Perimeter Firewall Design

By Archiveddocs

Archived: 2026-04-06 03:15:45 UTC

Updated : February 6, 2004

This module helps you to select a suitable firewall product for your organization's perimeter network. It presents the different classes of available firewalls and highlights their significant features. It also gives you guidance in determining your own requirements and helps you to select the most appropriate product.

Use this module to:

- Identify the features necessary in your perimeter firewall.
- Classify firewall products.

- Select the best firewall product for your perimeter firewall.

This module applies to the following technologies:

- Ethernet/IP-based firewall products

Before reading this module, you should have an understanding of the TCP/IP protocol, your own network architecture, and in particular the devices in your perimeter network. It would also be useful to find out what inbound traffic from the Internet can be considered valid and what is invalid.

The design guidelines presented in this module will help you select the features you need from your firewall, taking into account major considerations such as growth and cost. The module will also provide you with information on some of the most damaging intrusions so that you can determine which are most likely to occur in your environment and how intrusions can be prevented, not just by installing a firewall but, for example, by tightening up server configurations or discussing controls with your Internet Service Provider (ISP). This module also defines different classes of firewalls and using the design guidelines you should be able to select the most appropriate class of firewall to meet your requirements. From the knowledge provided in this module and the technical terminology, you should be able to discuss with firewall manufacturers the products they can provide and evaluate their suitability for your requirements.

Network intrusions from both internal and external users occur with increasing frequency, and protection from these intrusions must be established. Although a firewall offers protection for your network, it also costs money and creates an impediment to traffic flow, so you should look for one that is as cost effective and efficient as possible.

In an enterprise network architecture, there will generally be three zones:

- Border network

  This network faces directly onto the Internet via a router which should provide an initial layer of protection, in the form of basic network traffic filtering. The router feeds data through to the perimeter network via a perimeter firewall.

- Perimeter network

  This network, often called the DMZ (demilitarized zone network) or edge network, links incoming users to the Web servers or other services. The Web servers then link to the internal networks via an internal firewall.

- Internal Networks

  The internal networks link the internal servers, such as SQL Server and the internal users.

These networks are depicted in Figure 1.


Cc700828.networkcomponents-0-0(en-us,TechNet.10).gif

Figure 1. Enterprise Network Architecture

A firewall checks incoming IP packets and blocks those it believes are intrusive. Some blocking can be done by recognizing by default that certain packets are illegal, others by configuring the firewall to block them. The TCP/IP protocol was designed many years ago, without any concept of hacking or intrusion, and it has many weaknesses. For example, the ICMP protocol was designed as a signaling mechanism within TCP/IP, but this protocol is open to abuse and can lead to such problems as denial-of-service (DoS) attacks. A perimeter firewall can have a more restricted capability than an internal firewall, because incoming traffic is more limited since its legal destination is the Web server or other special services.

Many types of firewalls are available, differentiated partly by price, but also on features and performance. Generally, the more expensive the firewall, the more power and features it has. Later on this module, the firewalls are grouped into classes to differentiate them, but before selecting a firewall, you need to determine what your requirements are, taking the following considerations into account:

- **Budget**
- **Existing facilities**
- **Availability**
- **Scalability**
- **Features required**

What is the available budget? Every firewall in the environment should provide the highest possible level of service while remaining cost-effective, but be aware of the resultant damage to your business if the firewall is too restricted by cost. Consider the downtime costs in your organization if the service is suspended by a denial of service attack.

Are there existing facilities that can be used to save costs? There may already be firewalls in the environment that can be reused and routers that can have a firewall feature set installed. Your ISP can often implement firewall restrictions on your link, such as *rate-limiting*, i.e. limiting the rate at which certain packets are sent to you in order to reduce distributed denial of service attacks, DDoS, when your network is bombarded simultaneously by many other computers. Ask your ISP if they perform filtering according to RFCs 1918 and 2827.

Does the firewall need to be available at all times? If you are offering a public Web server facility when users may want to connect 24 hours a day, you need almost 100% uptime. With any firewall there is always a chance of failure, so you need to mitigate against that. The availability of a firewall can be improved by two methods:

- Redundant components

  Duplicating those components more likely to fail, such as the power supply, improves the resilience of the firewall, as the first component can fail with no effect on operations. Low cost firewalls usually do not have any redundant options, and adding resilience to your firewall adds to the cost without increasing the processing power.

