

Protecting the Distributed Enterprise

A White Paper prepared by SonicWALL, Inc.

SonicWALL, Inc.
1160 Bordeaux Drive
Sunnyvale, CA 94089-1209
1-888-557-6642
<http://www.sonicwall.com>

Overview

Since the rise of the Internet, much attention has been devoted to the Internet-related computer-security challenges facing large enterprises. The media, security vendor, consultants, and organizations like the Computer Security Institute all ensure that all IT managers and corporate executives understand the risks inherent in protecting against attacks.

The most common prescription against unwanted Internet access has been fortification of the enterprise network's main entrance against hackers— the publicly known servers and services providing Web, telnet, and ftp access for customers. High-end solutions, such as Check Point Software's FireWall-1, are now firmly and properly established at the main entrances to the enterprise network. But that is not enough. Although the front door may be fortified and monitored, other entrances that may not be as well protected against stealthy attacks from without or within. Remote offices may not be protected at all, placing their own data and application availability at risk, and perhaps also providing an unguarded "back door" into the fortified headquarters network.

The technology used to protect alternative portals into an enterprise network and remote networks from external attack, and to isolate internal segments of a large network from internal threats, are the same as those which protect the main entrance: firewalls at portals, and virtual private networks between sites.

Denial of Service Attacks

Increasingly, the types of attacks facing enterprise networks are geared at denying access to essential services by legitimate users — often by crashing servers or routers, or by overwhelming the network with enough traffic to degrade service. This is a change from the past, when many attacks were carried out to gain access to valuable corporate data or to "joyride" high-end computing resources.

Such *denial of service* attacks may be initiated either outside the firewall, or from within the enterprise. The subsequent disruption of servers and services may be also the actual goal of the attacker, or may be a prelude to a more sophisticated attack. For example, an automated system may attempt denial-of-service attacks on consecutive groups of IP addresses, and then ping them for a response. If there's no response from a downed server (and thus the denial-of-service attack was apparently successful), the site may be targeted for more sophisticated hacking at a future date.

Hacking is not limited to trying to gain entrance via the Web server or ftp host. A technique known as "port scanning" occurs when attackers attempt to gain entry using often-unguarded IP ports reserved for system management. There are more than 65,000 IP ports. Many have clearly defined purposes, such as port 21 for ftp traffic, 23 for Telnet, 25 for SMTP, 80 for HTTP (Web) traffic, and 110 for POP3 traffic. However, other ports may be used for tasks. Port 389, for example, is often used for LDAP directory access, port 8080 for a test Web site, and port 20000 is widely

used to administer a Netscape LDAP server. It is essential that enterprises be protected from attacks aimed at any in-use IP port, not merely the most popular.

See appendix A for a description of several common denial-of-service attacks – which may not always be blocked by firewall products, or may not be guarded against at all within the enterprise.

Firewalls: An Overview

A *firewall* guards a data path. It is transparent to “good” packets and messages, and speeds them to their destination, while blocking “bad” messages. Firewall functions may be implemented in dedicated hardware, as firmware or applications running inside WAN routers, or as services running on a general-purpose server. There are three primary mechanisms used for implementing firewall features: packet filtering, applications-level proxies, and stateful inspection. Stateful inspection is the simplest as well as most robust, and best able to withstand new methods of attack.

Packet filtering scans a packet to determine where it originated, based on IP address. Thus, packet-filtering firewalls are prone to being compromised using IP spoofing, which involves altering an IP packet so the firewall thinks the packet has an internal, rather than external, source address and therefore grants it network access. Some protocols, such as ftp and DNS, can’t be safely passed through packet filters because those protocols require opening “holes” in the firewall – i.e., that their IP ports not be filtered – and this compromises security.

Application-level proxy servers protect the network by examining the application layers, to see if the application generating the packet has originated from a known and trusted application. Unfortunately, this upper level examination requires a great deal of processing power and often leads to an unacceptable performance penalty. In addition, each application type, such as HTTP, ftp, SMTP or POP3, requires the installation and configuration of a different application proxy, making support for new applications a problem. In addition, this approach requires the user to reconfigure their network settings to support the proxy. Finally, application proxies track only the application state, not packet or connection state, which may introduce security vulnerabilities.

