

OMB:Switzerland-QR14

Main	Get Support	Documentation	Tools	Activities	Performance	Middleware	Catch-all Services	Security
-------------	--------------------	----------------------	--------------	-------------------	--------------------	-------------------	---------------------------	-----------------

Inspire reports menu:	Home •	SA1 weekly Reports •	SA1 Task QR Reports •	NGI QR Reports •	NGI QR User support Reports
------------------------------	--------	----------------------	-----------------------	------------------	-----------------------------

Quarterly Report Number	NGI Name	Partner Name	Author
QR 14	NG-CH	Switzerland	Sigve Haug

1. MEETINGS AND DISSEMINATION

1.1. CONFERENCES/WORKSHOPS ORGANISED

Date	Location	Title	Participants	Outcome (Short report & Indico URL)	
30.9-01.10	CSCS, Lugano (CH)	GridKa Cloud T1-T2 yearly face to face	25	ATLAS German cloud sites' technical solutions discussed. Direct contact between CSCS admins and ATLAS operation experts, [1]	-

1.2. OTHER CONFERENCES/WORKSHOPS ATTENDED

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
2013-09-16 - 20	Madrid	EGI Technical Forum Madrid	4	
2013-10-23	Lausanne	HPC-CH Forum	Michael Rolli, Nico Faerber	

1.3. PUBLICATIONS

Publication title	Journal / Proceedings title	Journal references <i>Volume number</i> <i>Issue</i> <i>Pages from - to</i>	Authors <i>1.</i> <i>2.</i> <i>3.</i> <i>Et al?"</i>
Grid Site Testing for ATLAS with HammerCloud	CHEP2013 Proceedings		Johannes Elmsheuser, Ludwig-Maximilians-Universitaet Muenchen Federica Legger, Ludwig-Maximilians-Universitaet Muenchen Ramon Medrano Llamas, CERN Gianfranco Sciacca, Universitaet Bern Daniel Colin van der Ster, CERN

2. ACTIVITY REPORT

CSCS'''

2.1. Progress Summary

Smooth operation in terms of network and storage.

2.2. Main Achievements

(by end of this month) Complete reinstallation of the whole compute cluster. Migrated to EMI-3 on SL6 and SLURM, this includes 4 CREAM-CEs, 2 ARC-CEs, 78 WNs and 1 APEL server.

Planning dCache upgrade to 2.6 by mid-november and becoming fully SHA-2 compliant.

PSI

1. Progress Summary

* Expanded SE by 360 TB raw storage by adding a NetApp E5460, same hardware as our SGI SI5500. Upgraded and homogenized all Firmware.

* Decision to continue our Solaris X4500 and X4540 for next one or two years based on availability of replacement parts (decommissioned machines from our Tier-2 at CSCS).

* Reinstalled a Sollaris network boot and install service (Jumpstart) to ease reinstallation of our Solaris10 machines.

* Complete phasing out of the disks that had troubled us in the X4540 machines over the last two years (frequent failures) by using replacements from the decommissioned machines from our Tier-2. Reinstallation of all X4540 servers.

2. Main Achievements

* SE expansion (360 TB raw)

* Ensuring a continued use of our existing aging HW by securing replacement parts and providing an adequately resilient Solaris infrastructure for fast reinstallations.

3. Issues and Mitigation

UZH Nothing reported

UNIBE-ID

1. Progress Summary

Stable operations with minor issues reported below

2. Main Achievements

- UBELIX relocation: At end of August the whole cluster was relocated to a new server room at von Roll complex with only minor problems. After 8.5 days of downtime the cluster was fully operational again. Though there were no hardware defects right after the relocation, over the next few weeks three harddisks and one mainboard broke.

3. Issues and Mitigation

- ARC CE usage recrd registration to the smscg database stopped working due to duplicated job ids occuring in the smscg db. After deduplication with a small batch script the usage records delivery works again and the accumulated use records where delivered too.
- At the new cluster location brand new switch from Brocade were installed. Since the relocation we are occasionally facing very short network link downs on two of those switches. Cases at Brocade are opened and the problem is investigated.

UNIGE-DPNC

Progress Summary

Stable operations...

Main Achievements

- HW delivery (4 x IBM x3630 M4, 43 TB for data)
- Got old free TDAQ CPU (35 x DELL PE 1950 MKIII)
- Cleanup in the SE, 1st since Jan 2013, 45%
- New users from AMS and CTA

UNIBE-LHEP

Quite stable operation of the production cluster (ce.lhep).

Preparation for migration from CentOS5 to SLC6 and expansion with nodes obtained from CERN/ATLAS

A new cluster with ~1500 cores (ce01.lhep) has been commissioned and operation stabilised. Some outstanding issues (details below)

* Main achievements

- ce01.lhep cluster in full production for ATLAS. Commissioning of the Infiniband local area network
- ce.lhep cluster has been operated with reasonable stability until early October. Shutdown for expansion and migration from SLC5 to SLC6. Added nodes from CERN/ATLAS, complete rationalised re-cabling (power and network).
- ce.lhep decommissioned and re-installed as ce02.lhep (SLC6.4, ROCKS 6.1 Front-end), added to GOCDB. ATLAS SLC6 WN image prepared. ROCKS images for Lustre nodes (MDS, OSS) prepared, with Lustre 2.1.6. Ready for mass-install.
- Enabled t2k.org VO on our Storage Element

* Issues and mitigations

Recurring problems on the ce01.lhep cluster:

- Frequent NIC lock-ups on lustre OSS nodes causing Lustre to hang and consequent cluster downtimes. Solution: switch LAN from TCP to Infiniband
- CVMFS cache/partition full issue causes WN's to become black holes. Mitigation: manual cache clean-up executed from time to time. Foreseen solution: will re-install with CVMFS 2.1.5 which is said to resolve the

bug

- NFS v4 new defaults caused the inability of the ATLAS software manager jobs to validate the CVMFS software deployment. Mitigation: run on ATLAS request the validation manually from time to time. Solution: identified and corrected the setting causing the issue
- One off failure of one PDU: the LAN switch for the ce01.lhep cluster affected (no redundant PSU): cluster hanging until power recovered

Problems affecting ce.lhep production cluster:

- Cooling instabilities have caused at least once a full cluster shutdown. A second instance saw only part of the WNs to shutdown spontaneously. Recovery implies re-installation
- Some obscure issues with ROCKS and a Lustre build against the latest kernel available have delayed mass-install of the cluster. Issue: a critical ARC bug causes the services to stop processing jobs upon a Data Staging failure. Mitigation: none. Rely on Nagios email notification from the PPS EGI Nagios service to catch failures and react by restarting the services

References

- [1] <https://indico.cern.ch/conferenceDisplay.py?confId=261676>

Article Sources and Contributors

OMB:Switzerland-QR14 *Source:* <https://wiki.egi.eu/w/index.php?oldid=61480> *Contributors:* Krakow, Shaug