Global Information Assurance Certification Paper

Distributed Firewall

Daniel Wan

(GSEC Practical Assignment Version 1.2c)

Introduction

Conventional firewalls serve as the sentry between trusted and untrusted networks. They rely on restricted topology and controlled network entry points to enforce traffic filtering. The key assumption of this model is that everyone on one side of the entry point – the firewall – is to be trusted, hence they are protected, and that anyone on the other side is, at least potentially, an enemy.

The expanded Internet connectivity makes conventional firewalls obsolete. Furthermore, end-to-end encryption, and other protocols are also threats to this type of firewall. To address these shortcomings, the concept of "a distributed firewall" has been proposed.

In this paper we will discuss the drawbacks of traditional firewalls, along with the needs, concepts and some implementations of distributed firewalls.

Conventional Firewalls and Their Problems

A *firewall* is a collection of components, interposed between two networks, that filters traffic between them according to some security policy [1]. Conventional firewalls depend on the topology restriction of the networks. The controlled entry point – the firewall – divides the networks into two parts, internal and external networks. Since the firewall cannot filter the traffic it does not see, it assumes that all the hosts on the internal networks are trusted and all the hosts on the other side (external) are untrusted.

This model works quite well when networks comply with the restricted topology. But with the expansion of network connectivity, such as extranets, high speed lines, multiple entry points, and telecommuting, this model faces great challenges:

- Firewalls do not protect networks from internal attacks. Since everyone on the internal networks is trusted and the traffic within these trusted networks is not seen by the firewall, a conventional firewall cannot filter internal traffics, hence it cannot protect systems from internal threats. For traditional firewalls, the only way to work around this is to deploy multiple firewalls within the internal networks, i.e. divide the network into many smaller networks, and protect them from each other. Since different policies have to be applied on these firewalls, both the load and complexity of administration increases.

- The vastly expanded Internet connectivity makes this model obsolete. The extranets and the telecommuters from outside are allowed to reach all or part of internal networks. Meanwhile, the telecommuters' computers that use the Internet for connectivity need protection when encrypted tunnels are not in place, especially as cable modems and DSL become more available and affordable. Currently, most such

telecommuters connect to organizations' internal networks through VPN tunnels. If they also use the same VPN tunnel for generic Internet browsing purpose, it is not only inefficient (causes the "triangle routing"), but may also violate the organization's guidelines. If they don't use the VPN channel, they either open a security hole or add the workload of maintaining numerous personal firewalls that are at different location.

- End-to-end encryption is another threat to the firewall. There are so many external Web proxies, such as www.anonymizer.com, out there on the Internet. Users can easily setup an end-to-end encryption tunnel between their desktop within the organization's internal network to a machine outside. Since the firewall does not have the key to look into the encrypted package, it cannot filter properly according to the security policy. By doing so, insider users can bypass the destination restriction and hide the traffic. If this channel is controlled by malicious hackers, it is almost impossible to detect, because all the packages are encrypted!

- Some protocols are not easily handled by a firewall. Because the firewall lacks certain knowledge, protocols like FTP and RealAudio need application level proxies to manage through the firewall.

- A firewall is the single entry point. This is the place traditional firewalls enforce their policy and filter the traffic. It is also a single point of failure. If the firewall goes down for any reason, the entire internal networks are isolated from outside world. Although high availability option, such as hot standby firewall configurations exist, they are usually cost prohibitive.

- Firewalls tend to become network bottlenecks. Due to the increasing speed of networks, amount of data passing through, and the complexity of protocols firewalls must support (such as IPSec) they are more likely to be the congestion points of networks [3].

- Unauthorized entry points bypass the firewall security. It has become trivial for anyone to establish a new, unauthorized entry point to the network without the administrator's knowledge or consent. Various forms of tunnels, wireless, and dial-up access methods allow individuals to establish backdoor access that bypasses all the security mechanisms provided by traditional firewalls. While firewalls are in general not intended to guard against misbehavior by insiders, there is a tension between internal needs for more connectivity and the difficulty of satisfying such needs with centralized firewalls [3].

In order to solve these problems while still retaining the advantages of the conventional firewalls, Steven Bellovin, an AT&T researcher, proposed a "distributed firewall" [2].

