

Incident Report - October 13th, 2015

Incident #2015-106

Slow Response Times for Enterprise Applications

Summary

On October 13th, 2015 multiple services connected to one of the Dupuis Datacentre networks had intermittent slow response times. These connection issues lasted for a few seconds at a time but occurred every few minutes.

Impact

All customers of our Enterprise services would have experience sluggish performance intermittently. For example: You may have loaded a Moodle page, and received an error while seconds later after a page reload, the page would load successfully.

Root Cause

The root cause was determined to be a faulty stack interconnect cable in the DCP network stack. When the cable was failing it was flapping up and down causing traffic to be rerouted through the switch stack. Reseating the cable caused a switch master election reboot loop in DCP network switches 2-6. Pulling the power on one of the switches allowed the stack to recover from the reboot loop.

Resolution

Our Incident Response Process was followed with the Manager of iOPS taking the lead. He initiated an incident triage team who worked to diagnose and correct the problem. The team diagnosed the issue with the cable and the connector, both of which were replaced to correct the problem.

Communications (Internal)

The incident triage team were isolated to a meeting room to focus on the problem. The incident response lead kept the department updated via phone and email.

Communications (External)

- ITS Notices: *Update*: As of 11:00 am all systems running normally. Contact the Support Centre if you are having issues. 36666
- 2015-10-13 09:59:08: Queen's ITS staff are investigating issues of degraded performance affecting multiple services. An updated notification will be posted when new information is available.
- ITS has fixed the issue by replacing two bad cables. Services are now back up and stable. Time: 5:20 pm October 13th, 2015
- 2015-10-13 17:00:17: ITS is working on a known problem that is causing a multiple service outage. Systems should be back up and running in 15 min. This notification will be updated in 15 min.
- Twitter was also used twice to let users know of multiple service outages.

Lessons Learned

The logging of the DCP network stack was not shipping logs to the logging server. This meant the only way to view them was on the switch console. At the same time an interface on a server was flapping and putting out many messages to the log buffer which hid the switch stack link issue. This configuration has since been fixed.