- Duplicate devices

Duplicating the firewall device provides a totally resilient system, but again at a considerable cost, as it also requires totally duplicate network cabling and duplicate connectivity in the routers or switches to which the firewall connects. However, depending upon the type of firewall, it may also double the throughput to compensate. In theory, all firewalls from the smallest to the largest could be duplicated, but in practice you also need a software switchover mechanism which may not be present in the smaller firewalls.

What is the throughput requirement of the firewall? Throughput can be considered both in terms of bits per second and packets transferred per second. If it is a new venture, you may not know the throughput rates, and if the venture is successful, the throughput from the Internet could escalate rapidly. In order to handle the change, you need to select a firewall solution that can scale up as the throughput increases, either by adding more components to your firewall, or by installing another firewall in parallel.

Which firewall features are required? Based on risk assessments conducted against the services provided in your organization, you can determine which firewall features are required to protect the assets that provide the services. If VPNs (Virtual Private Networks) are required, then this will affect the design.

This section provides a summary of some of the better known system attacks, along with reasons for using a firewall service as the first line of defense.

The Internet is a haven for those who want to adversely affect organizations or steal trade secrets to gain competitive advantage. If you install a perimeter firewall and look at the log of intrusions, you will be surprised by the volume. Most of these intrusions are just probes to see if your machine responds and to find out what services you are running. This may seem innocuous, but if the attacker discovers your machine he may then attack your service, knowing what weaknesses it has.

In addition to providing protection from Internet-based attacks, sensitive information must be protected. Most organizations have sensitive information that should be protected from certain users on the internal network, including employees but also vendors, contractors, and customers. While a perimeter firewall is primarily there to protect against external intrusions, knowledgeable internal users may try to enter via the Internet.

Intrusion threats can take many forms, and describing them all here would serve only a limited purpose, because new ones are created on a daily basis. Some intrusions, such as pinging a server address, may seem harmless, but after discovering the presence of a server, the hacker might attempt a more serious attack. In other words, all intrusions should be considered potentially harmful. Some of the major intrusions are:

- Packet Sniffers

  A sniffer is a software application or hardware device that attaches to the LAN and captures information from Ethernet frames. The original intention of these systems was to troubleshoot and analyze Ethernet traffic or to delve deeper into the frames to examine individual IP packets. Sniffers operate in promiscuous mode; that is, they listen to every packet on the physical wire. Many applications, such as Telnet, send user name and password information in clear text that can be read by sniffer products, and therefore a hacker with a sniffer could gain access to many applications.

Sniffing cannot be prevented by a firewall as a sniffer does not generate network traffic. There are various measures to counter sniffing, primarily by ensuring that strong encrypted passwords are used, but this is beyond the scope of this module.

- IP Spoofing

  IP spoofing occurs when the source address of an IP packet is changed to hide the identity of the sender. Because the routing operation within the Internet uses only the destination address to send a packet on its way and ignores the source address, a hacker can send a destructive packet to your system disguising the source without you knowing where it came from. Spoofing is not necessarily destructive, but it signals that an intrusion is at hand. The address may be outside your network (to hide the identity of the intruder) or it may be one of your trusted internal addresses with privileged access. Spoofing is typically used for denial of service attacks, which are described later in this module.

  It is possible to prevent IP spoofing by implementing either or both of the following mechanisms:

  - Access control

    Deny access to incoming packets from the Internet with a source address that is on your internal network.

  - RFC 2827 filtering

    It is important to ensure that no IP spoofing takes place on your outgoing traffic. Spoofed packets must originate on somebody's network; you want to be certain that your network is not being used as a source for spoofing. Therefore, you should prevent all outgoing traffic from your network that does not have a source address within your own allocation. Your ISP might also be able to drop spoofed traffic from your network by checking if the source address is one that belongs to your network. This technique is known as RFC 2827 filtering; contact your ISP for more information about how to implement it. Filtering outbound traffic has no benefit for you, but another network performing similar filtering could prevent a spoofed attack on your network. Most modern firewalls have the ability to prevent inbound IP spoofing.