The Distributed Firewall

A multitude of host-resident firewalls when centrally configured and managed makes up a distributed firewall [4]. In this architecture, the security policy is still defined centrally, but the enforcement of the policy takes place at each endpoint (hosts, routers, etc). The centralized policy defines what connectivity is permitted or denied. Then this policy is distributed to all endpoints, where it is enforced.

Three components are needed for distributed firewalls: (a) a *security policy language*; (b) a *policy distribution scheme*; and (c) an *authentication and encryption mechanism*, such as IPSec.

The *security policy language* describes what connections are permitted or prohibited. It should support credentials and different types of applications. After policy is compiled, it is shipped to endpoints. The *policy distribution scheme* should guarantee the integrity of the policy during transfer. This policy is consulted before processing the incoming or outgoing messages. The distribution of the policy can be different and varies with the implementation. It can be either directly pushed to end systems, or pulled when necessary, or it may even be provided to the users in the form of credentials that they use when they try to communicate with the hosts.

How the inside hosts are identified is very important. Conventional firewalls rely on topology. The hosts are identified by their IP addresses and network interfaces on the firewalls they are attached to, such as "inside", "outside", and "DMZ". This kind of structure is quite weak. Anyone with physical access to the internal network and get an internal IP address will be fully trusted, plus IP address-spoofing is not difficult at all. Since all the hosts on the inside are trusted equally, if any of these machines are subverted, they can be used to launch attacks to other hosts, especially to trusted hosts for protocols like *rlogin*.

It is possible that distributed firewalls use IP addresses for host identification. But a secure mechanism is more desirable. It is preferred to use certificate to identify hosts. IPSec provides cryptographic certificates. These certificates can be very reliable and unique identifiers. Unlike IP address, which can be easily spoofed, the digital certificate is much more secure and the ownership of a certificate is not easily forged. Furthermore, they are also independent of topology. Policy is distributed according to these certificates. If a machine is granted certain privileges based on its certificate, those privileges can be applied regardless of where the machine is physically located.

In this case, all machines have the some rules. They will apply the rules to the traffic. Since they have better knowledge of the connection (such as the state and the encryption keys, etc), they will make better judgment according to the policy. With a distributed firewall, the spoofing is not possible either, because each host's identity is cryptographically assured.

Advantages and Benefits

- **Topology Independence**

The most important advantage for distributed firewalls is that they can protect hosts that are not within a topology boundary. The telecommuters who use the Internet both generically and to tunnel in to a corporate network are better protected now. Before they either have to use the "triangle routing" to tunnel into organization's network for generic Internet traffic or are not protected when they are not tunneled, which is a security breach for the computer and the organization. With distributed firewalls, the machines are protected all the time, regardless of whether the tunnel is set up or not. No more triangle routing is needed.

- **Protection from Internal Attacks**
  After distributed firewalls abandon the topology restriction, hosts are no longer vulnerable to internal attacks. To the host, there is no more difference between "internal" and "external" networks. After a machine boots up, the policy is enforced on it for any inbound and outbound traffic. Also the hosts can be identified by the their encrypted certificates, this eliminates the chance of identity spoofing.

- **Elimination of the Single Point of Failure**
  A traditional firewall needs a single entry point to enforce policy. It not only creates the single point of failure but also limits the entire network's performance to the speed of the firewall. Multiple firewalls are introduced to work in parallel to overcome these problems; in many cases though, that redundancy is purchased only at the expense of an elaborate (and possibly insecure) firewall-to-firewall protocol [2]. With the deployment of distributed firewalls, these problems are totally eliminated. The performance, reliability, and availability no longer depend on one, or in some cases a group of, machine(s).

- **Hosts Make Better Decisions**
  More often that not, conventional firewalls don't have enough knowledge in terms of what a host intends. One example is end-to-end encrypted traffic, which can easily bypass the rules on conventional firewalls, since the firewalls don't have the necessary key. Also, many firewalls are configured that they will pass the TCP packets from outside world with the "ACK" bit set. Because they think these packets are the replies to the internal hosts who initialize the conversation. Sadly enough, it is not always true and the spoofed ACK packets can be used as part of "stealth scanning". Similarly, traditional firewalls cannot handle UDP packets properly, because they cannot tell if these packets are replies to outbound queries (and hence legal) or they are incoming attacks. In contrast the host that initializes the conversation or sends out the queries knows exactly what packets it is expecting, what is not, since it has enough knowledge to determine whether an incoming TCP or UDP packets are legitimate and it has the necessary key in the case of end-to-end encryption. Same thing is true for the protocols like FTP.

