

**The Redesign and Reimplementation of the Stevens Institute of Technology Campus
Network**

*Group 12
December 9, 2002*

Faculty Advisor
Prof. Sumit Ghosh

Group Members

*Philip Tan
Nicholas Evans
Grace Shabo
Hans Kim
Jay Grossman*

We pledge our honor that we have abided by the Stevens Honor System.

Table of Contents

I.	Abstract	1
	1. Acknowledgement	1
II.	Project Final Design Plan	
	1. Introduction	2
	2. Design Requirements	3
	3. System Design	6
	4. Financial Budget	22
	5. Project Schedule	23
III.	Summary	24
IV.	References	26

Appendices

Parts List

Addendum: Retrieving Supplemental Materials

I. Abstract

We present our final design for the next generation campus network for the Stevens Institute of Technology. The current infrastructure at Stevens is inadequate to scale to the future needs of the population and several existing components are key single points of failure that can render services unusable until their repair. Currently there are security issues on both the wireless TreeNET and the existing Ethernet based networks which are directly related to the core design and implementation of the network. We address these problems through the use of edge media deployment, migration of faculty to separate networks to better protect against foreign and local attack, increased use of cryptography, load balancing of existing services to increase redundancy and performance, and the addition of firewalls and intrusion detection in key areas of the network. We expect our solutions to allow the Stevens Institute of Technology to scale to the current and future needs of its population as well as accelerate its Internet-based course initiatives by providing services that are secure, reliable, and cost effective.

I-1. Acknowledgement

The group would like to thank Christopher Hose, Director of Networking, Telecommunication, and Video at Stevens Institute of Technology for his valuable information on the current infrastructure of the campus network. We would also like to thank Sumit Ghosh for advising us through the many stages of this design project. Without his feedback and advice none of our efforts would have come to fruition.

II. Project Final Design Plan

II-1. Introduction

The current Stevens network infrastructure, while in most cases is structurally and functionally sound, contains some flaws that limit its potential to deliver consistent high performance and secure access to the Stevens community. The current structure of the network, as described by Mr. Hose, consists of a DS3 with 15 Mbit commit bandwidth to the Internet, provided by Verizon Communications, LLC. In addition, a 155 Mbit OC3 connection, as part of the vBNS/Internet2 initiative, exists at Stevens. Any outgoing connection to another vBNS site is routed over this OC3 connection. The campus network is protected by a Cisco PIX firewall, installed in February 2002. All campus buildings are connected to the core routers in the computer center via fiber optic connections. The TreeNET wireless network provides an 11 Mbit connection to the Internet and other campus resources using the IEEE 802.11b protocol. The wireless network currently uses a 40-bit encryption scheme with a 40-bit WEP key, with various access points distributed around campus.

Shortcomings in the Stevens network infrastructure exist mainly in the lack of future expandability for multimedia and security of the wireless network. Currently, the core networking environment consists of a single switch with no real failsafe mechanism (hot fail-over backup switch). Also, certain servers on campus are severely overloaded with no failsafe mechanisms in place to deal with system failures. In addition, portions of the campus are still using extremely outdated networking technologies. These shortcomings will be addressed in our campus network redesign. The redesign will be

based on the "band-aid" approach previously discussed in our project proposal.

II-2. Design Requirements

The design of the new campus network will address the issues we have encountered including future expandability, core networking redundancy, server and/or service redundancy, wireless network security, and overhauling the hardware deemed extremely outdated¹.

Future Expandability

As student populations, on and off campus, increase, the need for a larger infrastructure to support them becomes more apparent. We are at the beginning of such a need. Stevens Institute of Technology is growing, its initiatives expanding. Therefore now is the time to begin expanding services to accommodate these new needs and desires. The most notable desire is Stevens's WebCampus initiatives. The faculty is expressing the desire to deliver its course content over the Internet to a greater degree. This includes a greater delivery of multimedia content including filmed lectures. The network, with its current capacity and configuration, will not be capable of supporting the high volumes of media rich traffic in addition to its current traffic for other services. Our new design employs techniques to augment the capabilities of the network such that it will be able to support future initiatives and general growth for the foreseeable future.

¹ "Outdated" hardware includes the token ring coax networking seen in some of the old labs in the Burchard building. This would most likely include replacing the supporting network structure of this coax setup as well.

Core Networking Redundancy

While the core networking infrastructure is adequate in terms of intracampus connections and Layer 3 switching (based on the ISO Reference Model of Open System Interconnection), the single core switch at the center of a star topology, located at the computer center, is a poor design. From a reliability standpoint, switch failure would cause a complete networking failure. There is no backup switch to compensate for the loss of the main switch. Currently the computer center uses an approach where portions of the existing switching infrastructure would be used as a backup system by juggling cross connects between the failed main unit (Cisco 6509) and the contingency units.

From a performance standpoint, a single switch is not very good either. It must take the entire load from every point on campus. Performance penalties, both throughput and latencies due to the amount of traffic passing through the switch per unit time, are inevitable. For a high speed media rich network elevated latencies are to be avoided.

The new network design would build more redundancy into the core switching system. At least two core switches would be incorporated. This way, failures of one switch can be dealt with efficiently and automatically with the backup switch. Also, performance can be increased if the load is shared between the two switches. Further optimization can be realized through traffic pattern analysis and redistribution of cross connects according to the conclusions gained through the analysis.

Server/Service Redundancy

Currently, the server known as "Attila" services the campus with terminal logins, e-mail, academic, club, and student web sites, and file storage space. Some of these sites include

high traffic entities such as the Registrar's web site and the Office of Residence Life. Clearly, when Attila goes down due to scheduled service or system failure, the entire campus is affected. A redistribution of these services, or even using multiple "Attila's" to serve the campus would provide the redundancy necessary to provide reliable service to the entire campus community.

In the new network design, Attila, currently a single SGI server², would be replaced with a rack of servers. This server cluster would also contain shared storage to allow for more quota disk space and ease of upgrading. Service redirection would be implemented to target specific traffic to the appropriate servers in the cluster. All of these servers would be load balanced under the existing "attila.stevens-tech.edu" DNS name, to provide a smooth transition for campus users.

