiB

## LHCb

**Total:** 960 TiB (pledge: 920TiB)

**Free:** 378 TiB




**T2 Space allocation**



# Last Issues


- *GPFS/DVS Issue*
  - **Issue:** DVS nodes experienced very high load, causing GPFS to hang and affecting heavily I/O operations -> Job efficiency
  - **Solution:** Some DVS configuration parameters had to be tuned
- *CMS filling completely dCache space*
  - **Issue:** CMS was filling completely the dCache storage causing write transfers to fail due to “no free space left” error
  - **Solution:** Since the “over allocation” was not under the CMS Central Operations control but more on the user part of the storage, along with PSI colleagues we managed to delete unused T3 data, old users home, test data and **some old/unused datasets**

# Last Issues - Fairshare



Consumed Cycles  
March and April


- The priority calculation did not generate a large enough spread based on share of the resource
- Resulted in a near even distribution of compute cycles
- Calculation adjusted to increase the emphasis on Fairshare for CHIPP jobs
  - Changes introduced mid-May
- Running workload immediately readjusted to work on balancing back to allocated shares



Consumed Cycles  
June, July and August

*Note: These charts are for specific periods to show the effects of the change, see accounting charts for overall delivery of cycles*


# Last Issues - CVMFS Thread Hangs



Recurring Pattern

CVMFS threads hang blocking the node which requires a node reboot to clear. The impact is more visible over a weekend. The situation is automatically detected and the node removed from service. Nodes are rebooted and returned to service. Impacts overall compute capacity.

Running Core Count per Account (atlas + cms + lhcb) and partitions (wlcg +...



Pattern Not Observed

New version of CVMFS installed on September 4<sup>th</sup>. Problem appears resolved, but more time is needed to declare it completely solved.

