

## Description:

Air Conditioning to the Lab 8 Machine Room failed over the holiday period resulted in a loss of services to MICE

## Notes

Computing Group does not provide out of hours support for any of its services  
heplnv150 runs the mousehole service  
heplnv151 runs the EPICS gateway  
heplnv152 runs the micewww service

## Impact:

All Linux based services were shut down for between 3 and 5 days depending on the machine. During the (out of hours) restart the system disk of one of the Mice Virtual Machine servers failed meaning most Mice service were unavailable until the Lab reopened on 2nd January and the archive data disk on the EPICS gateway was not restored until that machine could be reinstalled on 10th January.

## Timeline:

When	What
25/12/12 11:37	CAB received call from RAL security on 25th December at 11:37 AM informing me that the temperature alert for Lab 8 had triggered. KD had been informed, estates called and the back and corridor doors been opened. Checked the Ganglia Monitoring saw that the temperature was currently OK and falling. Some machines has automatically shut down at this point
25/12/12 c. 14:00	CAB Called RAL security, was informed that estates had attended and the temperature was now OK and the doors were closed. Was not in a position to verify that.
25/12/12 c. 17:00	CAB checked the ganglia monitoring to see the temperature rising again, took the decision to shutdown non core linux services including the Tier 2, and monitor the temperature for a while, which then appeared to be stable.
26/12/12 c. 14:00	CAB checked ganglia again on the afternoon of the 26th to see the temperature rising again, all linux services apart from the monitoring shutdown
28/12/12 c. 12:00	CAB Checked the status in ganglia to see that the temperature was returning to normal levels
30/12/12 c 15:00 to 17:00	CAB attended the lab and restarted all the Linux Servers. heplnm067 failed to restart, monitoring showed heplnv150, heplnv151 and heplnv152 still down
02/01/13	failure on heplnm067 investigated, heplnm068 restarted and heplnv152 recovered, heplnv150 and heplnv151 restarted on heplnm068
03/01/13	HN replaces the failed disk in heplnm067
10/01/13	RH Reinstalled heplnm067 - Archive Disk reexported to heplnv151

## Incident Details:

This is really two incidents, the air conditioning failure and the disk failure.

On the morning of 25th December the air conditioning in R1 Lab failed and the temperature started rising. Security we're alerted by the BMS and the room temperature reached a level that some servers are configured to shutdown (nn). Security opened the outside and corridor doors, and called KD, CAB and estates. KD and CAB determined that the room temperature was failing and decided no further action was necessary at that point. CAB called Security early in the afternoon and was informed that Estates had attended, the problem was fixed and the doors had been closed again. CAB checked the situation later in the afternoon to find the temperature rising again and too the that since the problem had reoccurred and CG were not in a position to frequently check the temperature in the Lab for the next few days that rather than call someone out for a second time on Christmas day, to shut down non essential systems as far as possible.

On the 26th December CAB checked the temperature of the room and found it rising towards 30 degrees again and took the decision to shut down the remaining Linux servers in Lab 8 apart from those needed to keep the monitoring up (heplnx010 and heplni001) although some machine probably hung during the shut down due to stuck [NFS](#) mounts

On the morning of the 28th Estates attended again and fixed the air conditioning. Computing Group were not informed. CAB was occasionally monitoring remotely and was going to contact estates via security that day

On then evening of Sunday the 30th December complaining about the disk in port 1 "operating outside specifications" or something like that "F1 to continue or F2 for Setup". Pressed F1 to continue the boot but was not able to log in at the console. Monitoring showed heplnm067, heplnv150, heplnv151 and heplnv152 down. Checked the other Mice physical machines, non were running a Xen kernel so assumed that all the VMs ran on heplnm067 so left it at that.

Once the lab reopened on 2nd January CAB with HN and CR went to investigate the problems. Heplnm067 was booted into rescue mode, and it was verified that /dev/sda was reporting smart errors, the disk was backed up across the network to a tar file. It was determined the heplnv150 and heplnv151 had been running on this host but that heplnv152 had been running elsewhere. Heplnm068 was checked and it was discovered that Xen was installed but not the default boot kernel. heplnm068 was restarted with the Xen kernel as the default and heplnv152 also came up. The configuration and data for heplnv150 and heplnv151 was copied to heplnm068 and the virtual machines started. HN reported the failed disk to Dell. The disk was replaced by HN the next day and quickly tested but it was not until 10/1/13 that the machine could be reinstalled and the final service brought back online (the [ArchiveData](#) disk on heplnv151).

Note: Machines in Lab 8 are monitor the room temperature and are configured to shut down progressively as the temperature reaches 28, 29 and 30 degrees. If the room temperature reaches 34 degrees the BMS will cut the power

## Analysis:

- The [HyperV](#) VM hosting cluster was not taken down during this incident.
- Mice services could have been left running (though at extreme risk) if Computing Group had been informed about the magnet tests and the implications for the Mice Nodes.
- heplnm067 took 7 days to reinstall after the disk was replaced due to a lack of CG effort (partly caused by the proximity to the holiday) RH returned to the Lab on 8/1/13 and CAB was on leave on 4/1/13 and offsite at a meeting on 7/1/13 and 8/1/13.

## Follow Up:

Action	Owner	Status
Verify the list of machines using the Lab 8 temp check configuration and their shutdown order	CAB	In Progress
Migrate the Mousehole, EPICs Gateway and <a href="#">MiceWWW</a> services to the <a href="#">HyperV</a> VM System	RH	In Progress - VM Blanks have been created for the services

-- [ChrisBrew](#) - 2013-01-09

Topic revision: r3 - 2013-01-17 - 09:43:29 - [ChrisBrew](#)

Copyright &© by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? [Send feedback](#)