Wireless Network Security

As discussed by Mr. Hose, the current wireless network implementation utilizes a non-rotating 40-bit WEP encryption key. Through the use of probing and sniffing of raw 802.11 frames on the wireless network, we successfully cracked this encryption key, within a few days. The current implementation requires that every wireless card on the network have its MAC address registered. In this way, the access points will not permit an unregistered MAC address from communicating. While providing some security, the current implementation is clearly lacking. Using the same network probing techniques described above, MAC addresses can be easily harvested. Combine this with the relative ease of cracking a 40-bit WEP key, and the security shortcomings become apparent.

² Currently Attila is an SGI Challenge machine with 6 CPUs, 1 GB RAM, and approximately 100 GB of disk space.

Legacy Hardware Replacement

Legacy hardware is still found in many places on the Stevens campus. The most notable are the antiquated systems in the sophomore physics lab on the fourth floor of Burchard. The lab contains Pentium 90 MHz machines running Windows 3.1, which are using token ring coax networking. Coax networking is extremely limited in performance and scalability, as compared to twisted pair Ethernet. As part of this project, all the legacy coax networking will be removed, including its supporting infrastructure. This will all be replaced by twisted pair Ethernet with supporting 10/100 Mbit switches.

II-3. System Design

For our final design we decided to pursue a modified band-aid approach. Considering that the core wiring infrastructure consists of fiber optic cross-connects between major buildings and the Computer Center's Datacenter we believe this to be the most cost effective solution. This solution will retain the bulk of this wiring infrastructure (aside from adding new fiber and copper wiring for the outdated areas of campus) and will allow Stevens to retain as much of their existing switching and routing hardware as possible. We have provided a network diagram to illustrate the service and basic network structure and will refer to this diagram to aide description of our final design. To describe our design we will start from the periphery of the network and work toward the core.

During our investigation of the existing network infrastructure we conducted several ping tests as well as monitoring of traffic to determine the potential loads a single user of the network can impose on the infrastructure. According to Chris Hose, the individual dorm networks consist of Cisco switches (we assume Catalyst 2900 or 3500

series variants) cross-connected together per floor and a gigabit cross-connect for the aggregate "switch" to the Computer Center's primary Catalyst 6509 in the datacenter. During our investigations we noted a bulk data transfer between two systems within Technology Hall. This transfer occurred between two computers on two different floors of "Tech", during evening ours (peak time) and, we believe, different switches. We noted a maximum transfer rate of about 10 megabits per second. Granted, part of this result may be due to system load on either end, but given the performance of the end systems involved (Athlon XP 1800+ @ 1.53 GHz) their ability to provide at least 50 megabits of throughput is certain. We performed other throughput tests that yielded better results, but in all of our testing the throughput between two systems in "Tech" on different switches was never greater than 35 megabits. Based on these tests we believe that the aggregate cross-connects may only be Cat5 copper. (See Appendix A for a sample of the results of the above investigation.)

The switching backplane on a Cisco Catalyst 2900 series device is between 4.8 and 8.8 gigabits per second (depending on the model), while the 3500 series boasts a 10.8 gigabit backplane. In either configuration the switch is not the limiting factor given the number of ports each model switch can accommodate. While the switch itself can have nearly all of its ports running at full-duplex 200 megabits without a perceivable loss in performance, the cross-connects aggregate that 4.8 or 9.6 gigabits of traffic down to either a 200 megabit or 2 gigabit full-duplex piece of copper or fiber, respectively. Considering that the bulk of traffic on the network would be accessing school-provided services (e-mail, Registrar, course web pages) or the Internet, and not reside within the ports on a single switch, a restrictive Cat5 cross-connect between aggregate switches is a

problem, even within a dorm. Intra-switch transfers will be limited by port speed and duplex, but Inter-switch transfers within a building are limited by their cross-connects' port speed and duplex. Therefore it is very important to augment these links, if they are not already, with fiber at 1 gigabit full-duplex. We are assured by the IT department that the uplinks between buildings and the Computer Center are fiber, but nothing was mentioned about the cross-connects within the buildings themselves.

While we are still on the periphery of the network we should discuss our media delivery schema for Stevens's WebCampus initiatives. In our opinion, the best Internet-based course would have streaming video of lectures and interactive course materials to engage the student and allow for reviewing at a later time. Moving course materials to the Internet allows Stevens to accommodate more students without a significant increase in its building infrastructure. This reduces costs while increasing revenue through increased numbers of tuition payers. While Stevens has the beginnings of a "WebCampus" through its online courses and WebCT system there is a lack of major media delivery across campus and on the web. It is this media delivery, mostly video, which will allow a student to view a lecture both live, and on-demand at a later time, perhaps in review for a final exam. Our solution to providing media delivery on campus is to use edge delivery schemes commonly used by major Internet media providers like Akamai. Essentially the media stream is recorded and encoded, or broadcast live to a "splitter", contained within the central datacenter. This splitter will then rebroadcast the stream to any number of media servers at the edge of the network near the viewers. As an example, consider there are 100 viewers in a building who wish to view a media stream. In the classic delivery system a central set of media servers would be present in a datacenter located outside the

example building. Each viewer would connect to the central media server and watch the video stream. During this scenario, if we examine the link between the viewers' building and the central media server we would see 100 separate connections transacting video data. In the edge delivery system there would be a media server within the building that the viewers can connect to. A central splitter would supply the video stream they view, but upon examining the link between the central location and the viewers building we see only one connection for the video stream. In effect we save 99 users of bandwidth, the local media server being the only "viewer" from the central media splitter. Considering that in any concentrated area of viewers on campus (a dorm for example) we could easily be accommodating more than 100 users, the benefits immediately become clear. Moving media to the edge of the network has the potential of saving extreme amounts of bandwidth on the cross-connects between buildings and the Computer Center. Lowering bandwidth usage on these connections will reduce load on the central switching infrastructure, allow other types of traffic to traverse the links, and reduce quality of service issues with the media being delivered. Diagram 2 visualizes the technique we prescribe in our design.

