

PROTECTING INFORMATION SYSTEMS WITH FIREWALLS: REVISED GUIDELINES ON FIREWALL TECHNOLOGIES AND POLICIES

Shirley Radack, Editor
Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology

Firewalls are essential devices or programs that help organizations protect their networks and systems, and help home users protect their computers, from hostile attacks, break-ins, and malicious software. Firewalls control the flow of network traffic between networks and between hosts that employ different security policies.

Firewalls were originally installed at the perimeter of networks, where hostile threats from external intruders could be detected and stopped. While these early firewalls provided some protection for an organization's internal systems, they could not recognize all instances and all forms of attack. For example, attacks sent from one internal host to another often did not pass through the network firewalls.

Networks are now often designed to provide protection at the network perimeter as well as at other network locations and to detect both external and internal attacks. Firewalls can now be used to restrict connectivity to and from internal networks that process personal information and carry out sensitive functions, such as accounting and personnel tasks. Firewalls can provide an additional layer of security by preventing unauthorized access to systems and information, and they can protect mobile devices that are placed directly onto external networks. To help organizations use today's firewall technology effectively, the Information Technology Laboratory of the National Institute of Standards and Technology (NIST) recently revised its guide to firewall technology and the development of firewall policies.

NIST Special Publication 800-41, Revision 1, Guidelines on Firewalls and Firewall Policy: Recommendations of the National Institute of Standards and Technology

Written by Karen Scarfone of NIST and Paul Hoffman of the Virtual Private Network Consortium, NIST Special Publication (SP) 800-41, Revision 1, replaces an earlier guide to firewalls that had been issued in 2002. The updated report provides an overview of firewall technology, and helps organizations plan for and implement effective firewalls.

The revised publication explains the technical features of firewalls, the types of firewalls that are available for implementation by organizations, and their security capabilities. Organizations are advised on the placement of firewalls within the network architecture, and on the selection, implementation, testing, and management of firewalls. Other issues covered in detail are the development of firewall policies, and recommendations on the types of network traffic that should be prohibited.

The appendices to the report contain helpful supporting material, including a glossary and lists of acronyms and abbreviations used in the text of the report. Also included in the appendices are listings of in-print and online resources that provide information on the effective use of firewalls as a component of a comprehensive approach to protecting information and information systems.

The revised guide to firewalls and firewall policies is available from the NIST Web page <http://csrc.nist.gov/publications/PubsSPs.html>.

Role of Firewalls in Network Communications

There are several types of firewall technologies, which can be most readily distinguished by which parts of network communications they can interpret. One way of accomplishing this is by referencing the four layers of network communications that are described by the Transmission Control Protocol/Internet Protocol (TCP/IP) interconnection standard:

- The **application layer**, the highest layer, sends and receives data for applications such as Domain Name System (DNS), Hypertext Transfer Protocol (HTTP), and Simple Mail Transfer Protocol (SMTP). The application layer itself is composed of layers of protocols, such as for message formats and message handling.
- The **transport layer** provides connection-oriented or connectionless services for transporting application layer services between networks, and can provide communications reliability services. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are commonly used at this layer.
- The **internet protocol** (or network) layer routes packets across networks using protocols including Internet Protocol version 4 (IPv4), IP version 6 (IPv6), Internet Control Message Protocol (ICMP), and Internet Group Management Protocol (IGMP).
- The **hardware** (or data link) layer is the lowest layer, controlling communications on the components of the physical network. The Ethernet protocol is a data link layer protocol.

The TCP/IP communications layers work together to transfer data between hosts. When a user wants to transfer data across networks, the data is passed from the highest layer through intermediate layers to the lowest layer, with each layer adding more information. The lowest layer sends the accumulated data through the physical network, and the data is then passed upwards by the receiving host through the communications layers to its destination.

Basic firewalls operate at one or a few layers, usually the lower layers, while the more advanced firewalls operate at all of the layers. Threats were previously most prevalent in the lower layers of network traffic, but now threats are common at the application layer. Firewalls are still needed to stop the significant threats at the lower layers of network

communications, but the firewalls that examine more layers can perform more comprehensive examinations. Firewalls that are effective in the application layer can potentially protect advanced applications and protocols, and provide services that are user-oriented. For example, a firewall that operates only in the lower layers usually cannot identify specific users, but a firewall with application layer capabilities can provide security services such as enforcing user authentication and logging events by specific users.

Firewall Technologies

Firewall technology is often combined with other technologies, such as routing and network address translation (NAT) capabilities. Firewalls may also include content filtering features and intrusion prevention technology. See Section 2 of the report for a discussion of the advantages and disadvantages of these firewall technologies.