- Denial-of-Service Attacks

  Denial of service (DoS) attacks are among the hardest to prevent. They differ from other types of attack in that they do not cause permanent damage to your network; instead, they try to stop the network functioning by bombarding a particular computer (either a server or a network device), or by degrading the throughput of network links to the point where performance is so abysmal it causes ill-will among customers and loss of business to the organization. A distributed DoS (DDoS) attack is an attack initiated from many other computers concentrating the bombardment on your system. The attacking computers have not necessarily initiated the attack themselves, but due to their own security vulnerabilities, they have allowed themselves to be infiltrated by a hacker who has directed them to send high volumes of data to your network, congesting either the link to your ISP or one of your devices.

- Application Layer Attacks

Application layer attacks are often the most publicized attacks, and usually exploit well-known weaknesses in applications, such as Web servers and database servers. The problem, particularly for Web servers, is that they are designed to be accessed by public users, who are unknown and cannot be trusted. Most attacks are against known deficiencies in the product, so the best defense is to install the latest updates from the manufacturers. The infamous Structured Query Language (SQL) Slammer worm affected 35,000 systems within a very short time of its release in January 2003. The worm exploited a known problem in Microsoft SQL Server™ 2000 for which Microsoft had already issued a fix four months earlier in August 2002, thus taking advantage of the fact that many administrators had neither applied the recommended update nor had adequate firewalls in place (which could have dropped packets destined for the port that the worm used). Note that a firewall is just a backstop in these situations; manufacturers recommend that upgrades should be applied to all products, particularly to prevent application layer attacks.

- Network Reconnaissance

  Network reconnaissance is the scanning of networks to discover valid IP addresses, domain name system (DNS) names, and IP ports prior to launching an attack. Although network reconnaissance is harmless by itself, discovering which addresses are in use can help someone launch a hostile attack. In fact, if you look at the logs for a firewall, you will find that most intrusions are of this nature; typical probes include scanning for listening transport control protocol (TCP) and user datagram protocol (UDP) ports, as well as for other well-known listening ports, such as those used by Microsoft SQL Server, NetBIOS, HTTP, and SMTP. All such probes seek a reply, which tells the hacker that the server exists and runs one of these services. Many of these probes can be prevented by the border router or by a firewall. Many services are present by default, but turn off any unrequired services, but turning off some of them may restrict your network diagnostics capabilities.

- Viruses/Trojan Horses

  Viruses generally cannot be detected by firewalls, as they are often embedded in email attachments. Traditional viruses tended to just damage the device that they had contaminated, but modern viruses often try to replicate and damage either other local machines or spread out onto the Internet by sending multiple emails with the virus attached. Many of these viruses install a Trojan Horse program on the contaminated device. A Trojan Horse program may not do any direct damage, but rather sends information from the device on which it is installed over the Internet to the hacker, who can then launch a targeted attack on that device, knowing what software it is running and where it is vulnerable. While the primary defense against viruses is always to maintain up-to-date anti-virus software on the device, the perimeter firewall may be useful in limiting the effectiveness of the Trojan Horse program.

A firewall is a mechanism for controlling the flow of IP traffic between two networks. Firewall devices typically operate at L3 of the OSI model, although some models can operate at higher levels as well.

Firewalls generally provide the following benefits:

- Defending internal servers from network attacks
- Enforcing network usage and access policies

- Monitoring traffic and generating alerts when suspicious patterns are detected

It is important to note that firewalls mitigate only certain types of security risks. A firewall does not usually prevent the damage that can be inflicted against a server with a software vulnerability. Firewalls should be implemented as part of an organization's comprehensive security architecture.

Depending on the features that a firewall supports, traffic is either allowed or blocked using a variety of techniques. These techniques offer varying degrees of protection, based on the capabilities of the firewall. The following firewall features are listed in increasing order of complexity:

- **Network adapter input filters**
- **Static packet filters**
- **Network address translation (NAT)**
- **Stateful inspection**
- **Circuit-level inspection**
- **Proxy**
- **Application layer filtering**

In general, firewalls that provide complex features will also support the simpler features. However, you should read vendor information carefully when choosing a firewall, because there can be subtle differences between its implied and actual capabilities. When selecting a firewall, you must inquire about the features and test it to ensure that the product can indeed perform according to specifications.