Both on-demand and live content can be accommodated in the edge delivery technique. Essentially the media splitters can rebroadcast a video stream, upon request, to a local server where the end user can then view it over the local dorm switching infrastructure. The central media splitter would pull on-demand content from a central file server which we will discuss later. Providing media delivery services in this way does increase cost somewhat, but the costs involved for the much larger amounts of bandwidth required to view all media from a central server far outweigh them.

For off-campus viewers a load balanced media server farm would deliver content through the routing infrastructure, over the Internet, and to their desktop wherever they are located. Depending on the number of off campus viewers the bandwidth to the Internet may need to be augmented, but we will discuss this more at a later time.

Before addressing the core network in our trip across the new Stevens network topology we should touch on the wireless "TreeNET". As we've mentioned before, the wireless access is fairly well distributed across campus. Lucent/Agere access points service most populated areas that do not already have wired access. Despite the good coverage on campus the security of this service is insufficient. While a prospective user should not expect security over a wireless medium, there are significant steps that can be taken to mitigate trivial interception of data by an unauthorized third party. We've proven that the current security configuration is not adequate, and now we will outline some steps to increase that security. In our design, 128-bit security is a must. The time and traffic required to crack a WEP key increase by an order of magnitude by adding the 88-bits of encryption. There is really no reason why TreeNET should be using only 40 bits of encryption as students are required to purchase gold standard cards which support 128-bit. The second issue we address in this design is the use of a random key. As we explained previously, WEP uses RC4 as a cipher. Certain character combinations yield statistically weaker keys, but there are several programs available to check the key strength. All keys should be randomly derived. Simply looking at the current 40-bit key reveals that its character combinations were not randomly chosen. We prescribe using a random character generator, preferably a `/dev/urandom` or other source on OpenBSD, in combination with a one way hash algorithm such as MD5 or SHA1 as the source of the

necessary characters for a key. This key should be put into service after it successfully passes the tests of the aforementioned integrity checking programs. Finally the WEP keys should be rotated at a minimum of once per semester, but preferably once a month. We realize this may impart some initial logistics problems until the user population becomes aware of the procedures set forth by the IT department. We feel that, in the end, they will appreciate the added security despite the increased effort required to maintain their connections.

We've finally made it to the core of the new network design. First let's address the underlying Layer 2/3 network. The existing network utilizes a relatively new Cisco 6509 Layer 3 switch which aggregates the fiber cross-connects from all over campus and then switches the logical subnet traffic from different segments. This Layer 3 switching is a significant improvement over traditional routing technology. It provides much higher throughput and greater flexibility in networking capabilities than traditional routing, kudos to the IT department for making the change. The only problem we can identify with this configuration is the lack of redundancy. While this series Cisco switch has some built in redundancy (i.e. multiple power supplies, multiple switch blades), if there is a large enough failure in the unit, the entire network will be taken down. The solution outlined by Chris Hose involved a juggling of fiber cables to other switches while the 6509 was repaired or replaced. With an increased reliance of both the student and faculty on the services provided by the Computer Center this is an insufficient solution as the downtime would impose a significant time window without service. We propose the obvious, the purchase of a secondary 6509 and its configuration as a hot fail-over device is required. Another benefit of purchasing the additional 6509 is the ability to split load

between the units. While these switches have a 256 gigabit backplane, any additional services utilized such as VLAN tagging or HSRP for hot-fail-over can reduce the performance of the switch and subsequently, during peak traffic volumes (especially in our new media rich network), can reduce the throughput or increase latency on the network. This also allows the Computer Center to analyze traffic patterns and optimize the location of segments in the switch fabric to keep higher volumes of traffic within a single 6509 instead of moving data across a cross-connect between the units.

Fault tolerance is an important issue addressed in our new network design. We believe there are no fault tolerant measures in place in the rest of the networking infrastructure. This includes the PIX firewall and the router used to access the Internet. We propose the purchase of an additional PIX firewall to act as a fail-over for the existing unit. There would be no need to load balance these units. Considering the current and proposed bandwidth, the Internet traffic is not great enough to make the PIX the bottleneck.

While we are discussing the connection to the Internet let us explain what our new network calls for in terms of bandwidth. We feel that the 15 megabit commit is sufficient for current Internet traffic, especially when combined with the traffic shaping measures the IT department recently put into place. However this bandwidth will not be sufficient if there are increases in off-campus students, or if on campus student population increases. Also, with the increase in WebCampus courses offered and the future inevitability of streaming lectures, 15 megabits just will not be enough to support the needs of the school. We propose two possible solutions, the first being to increase the commit bandwidth to the full DS3 rate of 45 megabits from Verizon. This will allow

more traffic through the pipe to the Internet, but will not solve fault tolerance issues or service issues. The better, but less cost effective, solution would be to bring in an additional DS3 from another provider, perhaps UUNet, Level 3, Sprint, MCI, GlobalCrossing, Cable and Wireless, etc. at some commit rate. The IT department could then use BGP to dynamically route Internet bound traffic. Attacking the problem this way allows continuous traffic to be diverted over either a less bandwidth costly link, over a better route path, or load balanced over both links. It allows dynamic route changes in the event of a link failure or topology change elsewhere on the Internet.

While we are discussing routing to the Internet we should describe how we go about this in our design. We have two Cisco 7507 routers, one primary, one fail-over, handling the routing on this network. By using the 7507 we can accommodate both DS3 links as well as the OC3 for the vBNS ATM link. We feel that using a lower model router, i.e. 7206VXR, would be pushing the limits of the device's capacity, therefore not allowing for future expansion. The raw amount of potential traffic (upwards of 245 megabits) combined with access lists and other services provided by the router would simply be too much for it to handle. Therefore we chose the more expensive 7507 as it offers over twice the performance of the 7206VXR and larger module capacity as well as better redundancy options. At this time we are unaware what the Computer Center's current routing situation is in terms of the devices they use, so a single 7507 may be employed right now. In that case we would do as we have done with the switching and firewall infrastructure and purchase another unit so we could implement hot fail-over ability between the units.

