January 2002

ITL Bulletin

ADVISING USERS ON INFORMATION TECHNOLOGY

GUIDELINES ON FIREWALLS AND FIREWALL POLICY

By John Wack, Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology

This *ITL Bulletin* discusses advances in firewall technology and outlines a number of issues involved in selecting the right kind of firewall for your organizational environment. It contains a series of recommendations for configuring and managing firewalls. The bulletin summarizes NIST Special Publication 800-41, *Guidelines on Firewalls and Firewall Policy*, which is available for download at http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf.

Firewall technology has improved substantially since it was introduced in the early 1990s, beginning with simple packet-filtering firewalls and advancing to more sophisticated firewalls capable of examining multiple layers of network activity and content. As the Internet has developed into the modern, complex network of today, Internet security has become very problematic, with break-ins, attacks, and web defacements now so commonplace as to be considered "part of the neighborhood." Thus, firewall technology is now a standard part of any organization's network security architecture, and even some home users on commercial dial-in and cable/DSL connections routinely employ personal firewalls.

Firewalls Now More Than Ever

All organizations connected to the Internet should use a firewall. Internet-borne attacks, ranging from direct intruder attacks to indirect attacks in the form of malicious active content in email or from web sites, are sufficiently prevalent that operating without firewall protection would be very dangerous. An increasingly important aspect of modern firewalls is their ability to filter on email and web content for viruses and malicious active content. Viruses are rampant; recent years have seen many outbreaks of viruses and worms that have caused major damage and losses to productivity.

Securing personal computers at home or remote locations is now as important as securing them at the office; many people telecommute or work at home and operate on organization or agency proprietary data. Home users dialing an Internet Service Provider (ISP) may have little firewall protection available to them, because the ISP has to accommodate potentially many different security policies. Therefore, personal firewalls have been developed to provide protection for remote systems and to perform many of the same functions as larger firewalls.

Newer Developments in Firewall Technology

The International Organization for Standardization's Open Systems Interconnect (OSI) model is a useful measuring stick for firewall capability. The OSI model describes seven layers of network protocols and functionality that apply to most networked systems. Early firewalls were capable of working at OSI layers 3 and 4, the Network and Transport layers, respectively. What this meant was that firewalls could examine fields in the TCP/IP packets containing the source and destination address, the protocol being used, and (sometimes) the application in use, e.g., SMTP (Simple Mail Transport Protocol) for email or HTTP (Hypertext Transfer Protocol) for web. Using these fields, early firewalls could make access control decisions (based on a policy) on TCP/IP packets, permitting or denying them, and restricting the packets and connections to specific systems.

Continued on page 2

ITL Bulletins are published by the Information Technology Laboratory (ITL) of the National Institute of Standards and Technology (NIST). Each bulletin presents an in-depth discussion of a single topic of significant interest to the information systems community. **Bulletins are issued on an as-needed basis** and are available from ITL Publications, National Institute of Standards and Technology, 100 Bureau Drive, Stop 8901, Gaithersburg, MD 20899-8901, telephone (301) 975-2832. To be placed on a mailing list to receive future bulletins, send your name, organization, and business address to this office. You will be placed on this mailing list only.

Bulletins issued since July 2000

- □ *Identifying Critical Patches with ICAT*, July 2000
- □ *Security for Private Branch Exchange Systems*, August 2000
- □ *XML Technologies*, September 2000
- □ *An Overview of the Common Criteria Evaluation and Validation Scheme*, October 2000
- □ *A Statistical Test Suite for Random and Pseudorandom Number Generators For Cryptographic Applications*, December 2000
- □ *What Is This Thing Called Conformance?* January 2001
- □ *An Introduction to IPsec (Internet Protocol Security),* March 2001
- □ *Biometrics– Technologies For Highly Secure Personal Authentication*, May 2001
- □ *Engineering Principles for Information Technology Security*, June 2001
- □ *A Comparison of The Security Requirements for Cryptographic Modules In FIPS 140-1 AND FIPS 140-2*, July 2001
- □ *Security Self-assessment Guide For Information Technology Systems*, September 2001
- □ *Computer Forensics Guidance*, November 2001

The early firewalls were often routers containing packet-filtering capability. As the Internet grew, so did the number of applications, the amount of traffic, and, unfortunately, the risk of attack from intruder activity, misconfigurations, and viruses and malicious code. Newer firewalls were developed to meet these needs and deal with the upper layers of the OSI model. The following types of firewalls are now available:

Boundary Router Packet Filter Firewall

An important point regarding packet filter routers is that their speed and flexibility, as well as capability to block denial-of-service and related attacks, makes them ideal for placement at the outermost boundary with an untrusted network. The packet filter, referred to as a *boundary router*, can block certain attacks, possibly filter unwanted protocols, perform simple access control, and then pass the traffic onto other firewalls that examine higher layers of the OSI stack.