Network adapter input filtering examines source or destination addresses and other information in the incoming packet, and either blocks the packet or allows it through. It applies only to incoming traffic and cannot control outgoing traffic. It matches IP addresses, port numbers for UDP and TCP, as well as the protocol of the traffic, TCP, UDP, and generic routing encapsulation (GRE).

For a perimeter firewall protecting a Web server, legal incoming traffic should only be able to access the Web server IP address and usually a limited range of port numbers, such as 80 for HTTP or 443 for HTTPS. Although the perimeter firewall should have this control, it should also be implemented in the border router.

Network adapter input filtering allows a quick and efficient denial of standard incoming packets that meet the rule criteria configured in the firewall. However, this form of filtering can easily be evaded, as it only matches the headers of the IP traffic, working on the basic assumption that the traffic being filtered follows IP standards and is not crafted to evade the filtering.

Static packet filters are similar to network adapter input filters in the sense that they simply match IP headers to determine whether or not to allow the traffic to pass through the interface. However, static packet filters allow control over outbound as well as inbound communications to an interface. Furthermore, static packet filters typically allow an additional function over the network adapter filtering, which is to check if the Acknowledged

(ACK) bit is set on the IP header. The ACK bit gives information on whether the packet is a new request or a return request from an original request. It does not verify that the packet was originally sent by the interface receiving it; it merely checks whether the traffic coming into the interface appears to be return traffic, based on the conventions of the IP headers.

This technique only applies to the TCP protocol and not the UDP protocol. Like network adapter input filtering, static packet filtering is very fast, but its capabilities are limited, and it can be evaded by specifically crafted traffic.

As with network adapter input filtering, static packet filtering should also be implemented on the border router in addition to the perimeter firewall.

In the worldwide IP address range, certain address ranges are designated as *private addresses*. These are intended to be used in your organization and have no meaning in the Internet. Traffic destined for any of these IP addresses cannot be routed through the Internet, so assigning a private address to your internal devices gives them some protection against intrusion. However, these internal devices often need to access the Internet themselves and so Network Address Translation (NAT) converts the private address into an Internet address.

Although NAT is not strictly a firewall technology, concealing the real IP address of a server prevents attackers from gaining valuable fingerprinting information about the server.

In stateful inspection, all outgoing traffic is logged in a state table. When the connection traffic returns to the interface, the state table is checked to ensure that the traffic originated from this interface. Stateful inspection is a bit slower than static packet filtering; however, it ensures that the traffic is allowed to pass only if it matches the outgoing traffic requests. The state table contains items such as destination IP address, source IP address, port being called, and originating host.

Certain firewalls may store more information in the state table than others (such as IP fragments sent and received). The firewall can verify that the traffic is processed when all or just some of the fragmented information returns. Different vendors' firewalls implement the stateful inspection feature differently; so you should read the firewall documentation carefully.

The stateful inspection feature typically assists in mitigating the risk posed by network reconnaissance and IP spoofing.

With circuit-level filtering, it is possible to inspect sessions, as opposed to connections or packets. Sessions are established only in response to a user request and may include multiple connections. Circuit-level filtering provides built-in support for protocols with secondary connections, such as FTP and streaming media. It typically assists in mitigating the risks posed by network reconnaissance, DoS, and IP spoofing attacks.

Proxy firewalls request information on behalf of a client. In contrast to the firewall technologies discussed above, the communication does not occur directly between the client and the server hosting the service. Instead, the proxy firewall gathers information on behalf of the client and returns the data it receives from the service back to the client. Because the proxy server gathers this information for one client, it also caches the content to disk or memory. If another client makes an identical data request, the request can be satisfied from the cache, resulting in reduced network traffic and server processing time.

For non-encrypted sessions, such as FTP read-only and HTTP sessions, a proxy firewall actually creates individual sessions with both the client and the server, so there is never a direct connection between the two. For encrypted sessions, on the other hand, the proxy server verifies that the header information conforms to the standards of Secure Sockets Layer (SSL) communication before allowing the traffic to pass. However, the proxy cannot inspect the data passing by, because it is encrypted end-to-end by the client and the server.

