


Paolo Franchini
University of Warwick

MICE computing infrastructure

23rd June 2015




Infrastructure


The computing infrastructure group is responsible for:

- the maintenance of **networking** and **computing** infrastructure in the MICE local (and remote) control room (hardware and OS)
- **GRID** services, including data archival
- the **Configurations Database**
- various MICE **web services**

Network

- micenet: secure virtual LAN administered by RAL networking
- Network stack in Rack Room 1 near MLCR
- Network stable and no major issues reported in last several months after the installation of the new network switch: *Hewlett Packard 5406R zl2* (dropouts with the old switch solved)
- Sufficient capacity (136 ports) to support the experiment completion
- External access through a gateway, using *ssh* key pairs





Monitoring

- All machines monitored using **NAGIOS**:
 - 80 machines and around 300 services
 - Our Nagios machine is monitored by the RAL PPD Nagios
 - I do not know who's monitoring the RAL PPD Nagios
 - ping, file system and space available, load, NTP synchronisation, zombie processes, etc.
 - Nagios is on MICENet (so you need a tunnel):
`http://nagios.micenet.rl.ac.uk/nagios`
 - Nagios will also be used to monitor the data movement chain
 - Please email me if you need to implement a check on a specific machine on MICENet



Monitoring

Nagios®

- General
 - Home
 - Documentation
- Current Status
 - Tactical Overview
 - Map
 - Hosts
 - Services
 - Host Groups
 - Summary
 - Grid
 - Service Groups
 - Summary
 - Grid
 - Problems
 - Services (Unhandled)
 - Hosts (Unhandled)
 - Network Outages
 - Quick Search:
- Reports
 - Availability
 - Trends
 - Alerts
 - History
 - Summary
 - Histogram
 - Notifications
 - Event Log
- System
 - Comments
 - Downtime
 - Process Info
 - Performance Info

Current Network Status

Last Updated: Tue Jun 23 00:13:22 BST 2015
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as *mice*

View Service Status Detail For All Host Groups
View Host Status Detail For All Host Groups
View Status Summary For All Host Groups
View Status Grid For All Host Groups

Host Status Totals

| Up | Down | Unreachable | Pending |
|----|------|-------------|---------|
| 80 | 4 | 0 | 0 |

All Problems All Types

| All Problems | All Types |
|--------------|-----------|
| 4 | 84 |

Service Status Totals

| Ok | Warning | Unknown | Critical | Pending |
|-----|---------|---------|----------|---------|
| 297 | 3 | 0 | 4 | 0 |

All Problems All Types

| All Problems | All Types |
|--------------|-----------|
| 7 | 304 |

Service Overview For All Host Groups


| MICE Webcams (camhalls) | | | | Mice Networked Devices (mice-devices) | | | | Mice Networked Devices (mice-devices-critical) | | | |
|-------------------------|--------|----------------------|---------|---------------------------------------|--------|----------------------|---------|--|--------|----------------------|---------|
| Host | Status | Services | Actions | Host | Status | Services | Actions | Host | Status | Services | Actions |
| camhall07 | UP | 1 OK | | ckovmon1 | UP | 1 OK | | HA_PLC01 | UP | 1 OK | |
| camhall08 | UP | 1 OK | | emrlvcrate | DOWN | No matching services | | HA_PT01 | UP | 1 OK | |
| camhall09 | UP | No matching services | | envmon1 | UP | 1 OK | | backup1 | UP | 1 OK | |
| camhall10 | DOWN | No matching services | | pabsorber | UP | 1 OK | | miceioc1 | UP | 1 OK | |
| camhall11 | UP | 1 OK | | sub25pqm | UP | No matching services | | miceioc2 | UP | 1 OK | |
| camhall12 | UP | 1 OK | | | | | | miceioc4 | UP | 1 OK | |
| camhall13 | DOWN | 1 CRITICAL | | | | | | miceioc5 | UP | 1 OK | |
| camhall14 | UP | 1 OK | | | | | | miceioc6 | UP | No matching services | |
| camhall15 | UP | 1 OK | | | | | | trackerdelay1 | UP | No matching services | |
| camhall16 | UP | 1 OK | | | | | | trackerdelay2 | UP | No matching services | |
| camhall17 | UP | 1 OK | | | | | | trackerdelay3 | UP | No matching services | |


