Revision: 2 Effective Date: TBD Status: Draft

Authors: Peter Hill, Kevin Miller, Russell Yount

Core and Edge Network Upgrades

1. Scope

This document addresses plans to upgrade the core and edge data networks maintained by Computing Services. It deliberately addresses the upgrade at a high level, as it is expected that specific deployment details will be documented in other forms.

2. Executive Summary

In broad strokes, the purpose of this upgrade is to modernize the data network. In doing so, we improve the capabilities at many levels (for example, the raw connection speeds available and the IP services offered). We improve management and monitoring capabilities, by standardizing the operating platform (Cisco IOS). Network security is boosted by capabilities of the edge network, including per-port machine limits to limit damage from common problems. The edge network will also enable the use of 802.1x to implement the next generation methods of restricting access to public outlets. Finally, we improve the resiliency of the network by implementing redundant routers on nearly every subnet across campus.

From a simple hardware perspective, we will be removing multiple generations of edge network devices from the network. These devices include shared Ethernet hubs first deployed in 1995, first generation Ethernet switches, and up through non-IOS based switches. For the core upgrades, we will be replacing components of the 6509 switches deployed between 1999 and 2001. Because of the concentration of services provided by the routers, their useful life is somewhat shorter than edge devices.

We have undertaken a detailed process to specifically identify the hardware scheduled for removal and the replacement hardware to install. This process involved the collection and aggregation of data crucial in the planning process, including reliable data regarding the number of active outlets, the type and termination of inter-closet and inter-building fiber cable (not previously recorded in database form), and current counts of deployed equipment. From this, a lightweight web-accessible database system[1] was designed to view and manage this data. In real-time, the system verifies the correctness of the specified network topology and alerts to any problems discovered. It calculates the auxiliary components necessary to construct the network, including fiber optics adapters and patch cables. The resulting reports reflect the project equipment needs to a high degree of accuracy.

The current estimated cost to fully implement the equipment replacement plans is $1.5 million. We anticipate some discounting of this estimate based upon trade-in or "coupon" programs available from some suppliers. As necessary, upgrades in some areas will

[1] http://km1.net.cmu.edu/n2004

be postponed. Of note, the plan addresses all edge networking provided by Computing Services, including residence halls. These areas are prime targets for postponement, though the upgrades would improve our ability to effectively manage these areas perhaps more so than on campus.

3. Core Network Upgrades

The plans for upgrades to the core network call for a reduction in the total number of 6509s in service (currently 21) while simultaneously adding redundancy to most areas of campus in terms of routing. A total of six "pods", or pair of routers, will be defined. One pod will closely mirror the existing core ("core0" and "core255") and inter-connect all other pods. One pod will be used for campus egress (that is, our Internet connections). One pod will be a "test" pod (similar to the use of "sysdev" in the lab today).

The remaining three pods will be the primary production pods. These pods will be connected to building aggregator switches (the latter of which is addressed as edge upgrades). The building aggregator will connect to both pod members (that is, two routers with nearly identical configurations). These routers will be running the Hot Standby Router Protocol (HSRP). If one router in a pod were to crash or be otherwise unavailable, the pod member would carry all building traffic. Thus, the network has added resiliency against router outages.

Members of a given pod will be located in different campus locations. This provides some resiliency against a mechanical outage of a single building. Further, in the case of the production pods, one of the members will be geographically close to the buildings aggregated on that pod. In some areas, this enables the use of alternate fiber to provide protection against an outage to a single fiber bundle[2]. Due to the best operating conditions on campus, however, we expect one member of each pod to reside in Cyert A84.

The core upgrades will involve the replacement of many components of the 6509s, including the supervisor (the router "brains"), fan tray, power supplies, and line cards. Many of these upgrades are necessitated by the supervisor upgrade, which refreshes the longevity of the unit and provides a number of new services and capabilities. The physical interface line cards will be upgraded to provide the greatest support for high-speed packet forwarding through the chassis.