Now we have reached the Computer Center itself, the core of the physical and

logical campus network. Above we covered basic redundancy upgrades, and now we will cover enhanced security measures and later, our expansion of service capacity. The current protection from the Internet is more than adequate. Cisco's PIX is an industry standard firewall. While we would have liked to see a combination of IPFilter³ and FreeBSD⁴ used for packet filtering services due to their low cost, good performance, and good security, we feel the purchase of the Cisco unit is a good alternative. The only prescription we have is to do a full audit from the Internet of the entire Class B address range Stevens owns and then adjust the access lists on the routers and PIX to better protect the campus. We were not able to do this ourselves and provide an exact list of security enhancements because we felt it would be seen as a hostile action by the Computer Center if detected and they may have taken action against us. This is one of the grey areas in the new design because it requires further investigation by an authorized party to gather the information needed to make an informed decision on the best solution. In our new design we add additional firewalls, rearrange some of the routing segments, and install an Intrusion Detection System ("IDS" from now on) to intercept would-be attackers.

First we describe our additional firewalls. We believe that even with increases in security on TreeNET any publicly accessible part of campus becomes suspect. Therefore, we think it is necessary to have a separate aggregate of switches that handle the public network jacks and the wireless access points which, before connecting to the primary Cisco 7509, must pass through a FreeBSD/IPFilter bridge-firewall. We chose FreeBSD/IPFilter for its low cost (hardware and time to configure), plain English filter

³ <http://www.ipfilter.org/>

⁴ <http://www.freebsd.org/>

construction, good performance (90Kpps @ 64bytes/packet)⁵, and exceptional filtering ability. The other benefit of this setup is that the unit is completely transparent to the network. It would have an IP address for management purposes, but the packet path traverses a logical bridge with no IP addresses on either interface of the bridge, therefore an attacker sees nothing from his or her point of view and cannot mount a successful attack. One might ask why we should bother having an intermediate firewall; why not just perform all filtering on the 6509? The answer is simple: while the 6509 is more than capable of filtering traffic from the publicly accessible sections of campus, it does not provide a layered security approach. If, for whatever reason, the access lists on the 6509 are compromised, or accidentally removed, there is no protection for the rest of the campus from attackers that can easily connect a mobile device to a network jack and mount an attack. With the addition of the bridging firewall an additional layer of security is in place which allows abstraction from the security policy of the rest of the campus. The idea is to have access lists on the bridging firewall as well as the 6509 so that in the event the bridging firewall stops filtering (system crash, user error, etc.) an attacker would still hit the access list "wall" on the 6509.

Another benefit of adding this bridging firewall between public access points (TreeNET, public network jacks) is that the traffic that flows over those connection points can be monitored with an IDS. This "node" of the IDS would connect to a greater intrusion system that covers other links to the Internet. We prescribe using Snort⁶, a freely available, high performance, and flexible system for detecting bad behavior on

⁵ This was conducted through our own testing using ping packets transmitted at the fastest rate possible. IPFilter was configured on a FreeBSD 4.6 system running on an Intel Celeron 900Mhz 1U server. IPFilter had an access list comprised of 300 static (non-stateful) rules.

⁶ <http://www.snort.org/>

networks. Our Snort based solution includes two sniffer nodes (the bridging packet filter described above and an additional unit) that gather and analyze traffic with various attack signatures and other protocol analyses. These forward sniffer nodes then transmit their data to a MySQL database, which is in turn accessed by an Apache web server, with PHP support, running a supplemental program called ACID⁷. ACID takes the data gathered in the MySQL database and interprets and displays it. The interface is very intuitive and it allows the user to see a near-realtime view of the potential bad traffic on the networks the sniffer nodes can see. To facilitate traffic gathering, a monitor port would need to be configured such that all traffic coming from the external routers into the network before they are filtered by the PIX can be seen by the sniffer's interface. An additional IP addressed interface would be needed on the sniffer so that communication between it and the MySQL server could be performed. In the case of the bridging firewall the IP address management interface would be able to communicate with the MySQL server. Examples of usage and the GUI can be seen on the Snort and ACID web sites provided in the footnotes. In addition, all of the solutions here are cost of hardware and time only. Snort, ACID, and MySQL are freely available programs, and the hardware requirements to support them are not great.

Finally, in our quest to augment security on the new network, we decided to segregate the academic and administration departments completely from the rest of the campus network. This involves using an additional FreeBSD/IPFilter bridging firewall and aggregating their cross connects before this new firewall. We would recommend that this device act as a third Snort gathering node so that the Academic and Administration departments could be protected from attackers inside the firewalls, i.e. students. Again we

⁷ <http://www.andrew.cmu.edu/~rdanyliw/snort/snortacid.html>

take the bridging approach because it is simple to install, requiring no changes in routing, and it is completely transparent to the attacker.

The last stop in our security augmentation approach is to actually encrypt VPN sessions. The VPN documentation posted on the Stevens IT web site⁸ explicitly instructs the user to deactivate encryption, and our tests concluded that all parts of the session (the initial connection, as well as the contents of the tunnel) are in plain text. We believe that Microsoft's VPN solution is the easiest to configure and maintain, but it may not be the best performing system around. Therefore we suggest load balancing of multiple servers to increase capacity as well as installation of cryptographic acceleration boards into each of the VPN servers. This allows the operating system to offload cryptographic functions to the board and significantly increase throughput and the amount of sustainable connections as well as reduce the requirements of the host system. These boards can be purchased for about \$300 each from varying vendors, but overall the solution is still far more cost effective than purchasing an IPSec based hardware solution from Cisco.