With *stateful packet inspection*, a firewall makes security decisions based on the origination of Internet sessions. A stateful-inspection firewall allows data coming from the untrusted side of the firewall (i.e., the Internet) only if it’s part of a session that was initiated by one of the users on the secure side, but will block all communications that are initiated from the Internet. Stateful packet inspection has the added benefit of being easy to manage, making it ideal for organizations who don’t have the technical resources for a packet filtering or proxy firewall.

When it is necessary or desirable to allow Internet sessions to originate from the untrusted side of the firewall, this can be accomplished by setting up *demilitarized zones* within the firewall, where packets addressed to certain IP ports

– such as 80 for Web traffic, or 23 for ftp sessions – are routed directly to a separate LAN segment, which contains the appropriate servers for those tasks.

Firewalls: Beyond the Headquarters Border

Enterprise gateway firewalls are complex and expensive products, and are actively managed by network administrators – often the same administrators who are managing LAN or WAN routers and connections. Leading enterprise firewalls include Firewall-1 from Check Point Software; the PIX Firewall, Centri Firewall, and IOS Firewall from Cisco Systems; Border Manager from Novell; and Raptor from Axent Technologies.

Interior and side-door firewalls require simplicity, transparency, and should be largely self-managing. For example, they should be able to detect and repel a wide variety of attacks, be able to handle high-bandwidth because they are often connected to 10Mbit/sec Ethernet or 100Mbit/sec Fast Ethernet network segments, rather than slower WAN links. They should be able to alert administrators when attacks are in progress, and when firmware or software upgrades are required.

It is often tempting to implement remote-office firewalls as services running on standard PC or Unix servers. Certainly, a general-purpose server and operating system *can* be used to implement an effective firewall – if it is proactively and professionally administered. However, this is not likely to be the case in remote offices or within end-user departments of a large enterprise – and the security of the firewall is only as strong as its weakest link. For example, the operating system itself must be *hardened*. In the case of Firewall-1 or Raptor running on top of Windows NT Server, hardening the operating system is a combination of installing the correct service and option packs on top of Windows NT, and keeping up to date with all Windows NT security flaws, registry settings, and Microsoft-supplied hot patches.

Remote Office Firewalls

Remote offices, particularly those beyond a few blocks or miles of a headquarters location, are often connected directly to the Internet via their own Internet Service Provider link, such as a T-1 or ISDN. This Internet connection is used for multiple purposes: making available public Web sites, allowing local employee access to the Internet, and local employee remote access to the remote office's IT resources. Increasingly, that Internet connection is also used to either supplement or replace a dedicated WAN connection to the headquarters location. The headquarters location, and very large remote offices, are typically protected by enterprise firewall products. However, due to the cost and complexity of typical firewalls, smaller offices, including those occupied by a single individual, are not so protected.

The vulnerability of those offices represents a real danger to the enterprise: the loss of password or other authentication stolen from a remote-office server may lead to a successful break-in at the headquarter location; if the remote

location is also connected via a dedicated and trusted WAN pipe to the headquarters, it may also represent a back door for hacker entrance. While successful denial-of-service attacks on remote offices would not have as severe an impact as such an attack on the headquarters, it would still cause lost productivity and revenues to the organization.

We believe that it is crucial to install firewall functionality in such locations. The main requirements should be high performance and effectiveness, transparency to the user and to the network, and lack of necessity for hands-on administration. However, there must be the capacity for proactive communication back from the firewall to the enterprise IT department should a break-in be attempted, and when software upgrades become available.

For remote offices, we recommend SonicWALL's SonicWALL PRO as the remote-office solution. Their Web-based administration (including monitoring and firmware upgrades) and proactive email communication and log transmissions, would satisfy the "keep it simple" needs of remote offices.

A second requirement should be for Virtual Private Network capability within those remote office firewalls, if they are connecting back to the headquarters office over the Internet, or if telecommuters wish to use the Internet to communicate with the remote office. A VPN provides a secure, encrypted path over the Internet, and the use of VPN should be required for accessing any non-public information over the Internet.

As VPN standards are still evolving, different vendors' implementations of IPSec are not always fully interoperable. Yet a good remote office firewall should be adaptable to support all of the leading enterprise VPN products. (See Figure 1.)

For VPN access, we also recommend the SonicWALL PRO, which supports 168 bit Data Encryption Standard (Triple-DES), 56 bit Data Encryption Standard (DES), and 56 bit ARC4 (ARC4).

