



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Achieving Defense-in-Depth with Internal Firewalls

Introduction

A single firewall at the Internet gateway is no longer sufficient. Currently there is a trend toward more and more outside access to the enterprise network by employees, partners, customers, and suppliers. In addition, attackers are becoming more sophisticated. A sound security perimeter today requires more than a single firewall connected at the Internet router. By segmenting the network with multiple firewalls, we can achieve the holy grail of network security – Defense-In-Depth.

Defense-in-Depth

Imagine your average bank. Think of the security features that you take for granted: a vault limits access to the cash; cameras record everything that happens; an alarm system can summon police quickly; and dye bombs disguised as money help identify the thief. You feel safe handing them your paycheck, not because of any one particular precaution, but because together they offer excellent safety for your hard-earned cash.

Now imagine another bank with only a single line of defense: a massive vault to protect your money. No matter how good that vault might be, would you feel as safe leaving your money there? What if an employee decides to flee the country with your retirement fund? What if they forget to lock it one night? What if someone knows how to pick the lock? Obviously, a vault alone is not enough, no matter how strong it might be. What this bank lacks is Defense-in-Depth. Rather than relying on any single security measure, a strategy of Defense-in-Depth assumes that any individual security precaution might fail, and has another line of defense ready [1].

Now think of a typical computer network. Which bank does it resemble? Many organizations still rely on the old concept of a single firewall at the Internet gateway to protect the entire enterprise network from attack. This may have been sufficient in the past, but this old technology can no longer deal with the realities of today's e-business environment [6].

A Moving Target

Many changes have taken place since the days when firewalls were first deployed to protect the internal network from attackers on the Internet. Network security will always be a difficult target to hit, due to the rapid changes needed in today's business world and

the swift pace of technological innovation by hackers. Firewalls have increased in sophistication to some degree, but the skill of the attackers has grown much faster. In addition, many of the assumptions that led to the deployment of these systems are no longer valid in the enterprise networks of today [2].

For example, most organizations once assumed that they could trust their employees, and firewalls were usually deployed solely to defend against external threats. However, in a report published by WarRoom Research, 61% of those responding to the survey reported an internal attack within the past 12 months [3]. Obviously, a single firewall at the Internet gateway can do nothing to prevent a temporary employee in accounting from attempting to gain inappropriate access to a server containing critical financial records.

Another assumption these systems relied upon was that only employees would be connected behind the firewall. Today's businesses are opening up their networks to partners, suppliers, and customers. Can we afford to trust all these new users of the network implicitly? As we make more and more connections into the internal network, the old model of a single network entry point, watched over by a single firewall, no longer makes much sense.

The rapid increase in network bandwidth has placed a strain on the ability of the firewall to inspect and log all connections. Whereas the Internet gateway was once assumed to be a point where all traffic would be passed through strict filters, and recorded in great detail, many firewall administrators are finding that the increased bandwidth requirements do not allow for this level of control. In many organizations, proxy servers have given way to state-aware packet filters, and the sheer volume of traffic has prompted some administrators to reduce the level of logging. Malicious activity may go unnoticed in these situations.

Finally, mobile users are now turning the most basic assumption of the traditional firewall system upside-down. Instead of being confined to a particular physical location, users can now be located anywhere. The rapid growth of Virtual Private Networks (VPN's) to support mobile users and telecommuters has extended the internal network to thousands of endpoints outside the organization's physical location. Encryption technology protects this traffic while it travels over the Internet, but this poses certain problems as well. The firewall cannot inspect the traffic until it has been decrypted, and therefore any encrypted traffic must be allowed inside the perimeter. If the VPN fails for any reason (compromise of the remote machine, failure of the VPN device, etc.) the firewall will have no way to protect critical assets from attack, since this VPN traffic is usually trusted as if it were from an internal user.