We have finally reached the area we feel holds the greatest potential for increases in performance and redundancy, the service infrastructure. Here we will cover what we plan to do with Attila as well as other services provided by the school. Our approach involves using load balancing services provided by F5 Networks' BigIP device to distribute service traffic among many backend servers. This will increase fault tolerance within service cluster groups, allow for higher throughput, ease servicing of those servers, and allow the service to scale to the needs of an ever-growing campus. Throughout this explanation, consult Diagram 1 for a visual representation.

Attila is seeing increases in usage as each new academic year starts. We have

reached the point where Attila's load is becoming too much for the server it is on (SGI Challenge, 6 MIPS processors). We consistently see load averages over 3 with surges as high as 33 ('top' output). A hardware failure or planned service on Attila results in students and faculty losing e-mail service, personal web site uptime, Registrar web page use, etc. There are simply too many services residing on the single system which has no backup unit that can be swapped in should Attila fail. Our solution involves converting Attila into a rack of servers each running a separate service and dynamically load balanced behind a Virtual IP address ("VIP" from now on) assigned to the F5 BigIP. We envision two servers per service in the beginning with the possibility to extend this to any number of servers needed to service the loads seen. The BigIP will be configured with a single VIP who's DNS name will be `attila.stevens-tech.edu`. We then configure the BigIP to redirect various services to their appropriate backend server. For instance, if we have two servers handling SMTP traffic for Attila, we instruct the BigIP to redirect port 25 traffic it sees on the Attila VIP to each of the servers `smtp1.stevens-tech.edu` and `smtp2.stevens-tech.edu`. We can then have the BigIP keep persistence on the connections which forces it to redirect the same source IP address to the same backend server each time that IP makes a request of the service. Doing this is good for web servers where session data would not be shared among multiple servers in the cluster. The BigIP allows for various methods of load balancing including least recently used, round-robin, least loaded, etc. The beauty of this approach is that when one of the SMTP servers needs to be serviced (program upgrades, etc.) that individual server can be taken out of the VIP load balancing rotation. The IT department then waits for any persistent connections to expire (can be monitored from the BigIP's interfaces), and then they are free to work on

⁸ <http://www.stevens-tech.edu/it/vpn.pdf>

that unit while the BigIP directs new connections to the other SMTP server in the rotation. Once the "down" server has its servicing complete it is simply added back into the SMTP rotation and the BigIP will send begin sending new connections to it as per the load balancing algorithm.

If a server in an individual service cluster crashed unexpectedly, the BigIP can detect that the device is no longer accessible and will dynamically direct new connections to the other good server(s) in the cluster. The BigIP's themselves are configured in hot fail-over configurations with a 2-3 second lag for fail-over. They share information about the state of their connection tables as well as their current hardware condition. The units share VIPs and MAC addresses on those "interfaces" such that, upon failure, they can immediately bring up the interface and begin servicing existing and new connections without a perceivable loss in connectivity to the end user. To handle the load of all the services we envision on campus, now and in the future, we believe two sets of BigIP units (four units total) would be required.

To facilitate moving Attila to a clustered service approach a central file storage mechanism would be required. The Computer Science department's Unix lab is very similar to our approach. They have each user's "home directory" stored on a central server (guinness.cs.stevens-tech.edu) which is then mounted via an NFS file sharing mount onto each desktop computer's Unix /home directory. This allows each user to access their data from any desktop without having to keep multiple copies of that data on each desktop. We would accomplish the same thing with our load balancing approach. The user's home directories would be stored on a Network Appliance Filer F840 network attached storage device. Each server in the load balanced cluster would then mount these

NFS exported directories locally. Essentially each service has access to the user's home directory in the same way Attila-native services do now. This allows us to reduce the cost of the cluster servers; since they do not need a large local storage capacity, we simply move everything to a central file storage device. In this case the Filer F840 is the best approach despite it being the most expensive piece of hardware in our design. We considered building our own central file storage unit at significant cost savings, but the F840 adds many features that ease administration. Its fault tolerance, automatic RAID rebuilding, scalability, and performance capabilities are far greater than anything that can be constructed from off the shelf components. These units are worth the cost. We start initially with a half terabyte, dual shelf unit, with an additional 5U controller device. Essentially the Filer, as it stands, utilizes 11U of rack space with an additional 3U required per addition of 250 gigabytes. This unit will act as shared storage for every service cluster in the Computer Center.

Using our SMTP service as an example of how the load balancing situation will work in the new datacenter, we will not outline any variation between it and the other services we provide. The rest of Attila's services will be bound under the `attila.stevens-tech.edu` DNS name we mentioned earlier and include: IMAP4 (SSL), POP3 (SSL), Shell Accounts (SSH), FTP, and HTTP for student web pages. Each of these servers need not be the same configuration. We envision utilizing dual processor Pentium III 1U, inexpensive rack servers for the majority of these services with exception to the shell account systems. Since those systems may involve users running scientific scripts or other processor intensive tasks we believe using more expensive dual or quad Pentium III XEON systems would be in order. The most important note about this configuration is

the cost. Even with the BigIP systems and the servers to replace Attila, the overall cost is still less than that of purchasing a new SGI system and migrating all of the content to it. The other scalability, ease of administration/service, and redundancy benefits become obvious and much more attractive when they cost less to implement.

Other services that we plan to load balance (see Diagram 1) are:

- Registrar's website and ES registration system (Registrar Farm)
- VPN servers (VPN Farm)
- Media splitters (RealMedia Splitter Farm)
- Off-campus media delivery servers (Off Campus RealMedia Farm)
- Webmail (Webmail Farm)
- Stevens FTP (FTP Farm)
- WebCT (WebCT Farm)
- HTTP Proxy servers (Proxy Farm)
- Netsrv (although this could be provided directly by the Filer F840)
- www.stevens-tech.edu (WWW Farm).