| MICE Un-Critical Devices (mice-devices-uncritical) | | | | MICE Hosts (mice-hosts) | | | | MICE Critical (mice-hosts-critical) | | | |
|--|--------|----------|---------|-------------------------|--------|----------|---------|-------------------------------------|--------|----------|---------|
| Host | Status | Services | Actions | Host | Status | Services | Actions | Host | Status | Services | Actions |
| miceac1 | UP | 1 OK | | cagateway1 | UP | 1 OK | | HX_INV01 | UP | 1 OK | |
| miceac2 | UP | 1 OK | | localhost | UP | 2 OK | | configdba | UP | 1 OK | |
| miceac3 | UP | 1 OK | | miceecserv2 | UP | 9 OK | | miceac1rr2 | UP | 1 OK | |
| miceac4 | UP | 1 OK | | miceisgateway | UP | 1 OK | | miceacq05 | UP | 1 OK | |

+ e-mail spam notifications

MICE Local Control Room

- Two **LCD TV screens** have been wall-mounted in MLCR:
 - one shows the real-time status of the various *EPICS controls* and monitoring subsystems and the *alarm handler*
 - the second one acts as a “*web whiteboard*” (MICE webcams, MICE Hall and ISIS status, plots from online reconstruction and the eLog)





MICE Remote Control Room


- A remote control room is going to be fitted up in the MICE office in the ATLAS building
- The machines will be connected to the MICENet in order to permit a quick access to the C&M and Online resources
- TV screens will show the status of the current run (online reconstruction plots, ...)



Spares and backups


- All the critical machines are relatively new and still under warranty
- `miceiocpc2`, `miceecerv2` and `miceiocpctk` (C&M servers for the EPICS client, software and tracker controls) have been replaced with new servers
- Backup machines ready in place
- Several Intel NUCs hot-swappable spares used as general purpose EPICS IOC machines





Automated data-mover

- **The goal:** process data and make them available for analysis as quickly as possible
- Data movement chain fully tested:
 - raw data from the DAQ is first compacted along with outputs from online reconstruction and online monitoring programs
 - data moved into `migestore`
 - data copied into tape (CASTOR)
- File compaction script runs automatically in the Run Control after the stop of each run
- Required a Nagios monitor/alarm



Batch processing

- Automated script submits jobs for reconstruction on the GRID
- Reconstruction performed on the fast-response **Tier-1** queue at RAL
- Reconstructed data moved on permanent storage and available on web
- Further (re)processing will be done on the **Tier-2** queues (Brunel and Imperial College)

- GRID proxy renewal for job submission is now manual
- Required a monitor as well




Batch processing

gfe02.grid.hep.ph.ic.ac.uk:8301/RECO/MAUS-v0.9.4/1/Step4/06100/

dCache

RECO / MAUS-v0.9.4 / 1 / Step4 / 07000

| Name | Size | Last Modified |
|-------------------|---------|------------------------------|
| 07078_offline.tar | 466338 | Tue Jun 23 00:12:22 BST 2015 |
| 07076_offline.tar | 478879 | Mon Jun 22 23:12:19 BST 2015 |
| 07075_offline.tar | 1434550 | Mon Jun 22 08:12:17 BST 2015 |
| 07074_offline.tar | 493655 | Sun Jun 21 06:12:23 BST 2015 |
| 07072_offline.tar | 896289 | Sun Jun 21 06:12:23 BST 2015 |
| 07071_offline.tar | 453030 | Sun Jun 21 01:12:29 BST 2015 |
| 07070_offline.tar | 422003 | Sun Jun 21 01:12:29 BST 2015 |
| 07067_offline.tar | 414440 | Sun Jun 21 00:12:49 BST 2015 |
| 07068_offline.tar | 459648 | Sun Jun 21 00:12:49 BST 2015 |
| 07066_offline.tar | 430814 | Sun Jun 21 00:12:24 BST 2015 |
| 07065_offline.tar | 980007 | Sat Jun 20 08:12:20 BST 2015 |
| 07050_offline.tar | 935118 | Fri Jun 19 08:12:17 BST 2015 |
| 07049_offline.tar | 441110 | Fri Jun 19 02:12:16 BST 2015 |
| 07044_offline.tar | 432096 | Fri Jun 19 01:12:33 BST 2015 |