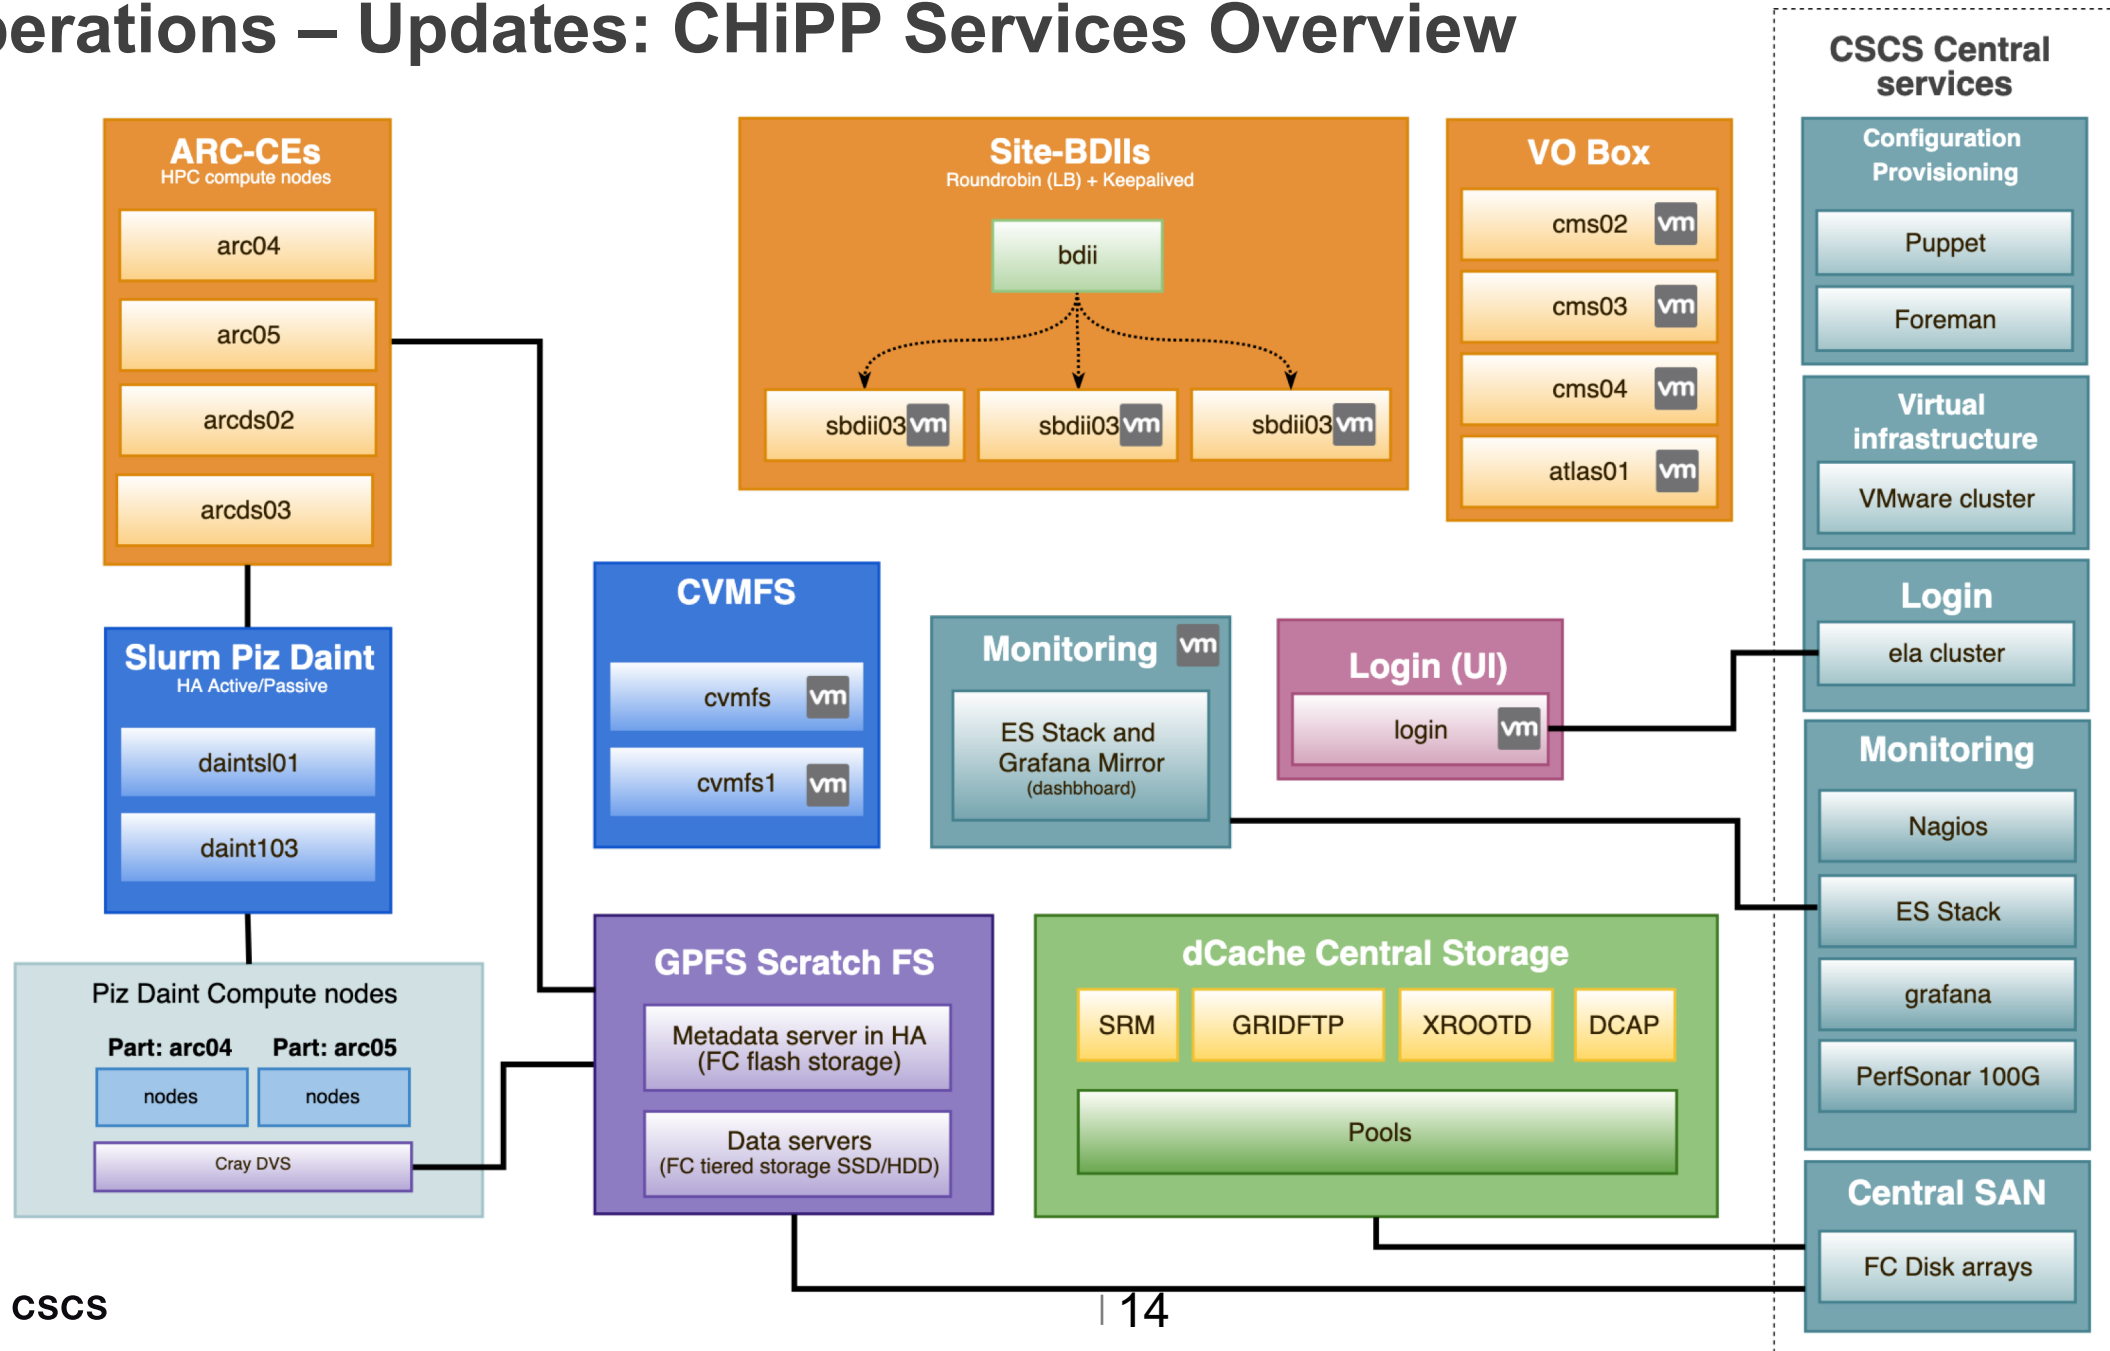
# Last Issues

- ARC hangs
- Atlas internal submission issues
- Yesterday dCache DB issue


# Operations

---

# Operations – Updates: CHiPP Services Overview




# dCache migration and deployment



# dCache Migration Performance


This is the bandwidth we were able to reach while migrating data from the old dCache location/nodes to the new ones






# dCache Architecture Evolution (Infrastructure)


## Past deployment




## 2019



## Since 2015



## Future (to be tested)



# Implemented Changes

- GPFS / Scratch

This year we did some important improvements to the GPFS infrastructure that is providing the scratch filesystem (all changes have been applied online)

- Moved from Infiniband to Ethernet 25G
- Moved the HDD pool to a new storage system (DDN SFA12k to DELL-EMC SC9k)
- Doubled the Fibre Channel bandwidth

Info: Since April 27<sup>th</sup>, **20PiB** have been **written** and **44PiB** have been **read** to/from the filesystem.

# Implemented Changes

## ■ Phoenix Shutdown



- After many years of successful service, Phoenix was permanently removed from service on March 4<sup>th</sup>, 2019. Phoenix began operation in 2007 and provided computational cycles for the Physics community for 12 years, ending in 2019.

## ■ Phoenix Nodes Distributed

- Phoenix continues to make a contribution to the Physics community after its decommissioning. The ATLAS (Bern) community were sent 114 compute nodes and the LHCb (UZH) community were sent 8 compute nodes from the Phoenix hardware.

## ■ Dashboard Enhancements

- Last fall, CSCS released an active dashboard to the VOs. A few minor enhancements (tweaks) were made post release, but since then, it has been running stable. Please let us know if there are additional needs that should be considered.