These services are load balanced in largely the same manner as we described with the Attila services. Some notable exceptions are VPN which does not require a shared file system. Complications may arise with the configuration of the proxy servers if a shared storage medium is used, but it is not advisable anyway since the additional latency of defaulting to the network attached storage may not be desirable, despite the exceptional speed of the Filer F840. Webmail may require persistence of connections depending on the software's design, as would the ES system. Since the backend of the registration software is obviously a database, all data is centralized; however, the access to that data

would not be granted unless persistence is activated. All HTTP based services generally should be load balanced with persistence, as HTTP will often establish many connections to transfer data. This behavior may affect connection tracking mechanisms of the software used on those web servers. As we mentioned before, Netsrv services could be provided directly by the Filer, but if finer grained access controls are needed front-end Samba⁹ based systems can be used. In our design we are not load balancing DNS because most end systems have functionality to use multiple DNS servers and the DNS system itself has some measures for fault tolerance. Taking the measures we have outlined in our design will allow the campus network to scale its services to the needs of the community while maintaining a cost-effective, redundant and secure medium for delivering those services.

II-4. Financial Budget

Here we show an estimated financial budget for the components required to build the network as we have designed it. All of the major components are accounted for in the table on the following page, and we have included a miscellaneous item entry to cover any expenses we did not anticipate during the course of the design. This entry will also cover any underestimates we make. References for the prices are annotated under the table.

⁹ <http://www.samba.org/>

Items	Qty.	Unit Price	Total
Additional Cisco 6509 Chassis ¹	1	\$ 22,000.00	\$ 22,000.00
Additional PIX Firewall ²	1	\$ 10,000.00	\$ 10,000.00
Cisco 7507 Routers ³	2	\$ 30,000.00	\$ 60,000.00
Switching Blades for 6509*	10	\$ 3,000.00	\$ 30,000.00
DS3 Cards for 7507*	4	\$ 5,000.00	\$ 20,000.00
ATM Cards for 7507*	2	\$ 5,000.00	\$ 10,000.00
F5 Networks BigIP ⁴	4	\$ 15,000.00	\$ 60,000.00
Network Appliance Filer F840 ⁵	1	\$ 115,000.00	\$ 115,000.00
Additional Filer F840 Storage Shelf ⁵	1	\$ 15,000.00	\$ 15,000.00
Generic Service Servers (Dual Pentium4 1.7Ghz, 256 Ram, 40GB HD.) ⁶	50	\$ 1,500.00	\$ 75,000.00
Specialty Shell Servers for "Attila" (Dual XEON, 2Gig Ram) ⁶	2	\$ 7,500.00	\$ 15,000.00
Avocent/Cyber KVM Controller Units ⁷	7	\$ 700.00	\$ 4,900.00
Intel Fiber Gigabit adapter (5 packs) ⁷	5	\$ 2,800.00	\$ 14,000.00
Wiring new 10/100 jacks in outdated campus areas ⁸	200	\$ 300.00	\$ 60,000.00
Miscellaneous Expenses	1	\$ 75,000.00	\$ 75,000.00
Total			\$585900.00

References

1. <http://shopper.cnet.com/shopping/search/results/1,10214,0-1257,00.html?tag=top&qt=cisco+6509&cn=&ca=1257>
2. <http://shopper.cnet.com/shopping/search/results/1,10214,0-1257,00.html?tag=top&qt=Cisco+PIX&cn=&ca=1257>
3. <http://shopper.cnet.com/shopping/search/results/1,10214,0-1257,00.html?tag=top&qt=Cisco+7507&cn=&ca=1257>
4. http://www.kernelsoftware.com/products/f5_-_hardware.html
5. <http://www.ftsi.fujitsu.com/services/solutions/federal/gsa/netapphard.html>
6. <http://www.offmyserver.com/cgi-bin/oms/home.html> (FORMERLY iXSYSTEMS.NET)
7. www.cdw.com
8. Estimate provided by Chris Hose.

* Our own overestimate.

II-5. Project Schedule

The project in the second semester will mainly consist of gathering more network statistics through the use of a mock network. This mock network will be used in conjunction with the results obtained from the first semester regarding an average student's network usage. This mock network will allow students and/or faculty to freely use the network while their usage is monitored to help provide a more accurate sample.

The Gantt chart at the end of the report will detail this work.

III. Summary

The goals of this project are to provide Stevens with a secure, reliable, and high performing network with the potential for expandability. Much of the core network infrastructure is sound, with a few problems that are quite serious. Our design will address the most serious of these problems and provide a new networking environment which improves the quality of service dramatically.

In terms of security, the wireless network will be enhanced through the use of a 128-bit encryption key, and the use of a random key generator to minimize “trivial” encryption keys. Also, the key will be rotated periodically to further reduce the chance of compromising the network. The VPN setup will be augmented with encryption turned on being the default setting. Cryptographic acceleration boards will be installed in the VPN servers to help reduce the overhead associated with encrypted transmissions.

In terms of reliability, Attila will be transformed into a rack of servers with its current services being distributed among the servers. This will reduce the risk of a complete loss of Attila’s services should a server fail. Also, such load distribution will increase performance and response time of Attila when the server load is at its peak. Other services will be augmented in a similar fashion. The single central switch at the center of the network’s star topology will be paired up with another switch to help prevent a complete loss of network switching and subsequently a complete network shutdown. The Cisco PIX firewall will also be paired up to create redundancy.

To boost network performance, the aforementioned redundancy will be implemented. Since the main wiring around campus is sound, implementing redundancy

in what exists will have a profound effect on network performance. Also, the parts of campus that are currently still in the Dark Ages of networking (those locations with Token Ring coax) will be upgraded to help increase performance.

Expandability will be addressed through investigations of increasing the current 15 Mbit commit to the Internet. Placing media servers in the dorms would allow streaming media to be delivered to entire buildings while only placing one load per building on the core media servers.

With all of the changes mentioned above, we feel that the Stevens network can fully meet or exceed current networking demands and can fully meet any foreseeable future need due to the school's growth both on and off campus.

