Firewall Strategies for Protecting Academic Resources

Brett J. L. Landry

Department of Management
University of New Orleans
2000 Lakeshore Drive
New Orleans, LA 70148
Tel: 504-280-3226
E-mail: Brett.Landry@uno.edu

Sathi Mahesh

Department of Management
University of New Orleans
2000 Lakeshore Drive
New Orleans, LA 70148
Tel: 504-280-6964
E-mail: smahesh@uno.edu

M. Scott Koger

University Computing and Communications
University of New Orleans
2000 Lakeshore Drive
New Orleans, LA 70148
Tel: 504-280-1058
E-mail: Scott.Koger@uno.edu

ABSTRACT

There are many myths surrounding firewalls at universities. Many IT staff and faculty feel that it will block what content they view or will bring the entire campus network to its knees. While poor implementations may do this, a properly installed firewall can provide secured access, better uptime for server resources, and faster overall access. The goal of this paper is to dispel some common myths regarding campus firewalls and to provide a road map for campuses to protect their academic resources.

INTRODUCTION

With the recent Blaster and Sobig worms that devastated campus networks in the Fall 2003, universities need to protect their campus resources. Corporate IT resources are widely protected by firewalls with once recent study showing nearly 100 percent using firewalls to protect their systems (McAdams, 2004). One way to protect campus IT resources is to place firewalls between the Internet and campus systems and between campus systems and academic computing resources While the need for firewalling campus systems from the Internet is easily understood, the need to firewall academic computing resources and the internet from campus systems is less well understood. Campus systems are powerful workstations typically managed by fiercely independent researchers who are skilled in their application domain and its software but are weak in systems security and often unwittingly allow their machines to be hijacked by attackers. A broad overview of campus systems security and a ranking of the applicability of a wide range of security technologies is presented in a report by Educause (Misra, et.al., 2004). This report finds that in 2003, use of firewalls declined steeply with the level of fundamental research performed. While 83% of baccalaureate institutions reported the use of some type of firewall, only 40% of universities with extensive doctoral research used firewalls. Firewalling academic computing resources accomplishes three functions: it restricts attacks from outside and inside of campus, it can restrict the types of traffic allowed to access the academic resources, and it can disallow infected machines using precious bandwidth to attack academic resources.

On February 7 and 8th 2000, university computers were used to attack major web sites including Yahoo!, Amazon.com, eBay, CNN.com, and ZDNet. The attack is called a Denial of Service (DoS) and happens when a web site is overwhelmed my traffic and is effectively taken down. While the site is not actually broken into, the net effect is a disruption of services. It is important to note that while, the DoS traffic originated from university PCs, the attackers were actually elsewhere. The university PCs had been hacked into and then hijacked as ‘Zombies’ to further launch a coordinated attack (Dube, 2000), (Brown, 2000).

In a 2002, a panel of academics, practitioners and government officials met at Stanford to unveil the National Strategy to Secure Cyberspace. The plan called for many of the things we all know but do not do. These include the enforcement of strong passwords, the use of firewalls, updated virus software, and systematic evaluation of software vulnerabilities. Richard Clarke, chairman of the President’s Critical Infrastructure Protection Board used the analysis of a crime infested neighborhood to describe the problems with cyber security. According to Clarke “It is as if we are in a town where houses are continually being burglarized. The police investigate all the crimes and find that there are no locks on the doors. Should we just go out and get more cops, or should we put locks on the doors?” (Dominik, 2002). Universities are still subject to increasingly sophisticated attacks and need to be made more secure (Brown, 2004).

Why are Universities Vulnerable?

According to Randy Katz, a professor at Berkley, universities do not setup firewalls because outside access is needed to university networks (Lasher, 2000). According to Marilyn McMillian, NYU’s Chief Information Technology Officer, NYU had firewalled their network to protect from outside attacks and that it was indeed possible with current (2000) technology to segment and protect resources (Lasher, 2000). The choice is whether to be reactive or proactive. Georgia Tech took a proactive step in 2003, when they deployed Intrusion Detection Systems to dynamically self-regulate and self-heal the network. However at the time of the article, the university had not deployed firewalls and saw this as a future direction.

The campus network security problem is not limited to the United States. In fact, any campus environment with powerful computer systems and good network access becomes a victim of malicious attacks. Tsinghua University in China experienced increased external attacks from

Internet worms and viruses, and far more and unanticipated internal attacks as well after rolling out a Gigabit Ethernet network (Force, 2004).