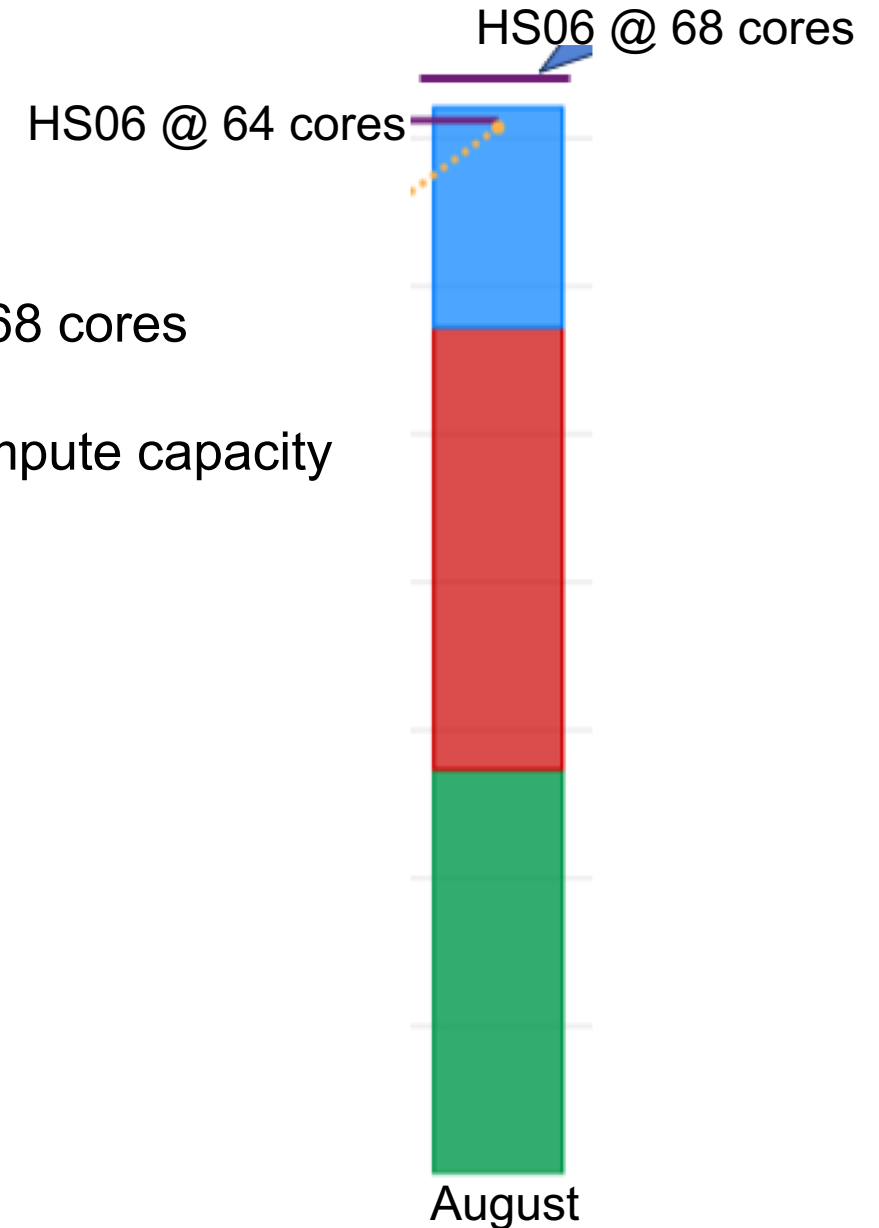
# Implemented Changes

- PerfSonar (100Gb)
- Services reinstalled on new Puppet
- New ARC for Daint
- *QuoVadis CA: CSCS Application approved and new procedure to request host certificates successfully tested.*

# Ongoing Activity

## ■ 68 Core Experiment (*early results*)

- Compute nodes were configured to use 64 cores
- On August 1<sup>st</sup>, the compute nodes were adjusted to run at 68 cores
- Testing for delivery of cycles as well as stability
- The gap to the 68 core line is most probably due to lost compute capacity from CVMFS thread hangs (*problem may be resolved*)
- The early results are promising, but more time is needed



# Future Operations

1. Upgrade ARC for Daint
2. Sarus
3. dCache Upgrade to 4.2/5.2

Daint update:

- CLE update (November '19)

Center-wide outage (Sept '19, Jan '20)

Increase capacity as planned (April '20):

- KHS06 161(179 raw)
- dCache 5.3 PB

# Dissemination

- HEPiX Spring/Autumn (Dino)
- NorduGrid (Dino)
- dCache Workshop (Dario)
- CERN OpenLab Meetings Geneva/Lugano (Miguel, Pablo)

# Monthly Operation Meeting

How to improve and make this meeting more effective?

- Track all issues and downtimes (it requires everyone at the meeting)
- Accounting figures (crosscheck)
- Suggestions?






**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich



**Thank you for your attention.**