UW-SUPERIOR TECHNOLOGY SERVICES

Post Incident Review

November 2014 Network Outage

Gigi Koenig, Tom Janicki, Joe Kmiech and Jim Rink

12/1/2014

The purpose of this report is to provide a functional description of major system outages. Technical details can be found in the Footprints ticketing system. Contact any IT manager for additional information or clarification.

Contents

UW-Superior Technology Services Post Incident Report

Incident Summary

On Wednesday, November 5th following a 10:30 pm scheduled electrical upgrade in the data center, the campus' two primary core switches and the primary and secondary firewalls failed to start normally on reboot due to significant hardware failures leaving the campus with no network services. New hardware was shipped to the campus and the majority of systems were restored on Thursday, November 6th by 9 pm with the remainder of systems restored by 2:30 pm on Friday, November 7th.

Incident Details

In the course of performing routine maintenance on the uninterruptable power supply (UPS) which feeds the primary campus data center, technicians from Eaton Corporations found a bad output breaker that needed to be replaced as soon as possible. The repair required power to be shut down to the datacenter and the main and intermediate distribution facilities in Swenson Hall. Since core switching equipment is housed in this location, all campus technology services would be offline.

In preparation for the scheduled outage:

- Technology Services notified the campus, our Internet Services Provider (ISP), Community Area Network (CAN) members, and Governmental Internet Service Providers with equipment located in this facility of the scheduled outage.
- There were limited external power supplies available and it was decided to move our ISP to maintain our Domain Naming System (DNS) presence and external partners such as governmental ISP's and CAN members to these power sources.
- All devices were prepared for the shut down and backups of critical configuration files were validated.
- The emergency website hosted off campus was updated and plans were in place to forward queries for www.uwsuper.edu to emergency.uwsuper.edu.
- Eaton technicians estimated that the breaker replacement would take twenty minutes, so a decision was made to make no changes to DNS records.

Individual systems were powered down per standard shut down procedures beginning at 10:30 pm, and at approximately 11:00 pm the main power was shut down and the breaker was successfully replaced as scheduled. Power was restored at approximately 11:20 pm and standard start-up procedures were initiated. During start-up, two redundant primary core switches (which are the heart of the campus network) and the primary and secondary firewalls (security systems that help prevent malicious attacks on our network) failed to start normally.

Tom Janicki, Director of Network infrastructure and technicians Jerod Boisjoli, Brian Hood, and Ed Knudsen worked to restore services until approximately 12:00 am when the decision was made to connect to a system outside of the campus firewalls and initiate a support case with Cisco Systems.

A Cisco technician believed the issue to be hardware related and asked for serial numbers to start a “Return Merchandise Authorization” (RMA) to replace suspected hardware. IT staff left the data center at 2:30 am on Thursday. At 4:00 am Tom Janicki was following up on the RMA and discovered that the parts had not been ordered. It was later discovered to be the result of a shift change miscommunication at Cisco Systems.

At 6:20 am Tom returned to campus to notify offices of the outage. Application Services Director Jim Rink was in the office and took over coordinating campus communications to allow Tom to focus on technical issues. A call was placed to Lynne Williams the director of campus communications who used the campus emergency alert system (RAVE Alerts) to notify the campus of the situation and Lynne also briefed senior leadership. IT communication coordination was subsequently handed over to the Director of Technology Support Services, Joe Kmiech and his staff at 7:30 am.

Resolution / Restoration of Services

Tom contacted the regional Cisco sales representative at 6:53 am who in turn contacted the regional post-sales engineer to push up the priority of the case and aid in troubleshooting. A three way conference call was established and it was determined that four components from one core switch and one component from the other core switch needed to be replaced. At approximately 10:30 am it was decided the parts could be sent via courier from the Twin Cities and were expected to arrive within four hours.

A second case regarding the firewalls was being worked simultaneously by Tom and technician Bob Iverson. They worked with Cisco to bring up the primary firewall and decided to leave the secondary firewall offline so that work could continue on troubleshooting the core switches. The image files on the firewalls were discovered to be corrupt; however, the configuration files were intact.

At approximately 2:00 pm, the replacement components arrived on campus and installation began. The first core switch was brought online. Because the second core lost a supervisor engine, the configuration would need to be restored. After the first core was back online, Ed Knudsen began restoration on the physical and virtual servers, which allowed access to the configuration file for the second core switch and restored it to an operational state.

By 9:30 pm we believed we had approximately 90% of all service online, however, it was discovered later that wired service in labs, Swenson Hall, JDH Library and Barstow as well as credit services and guest wireless access were still impacted by the outage. At approximately 12:00 pm on Friday, Tom was called back in and he restored the remaining services by 2:30 pm.

In reviewing the issues that extended into Friday, it was determined that during recovery of systems, some cables had been plugged into the wrong ports and some virtual interfaces were administratively left down.

