



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETHzürich



LHConCRAY status update

CHIPP/CSCS F2F meeting

Miguel Gila, CSCS

January 27, 2017

LHConCRAY status update



- Updates/changes since last meeting
- System Statistics
- Technical conclusions



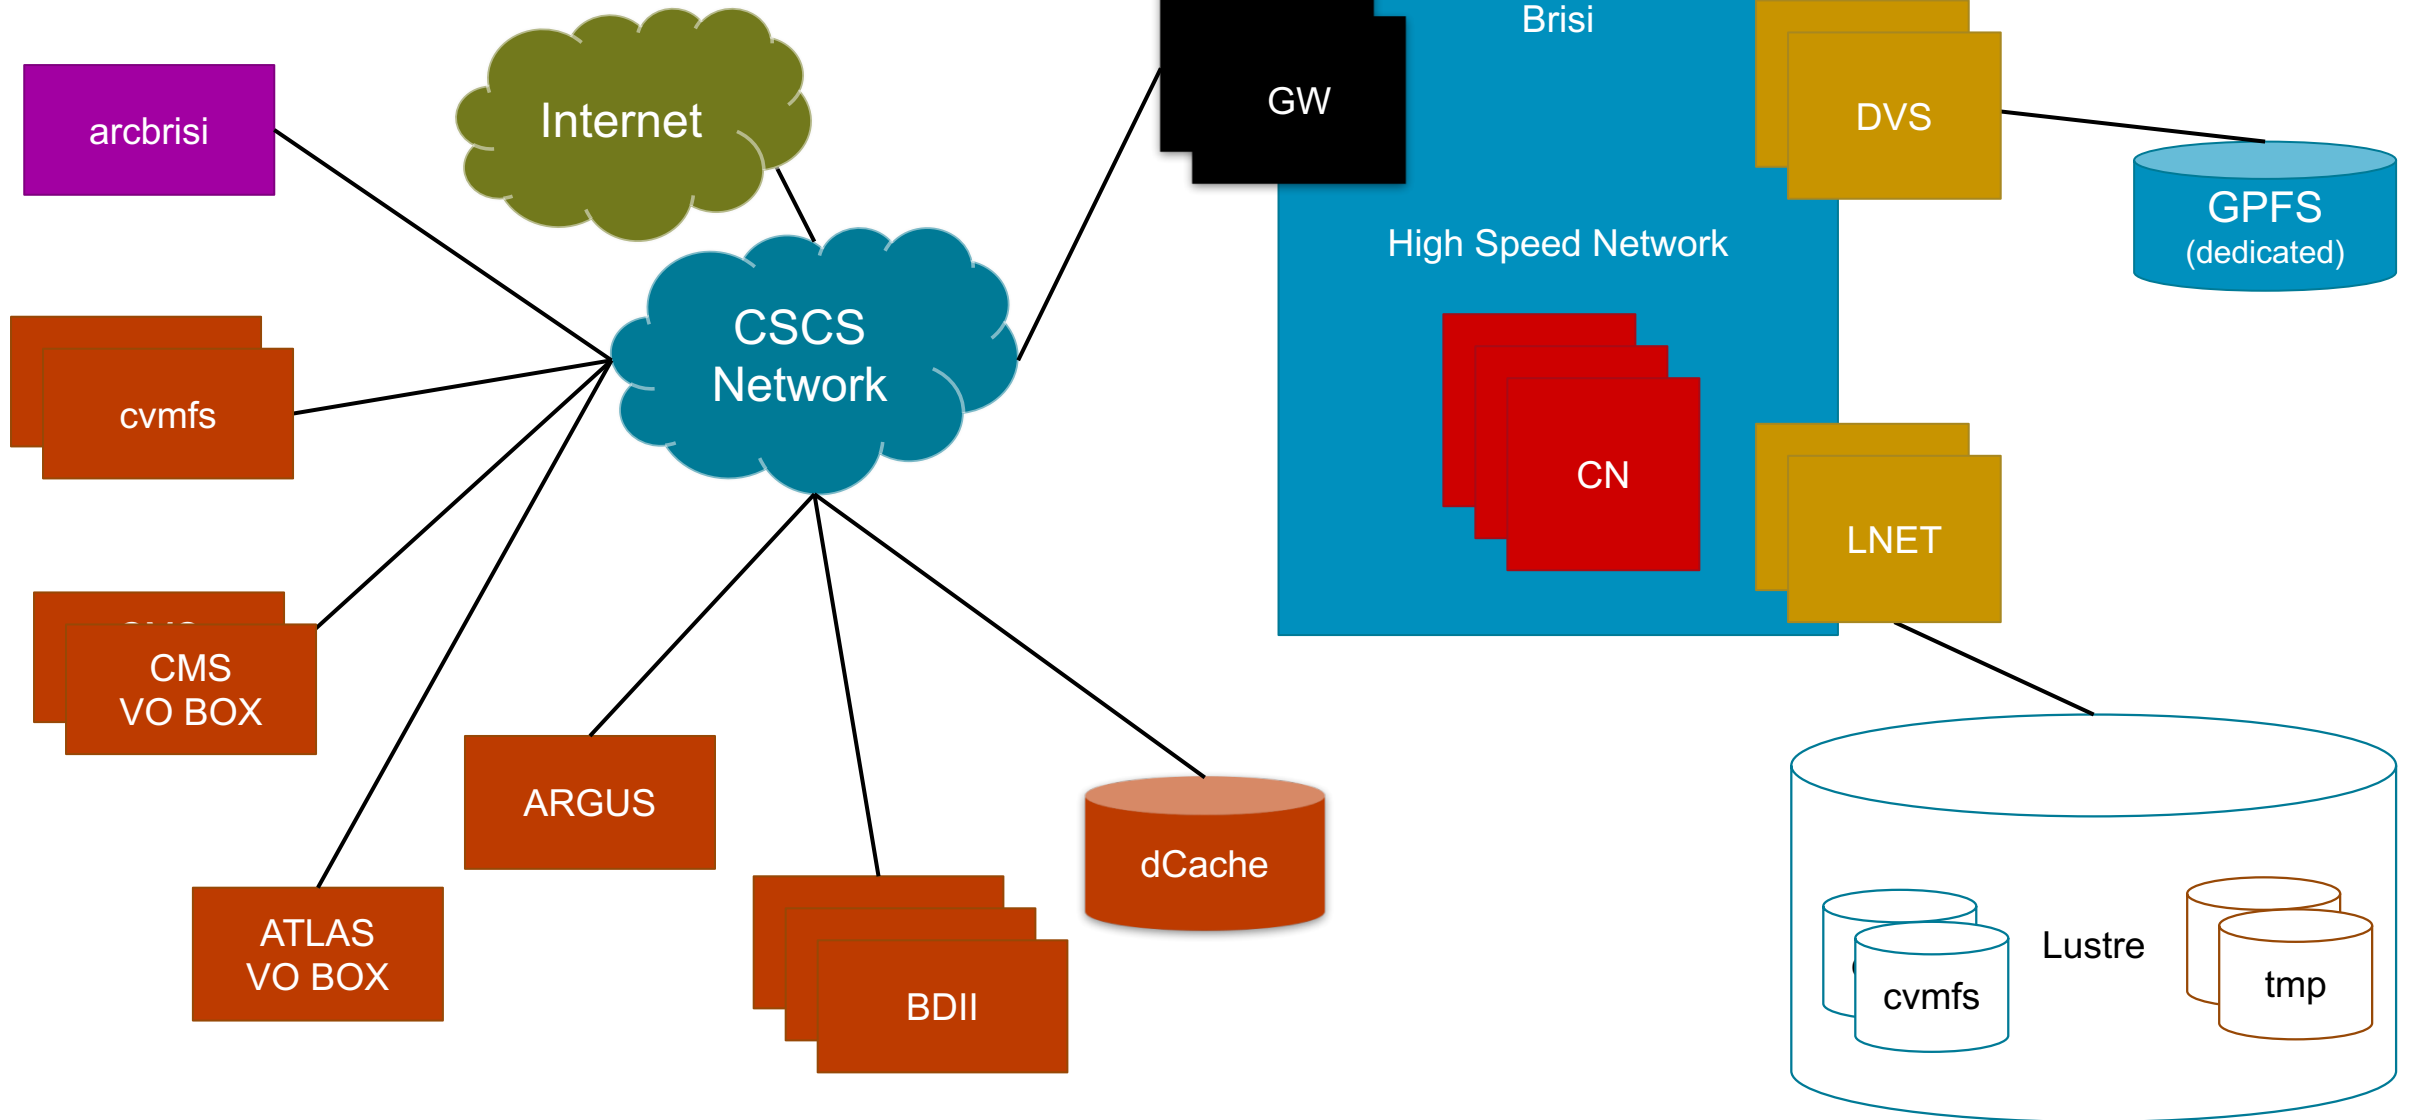
CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Updates/changes since last meeting