The figure below shows a packet filter firewall used as a boundary router. The router would accept packets from the untrusted network connection, which typically would be another router owned or controlled by the Internet Service Provider (ISP). The router would then perform access control according to the policy in place, e.g., block SNMP, permit HTTP, etc. It would then pass the packets to other more powerful firewalls for more access control and filtering operations at higher layers of the OSI stack. The figure also shows an internal, less trusted network between the boundary router and an inner firewall, sometimes referred to as the external DMZ (DeMilitarized Zone) network.

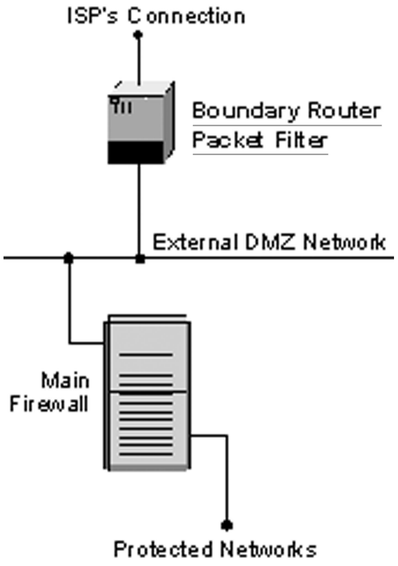

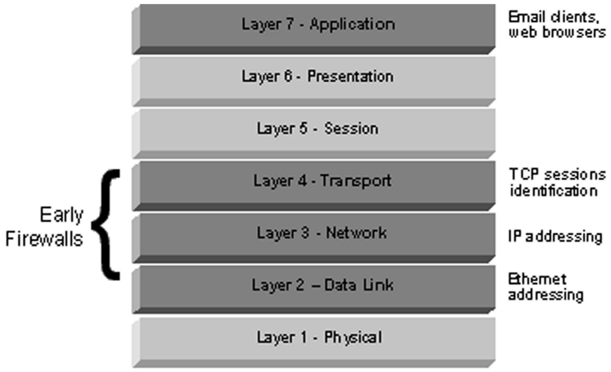

OSI Layers Handled by Modern Firewalls

Stateful Inspection Firewalls

Stateful inspection evolved from the need to accommodate certain features of the TCP/IP protocol suite that make firewall deployment difficult. When a TCP (connection-oriented transport) application creates a session with a remote host system, a port is also created on the source system for the purpose of receiving network traffic from the destination system. According to the TCP specifications, this client *source port* will be some number greater than 1023 and less than 16384. According to convention, the destination port on the remote host will likely be a “low-numbered” port, less than 1024. This will be 25 for SMTP, for example. To permit traffic to return from the destination port to a number of higher-numbered source ports, early firewalls (and routers in general) needed to permit return traffic to any source ports numbered higher than 1023. This left a large window of vulnerability open for a variety of attacks; one attack involved running a server on an internal system at a high-numbered port. This would then permit outside systems to connect freely to the internal system’s server, since it would be permitted by the firewall’s policy.

Stateful inspection firewalls address this problem, as well as some other inherent security problems in the TCP/IP protocol suite, by maintaining a database of connections and associated ports. In essence, they maintain a record of the state of each connection, and they make access control decisions based on that state. If a packet arrives at the firewall destined for a high-numbered source port, the firewall determines if there is a current connection associated with that port and subsequently passes or blocks the packet according to the policy.

Application-Proxy Gateway Filtering

These firewalls are fundamentally different from packet filtering routers in that they do not provide any inherent routing capability built into the firewall. Instead, software applications known as application-proxies pass or route the traffic, e.g., a web application-proxy routes web traffic, an email application-proxy routes email traffic. Each individual application-proxy interfaces directly with the firewall access control ruleset to determine whether a given piece of network traffic should be permitted to transit the firewall. In addition to the ruleset, each proxy agent has the ability to require authentication of each individual network user. This user authentication can take many forms, including passwords, authentication tokens, and biometric devices.