One of the problems is that university networks are not routinely audited and when they are they are done so, it is by auditors who mainly focus on basic items such as password resets. A 2004 audit at the Kingsville campus of Texas A&M revealed that the last time the university security policies had been updated was in 1996. Unfortunately, further audits revealed that universities in Florida, Texas, and New York had similar problems. Universities that are audited and discover that there are vulnerabilities are frequently reluctant to share that information because they are afraid it will announce their vulnerabilities to hackers (Foster, 2004). The challenge of securing a university network is further complicated by the fact that universities often deploy the latest technology, often install software undergoing testing, and have rapidly evolving networks that change considerably. Unlike the business environment in which new systems are installed after they are proven reliable and have been reasonable well tested, university researchers install and even build radically innovative and untested systems.

Many educators are fearful that technologies such as firewalls will make university networks more restrictive and infringe on academic freedoms. Much of this has to do with the manner in which firewalls were initially deployed. Originally, firewalls were deployed to prevent access to and from certain locations. Therefore, a researcher might not be able to communicate with colleagues in China because administrators had restricted network access from Asia due to the number of viruses that are created there (Carenevale, 2003).

According to Rodney Petersen, a project coordinator for the joint Educause and Internet2 security project, tighter security on Internet2 networks will give academics' academic freedom because it will ensure more reliable networks (Carenevale, 2003). Certain parts of the Internet such as e-mail have gone from being a research project to a necessary utility for business processes on campus. An unstable system for mission critical tasks is a greater problem than perceived restraints on unfettered access to the net. Another problem is the lack of centralized control over IT resources and administration in universities. According to Stanley J. Yuraitis, director of computer information systems Kingsville campus of Texas A&M it is difficult to dictate what the College of Engineering does in terms of computer standard because the college operates it's own IT support functions (Foster, 2004). For one university that wished to remain anonymous, a compromised server exposed credit card numbers and as a result they were no longer able to process Visa credit card payments (Foster, 2004). Due to the mobile nature of faculty and students at universities is very unlikely that all campus vulnerabilities can be eliminated.

Since centralized control is difficult in the university environment, the firewalls in the university need to be administered by the system managers of individual units. Administrators of systems on campus need to be trained in managing their firewalls. A good human computer interface is necessary to get these administrators to involve themselves in setting and managing firewalls. Unfortunately most state of the art security system development has not focused on a friendly human computer interface. Criteria for a successful human-computer interface are proposed by Johnston (Johnston, Eloff, Labuschagne, 2003).

New technologies in firewalls have moved beyond static rules for stateful inspection to reactive firewalls that adapt their rules to adverse situations (Hunt, Verwoerd, 2003). Intelligent new technologies under development, it may become even easier for unit administrators to manage their firewalls effectively.

MYTHS AND MISCONCEPTIONS ABOUT CAMPUS FIREWALLS

Packet Filtering and Access Control List

One common misconception about firewalls is that they are the same thing as packet filtering. Packet filtering involves restricting all traffic based on a given port. Packet filtering can also places increased load on switches and routers campus networks. Access Control Lists (ACL) have been called 'the poor man's firewall' and do not offer the granularity for most users and often slow down network equipment. ACLs also put additional processing load on network edge switches. Most environments with ACL's and packet filter rules are not granular enough support all of a network's many needs. The other problem is that the filters based on a port address have no way of inspecting the ports to see if the traffic that is being passed contains attacks.

Firewalls Block Port 80 and Kill All Web Traffic

There's a common misconception that firewalling campus resources will block all of a particular web traffic type (Bryan, 2002). This misconception stems from the fact that many individuals confuse packet filtering with firewall techniques. Basically, this is a falsehood that training and education can overcome. Treat all campus and Internet connections as potential threats. The idea here is simple. If you do not have control over the network and the machines on that network, security can be compromised at any time. Thomas Ptacek (2003) has coined the term 'virtual perimeters' to describe this process of segmenting your network around departments. We know how to build perimeters between the Internet and the internal network. It's time to figure out how to build perimeters between the different business user groups on the network. Basically this involves using the same techniques for segmenting from the Internet and applying to inside the network. Firewalls can also limit the number of sessions from a given host to prevent a compromised PC or disgruntled user from launching DoS attacks. The firewall protects the network, without disabling the user long term.

UNO'S FIREWALL EXPERIENCE

The Department of Management at the University of New Orleans installed six servers as part of a statewide IT initiative for teaching. These servers were being used for hands on experience in teaching IT topics. Since the servers were initially not protected by a campus perimeter firewall, they were frequently attacked from the Internet. Additionally, the network infrastructure at the time did not provide a means of protection from local attacks either. These attacks made the servers frequently unavailable and imposed a burden on the system administration staff. In fact the servers were rebuilt or restored after being compromised on more than one occasion.