The advantages of a proxy server over the firewall technologies discussed above include:

- No direct connections between client and server

  The client and server do not usually make direct connections to each other; even if they do (such as with SSL), protocol header and traffic inspection is performed.

- The server can cache the content of frequently requested sites

  Caching saves bandwidth and prevents unnecessary requests from exiting the environment.

- Validation of protocols that pass through it

  In addition to validating the port number through which the communication travels, proxy servers also validate the protocols that pass through them. The most typical protocols that are inspected are FTP download only, HTTP, SSL, and some text messaging services (such as text only, no video, audio, or file transfers).

- Can be configured to forward requests based on a user's ID

  Proxy servers can often be configured to forward requests based on user ID (that is, restrictions can be set only for certain users), rather than just source IP, port, or protocol.

The main drawback to a proxy server is that it requires much more processing power to perform protocol inspection. However, processing power is increasing all the time, so this is becoming less of an issue. Still, proxy servers do not have the throughput of a stateful or packet filtering firewall. Arguably, the added benefits of protocol inspection are necessary in a world where high-speed networks abound for home users and where Internet connectivity is becoming increasingly available to non-trusted nodes that are connected by ISPs with little or no legal obligation to provide trusted Internet services.

The proxy feature typically assists in mitigating the risk posed by network reconnaissance, DoS, IP spoofing attacks, virus/Trojan horse, and some application layer attacks.

The most sophisticated level of firewall traffic inspection is application-level filtering. Good application filters allow you to analyze a data stream for a particular application and provide application-specific processing, including inspecting, screening or blocking, redirecting, and modifying data as it passes through the firewall.

This mechanism is used to protect against things like unsafe SMTP commands or attacks against internal Domain Name System (DNS) servers. Third-party tools for content screening, such as virus detection, lexical analysis, and site categorization, can usually be added to your firewall.

An application layer firewall has the ability to inspect many different protocols, based on the traffic that passes through it. Unlike a proxy firewall, which usually inspects the Internet traffic (such as HTTP, FTP download, and SSL) the application layer firewall has much greater control over the way that traffic travels through the firewall. For example, an application layer firewall is capable of allowing only the UDP traffic that originates inside the firewall boundary to pass through. If an Internet host was to port scan a stateful firewall to see if it allowed DNS traffic into the environment, the port scan would probably show that the well-known port associated with DNS was open, but once an attack is mounted, the stateful firewall would reject the requests, because they did not originate internally. An application layer firewall might open ports dynamically, based on whether or not the traffic originates internally.

The application layer firewall feature assists in mitigating the risks posed by IP spoofing, DoS, some application layer attacks, network reconnaissance, and virus/Trojan horse attacks. Drawbacks of an application layer firewall are similar to the proxy, in the sense that it requires much more processing power and is typically much slower at passing traffic than stateful or static filtering firewalls. The most important consideration when using an application layer firewall is determining what the firewall is capable of doing at the application layer.

The application layer feature ensures that the traffic being passed over a port is appropriate. Unlike a packet filter or stateful inspection firewall that simply looks at the port and at the source and destination IP addresses, firewalls that support the application layer filtering feature have the ability to inspect the both data and the commands being passed back and forth.

Most firewalls that support the application layer feature only have application layer filtering for clear text traffic, such as a proxy-aware messaging service, HTTP, and FTP. It is important to keep in mind that a firewall which supports this feature can govern traffic going in and out of the environment. Another advantage of this feature is the ability to inspect DNS traffic as it goes through the firewall to look for DNS-specific commands. This additional layer of protection ensures that users or attackers cannot conceal information in allowed types of traffic.

If your organization has an online store, which collects credit card numbers and other personal information about customers, it is prudent to take the highest level of precautions in protecting this information. In these cases, it is essential that this type of high security data is encrypted between the user's PC and your Web servers, using the Secure Sockets Layer (SSL) protocol.