System schematics (before)




Changes since last meeting (1/2)

- GPFS over DVS
 - First time GPFS is mounted over 2 DVS with 40GbE devices across IB/eth network bridge
- CVMFS update
 - CVMFS updated to latest release, built in-house for CLE 6.0
- Complete ARC submission script within Shifter
 - affecting ATLAS mostly due to how ARC executes the RTEs
- Detected ATLAS over utilization of memory
 - This led to certain jobs being killed by Slurm
 - Moved to PSS instead of RSS to account for memory and artificial increase of memory requirement

Changes since last meeting (2/2)

- Network performance of VMs
 - First 10GbE enabled VM at CSCS
 - 40GbE/FDR IB LNET router
- 'Large scale' ATLAS test
 - ~1000 core utilized by ATLAS, ~140 by CMS and ~140 by LHCb
 - 30core/node utilized (64GB/RAM per node)
- Datamover for ARC: currently a physical node is configured to run stage in/out tasks.
 - At the moment, useful only for ATLAS.
 - Could extend this to a small mini-cluster of 3 or 4 machines.
 - First dedicated data mover service for ATLAS on 40GbE/FDR IB

Final system (simplified schematics)



IB FDR 56Gbps via IB/ETH bridge (up to 80Gbps)

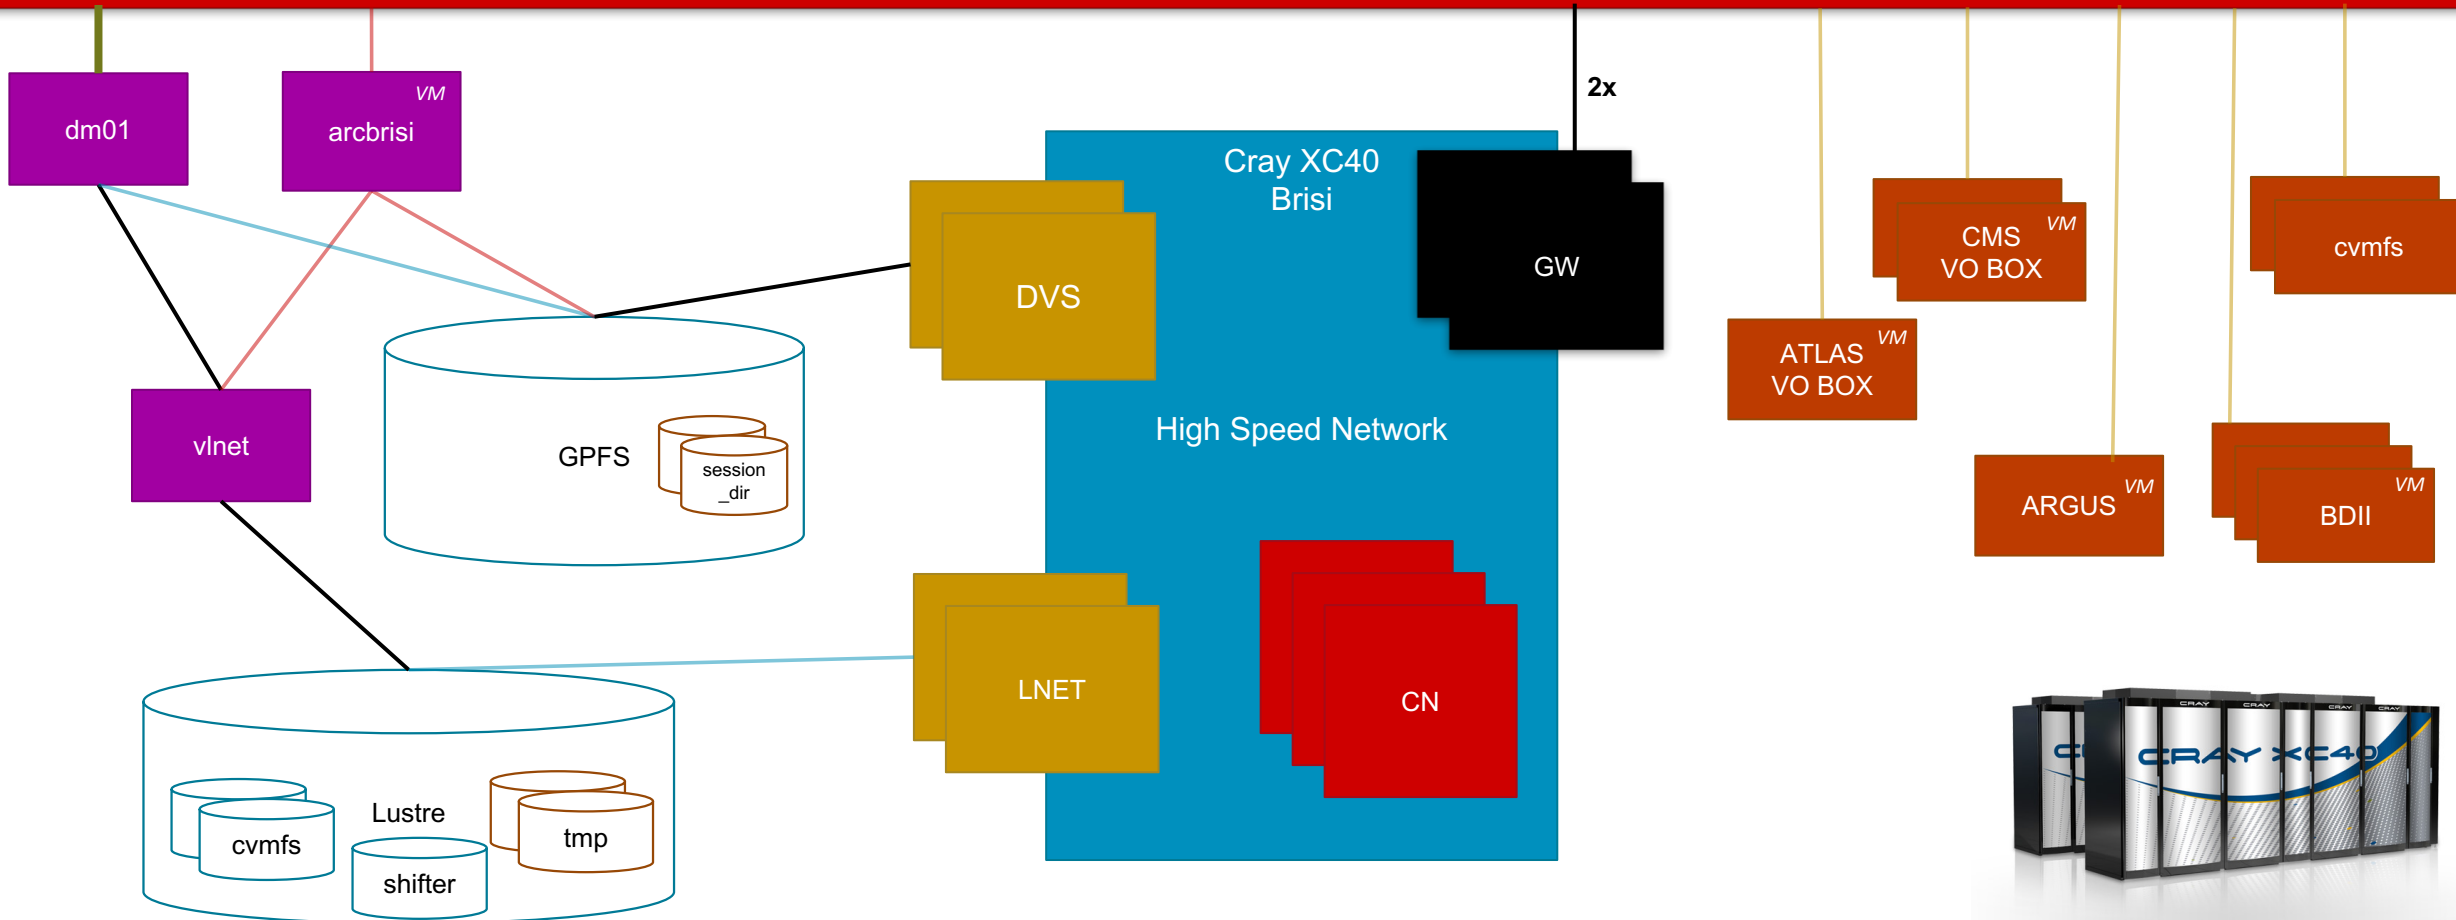
40GbE

10GbE

IB FDR 56Gbps

1GbE

CSCS Network/Internet





CSCS

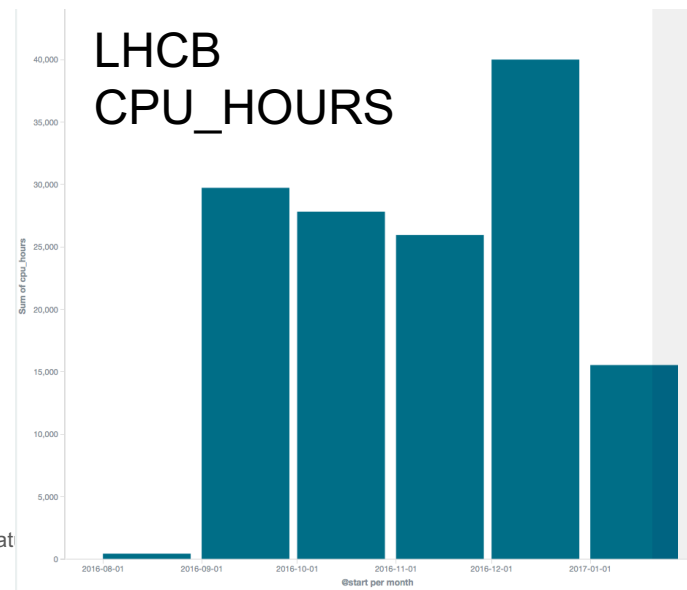


Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

System Statistics

Some statistics in the last 6 months


- ATLAS, CMS, LHCb running production jobs
- 193.655 jobs since 01 Aug 2016
 - 125.580 ATLAS jobs
 - 45.941 CMS jobs
 - 11.853 LHCb jobs
- 574.516 CPU hours since 01 Aug 2016
 - 279.026 ATLAS CPU hours
 - 155.944 CMS CPU hours
 - 139.545 LHCb CPU hours



LHConCRAY stat



More statistics

CHIPP Accounting for Brisi



Lustre performance since Dec 1 2016

ATLAS datamover
and arcbrisi hung



10 nodes



Lustre shared with other users

ATLAS Large scale test (~1000 cores)


Some ATLAS results (Dec 12 – Jan 10)




WallClock Efficiency based on success/all accomplished jobs
29 Days from 2016-12-12 to 2017-01-10



WallClock HEPSPROC6
4 Weeks from 2016-12-12 to 2017-01-10



CSCS-LCG2_MCORE - 11.41% (43,068)
CSCS-LCG2-HPC_MCORE - 2.23% (8,429)