Limitations

Like every other technology, distributed firewalls have their limitations too.

One issue with distributed firewalls is that they cannot protect legacy applications. This is because the old protocols do not understand strong cryptography used in a distributed firewall. This problem can be solved by combining the distributed firewall with conventional firewalls during implementation.

Like conventional firewalls, distributed firewalls are vulnerable to some attacks too, such as the "smurf" [8] -- one of many DoS attacks. As a matter of fact, neither form of firewall offers an effective defense.

Intrusion detection is harder to achieve on distributed firewalls. Modern firewalls can detect the attempted intrusions. In a distributed firewall, it is not a problem for a host to detect the intrusions, but the collection of data is more problematic, especially at times of poor connectivity to the central site, or when the either site (host or central site) is under attacks such as DoS.

Implementation

Proposed by Steven M. Bellovin in November 1999, distributed firewalls are still in their infancy. But we see that several institutes, and software/hardware manufacturers are working towards to making distributed firewalls a reality.

A project supported by DARPA was fulfilled at the University of Pennsylvania in 2000 [3]. In this project, a prototype of distributed firewall was constructed. OpenBSD was chosen as the operating system, because it was an attractive platform for developing security applications with well-integrated security features and libraries. Keynote was chosen as the security policy language. It was also used to send credentials over an untrusted network. IPSec was used for traffic protection and user/host authentication. This was a concept-prove implementation. Other improvements, such as moving the policy daemon to the kernel and adding IP filters, etc., needed to be done on this prototype firewall.

Currently, Network-1 Security Solutions Inc. is offering a commercial host-resident firewall, CyberwallPLUS, on the windows platform. This host-resident firewall includes personal firewalls for remote users, firewall agents for workstations, and application-server resident firewalls. It's very similar to distributed firewall. Actually, when multiple host-resident firewalls are centrally configured and managed, it indeed is a distributed firewall [4]. But CyberwallPLUS still has its challenges currently -- it cannot collect reports centrally [6].

When distributed firewalls inherit the workload of conventional perimeter firewalls, each host picks up some extra work namely policy enforcement and data encryption. This extra work does not necessarily slow down the hosts. A lot of them can be off-loaded from CPU to NIC. 3Com announced a new firmware add-on for its network interface cards earlier this year. The NICs can boost the speed of nodes that send encrypted data by off-loading encryption algorithms and other packet processing tasks from a PC's

processor to the NIC. Some of these cards even have firewall functions embedded in firmware [7].

Driven by the tremendous demands and fueled by the works from different resources, it is certain that we will see some mature distributed firewall products on the market soon.

Conclusion

With the increasing of line speed, connectivity, and complexity of protocols, conventional firewalls can no longer handle their purpose adequately. A new concept of firewall, distributed firewall, was introduced. Distributed firewall retains the advantages of conventional firewalls, while solving many of their problems. We can expect the to see more and more distributed firewall products in the future.

References:

[1] W. R. Cheswick and S. M. Bellovin. *Firewalls and Internet Security: Repelling the Wily Hacker.* Addison-Wesley, 1994

[2] Steven M. Bellovin, *Distributed Firewalls,* http://www.research.att.com/~smb/papers/distfw.html

[3] Sotiris Ioannidis, Angelos D. Keromytis, Steve M. Bellovin, and Jonathan M. Smith, *Implementing a Distributed Firewall* http://www.cis.upenn.edu/~angelos/Papers/df.pdf

[4] Avi Fogel, Pushing Security to Network Endpoints, http://www.nwfusion.com/news/tech/2001/0122tech.html

[5] http://www.network-1.com

[6] Avi Fogel, http://www.nwfusion.com/archive/2000/99612_06-26-2000.html

[7] Phil Hochmuth, *3Com improves NIC security features* http://www.nwfusion.com/news/2001/0409infra.html

[8] CERT® Advisory CA-1998-01 Smurf IP Denial-of-Service Attacks http://www.cert.org/advisories/CA-1998-01.html