It is important to distinguish those cases where the application layer feature is used in conjunction with SSL. SSL is encrypted, and the firewall cannot understand the protocol commands because they are located within the encrypted packet. Each firewall that supports the application layer feature handles this differently, so it is important to read the fine print of the documentation for whichever firewall you choose.

The problem is that no device is supposed to be able to inspect data once an SSL session is established and the encryption is negotiated. For example, a client using a firewall that supports the proxy-type application layer feature requests the firewall to initiate a connection to a secure Web server on its behalf. The firewall and the server do the initial setup of the TCP connection, and the firewall hands over the connection to the client to set up the encryption with the server. After the connection is handed over to the client, the firewall no longer has the ability to inspect the data.

When the application layer feature is used to expose Internet services publicly, the following options are available:

- Terminating the SSL traffic at the firewall

  This allows the firewall to inspect incoming SSL connections for legitimate Web traffic and to discard traffic as the firewall decrypts the data for the Internet service.

- Regenerating SSL traffic from the firewall to the exposed Web service

  This is particularly helpful if basic credentials (such as clear text user name and password) are used within the SSL tunnel. Individuals who can sniff traffic between the internal interface of the firewall and the published Web service cannot get at the traffic because it is re-encrypted.

- Allowing the SSL traffic to pass through the firewall to the back-end server

  This is essentially the reverse approach of the SSL connection between the internal client and the external server.

These options provide numerous ways of controlling how far an encrypted session can be allowed to tunnel into an environment. In general, the closer you can keep encrypted traffic to the edge of your environment the better, because nothing in between can really see what is inside that tunnel.

This section presents a number of firewall classes, each of which provides certain features. Specific firewall classes can be used to respond to specific requirements in the IT architecture design.

Grouping firewalls into classes allows for the abstraction of the hardware from the requirements of the service, so that service requirements can be matched against class features. As long as a firewall fits into a specific class, you can assume it supports all the services of that class.

The various classes are as follows:

- Personal firewalls
- Router firewalls
- Low-end hardware firewalls
- High-end hardware firewalls
- Server firewalls

It is important to understand that some of these classes overlap; this is by design. The overlap allows one type of firewall solution to span multiple classes. Many classes can also be served by more than one hardware model from the same vendor, so that an organization can select a model that best suits their needs both now and in the future.

Apart from the price and feature set, firewalls can be classified on the basis of performance (or throughput). However, many manufacturers do not provide any throughput figures for their firewall. Where they are provided (usually for hardware firewall devices), no standard measurement process is followed, which makes comparisons between manufacturers difficult. For example, one measure is the number of bits per second (bps), but as the

firewall is actually passing IP packets, this measure is meaningless if the packet size used in measuring the rate is not included.

The following sections define each firewall class in detail.

A personal firewall is defined as a software service that provides a simple firewall capability for a personal computer. As the number of permanent Internet connections (as opposed to dial-up connections) has grown, the use of personal firewalls has increased.

Although designed to protect a single computer, a personal firewall can also protect a small network, if the computer on which it is installed is sharing its connection to the Internet with other computers on the internal network. However, the performance of personal firewall software is limited and it degrades the performance of the personal computer on which it is installed. The protection mechanisms are usually less effective than a dedicated firewall solution, because they are usually restricted to blocking IP and port addresses, although generally speaking a lower level of protection is needed on a personal computer.

Personal firewalls may be supplied with an operating system or at a very low cost. They are suitable for their intended purpose, but because of their restricted performance and functionality, they should not be considered for use in an enterprise, even in small satellite offices. They are, however, particularly suitable for mobile users on laptop computers.

Personal firewalls vary tremendously in their capabilities and price. However, lack of a specific feature, especially on a laptop, might not be of great importance. The following table shows the features commonly available in personal firewalls.

Table 1. Class 1 - Personal Firewalls

Firewall Attribute	**Value**
Basic features supported	Most personal firewalls support static packet filters, NAT, and stateful inspection, while some support circuit-level inspection and/or application layer filtering.
Configuration	Automatic (manual option also available)
Block or allow IP addresses	Yes
Block or allow protocol or port numbers	Yes

Firewall Attribute	Value
Block or allow incoming ICMP messages	Yes
Control outgoing access	Yes
Application protection	Possibly
Audible or visible alerts	Possibly
Log file of attacks	Possibly
Real-time alerts	Possibly
VPN support	Typically no
Remote management	Typically no
Manufacturer support	Varies widely (depends on the product)
High-availability option	No
Number of concurrent sessions	1 to 10
Modular upgradeability (hardware or software)	None to limited
Price range	Low (free in some cases)

Personal firewalls offer the following advantages and disadvantages.