completed app-successful app-failed site-failed abort



CSCS-LCG2 CSCS-LCG2_MCORE CSCS-LCG2-HPC_MCORE ANALY_CSCS CSCS-LCG2-HPC ANALY_CSCS-HPC

Maximum: 52,115 , Minimum: 0.00 , Average: 31,688 , Current: 14,282



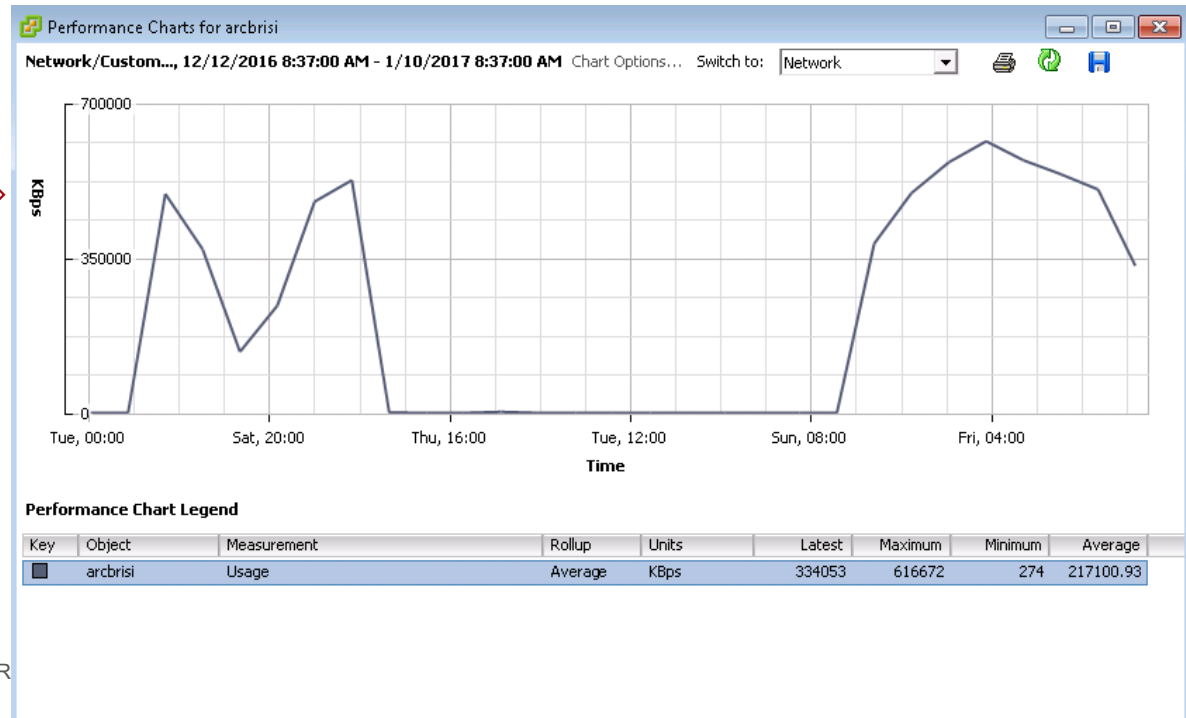
2_MCORE (43,068) ANALY_CSCS-HPC (26,608)

15 /s

But there were a few hick-ups (1/2)

- ATLAS tends to eat all the bandwidth of an ARC server in native mode, as it does most staging before the job gets submitted to Slurm
 - This caused arcbrisi to die miserably a few times due to high load and network resources exhaustion.
 - 1GbE not sufficient! → moved to 10GbE VM ... still not sufficient!
 - Traffic shaping is difficult, as gridftp transfers can happen in multiple ports... and it's not the only protocol used (http, xrootd). Identify which VO is doing what is difficult.

5 Gbps max on average, instant peaks reaching ~9Gbps



But there were a few hick-ups (2/2)

- GPFS via DVS in parallel mode showed some problems with cache consistency
- This affected some things that were mounted off the DVS filesystem:
 - Shifter images cannot be stored on any DVS mounted filesystem without cache consistency → kernel panic
 - perNodeCache /tmp → kernel panic
- Quickly rolled back to
 - GPFS via DVS in serial mode for session directory (job's home)
 - perNodeCache /tmp, cvmfs and swap file in Lustre
- This allowed ATLAS tests to start on Dec 12 2016
- Need to understand if this was due to a misconfiguration or an actual problem
 - The good thing is that we can change settings on compute nodes live, just need to drain them



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Technical conclusions

A few technical conclusions after this 6 months

- Running production ATLAS, CMS and LHCb seems technically possible
- ARC Data staging and bandwidth usage is significant, needs to be properly accounted for when designing the architecture.
 - Either using 10GbE-enabled VMs with proper 40GbE Inet routers...
 - Or go all physical with direct IB (FDR) or 40GbE connections to storage
 - ARC maxdelivery/maxprocessor to set based on network performance
- GPFS via DVS in cluster parallel mode not working well (in CLE 6.0.UP01)
 - Serial mode is sufficient for session directories... but would like to test more the cluster parallel mode
 - Read-only load-balanced might be an option for Shifter images.
- Small Lustre Sonexion seems sufficient for ~1000 cores
 - Could investigate GPFS with more DVS servers
 - Could investigate DataWarp (using SSDs to cache access to a large Lustre Sonexion)

Extra technical notes


- ARC scripts have suffered massive changes
 - Version installed on arcbrisi diverges from standard ARC codebase
- Some of these modifications can be used by the community
 - Would like to share them with ARC developers



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you for your attention.