The 6500s continue to be the platform upon which we provide advanced layer 3 services. In upgrading to the latest hardware available, a number of additional services are enabled, including:

- redundancy ("Hot Standby Routing Protocol")
- unicast reverse path forwarding (URPF; source address verification for greater network security)
- blackhole filtering with URPF (using unique strategies developed at CMU to quickly respond to network problems)
- NBAR application identification to classify and queue traffic
- IPv6 support in hardware
- per-host quality of service policing

Except for HSRP redundancy, these services cannot be provided using our existing hardware (or would severely limit the router's forwarding capabilities). Furthermore, the

[2] Wherever possible, we selected diverse fiber optic routes for redundant connections.

upgrades enable the core to support 10Gbps connections as needed to support campus connectivity.

4. Edge Network Upgrades

The edge of our network continues to languish as shared hubs continue to provide hundreds of connections to user outlets. These devices are unsupported by the manufacturer and were first deployed on campus in 1995. Many have begun to reach permanent end of life by way of hardware failure, and we anticipate the number only accelerating with time.

The edge network suffers from a lack of uniformity in construction, especially with the diverse array of connections. Our edge networks are providing both shared Ethernet connectivity to simple email users and Gigabit Ethernet to researchers exchanging considerable data sets with other institutions. Eliminating the trailing edge (shared hubs) and advancing the leading edge (switched Fast and Gigabit Ethernet) will enable us to continue managing the network to provide these diverse services. This diversity in the range of services has led to problems in the past where high-bandwidth users would negatively impact the performance of other hub-based users.

The overall goals of the edge upgrade are to:

- Eliminate non-IOS equipment (IOS has been the platform of choice for some time)
- Provide switched Fast Ethernet connections to every outlet
- Uplink each data closet at Gigabit Ethernet speed

User Outlet Switches

We have standardized on the Cisco 2950 series of switches for most user outlets. These managed switches provide 10/100 connections at a reasonable price point. They support the features we deem crucial for the next generation of network, including:

- Vlans and Vlan Trunking
- Port-based MAC limits (to limit damage from network bridging)
- 802.1x port-based authentication
- Configuration interface similar to other equipment (Cisco IOS)
- Secure/encrypted management access

The vlan capabilities of modern enterprise switches alone provide many opportunities for efficiency savings and flexibility in network design. The network topology can be simplified, making it easier to add redundancy where needed and aiding quick troubleshooting of problems. Additionally, vlans are one component used by voice-over-IP to protect and separate this traffic.

Distribution Switches

We standardized on the Cisco 3750 switch series for the network distribution layer. This encompasses primarily aggregation devices, such as those aggregating end-user switches in a network closet. The 3750 line has a unique and well-designed stacking capability, enabling "pay as you go" expansion capabilities. Furthermore, it provides high-density Gigabit Ethernet (GbE) ports at a reasonable price point.

One goal in designing the edge network was to minimize the number of devices between edge outlets and the router. In many cases today, devices are connected in a

cascading fashion. This can lead to multiple device hops to reach the router. Our approach is to install aggregation devices where needed (for example, in satellite data closets). In most cases, there will be no more than 3 hops from router to edge outlet (building aggregator, closet aggregator, edge switch).

The 3750s are also the platform of choice for providing GbE connections to end-user outlets, where desired. Our plans call for the replacement of the 6509 in Cyert A100 with two stacks of 3750 switches, providing GbE speeds to every host in the machine room. These stacks would be redundantly connected to a pod, as with other building connections.

5. Summary

The network upgrade plans call for substantial changes that will revitalize the campus network. Importantly, we will eliminate a trailing edge of equipment and position the network to enable leading edge technologies, including 10 Gigabit Ethernet, IPv6, and 802.1x authentication.

5. Revision History

Revision: 2	Date: 3/19/2004	Author: Kevin Miller
Changes: Added network diagram		

Revision: 1	Date: 3/18/2004	Author: Kevin Miller

Appendix A: Diagram of Core Network Design

The blocks for Wean, Baker/Porter, and CFA represent building aggregators as they would be connected to the core.