References

1. Stevens, Richard. TCP/IP Illustrated Volume 1: The Protocols. Addison Wesley, 1994.
2. Garfinkel, Simson; Spafford, Gene. Practical Unix & Internet Security. O'Reilly & Associates, Inc., 1996
3. Slatter, Terry; Burton, Bill. Advanced Routing in Cisco Networks. McGraw Hill, 1999.
4. Scambray, Joel; McClure, Stuart; Kurts, George. Hacking Exposed Second Edition. McGraw Hill, 2001.
5. Comer, Douglas E. Internetworking with TCP/IP: Principles, Protocols, and Architectures. Prentice Hall, 2000.
6. Callaghan, Brent. NFS Illustrated. Addison Wesley, 2000.
7. Perlman, Radia. Interconnections Second Edition: Bridges, Routers, Switches, and Internetworking Protocols. Addison Wesley, 2000.
8. Hulton, David. Practical Exploitation of RC4 Weaknesses in WEP Environments. [ONLINE]
<http://www.dachb0den.com/projects/bsd-airtools/wepexp.txt>
9. Fluhrer, S.; Mantin, I.; Shamir A. Weaknesses in the Key Scheduling Algorithm of RC4. [ONLINE] <http://citeseer.nj.nec.com/486392.html>
10. Chase, Jeffrey S.; Gallatin, Andrew J.; Yocum, Kenneth G. End-System Optimizations for High-Speed TCP. [ONLINE] www.cs.duke.edu/ari/publications/end-system.pdf

Appendix A

The following shows part of the output of the “tcptrace” program used to monitor a student’s network usage over a 24-hour period. The sample output is approximately a 15-minute span. Of particular note is the “Throughput” field at the bottom.

TCP connection 26:

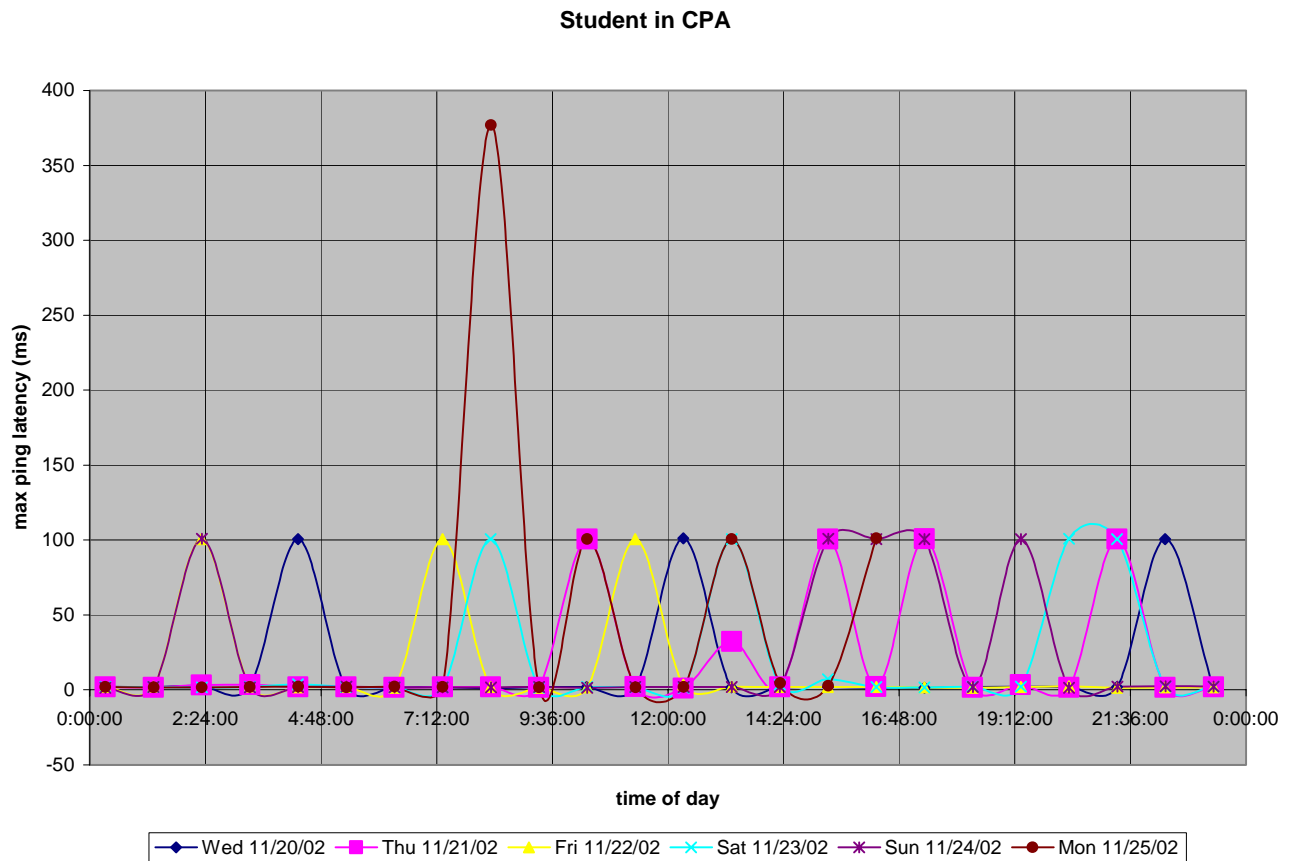
```
host ay:      hkim5.u02.stevens-tech.edu:4776
host az:      somestevensuser:139
complete conn: yes
first packet: Wed Nov 20 18:50:37.835377 2002
last packet:  Wed Nov 20 19:06:27.364074 2002
elapsed time: 0:15:49.528696
total packets: 926281
filename:     4-no-nbt-broadcast

ay->az:                az->ay:
total packets:         315944          total packets:         610337
ack pkts sent:         315943          ack pkts sent:         610337
pure acks sent:        296776          pure acks sent:         48
sack pkts sent:         0              sack pkts sent:         0
max sack blks/ack:     0              max sack blks/ack:     0
unique bytes sent:     1305661         unique bytes sent:     870008205
actual data pkts:      19166          actual data pkts:      610287
actual data bytes:     1306582         actual data bytes:     870014312
rexmt data pkts:       15             rexmt data pkts:       39
rexmt data bytes:      921            rexmt data bytes:      6107
zwnd probe pkts:       0              zwnd probe pkts:       0
zwnd probe bytes:      0              zwnd probe bytes:      0
outoforder pkts:       0              outoforder pkts:       0
pushed data pkts:      19166          pushed data pkts:      19170
SYN/FIN pkts sent:    1/1            SYN/FIN pkts sent:    1/1
```