In most networks today, if an attacker does manage to break through the outer security perimeter, the game is up for the network administrator. The "soft and squishy" interior of the average corporate network provides very little challenge for any knowledgeable attacker. More often than not, simply using an IP address from the interior network is all that is required to gain full access to some services, since many network administrators

avoid using passwords on the internal network. Even when passwords are required, sniffing traffic inside the network for even a brief period of time will usually reveal several username/password pairs. Once an attacker can masquerade as an employee, it becomes very difficult to distinguish their traffic from ordinary traffic. How many administrators closely monitor what every employee is doing on the network all day?

Strength In Numbers

When we consider all these issues, it is apparent that we cannot rely upon a single firewall to protect the enterprise network. How can we give the users of our network the same level of confidence they feel at the local bank? A good Defense-in-Depth strategy involves many different technologies, such as Intrusion Detection, Content Filtering, and Transport Layer Security. The single most important element, however, is a system of internal firewalls. Proper deployment of these devices can address all of the concerns raised above:

- Employees will not have unrestricted access to the entire network, and their activity can be monitored.
- Partners, customers, and suppliers can be given limited access to whatever resources they require, while maintaining isolation of critical servers.
- Critical servers can be closely monitored when they are isolated behind an internal firewall. Any malicious activity would be much easier to detect, since the firewall has a limited amount of traffic passing through it.
- Remote users can be restricted to certain portions of the network, and VPN traffic can be contained and easily monitored.
- A security breach in one segment of the network will be limited to local machines, instead of compromising the security of the entire network.

With a system of internal firewalls in place, we can come much closer to our ideal network. Instead of an all-or-nothing security posture, we can achieve Defense-in-Depth by forcing an attacker to penetrate multiple layers of security to reach mission-critical servers.

The Next Step

If we extend the idea of internal firewalls to its logical conclusion, we can envision a system that places a firewall on every network device. Such systems, known as “Distributed Firewalls,” are still in the early stages of development, but they are beginning

to attract the attention of network managers who have seen their perimeter firewalls fail to stop an attack. One example of the need for such systems is the recent Microsoft attack, where an attacker gained access to an employee's home machine, then exploited the VPN connection into the corporate network [4]. Since the access appeared to be originating from an employee, the perimeter firewall was not effective in stopping the attack. A system of distributed firewalls could have stopped the user's home machine from being compromised, or at the least it would have limited the damage done. In a paper written for AT&T labs, Steven Bellovin describes a system where "...policy is still centrally defined; enforcement, however, takes place on each endpoint. We thus retain the advantages of firewalls while avoiding most of the problems..." [5]. In effect, a central policy server knows what connections should be allowed to any given machine, and it pushes this configuration out to a host-resident firewall on that device. An attack cannot really penetrate the security perimeter, since the security perimeter is everywhere. Obviously, the rule set on the central policy server could be quite complex, since it would have to define exactly what traffic is allowed to each device on the network. In most cases, however, no inbound access would be required to workstations, and outbound access could be limited to a few protocols, such as HTTP, IMAP, etc. This would simplify the creation of such a policy, and it is possible that creating it would not be much more difficult than configuring a traditional firewall. With this type of system deployed, an attacker's job would become much more difficult. Today, the attacker need only focus on finding a gap in the security perimeter of an organization. Once inside the network, there is very little to stop the attacker from doing whatever he or she pleases. But with a distributed firewall in place, the attacker does not hit the jackpot if a single machine is compromised. Other servers not associated with the compromised machine are just as secure as they always were. With distributed firewalls, we can finally achieve true Defense-in-Depth. Unfortunately, this technology is not yet ready for use. It may be some time before stable and easy-to-use distributed firewall software is available. In the meanwhile, we can deploy internal firewalls as a step towards this goal. In most cases, much of the benefit of a distributed firewall system can be achieved with internal firewalls. On paper, the concept of segmenting the enterprise network with internal firewalls seems straightforward; in the real world, however, there are often a number of challenges to be overcome.