Application-proxy firewalls provide a high level of security because they are able to examine the application traffic itself. For example, a web application-proxy can filter the traffic for JavaScript or other active content if so desired. An email application-proxy can filter on MIME (Multipurpose Internet Mail Extensions) attachments and pass or deny them according to the policy. This results in a finer level of access-control for each application and connection.

Hybrid Firewalls

Recent advances in network infrastructure engineering and information security have caused a "blurring of the lines" that differentiate the various firewall platforms discussed earlier. The main result of these advances is that it is now common to see many application-proxy gateway firewalls with basic packet filter functionality. Likewise, many packet filter or stateful inspection packet filter firewall vendors have implemented basic application-proxy functionality to offset some of the weaknesses associated with their firewall platform. In most cases, packet filter or stateful inspection packet filter firewall vendors implement application proxies to provide improved network traffic logging and user authentication in their firewalls.

Nearly all major firewall vendors have introduced hybridization into their products in some way, shape, or form, so it is not always a simple matter to decide which specific firewall product is the most suitable for a given application or enterprise infrastructure. Hybridization of firewall platforms makes the pre-purchase product evaluation phase of a firewall project important. Supported feature sets, rather than firewall product classification, should drive the product selection. For example, some sites may be more interested in stateful filtering and speed, and not care about email and web application-proxy capability. Choosing a firewall with application-proxy capability might reduce the throughput of the firewall, therefore going with the stateful firewall, which has less overall capability, would be more appropriate.

Who we are
The Information Technology Laboratory (ITL) is a major research component of the National Institute of Standards and Technology (NIST) of the Technology Administration, U.S. Department of Commerce. We develop tests and measurement methods, reference data, proof-of-concept implementations, and technical analyses that help to advance the development and use of new information technology. We seek to overcome barriers to the efficient use of information technology, and to make systems more interoperable, easily usable, scalable, and secure than they are today. Our web site is http://www.itl.nist.gov/.

Host-Based Firewalls

Firewall packages are available in some operating systems such as Linux or as add-ons; they can be used to secure the individual host only. This can be helpful for using with internal servers; for example, an internal web server could be placed on a system running a host-based firewall. This carries several advantages, including the following:

- The server application is protected better than if it were running alone; internal servers should be protected and should not be assumed to be safe from attack because they are behind a main firewall.
- A separate firewall and subnet isn't necessary for securing the server; the host-based firewall performs these functions.

Host-based firewall packages typically provide access-control capability for restricting traffic to and from servers running on the host, and there is usually some limited logging available. While a host-based firewall is less desirable for high-traffic, high-security environments, in internal network environments or regional offices they offer greater security usually at a lower cost. A disadvantage to host-based firewalls is that they must be administered separately, and after a certain number, it becomes easier and less expensive to simply place all servers behind a dedicated firewall configuration.

Personal Firewalls and Firewall Appliances

Personal firewalls are typically implemented in one of two configurations, the first being a software product that runs on an individual personal computer. Such personal firewalls are installed on the system they are meant to protect; usually they do not offer protection to other systems or resources. Likewise, personal firewalls do not typically provide controls over network traffic that is traversing a computer system – they only protect the computer system on which they are installed. Personal firewalls, configured appropriately, can be very effective at protecting personal computers and laptops, and NIST recommends that all laptop and home users install and use a personal firewall product. Vendors now offer a variety of personal firewall software products such as Network ICE"s BlackICE Defender, Symantec's Personal Firewall, Zone Labs' ZoneAlarm*, and many others.

The second configuration is called a *Personal Firewall Appliance*, which is in concept more similar to that of a traditional firewall. In most cases, personal firewall appliances are designed to protect small networks such as networks that might be found in home offices or very small businesses. These appliances usually run on specialized hardware and integrate some other form of network infrastructure components in addition to the firewall itself, including the following:

- Broadband WAN Routing,
- LAN Routing (dynamic routing support),
- Network hub,
- Network switch,
- DHCP (Dynamic Host Configuration Protocol) server, and
- Network management (SNMP) agent.

Incorporating these infrastructure components into a firewall appliance allows an organization to deploy effective solutions consisting of a single piece of hardware.