The older firewalls were primarily **packet filter firewalls**. These are routing devices that provide access control capabilities for host addresses and communication sessions. Also known as *stateless inspection firewalls*, these devices do not keep track of the state of each flow of traffic that passes through the firewall; as a result, they cannot associate multiple requests within a single session to each other. Packet filter technology is still employed in most modern firewalls, along with other firewall methods. Unlike more advanced filters, packet filters do not analyze the content of packets. Their access control functionality is governed by a set of directives referred to as a *ruleset*. Packet filtering capabilities are built into most operating systems and into devices capable of routing, such as a network router that employs access control lists. Packet filters operate at the network layer, and can filter both inbound and outbound traffic.

Firewalls with **stateful inspection** functions improve on the capabilities of packet filters by tracking the state of connections and by blocking packets that deviate from the expected state. These firewalls have greater awareness of the transport layer; packets are intercepted at the network layer and inspected for adherence to an existing firewall rule, as packet filters do, but stateful inspection firewalls also keep track of each connection in a state table that contains information such as source IP address, destination IP address, port numbers, and connection state information.

Application firewalls add a *stateful protocol analysis* capability. Some vendors refer to this feature as *deep packet inspection*. Stateful protocol analysis improves upon the standard stateful inspection by providing basic intrusion detection technology to analyze protocols at the application layer and identify suspicious events. These firewalls can allow or deny access based on how an application is running over the network. An application firewall can determine if an email message contains a type of attachment that the organization does not permit, determine if protocols are being used incorrectly, block connections that are not allowed, and allow or deny access to Web pages. Firewalls with both stateful inspection and stateful protocol analysis features provide extensive capabilities to detect and prevent attacks, but they are not complete intrusion detection and prevention systems (IDPSs).

A firewall that has an **application-proxy gateway** capability combines lower-layer access control with upper-layer functionality. These firewalls contain a proxy agent that acts as an intermediary between two hosts communicating with each other. The proxy agent never allows a direct connection between the hosts. Each successful connection attempt results in the creation of two separate connections; one connection is between the client and the proxy server, and another connection is between the proxy server and the destination address. From the perspective of the two hosts, the connection appears to be direct. Because external hosts only communicate with the proxy agent, internal IP addresses are not visible to the outside world. The proxy agent interfaces directly with the firewall ruleset to determine whether specific network traffic should be allowed to pass through the firewall.

Some proxy agents can require authentication of each individual network user. This authentication process can include user identification (ID) and password, hardware or software token, source address, and biometrics. Like application firewalls, the proxy gateway operates at the application layer and can inspect the actual content of the network traffic.

Dedicated proxy servers retain proxy control of traffic, but they usually have much more limited firewalling capabilities than application-proxy gateway firewalls. Many dedicated proxy servers are application-specific, and some actually perform analysis and validation of common application protocols. These servers are usually deployed behind traditional firewall platforms. The proxy server can filter or log incoming traffic forwarded by the main firewall, and then forward the traffic to internal systems. A proxy server can also accept outbound traffic directly from internal systems, filter or log the traffic, and pass it to the firewall for outbound delivery. Lately, the use of *inbound* proxy servers has decreased considerably because they must mimic the capabilities of the main server that they are protecting. Most proxy servers now in use are *outbound* proxy servers.

Firewall devices may have to encrypt and decrypt specific network traffic flows between the protected network and external networks. This function is accomplished through **virtual private networks** (VPNs), which use additional protocols to encrypt traffic and provide user authentication and integrity checking. Common VPN architectures are gateway-to-gateway and host-to-gateway. Gateway-to-gateway architectures connect multiple fixed sites over public lines through the use of VPN gateways. An example of this task is the connection of branch offices to an organization's headquarters. The host-to-gateway architecture provides a secure connection to the network for individual users, usually called *remote users*, who are physically located outside of the organization, such as at home or in a hotel.

Gateway-to-gateway and host-to-gateway VPN functionality is often part of the firewall itself. Placing it behind the firewall would require the VPN traffic to be passed through the firewall while encrypted, thus preventing the firewall from inspecting the traffic. The

organization's VPN policy can specify the resources that users and groups are authorized to access.

Firewalls at the edge of a network may have to perform client checks for incoming connections from remote users and allow or disallow access based on those checks. This checking process, commonly called **network access control** (NAC) or **network access protection** (NAP), allows access based on the user's credentials and the results of the checking process, assuring that the user's computer complies with organizational policy concerning the use of patches, security software, and configuration settings.