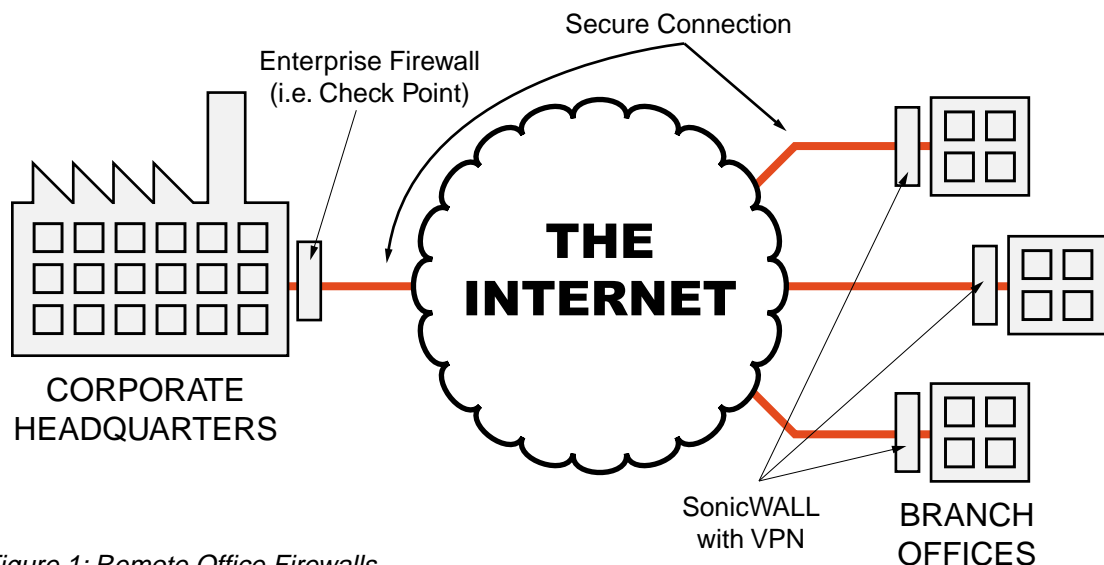


Figure 1: Remote Office Firewalls

Interior Enterprise Firewalls

A large percentage of successful hacks and attacks originate from within the enterprise network, and thus are not shielded by the enterprise border firewall. In those situations, two goals should be to minimize the exposure of enterprise IT assets to harm, and to limit the scope of a possible attack. A second goal would be to speedily identify that an attack is in progress, and if possible to determine the source of the attack.

One way of protecting the enterprise assets is to protect existing local area network segments with individual firewalls. A high-speed, transparent firewall should allow legitimate traffic to pass between network segments, using the same stateful inspection methodology, while blocking unwanted traffic, or packets which attempts to use inappropriate IP ports.

If the enterprise network is based on multi-protocol routers with built-in firewall capability, or with accessory firewalls blades, that may be a natural location to store the interior firewall. However, IP-only networks are increasingly switched, not routed, to improve performance and to reduce cost and complexity.

In such situations, an affordable solution may be to protect each segment connecting to an enterprise backbone switch with its own firewall appliance. (See Figure 2.)

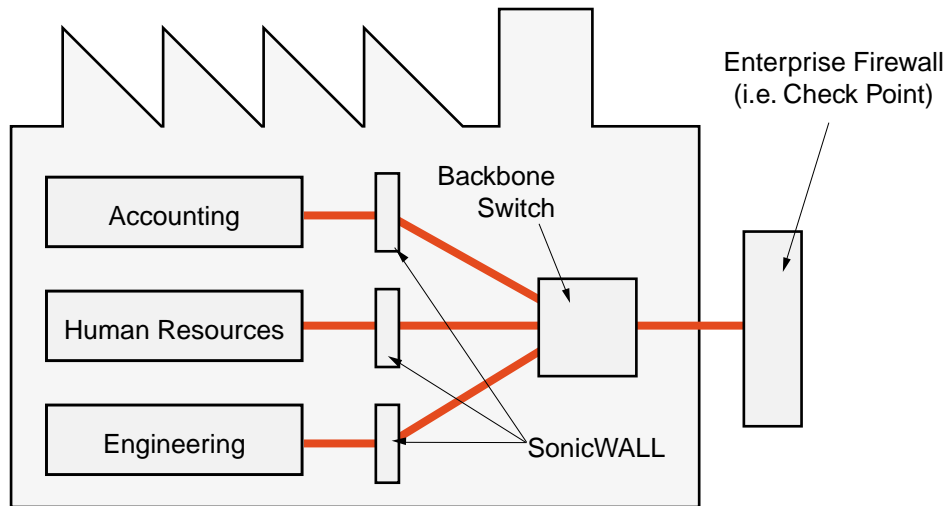


Figure 2: Interior Firewalls

As each LAN segment is typically running at 100Mbit/sec Fast Ethernet, a 10/100 product such as SonicWALL's SonicWALL PRO can provide transparent protection for each segment; its Web-based management interface would allow network administrators to monitor or configure the firewalls as needed.