The advantages of personal firewalls include:

- Inexpensive

  When only a limited number of licenses are required, personal firewalls are an inexpensive option. A personal firewall is integrated into versions of the Microsoft Windows XP operating system. Additional products that work with other versions of Windows or other operating systems are available for free or at limited cost.

- Easy to configure

  Personal firewall products tend to have basic configurations that work out-of-the-box with straightforward configuration options.

The disadvantages of personal firewalls include:

- Difficult to manage centrally

  Personal firewalls need to be configured on every client, which adds to the management overhead.

- Only basic control

  Configuration tends to be a combination of static packet filtering and permission-based blocking of applications only.

- Performance limitations

  Personal firewalls are designed to protect a single personal computer. Using them on a computer that serves as a router for a small network will lead to degraded performance.

Routers usually support one or more of the firewall features discussed previously; they can be subdivided into low-end devices designed for Internet connections and high-end traditional routers. The low-end routers provide basic firewall features for blocking and allowing specific IP addresses and port numbers, and use NAT to hide interior IP addresses. They often provide the firewall feature as standard, optimized to block intrusions from the Internet, and while they need no configuration, they can be refined with further configuration.

High-end routers can be configured to tighten up access by barring the more obvious intrusions, such as pings, and by implementing other IP address and port restrictions through the use of ACLs. Additional firewall features may be available, which provide stateful packet filtering in some routers. In high-end routers, the firewall capability is similar to that of a hardware firewall device at a lower cost, but also with a lower throughput.

Table 2. Class 2 - Router Firewall

Firewall Attribute	Value
Basic features supported	Most router firewalls support static packet filters. Lower-end routers typically support NAT. Higher-end routers may support stateful inspection and/or application layer filtering.
Configuration	Typically automatic on lower-end routers (with manual options). Often manual on higher-end routers.
Block or allow IP addresses	Yes
Block or allow protocol/port numbers	Yes
Block or allow incoming ICMP messages	Yes
Control outgoing access	Yes
Application protection	Possibly
Audible or visible alerts	Typically
Log file of attacks	In many cases
Real-time alerts	In many cases
VPN Support	Common in lower-end routers, not as common in higher-end routers. Separate dedicated devices or servers for this task are available.

Firewall Attribute	Value
Remote management	Yes
Manufacturer support	Typically limited in lower-end routers and good in higher-end routers.
High-availability option available	Low End: No - High End: Yes
Number of concurrent sessions	10 - 1,000
Modular upgradeability (hardware or software)	Low End: No - High End: Limited
Price range	Low to High

Router firewalls offer the following advantages and disadvantages.

The advantages of router firewalls include:

- Low cost solution

  Activation of an existing router firewall feature may not add any cost to the price of the router, and it requires no additional hardware.

- Configuration can be consolidated

  Router firewall configuration can be accomplished when the router is configured for normal operations, thereby minimizing the management effort. This solution is particularly suitable for satellite branch offices, since network hardware and manageability are simplified.

- Investment protection

  Router firewall configuration and management is familiar to the operations staff, so no retraining is required. Network cabling is simplified, because no additional hardware is installed, which also simplifies network management.

The disadvantages of router firewalls include:

- Limited functionality

  In general, low-end routers only offer basic firewall features. High-end routers typically offer higher-level firewall features, but may need considerable configuration, much of which is done through the addition of controls that are easily forgotten, making it somewhat difficult to configure correctly.

- Only basic control

  Configuration tends to be a combination of static packet filtering and permission-based blocking of applications only.

- Performance impact

  Using a router as a firewall detracts from the performance of the router and slows the routing function, which is its primary task.

- Log file performance

  Use of a log file to catch unusual activities can seriously reduce the performance of the router, especially when it is already under attack.