As a result, In January of 2003, IT professors began the process of installing a stateful inspection firewall to protect the application servers from on and off campus threats. Since a significant number of compromised machines were on campus, the firewall had to protect the servers from attacks originating from inside the campus systems and from across the world. There was a widespread belief on campus that this approach would not work on campus. Some of the reasons cited focused on the need for the servers to be accessible for students and faculty from home and the wide range of software supported on the servers. However, after defining what access was truly necessary for the systems to be successfully deployed for home and lab access, a set of appropriate firewall rules were created. Following the enforcement of the firewall rules, the servers ran faster, and with zero attacks. Rather than slowing down access, the machines delivered access to students and faculty more speedily since traffic was limited to only valid requests. Students and faculty were able to access the servers from on and off campus, and the servers were able download Microsoft patches and virus pattern files. During the peak of the Blaster and Sobig worm attacks in the Fall 2003, the servers that were behind the firewall were untouched and unaffected because the firewall prevented all unwanted access.

SUMMARY

Since university campuses have some of the most powerful machines linked to the open Internet, it is imperative that campus network administrators develop a security model that will offer the necessary protection from outside attack, and from internal machines being used for illegal activity. Firewalling Academic Resources is not a complete solution, but it is a first step. There may be many reasons why campuses cannot firewall their entire campus. Realizing this to be an issue, departments can firewall their own servers to protect them from attacks. The most important reason for doing this is that academic resources are mission critical devices. Classes that depend on these resources so they must be dependable and reliable.

Another aspect of firewalling academic resource segments on campus is that individual unit administrators control access to their resources. This matches the distributed computing environment on university campuses. The implementation of a segment firewall forces unit administrators to determine what systems are running on their own machines, to monitor traffic to and from their segments and become more aware of security needs for network access. We recommend that campus network administrators select a suitable academic resource segment firewall product and offer training to interested unit administrators. This approach will reduce attacks and lead to improved efficiency on the network. In addition, servers controlled by these units will not be easily hijacked by attackers. Rules for firewalls are set by individual units and the campus wide firewall can enforce commonly accepted firewall standards while allowing unit administrators to enforce a second, higher level of security. This multi-tiered security matches the administrative model of university campuses and will provide enhanced security.

REFERENCES

Brown, A. (2000). Stanford Computers Used in Web Attacks, *The Stanford Daily,* February 14, 2000.

Bryan, J. (2002). Internal memo, Lack of UCC Responsiveness to COBA, Nov. 20, 2002

Carnevale, D. (2003). Awareness of computer-security threats is still inadequate, report warns. *Chronicle of Higher Education*, 11/14/2003, v. 50(12), pp. 30-33

Dominik, M. (2002). Officials unveil cybersecurity plan. *The Stanford Daily*, Sep. 19, 2002

Dube, J. (2000). Zeroing In. *ABC NEWS* http://more.abcnews.go.com/sections/tech/dailynews/webattacks000214.html

Force (2004). Tsinghua University Expands Campus Network with Force10 E-Series Switch/Routers, *http://www.force10networks.com/applications/profiles-tsinghua.asp*, Dec 1, 2004

Foster, A. L. (2004). Insecure and Unaware. *The Chronicle of Higher Education.* 5/7/2004, v. 50(35), pp. 33-36

Hunt, R, Verwoerd, T. (2003). Reactive Firewalls: A New Technique, *Computer Communications*, v. 26(12), pp. 1302-1318

Johnston, J., Eloff, J.H.P., Labuschagne, L (2003). Security and the human computer interface, *Computers and Security,* v. 22(8), pp. 675-685

Lasher, M. (2000). Universities Learn from Web Attacks.

McAdams, A. (2004) Security and Risk Management: A Fundamental Business Issue, *Information Management Journal,* Jul/Aug 2004, Vol. 38(4), pp. 36-43

Misra, C., Schulman, M., St. Sauver, J., and Suess, J. (2004) Effective Practices and Solutions in Security, *EDUCAUSE*, http://www.educause.edu/1246

Ptacek, T. (2003). 10 Tips for Improving Security Inside the Firewall. *Computerworld*. Nov. 13, 2003

Roberts, P. (2004). Attacks at Universities Raise Security Concerns, *Network World Fusion*, 4/14/2004