Configurations Database

- Infrastructure group take care of:
 - CDB servers
 - table implementations
 - API (Python and C)
- Beamline C-API (aka *Run Control API*) developed (90% done, pending one minor get function)
- Available on launchpad
- Requires further tests to make sure things work




New Jenkins v1.6 installation

- New installation on a SL 6.4 PPD machine 157
- All the users jobs copied in this machine
- Tests undergoing <http://heplnv157.pp.rl.ac.uk/>
- Soon will replace the old installation

The screenshot shows the Jenkins web interface for MICE. The top navigation bar includes the Jenkins logo, a search bar, and the user name 'Paolo Franchini' with a 'log out' link. The left sidebar contains navigation options: 'New Item', 'People', 'Build History', 'Manage Jenkins', 'Credentials', and 'My Views'. Below the sidebar, there are sections for 'Build Queue (1)' and 'Build Executor Status'. The 'Build Queue' section shows a job 'MAUS_per_commit_third_party_heplnm071'. The 'Build Executor Status' section shows two executor groups: 'master' with 4 idle executors and 'brunei_test_box' with 4 idle executors. The 'heplnm071' executor is shown as 'offline'. The main content area displays a table of build jobs with columns for 'Name', 'Last Success', 'Last Failure', and 'Last Duration'. The table lists various jobs, including 'MAUS_clurga_release', 'MAUS_clurga_release_candidate', 'MAUS_clurga_step4', 'MAUS_clurga_trunkdev', 'MAUS_per_commit_third_party_heplnm071', 'MAUS_release_sl64', 'MAUS_rogers_online', 'MAUS_rogers_sl64', 'MAUS_rogers_xboa', 'MAUS_per_commit_third_party_sl64', 'MAUS_full_install', 'MAUS_karadzkov', 'MAUS_merge_uchids', 'MAUS_per_commit_third_party', 'MAUS_per_commit_third_party_heplnv157', and 'MAUS_release'.

| S | W | Name | Last Success | Last Failure | Last Duration |
|---|---|---|-----------------------------------|-----------------------------------|---------------|
| | | MAUS_clurga_release | N/A | N/A | N/A |
| | | MAUS_clurga_release_candidate | N/A | N/A | N/A |
| | | MAUS_clurga_step4 | N/A | N/A | N/A |
| | | MAUS_clurga_trunkdev | N/A | N/A | N/A |
| | | MAUS_per_commit_third_party_heplnm071 | N/A | N/A | N/A |
| | | MAUS_release_sl64 | N/A | N/A | N/A |
| | | MAUS_rogers_online | N/A | N/A | N/A |
| | | MAUS_rogers_sl64 | N/A | N/A | N/A |
| | | MAUS_rogers_xboa | N/A | N/A | N/A |
| | | MAUS_per_commit_third_party_sl64 | N/A | N/A | N/A |
| | | MAUS_full_install | 1 day 17 hr - #17 | 6 days 8 hr - #16 | 5 hr 4 min |
| | | MAUS_karadzkov | 10 days - #2 | N/A | 6 hr 57 min |
| | | MAUS_merge_uchids | 11 hr - #5 | 7 days 3 hr - #4 | 3 hr 17 min |
| | | MAUS_per_commit_third_party | 2 days 13 hr - #6 | N/A | 7 hr 15 min |
| | | MAUS_per_commit_third_party_heplnv157 | 2 days 13 hr - #2 | 9 days 13 hr - #1 | 7 hr 0 min |
| | | MAUS_release | 11 days - #1 | N/A | 5 hr 40 min |



Conclusions

- The computing infrastructure is in good shape
- The **data mover** is reliable and fully integrated in the Run Control
- **Batch reconstruction** worked well in last weeks of run
- Old machines have been replaced and **backup** machines are in place
- Need a better monitor for the *file compaction* → *data movement*
→ *job submission* chain