

Incident Report - May 10th, 2015

Incident #2015-80

Summary

On Sunday May 10th 2015, a production shared storage unit was upgraded to the latest software revision. The upgrade proceeded with no known failures. However, shortly after the upgrade completed the production server cluster began to lock up. This resulted in all virtual servers being unavailable. From 12:00 pm until 8:00 pm there were various services down as the failure cascaded from each of the production servers.

Impact

All production enterprise services were affected at various times throughout the 12:00 pm to 8:00 pm window. They were either inaccessible or sluggish.

These services included, but were not limited to SOLUS, PeopleSoft Finance and Human Resources, MyQueensU, the Queen's University website (www.queensu.ca), Moodle and Single Sign-on.

Root Cause

The problem is related to the ability of our fibre switches to recover from unexpected disconnections. Ultimately it resulted in a lock of all I/O operations.

Resolution

Each server had to be individually rebooted. Once rebooted they returned to normal. Firmwares on our production server cluster were also updated to prevent this issue in the future.

Communications (Internal)

The upgrader contacted his manager via telephone and they worked over the phone to correct the issues.

An email was sent on May 15th, 2015 to our hosting customers describing the issue and explaining our plan to ensure it doesn't happen again.

ITSP Communication (External)

Multiple notices were posted:

"2015-05-10 13:53:18

*Title: Multiple Service Outage *Update**

We are currently experiencing an outage of a number of services. IT Services is investigating and will update once we have more information."

Previous Description (2015-05-10 14:14:21):

Update

An interruption in IT Services virtual environment caused a number of virtual servers to become unavailable. This impacted services including OUR and Careers.

The problems have been resolved, but IT Services will continue to investigate into the cause.

Previous Description (2015-05-10 20:40:50):

Update

It was determined that there would have been degraded services across many services. It appears that they would have been slower than they are typically but remained up. This would have been between the hours of 12pm and 8:30pm today.

ITServices is working the vendor of our storage system to determine root cause of the issue.

Lessons Learned

Future storage upgrades will be expected to cause issues. We will consider earlier morning changes and set expectations that there may be extended outages. We will also pro-actively put systems into maintenance mode and reboot them in sequence after the upgrade takes place.

In the longer term, we will heavily consider non fibre-channel technology such as iSCSI.