Unified threat management (UTM) systems combine multiple security features into a single system, including a firewall, malware detection and eradication, and sensing and blocking of suspicious network probes. This approach may reduce the complexity of setting and maintaining policies on all of the systems that are deployed at the same location on a network, but it requires that the UTM have all the desired features to meet all security objectives.

Web application firewalls are installed in front of the Web server to detect the placement of malicious software on the computer of users who are accessing information on the Web, or to deter attempts to solicit private information from users. These firewalls, which are considered to be different from traditional firewalls, are a relatively new technology, with changing capabilities.

Firewalls for **virtual infrastructures** are another new area of firewall technology. These firewalls monitor *virtualized networking*, which allows more than one operating system running on a single computer simultaneously to communicate with each other as if they were on a standard Ethernet. Network activity that passes directly between virtualized operating systems within a host cannot be monitored by an external firewall. Some virtualization systems offer built-in firewalls or allow third-party software firewalls to be added as plug-ins.

Host-based firewalls for servers and **personal firewalls** for desktop and laptop personal computers (PCs) provide an additional layer of security against network-based attacks. These firewalls are software-based, residing on the hosts that they are protecting and monitoring and controlling the incoming and outgoing network traffic for a single host. They can provide more granular protection than network firewalls to meet the needs of specific hosts. Host-based firewalls are available as part of server operating systems. Many host-based firewalls can also act as intrusion prevention systems (IPSs) that, after detecting an attack in progress, take actions to thwart the attacker and prevent damage to the targeted host.

NIST Recommendations to Organizations on the Use of Firewalls and Firewall Policies

NIST recommends that organizations improve the effectiveness and security of their information systems by taking the following actions:

- **Create a firewall policy that specifies how firewalls should handle inbound and outbound network traffic.**

Firewall policies, which are described in Section 4 of the report, define how organizational firewalls should handle inbound and outbound network traffic for specific IP addresses and address ranges, protocols, applications, and content types based on the organization's information security policies. Organizations should conduct risk analyses to identify the acceptable types of network traffic and how they must be protected, specifying the types of traffic that can pass through the firewall and under what circumstances. For example, an organization might permit only necessary Internet Protocol (IP) protocols to pass, appropriate source and destination IP addresses to be used, particular Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports to be accessed, and certain Internet Control Message Protocol (ICMP) types and codes to be used. Generally, all inbound and outbound traffic not expressly permitted by the firewall policy should be blocked because such traffic is not needed by the organization. This practice will reduce the risk of attack and can also decrease the volume of traffic carried on the organization's networks.

- **Identify all requirements that should be considered when determining which firewall to implement.**

In planning for and selecting firewalls, organizations should determine the network areas to be protected, and consider the types of firewall technologies that will be most effective for protecting the required types of traffic. Organizations should also take into consideration firewall performance issues and the integration of the firewall into the existing network and security infrastructures. The design of the firewall should take into account any requirements relating to physical environment and to personnel, and to future needs, such as plans to adopt new IPv6 technologies or virtual private networks (VPNs).

- **Create rulesets that implement the organization's firewall policy while supporting firewall performance.**

Firewall rulesets should be as specific as possible for the network traffic that will be controlled. Organizations should determine the types of network traffic that are required, including the protocols that the firewall itself may need to use for management. The details of creating rulesets will vary according to the type of firewall and the specific firewall product, but the performance of many firewalls can be improved by optimizing firewall rulesets. For example, some firewalls check traffic against rules in a sequential manner until a match is found; for these firewalls, rules that have the highest chance of matching traffic patterns should be placed at the top of the list wherever possible.

- **Manage firewall architectures, policies, software, and other components throughout the life of the firewall solutions.**

The types of firewalls to deploy and their positions in the organization's networks can affect the security policies that the firewalls can enforce. Policy rules should be monitored and changed as the organization's requirements change, such as when new applications or hosts are implemented within the network. The performance of firewall components should be monitored to enable potential resource issues to be identified and addressed before performance is adversely affected. Logs and alerts should also be continuously monitored to identify threats, both successful and unsuccessful. Firewall rulesets and policies should be managed through formal change management control processes to avoid any impact on security and on business operations. Ruleset reviews or tests should be performed periodically to ensure continued compliance with the organization's policies. Firewall software should be patched as vendors provide updates to address any discovered vulnerabilities.

Related Publications

For information about NIST standards and guidelines related to the use of firewalls, as well as other security-related publications, see NIST's Web page <http://csrc.nist.gov/publications/index.html>.

Disclaimer

Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by NIST nor does it imply that the products mentioned are necessarily the best available for the purpose.