Reality Check

Most enterprise networks in use today were built on the old assumptions outlined earlier. Many networks are not segmented into autonomous networks, and do not have well-defined internal boundaries between departments. Obviously, there would be some challenges to be overcome before internal firewalls could be deployed in the typical corporate environment. The biggest issue is the widespread use of "Mesh" networks. Many organizations treat the entire enterprise WAN as a single security zone, and as a result any machine within the WAN has some form of access to all other machines. There might be some type of password protection on various servers – such as accounting systems – but in general any machine can connect to any other machine. This is

sometimes called a Mesh network, due to the way that the network connections would appear on a diagram. An example might be a large enterprise using the Windows “Single Domain” model. This type of network is incompatible with both Defense-in-Depth principles and with internal firewalls. In order to successfully implement internal firewalls, the security administrator must be able to limit connections to servers and workstations so that only those who need access to a particular device have the ability to connect. The typical Windows Single Domain model will be incompatible with this goal due to the requirement for NetBIOS traffic to pass between machines. All machines in the domain must access certain services - such as WINS – using the NetBIOS protocol. Since this protocol offers no authentication or integrity capabilities, allowing this protocol to pass through the internal firewalls would defeat the purpose of installing such firewalls in the first place. Blocking this traffic, however, means that a Windows Domain cannot traverse a firewall. As a result, we must redesign the network. Each department must be their own domain, or we must deal with a host of problems that will result from blocking NetBIOS traffic, such as Network Neighborhood problems, authentication problems, and communication between domain controllers. Migrating to a pure Windows 2000 environment might make the task much easier, since it should be possible to use DNS exclusively for name resolution, and the more flexible domain controller structure and Organizational Unit architecture should make it possible to design a network that is compatible with internal firewalls. Unfortunately, it will be some time before most network administrators can upgrade or replace all their legacy clients.

In addition to the problem of redesigning mesh networks, a few other issues may also need to be addressed. Fortunately the remaining problems are small ones by comparison. Applications must be well behaved – i.e., they should not use random ports or make many separate connections to a client. Additionally, administrators must define exactly what type of traffic is allowed to each machine, and create a firewall policy that allows the required traffic to flow. Knowing the exact nature of the traffic used by a particular application can be difficult if the application is not well documented, or if the system is at a remote site.

With all these issues to be overcome, it is apparent that it may take some time before distributed firewall systems are ready for wide scale deployment. Even internal firewalls are often problematic, since they may require additional subnets to be created, or for machines to be moved from one network to another. The biggest obstacle is the inertia of users who are accustomed to the old model of complete access. Security always has a price, and from the perspective of the end users, the price is the ease-of-use that was afforded by the lack of strong internal access controls. However, with proper education of the end-users, good design, and careful planning, segmented networks can be successfully deployed in most environments.

Conclusion

Distributed firewalls offer a promising solution to the limitations of a traditional single

firewall security perimeter, but it is unlikely that conventional firewalls will be replaced anytime soon. In the meanwhile, we can realize many of the benefits of such a system by adding internal firewalls to our enterprise networks. By segmenting the network into several pieces, and controlling the access between these segments, we can achieve a good measure of Defense-in-Depth, and make the attacker's job much more difficult.

References

- [1] Dr. David S. Alberts, Defensive Information Warfare, August 1996. URL: <http://www.ndu.edu/ndu/inss/books/diw/ch15.html>
- [2] Wei Li, Distributed Firewall, December 5th, 2000. URL: <http://www.cs.helsinki.fi/u/asokan/distsec/documents/li.ps.gz>
- [3] Dov Herdan, Only Adaptive Security Can Stop The Hackers, October 19th, 1998. URL: <http://www.warroomresearch.com/MediaPresenSpeak/TechWeb.asp>
- [4] Deborah Radcliff, Feature: Firewalls Reach Out, March 26th, 2001. URL: <http://www.nwfusion.com/net.worker/news/2001/0326firewalls.html>
- [5] Steven M. Bellovin, Distributed Firewalls, November 1999. URL: <http://www.research.att.com/~smb/papers/distfw.html>
- [6] Tony Harrington, Shoring Up Firewalls For the 21st Century, March 16th, 2000. URL: <http://www.vnunet.com/Features/600710>