Conclusion

Enterprise IT, networking, and WAN managers have expended much time and effort, as well as dollars, fortifying the main entrances to the enterprise network: its gateway to the Internet. High-end, high-performance firewalls, operating in conjunction with routers and other WAN access devices, provide adequate security against most frontal attacks.

Paradoxically, other critical network locations, such as the Internet connections to remote offices or telecommuters' homes, remain unguarded. The interior of the enterprise LAN is often considered a "trusted" zone, and thus is largely unprotected against attacks originating inside the trusted space. Internet connections between offices are often normal HTTP sessions, without the benefit of additional VPN security.

All three of those security holes can be plugged by the deployment of low-cost, high-performance firewall appliances. A well-defined enterprise security plan would take the availability of such devices into account.

APPENDIX A: TYPICAL DENIAL-OF-SERVICE ATTACKS

There are three types of Denial of Service attacks: those that exploit bugs in a TCP/IP implementation, those that exploit weaknesses in the TCP/IP specification, and brute-force attacks that flood a network with useless data, denying system resources to legitimate traffic.

Implementation Bugs

Ping of Death and Teardrop attacks exploit bugs in the TCP/IP implementations of various computer and host systems.

Ping of Death uses a “ping” utility to create an IP packet that exceeds the maximum 65,536 bytes of data allowed by the IP specification. The oversize packet is then sent to an unsuspecting system. Systems may crash, hang, or reboot.

Teardrop attack exploits weaknesses in the reassembly of IP packet fragments. As data is transmitted through a network, IP packets are often broken up into smaller chunks. Each fragment looks like the original IP packet except that it contains an offset field that says, for instance, “This fragment is carrying bytes 200 through 400 of the original (non-fragmented) IP packet.” The Teardrop program creates a series of IP fragments with overlapping offset fields. When these fragments are reassembled at the destination, some systems will crash, hang, or reboot.

TCP/IP Weaknesses

Weaknesses in the TCP/IP specification leave it open to SYN Flood and LAND attacks. These attacks are executed during the handshake that initiates a communication session between two applications. Under normal circumstances, the application that initiates a session sends a SYN (synchronize) packet to the receiving application. The receiver sends back a SYN-ACK (acknowledgment) packet, and then the initiator responds with an ACK (acknowledgment). After this handshake, the applications are set to send and receive data.

SYN Attack floods a targeted system with a series of SYN packets. Each packet causes the targeted system to issue a SYN-ACK response. While the targeted system waits for the ACK that follows the SYN-ACK, it queues up all outstanding SYN-ACK responses on what is known as a backlog queue. SYN-ACKs are moved off the queue only when an ACK comes back or when an internal timer (which is set at relatively long intervals) terminates the three-way handshake. Once the queue is full, the system will ignore all incoming SYN requests, making the system unavailable for legitimate users.

In a LAND Attack, hackers flood SYN packets into the network with a spoofed source IP address of the targeted system. This makes it appear as if the host computer sent the packets to itself, making the system unavailable while the target system tries to respond to itself.

Brute-Force Attacks

A bandwidth attack, such as a Smurf attack, targets a feature in the IP specification known as direct broadcast addressing, to quickly flood the target host or network with useless data.

A Smurf hacker floods a router with Internet Control Message Protocol (ICMP) echo request packets (pings). Since the destination IP address of each packet is the broadcast address of the network, the router will broadcast the ICMP echo request packet to all hosts on the network. If there are numerous hosts, this will create a large amount of ICMP echo request and response traffic.

If a hacker chooses to spoof the source IP address of the ICMP echo request packet, the resulting ICMP traffic will not only clog up the “intermediary” network, but will also congest the network of the spoofed source IP address, known as the “victim” network. This flood of broadcast traffic consumes all available bandwidth, making communications impossible.