Technical Lessons Learned/Action Items

The Positives

- The replacement cost for the equipment that failed would have been approximately $170,000, but was covered under service agreements established with the vendor.
- The team work and commitment was outstanding. Our technicians showed themselves to be skilled and dedicated IT professionals that have well defined lines of responsibility.
- An on-gong commitment by the IT unit toward continuous professional development paid off during the incident. Unlike large IT shops where there are dedicated resources for specific functions; UW-Superior IT staff must be cross-trained to take on a variety of technical functions.
- Tom Janicki showed exceptional leadership, a high level of technical skill and a creative and exceptional use of the resources at his disposal to resolve this issue as quickly as possible.
- Communication to the campus through the RAVE system and actual conversations was timely. Progress reports to the campus through the day were positive and informative.

Improvements

The following have been identified as areas for improvement or investigation. Those responsible for action items are identified in *italics*:

- The electrical component -- single point of failure
    - The electrical component that required replacement, while not the cause of this incident was the catalyst and represented a single point of failure. The cost of redundancy for this component is significant but the *infrastructure staff will* complete a risk assessment to determine if we should pursue redundancy or look for lower cost alternatives.
- Hardware known issue
    - Cisco engineers informed us that there was a known power issue for the specific equipment that failed, however; the method by which Cisco communicates these types of issues is not direct. *The infrastructure team will* suggest to Cisco that there be more direct notifications and internally, will establish new procedures for technical documentation review both on-going and prior to major upgrades.
- Scheduling down-time
    - Campus expectations on the network is 24/7/365 however, there are times when reliance on services is heightened such as during early registration and advisement. The *IT management team will* broaden our review of scheduled maintenance to minimize campus impact where feasible going forward.
- Establish backup relationships (MOA's) with like sized UW campuses and availability of surplus equipment at other UW Campuses that could be made available for use.

    - *The infrastructure Director will* work on creating backup relationships (MOA's) between UW-S and other schools such as Eau Claire or Stout and put out a call for access to surplus equipment from other Universities.
- LDAP Authentication
    - There are many systems on our campus that are hosted externally but rely on authentication to the campus for access. *The infrastructure team will* look into the feasibility of establishing alternate authentication options.
- Double check critical work -- Two sets of eyes
    - *Infrastructure staff will* make a procedural change on major component restorations that require double checks on critical work.
- Hardcopy documentation
    - We found that some of our documentation and resource listings were only available electronically. *IT managers* will publish, maintain and store hardcopy resources in our local copies of the COOP manuals or in other logical locations.

Communication Lessons Learned/Action items

Overall Incident communication was very good. The following is a recap of what worked and what could be improved:

Initial Issue Communication

Good:

- We established an initial IT incident coordinator to manage communications for IT and to work with the campus communications officer who invoked procedures for notifying senior campus leadership and the campus at large. The use of the campus emergency communications software (Rave Alerts) was appropriate to the situation.
- Use of signage on building entry doors proved to be a very effective communication tool.
- Several IT staff members work a year round flexible schedule which provided resources to assist with incident communications before standard work hours.

Improvement: The following have been identified as areas for improvement with those responsible for action items identified in *italics*:

- While there was staff on campus prior to standard business hours, they were not aware of the incident before they arrived to work and for a short period after they had arrived. All IT managers were not personally notified of the incident as it was discovered. *IT managers will* empower and encourage all IT staff to attempt to get in contact with all IT managers at the first sign of a significant issue with critical systems.
- IT managers need easy access to hard copy, regularly updated contact information for key staff members and external resources that may need to be contacted during a major outage. *IT managers will* work internally and with the campus communications officer to make this information available.
- There was some minor confusion on who was the “campus officer of the day”. *Lynne Williams* will follow up on that issue.

Campus Notifications

Good: Overall, notification to the campus was organized and was well received by campus constituents who contacted us. Key successes include:

- Signage on building entry doors was an effective method of communicating to staff and students.
- The continued use of the RAVE Alerts system was appropriate for this issue and we used the “texting” method to send the communication to the campus prior to the start of standard business hours.
- Social media outlets were successfully used as a method of campus communications.

- UW-System was notified of our situation in a timely manner.

Improvement: The following have been identified as areas for improvement with those responsible for action items identified in *italics*:

- We were unable to locally update or access the campus emergency web page to add and update content. *IT managers will* meet with the Campus Web Master to find a solution that allows us to easily and securely update the page from non-networked locations
- We could not locate instructions on how to broadcast a message to all extensions through the voice mail system. *IT management will* find the instructions and will make sure that they are included with COOP documentation going forward.
- A non-LDAP dependent method to access the RAVE administrator's site needs to be established. Lynne Williams established this during the incident *and will* communicate the procedures to those that need that access.
- There needs to be a checklist of all available communication methods in the COOP manual. The *IT management team and Lynne Williams* will work with the COOP coordinator to have this incorporated into the manual.
- The campus should encourage more campus community members to participate in the RAVE alerts program and official campus social media sites. The *IT managers will work with Lynne Williams* to promote these as alternate methods of communications.

This report was submitted by the IT management team on 12/1/2014. Please contact any team member listed below for additional information.

Gigi Koenig
Vice Chancellor of Administration and Finance
gkoenig1@uwsuper.edu

Tom Janicki
Director of Infrastructure Services
tjanicki@uwsuper.edu

Joseph Kmiech
Director of Technology Support Services
jkmiech@uwsuper.edu

Jim Rink
Director of Application Services
jrink@uwsuper.edu