

OMB:Switzerland-QR13

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 13	NG-CH	Switzerland	Sigve Haug

1. MEETINGS AND DISSEMINATION

1.1. CONFERENCES/WORKSHOPS ORGANISED

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
				-

1.2. OTHER CONFERENCES/WORKSHOPS ATTENDED

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
2013-05-27	Berlin	7th dCache WS	Fabio Martinelli	Martinelli presented a local solution for establishing advisory quotas for users and groups (q.v. progress summary)
2013-06-03/04	CSCS Lugano	Infiniband Foundation Course	Michael Rolli	

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>

2. ACTIVITY REPORT

CSCS

2.1 Progress Summary

* Smooth operation at CSCS-LCG2 despite some network problems.

2.2 Main Achievements

* Upgraded all WNs from SL5.7 to SL6.4 and moved away from Mellanox IB stack.

* Repartitioned SE (dCache) to provide LHCb with ~100TB of storage capacity (taken 50% from ATLAS and 50% from CMS).

* Found some serious problems on the Infiniband-Ethernet bridges that

may be due to a firmware bug on the bridges. As a consequence, MTU has been reduced on all the network to 1500.

Working to find a solution as soon as possible.

* All VMs are now running on new hardware and have been installed with SL6.4, with the exception of cmsvobox due to some services being not ready yet in SL6.

PSI

2.1 Progress Summary

- finished procurement of storage element extension: 360 TB of raw disk space in form of another NetApp E5400 system
- advisory quotas for users and groups. Space calculations implemented as PostgreSQL stored procedures. Integrated with LDAP based user and group information. Alarming implemented through Nagios.
- Migration of all virtual machines to PSI's new VMware cluster.
- Middleware Upgrades
 - all WN to emi-wn-2.6.0-1_v1
 - all UI to emi-ui 2.0.2-1.e15

2. Main Achievements

- implementation of advisory dCache quotas for managing our 50+ users.

3. Issues and Mitigation

- We still are running 6 SUN X4500 and 5 SUN X4540, and the X4540 are suffering from frequent disk failures.
- ZFS RaidZ2 + 3 spares has managed to cover us adequately for 3 years, so we never experienced a real data loss, even though the frequent disk exchanges and repairs are associated with a large time investment by the operator.
- we managed to get an Oracle maintenance contract which will allow us to continue running this HW for about one year in reasonable safety. Then the machines will be decommissioned.
- VMware virtual machine instabilities: We suffered from stuck operating systems on machines with high I/O for some time. After we disabled the integrated VMware/NetApp snapshots and just relied on pure NetApp file system level snapshots, the problem did not occur again.

UZH Nothing to report

UNIBE-ID

1. Progress Summary

Stable operations with minor issues reported below

2. Main Achievements

- successful readdressing of all publicly available servers to a new subnet during a scheduled maintenance down
- setup 12 new WN
- detailed planning for the relocation of the cluster at end of august

- take over of the new server room and setup of infrastructure
- preparation of material for the new setup
- installation of new Intel Compiler versions (C and Fortran)

3. Issues and Mitigation

- some minor hardware failures with subsequent replacement
- mitigation and fix for CVE-2013-2094
- 2h cluster shutdown due to cooling outage of external supplier. Afterwards gradually startup of the cluster over some hours.

UNIGE-DPNC

2.1. Progress Summary

Stable operations...

2.2. Main Achievements

- First batch worker node running SLC6 put in operation.
- Hardware procurement for the 2013 upgrade is under way.
The order is out. We will replace disk servers in the Storage Element running Solaris (six machines, 96 TB net) with new machines running Linux (IBM x3630 M4, 4 machines, 172 TB total).
- CernVM file system was set up. We use NFS to deploy it to the worker nodes and login machines.
- Adaptation of operational procedures, especially the cleanup of "dark data", to the new version of the ATLAS Distributed Data Management software "Rucio".

2.3 Issues and mitigation

- Chasing problems of occasional batch jobs, which don't run in the SLC6.
- A few disk failures on IBM x3630 M3 and Sun X4500+ machines => disks replaced by cold spares (in case of IBM later replaced under warranty).
- Memory errors on one Sun X4540 disk server => new memory bought and installed.
- Failure of hardware raid on one IBM x3630, due to overheating => machine repaired under warranty. No data loss, but no access to data for 5 days.
- Automatic cleanup of /tmp is affecting very long jobs. Files are removed after 5 days. We still don't understand why. A work-around solution was set up for the user concerned.

UNIBE-LHEP

* Main achievements

- Mostly stable operation, short down times due to critical ARC bug.
- Added 300TB to DPM SE (total capacity now 500TB)
- Upgraded all DPM disk servers to SLC6.3 and emi-dpm_disk.x86_64 0:1.8.6-1.e16 (EMI-2)
- Updated DPM head node to emi-dpm_mysql.x86_64 0:1.8.6-1.e15 (EMI-2, Security update from 1.8.4-1.e15)
- New ARC CE deployed with SLC6.3 and nordugrid-arc-2.0.1-1.e16.x86_64 (EMI-2)

- Completed most of the ground work to put online a new cluster with 1500 cores

* Issues and mitigations

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

Article Sources and Contributors

OMB:Switzerland-QR13 *Source:* <https://wiki.egi.eu/w/index.php?oldid=58475> *Contributors:* Krakow, Shaug