



Paolo Franchini
University of Warwick

MICE computing infrastructure

MICE CM 43

30th October 2015






Infrastructure


The computing infrastructure group is responsible for:

- **Network, computing, monitoring, spares and backups**
- **Data archival and data processing**
- **Configurations Database**
- **Web services**

Network

- micenet: secure virtual LAN administered by RAL networking
- Network isolation test performed in August:
 - no effects in the Control Room functionality
 - DAQ and C&M worked fine
- Crash of the DNS/DHCP machine (failure of the RAID system) solved in 4 hours after disks replacement and machine rebuild:
 - 2 DNS failovers in place
 - DHCP lease time extended to 1 week





Monitoring


- All machines monitored using **NAGIOS**:
 - 92 machines and around 400 services
 - ping, ssh connectivity, file system and space available, load and memory, NTP synchronisation, zombie processes, network connection, etc.
 - custom checks for the datamover machine (proxy validity, CASTOR upload and reconstruction backlog)
 - different contact groups defined for different subsystems (target, online, daq, etc.)
 - external check of the Nagios machine done by the PPD Nagios system
 - snapshot available on <http://micewww.pp.rl.ac.uk/nagios/>

Spares

- Most of the critical machines are relatively new and still under warranty
- Several machines have been replaced with new servers
- Backup machines and mirrors machine ready in place
- Wide use of Intel NUCs:
 - general purpose EPICS IOC machines
 - several hot-swappable spares
- Laptops running EPICS available for the hall




- Many spare 3.5" and 2.5" 1 TB hard drives are available
- New 2 TB NAS backup drive has been installed replacing the 3-year-old one, used for the daily backup



Backups (1)

- *MICENet* machines backup:
 - 2TB RAID NAS driver in the MICE Control Room
 - daily back of sensible directories of the machines (datamover, tracker, onrec, micestore, EPICS client servers, etc.)
- *EPICS repository and Archive data* mirroring:
 - 2 mirror machines (`miceecserv1` and `miceecserv2`)
 - 1 PPD mirror machine for external read only access (`hep1nv154`)
- *Redmine and CHEESE* backup:
 - daily backup on a PPD RAID machine



Backups (2)

- Backup on the GRID:
 - Archive data, calibration data, hall probes, etc.: to be refined
- *eLog*:
 - Automatically mirrored on a MICE net machine
 - daily and weekly snapshots on a PPD machine
- *CDB*:
 - New slave running on a PPD machine




Automated data-mover

- File compaction script run integrated in the Run Control
- NAGIOS controls implemented:
 - VOM proxy validity
 - Tarball copy and verification
 - Castor upload backlog
 - Reconstruction on the GRID

cornigaba	ssh	OK	10-29-2015 20:01:31	9d 3h 33m 30s	1/3	SSH OK - OpenSSH_4.3 (protocol 2.0)
datamover2	DM: Castor upload backlog - Step 2 only	OK	10-29-2015 20:01:28	9d 3h 38m 33s	1/3	OK - Castor upload backlog (step 2 ONLY) #files= 1 within the OK limit =3
	DM: Full chain Castor upload backlog - Step 1 and 2	OK	10-29-2015 20:02:44	9d 3h 37m 17s	1/3	OK - Full chain Castor upload backlog (Step 1 and 2) #files= 0 within the OK limit =3
	DM: Tarball copy and verification backlog - Step 1	OK	10-29-2015 20:00:16	9d 3h 34m 45s	1/3	OK - Tarball copy & verification backlog (Step 1) 0 within the OK limit =3
	DM: Tarball verification	OK	10-29-2015 20:01:32	9d 3h 33m 29s	1/3	OK - Tarball verification failures #files= 0 within the OK limit =0
	DM: VOMS proxy probe	OK	10-29-2015 20:01:29	9d 3h 38m 32s	1/3	OK - 1021 minutes before proxy expires (above the OK limit 360)
	Filesystem free on sda1	OK	10-29-2015 20:02:45	9d 3h 37m 16s	1/3	DISK OK - free space: 4607 177 MB (75% in use=99%)

- A backup machine for the datamover is being installed




Batch processing

Thanks of the recent speedup of MAUS, an attempt has been made to perform the online reconstruction on a local machine located in the MLCR


- a dedicated machine has been purchased: `miceoffrec01`
- the machine is in place
- . . . installation in progress . . .

Configurations Database




- Write access to the CDB is possible only from the MLCR
- A viewer is available as a Java webservice
`http://cdb.mice.rl.ac.uk/`
- Two standby services hosted on PPD machines. Waiting for a final migration
- A pre-prod CDB has a write access
- Lot of work done by Janusz and Pierrick to implement the C-API for the beam line and for the cooling channel

→ We still need to test the master/slave failover procedure




Web services

- `mice.iit.edu`: primary MICE website
 - `micewww.pp.rl.ac.uk`: Wiki, working group pages
 - `cdb.mice.rl.ac.uk`: Web interface to CDB
 - `test.mice.rl.ac.uk`: Jenkins test server for MAUS
 - `reco.mice.rl.ac.uk`: reconstructed data
 - `repo.mice.rl.ac.uk`: MICE repository
 - *SSH bastion*: Gateway to access MLCR machines
 - *EPICS gateway*: Remote read-only access for C&M
-
- No major issues in the last months
 - Machines upgrades and migrations are undergoing



Remote Control Room

- A kind of remote control room has been fitted up in the MICE office in the ATLAS building
- The machines will be isolated from the MICE Network but will be able to run the C&M mimics in read only mode
 - pod moved to PPD network
 - A gateway will be used to feed the EPICS mimics



Conclusions

- The computing infrastructure is in good shape: no major issues have prevented the normal activities
- The data mover is reliable and fully integrated in the Run Control; a failover machine is being installed
- Batch reconstruction worked well and the reconstruction chain is monitored under Nagios. A local reconstruction is a valid alternative
- New machines and backup machines are in place, while few PPD machines replacements need to be finalized
- We still to test failovers:
 - CDB: master/slave
 - EPICS mirror servers
 - primary IOC servers