ITL Bulletins Via E-Mail

We now offer the option of delivering your ITL Bulletins in ASCII format directly to your e-mail address. To subscribe to this service, send an e-mail message from your business e-mail account to listproc@nist.gov with the message **subscribe itl-bulletin**, and your name, e.g., John Doe. For instructions on using listproc, send a message to listproc@nist.gov with the message **HELP**. To have the bulletin sent to an e-mail address other than the From address, contact the ITL editor at 301-975-2832 or elizabeth.lennon@nist.gov.

Although personal firewalls and personal firewall appliances lack some of the advanced, enterprise scale features of traditional firewall platforms, they can still form an effective piece of the overall security posture of an organization. In terms of deployment strategies, personal firewalls and personal firewall appliances normally address the connectivity concerns associated with telecommuters or branch offices. However, some organizations employ these devices on the organizational intranet, practicing a defense-in-depth strategy. Personal firewalls and personal firewall appliances can also be used to terminate VPNs (Virtual Private Networks): many vendors currently offering firewall-based VPN termination also offer a personal firewall client as well.

Suggestions for Deployment

Very often, a firewall environment will consist of several firewalls and related systems, operating in tandem to provide a firewall capability for the internally and externally accessible systems of the organization. In the diagram below, there are three firewalls working together, a boundary router and two larger firewalls. They create two DMZ networks, the exterior one being used to locate externally accessible servers. The boundary router can filter traffic to these systems and provide basic protection, and the main firewall prevents any external traffic from entering the protected network (unless explicitly permitted, e.g., from the dial-in server).

The figure also shows network and host-based intrusion detection systems (IDSs) being used throughout the networks and on servers. These systems provide a check on the proper implementation of the firewall policies as well as provide some capability to recognize intrusion attempts.

The figure also shows a VPN server integrated in the main firewall, as well as the dial-in server located external to the main firewall. Both of these provide a means for external access to the internal networks and therefore should be located such that their traffic passes through the firewalls and is subject to all policy checks and controls.

In this diagram, an email server is located on an internal DMZ network, partitioned from direct external and internal access and so that it could filter incoming and outgoing mail for viruses in attachments and embedded malicious content. It is important to note here that the firewall environment can filter both incoming and outgoing traffic – filtering outgoing traffic for viruses, for example, helps protect the organization from accidentally spreading viruses.

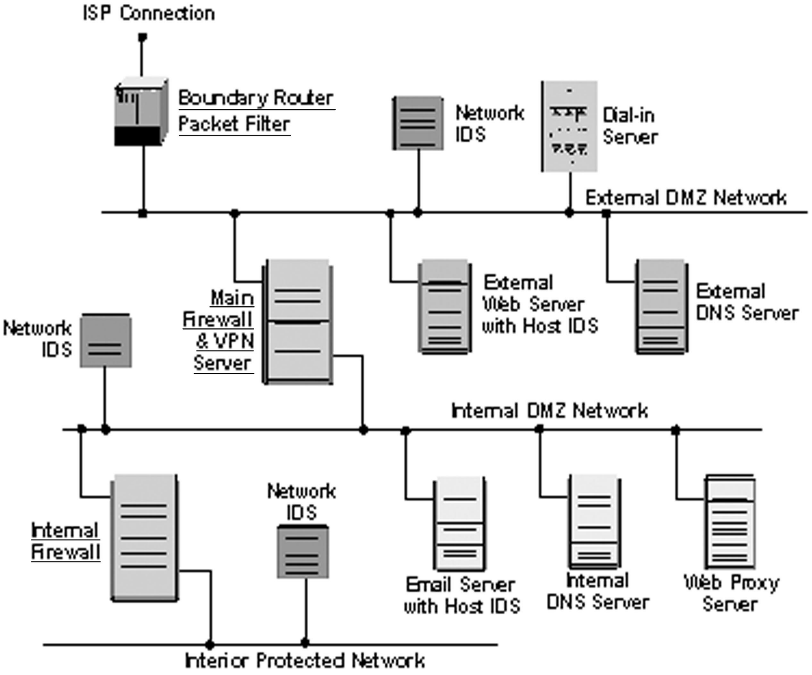

NIST Special Publication 800-41 suggests various methods for deploying firewalls. Here are some selected excerpts:

Keep It Simple. The KISS principle is something that should be first and foremost in the mind of a firewall environment designer. Essentially, the more simple the firewall solution, the more secure it likely will be and the easier it will be to manage. Complexity in design and function often leads to errors in configuration.