req 1323 ws/ts:	Y/N	req 1323 ws/ts:	Y/N
adv wind scale:	2	adv wind scale:	0
req sack:	Y	req sack:	Y
sacks sent:	0	sacks sent:	0
urgent data pkts:	0 pkts	urgent data pkts:	0 pkts
urgent data bytes:	0 bytes	urgent data bytes:	0 bytes
mss requested:	1460 bytes	mss requested:	1460 bytes
max segm size:	294 bytes	max segm size:	1460 bytes
min segm size:	39 bytes	min segm size:	1 bytes
avg segm size:	68 bytes	avg segm size:	1425 bytes
max win adv:	256960 bytes	max win adv:	5840 bytes
min win adv:	255500 bytes	min win adv:	4381 bytes
zero win adv:	0 times	zero win adv:	0 times
avg win adv:	12247 bytes	avg win adv:	5114 bytes
initial window:	72 bytes	initial window:	4 bytes
initial window:	1 pkts	initial window:	1 pkts
ttl stream length:	1305661 bytes	ttl stream length:	870008205 bytes
missed data:	0 bytes	missed data:	0 bytes
truncated data:	425564 bytes	truncated data:	841945977 bytes
truncated packets:	18811 pkts	truncated packets:	609621 pkts
data xmit time:	949.527 secs	data xmit time:	949.527 secs
idletime max:	120193.1 ms	idletime max:	120242.2 ms
throughput:	1375 Bps	throughput:	916253 Bps

Appendix B


The following graph shows the ping statistics of a student's computer in Castle Point Apartments (CPA). A machine running FreeBSD 4.7 was set up in Technology Hall (Tech) and pinged various machines around campus, for a span of approximately one week. Each ping was performed hourly, with 1000 packets being sent out during each ping operation. As can be seen, the maximum ping latency, for the most part, was quite low. This leads to the conclusion that the wiring around campus is sound.



Appendix C

The following graph shows the ping statistics of Attila. The ping operation is the same as described in Appendix B. As can be seen from this graph, the maximum ping latency was extremely erratic, and extremely high. Of particular note are the times at which the latency is the highest. Pings to Attila had the highest latencies at early hours, between 2 AM and 7 AM, the times when the campus's network traffic would be the least. Since the wiring from Tech to CPA is essentially the same as the wiring from Tech to the computer center where Attila is housed, we were led to believe from this graph that Attila is severely overloaded, thus requiring an upgrade.

Attila Ping Statistics



Parts List:

The following is a list of parts we will require for next semester in order to build our simulation network. This list is subject to change depending on the approach we decide to take in this simulation.

Part	Quantity
Cisco 6509 Representation (Cisco 3500 or 2900 series switches)	2
Routers (Pentium Processor computers)	2
PIX Firewall Representation (Pentium Processor computer)	2
FreeBSD/IPFilter Firewall (Pentium Processor computer)	2
BigIP Representation (Pentium Processor computer)	2
Attila Representation (Pentium Processor based PC)	5
VPN (Pentium Processor Based PC)	1
WWW Farm (Pentium Processor based PC)	2
Network Appliance Representation (Pentium Processor based PC)	1
IDS (PC)	1
DNS (PC)	1
Aggregate switches (Inexpensive 8-16 port unmanageable switches)	4-8
Wireless Network (Lucent/Agere Access Point with several client cards)	1
Cabling (Category 5 Patch Cords, Crossover cables)	25-50
Clients (Laptop PC's)	Depends on Scale

Simulating the network will be tricky considering we will not have access to the exact equipment used in the real-life version. We will try to simulate performance specifications to as close a scale parameter as we can. The Cisco 6509 will be the hardest part of the network to simulate since the Cisco 3500 series switches are not capable of performing Layer 3 switching we may be required to replace them with Pentium PC's that represent the switch and simply route traffic between segments. Their performance will need to match the switch to scale. We are planning on providing as much of the hardware as physically possible before turning to other sources for hardware on loan. Our requirements above are for approximately 19 to 21 Pentium class PC's. We believe that

collectively we can supply about 10-15 PC's to represent the core components of this network design, the rest we will need to purchase or borrow from the ECE Department if at all possible (Professor Ghosh mentioned that there may be some systems in the ECE Department we can use). The systems we purchase, if any, would not be new computers, but very inexpensive (\$100-125) used systems. To effectively simulate the design we will need to represent each major structure in the network. Therefore we must represent the Internet-bound routing, the firewalls, the core switching infrastructure, at least one service farm (WWW Farm), Attila with service redundancy for at least one service (WWW) and provision for at least one other Attila-based service, the Net App Filer, VPN, the Wireless network equipment, IDS, DNS/DHCP, and the Load balancers. The client (student/faculty) computers will be our own laptop PC's as well as those of our peers who can assist us in the simulation at the time of the simulation. While we require a larger number of PC's we do not require the same amount of monitors, as most systems will be UNIX based, we can simply use shell accounts over the network or serial line consoles to login and work with each system.

As stated before, this list is not final, and not necessarily reflective of our simulation plans. Over the course of the coming weeks we will determine the best course of action after conferring with Professor Ghosh and Professor McNair. We will then proceed with simulating the network, or it's substructures, based on those conferences.

Addendum: Retrieving Supplemental Materials

There are two supplemental diagrams as well as a Gantt chart that may assist the reader. Due to their large size we are not able to embed them into this document.

Therefore we have posted them on our official FTP site:

ftp://sd.syphen.net/pub/final_supplemental/

Login using your FTP client's standard anonymous login procedures or simply click the link in Acrobat Reader. Any problems regarding file downloads, or if you wish to download a copy of the original Visio documents, should be directed to Nick Evans at nevans@stevens-tech.edu.