Use Devices as They Were Intended to Be Used. Using network devices as they were primarily intended in this context means do not make firewalls out of equipment not meant for firewall use. For example, basic routers are meant for routing; their packet filtering capability is not their primary purpose and relying solely on them for all firewall capability is not wise. In many cases, hybrid firewalls and firewall appliances are better choices simply because they are optimized to be firewalls first and foremost.

Create Defense in Depth. Instead of relying solely on the firewall to protect the network, use additional firewalls internally to protect sensitive systems as needed, e.g., financial systems. Some systems have a firewall capability built-in that can be enabled. Keep internal systems up to date with security patches and proper configuration. Make it more difficult for an intruder or for an insider to attack internal systems if the firewall fails or use the network as a platform for attacking other sites.

Pay Attention to Internal Threats. Attention to external threats to the exclusion of internal threats leaves the network wide open to attack from the inside. While it may be difficult to think of your work colleagues as posing a potential threat, consider that an intruder who gets past the firewall somehow could now have free reign to attack internal or external systems. Therefore, important systems such as internal web and email servers or financial systems should be placed behind internal firewalls or in DMZ environments.

Policy Guidelines and Recommendations

NIST Special Publication 800-41 contains numerous policy guidelines and recommendations for configuring and operating firewalls. The main recommendations are as follows:

Use firewalls to secure Internet connections and connections to other networks. At remote locations, use personal firewalls and firewall appliances to secure connections to the Internet and Internet Service Providers.

Examine carefully which firewall and firewall environment is best suited to your needs. Assistance is available from a number of commercial sites that deal with firewall selection and analysis including www.icsa.net/ (International Computer Security Association); a list of evaluated products for use in federal agencies is maintained by the National Information Assurance Center at http://csrc.nist.gov/niap.

Create a strong firewall security policy. A general risk assessment and cost-benefits analysis should be performed on the network applications required by the organization. This analysis should result in a firewall policy containing a list of the network applications and the methods that will be used to secure the applications. From there, a firewall environment can be created that best suits the organization's needs.

Audit the firewall and its policies at least quarterly. The firewall needs to be tested to ensure it is configured according to the policy and to ensure the policy is sufficiently rigorous. Patches and vulnerabilities need to be addressed as soon as possible.

Address inherent vulnerabilities in TCP/IP. Many successful attacks simply exploit commonly known vulnerabilities that can be largely mitigated by boundary routers. Stateful inspection firewalls also can correct features of TCP/IP that otherwise can leave large holes for intruders to exploit.

Employ filtering at the firewall for viruses and active content. Firewalls can scan for viruses and malicious code in email attachments and embedded content. Depending on the policy, firewalls can block certain types of attachments such as executable files. Firewalls can also block web-based active content, if desired, including JavaScript, Java™, and ActiveX® controls. Scanning can be done for both inbound and outbound traffic.

Separate externally accessible systems from private networks. Use firewalls to cordon off public web servers, directory servers, or other servers accessible to the public by placing them on DMZ networks. Use intrusion detection on those networks and systems to detect intruder activity.

Ensure the firewall is well managed. Firewalls need close supervision: logs must be read, adjustments must be made, vulnerabilities checked, equipment maintained, all of which can require significant amounts of time. Therefore, firewall administrators need to be assigned and given adequate time and training to do their jobs well.

Monitor incident response team reports and security websites for information about current attacks and vulnerabilities. Update the firewall policy as necessary. A formal process should be used for managing the addition and deletion of firewall rules.

Stay current with Internet security information. Technology, applications, and associated threats change daily. New viruses and worms can infiltrate organizations within hours and seriously harm productivity. Product patches are often released, security vulnerability reports are issued frequently, and administrators and management must keep up with the pace of change and ensure that their firewalls are configured properly.

Lastly, don't rely exclusively on the firewall. Internal security must still be a top priority. Internal systems must be patched and configured in a timely manner.

*Disclaimer: *Certain commercial products are described in this document as examples only. Inclusion or exclusion of any product does not imply endorsement or non-endorsement by NIST or any agency of the U.S. Government. Inclusion of a product name does not imply that the product is the best or only product suitable for the specified purpose.*
Java™ is a trademark of Sun Microsystems, Inc., in the United States and other countries.

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
100 Bureau Drive, Stop 8900
Gaithersburg, MD 20899-8900