At the low end of the hardware firewall market are Plug-and-Play units, which require little or no configuration. These devices often incorporate switch and/or VPN functionality as well. Low-end hardware firewalls are targeted at small businesses and for internal use in larger organizations. They generally offer static filtering capabilities and basic remote management functionality. Devices from larger manufacturers may run the same software as their higher-end counterparts, providing an upgrade path should one be required.

Table 3. Class 3 - Low-end Hardware Firewall

Firewall Attribute	**Value**
Basic features supported	Most low-end hardware firewalls support static packet filters and NAT. May support stateful inspection and/or application layer filtering.
Configuration	Automatic (manual option also available)
Block or allow IP addresses	Yes
Block or allow protocol/port numbers	Yes

Firewall Attribute	Value
Block or allow incoming ICMP messages	Yes
Control outgoing access	Yes
Application protection	Typically not
Audible or visible alerts	Typically not
Log file of attacks	Typically not
Real-time alerts	Typically not
VPN Support	Sometimes
Remote management	Yes
Manufacturer support	Limited
High-availability option available	Typically not
Number of concurrent sessions	> 10 - 7500
Modular upgradeability (hardware or software)	Limited
Price range	Low

Low-end hardware firewalls offer the following advantages and disadvantages.

The advantages of low-end hardware firewalls include:

- Low cost

  Low-end firewalls can be purchased inexpensively.

- Simple Configuration

  Almost no configuration is required.

The disadvantages of low-end hardware firewalls include:

- Limited functionality

  In general, low-end hardware firewalls only offer basic firewall functionality. They cannot be run in parallel for redundancy.

- Poor throughput

  Low-end hardware firewalls are not designed to handle high-throughput connections, which may cause bottlenecks.

- Limited manufacturer support

  As these are low cost items, manufacturer support is usually limited to e-mail and/or a Web site.

- Limited upgradeability

  Usually there can be no hardware upgrades, though there are often periodic firmware upgrades available.

At the high end of the hardware firewall market, there are high-performance, highly resilient products, which are suitable for the enterprise or service provider. These usually offer the best protection, without reducing the performance of the network.

Resilience can be achieved by adding a second firewall running as a hot standby unit, which maintains the current table of connections through automatic stateful synchronization.

Firewalls should be used in every network connected to the Internet, because intrusion happens constantly; DoS attacks, theft, and data corruption are being attempted all the time. High-end hardware firewall units should be considered for deployment in central or headquarters locations.

Table 4. Class 4 - High-end Hardware Firewall

Firewall Attribute	Value
Basic features supported	Most high-end hardware firewalls support static packet filters and NAT. They may support stateful inspection and/or application layer filtering.
Configuration	Typically manual
Block or allow IP addresses	Yes
Block or allow protocol/port numbers	Yes
Block or allow incoming ICMP messages	Yes
Control outgoing access	Yes
Application protection	Potentially
Audible or visible alerts	Yes
Log file of attacks	Yes
Real-time alerts	Yes
VPN support	Potentially
Remote management	Yes
Manufacturer support	Good

Firewall Attribute	Value
High-availability option available	Yes
Number of concurrent sessions	> 7500 - 500,000
Modular upgradeability (hardware or software)	Yes
Price range	High

High-end hardware firewalls offer the following advantages and disadvantages.

The advantages of high-end hardware firewalls include:

- High performance

  Hardware firewall products are designed for a single purpose and provide high levels of intrusion-blocking together with the least degradation of performance.

- High availability

  High-end hardware firewalls can be connected together for optimal availability and load balancing.

- Modular systems

  Both hardware and software can be upgraded for new requirements. Hardware upgrades may include additional Ethernet ports, while software upgrades may include detection of new methods of intrusion.

- Remote management

  High-end hardware firewalls offer better remote management functionality than their low-end counterparts.

- Resilience

  High-end hardware firewalls may have availability and resilience features, such as hot or active standby with a second unit.

- Application layer filtering

Unlike their low-end counterparts, high-end hardware firewalls provide filtering for well-known applications at the L4, L5, L6, and L7 layers of the OSI model.

The disadvantages of high-end hardware firewalls include:

- High cost

  High-end hardware firewalls tend to be expensive. Although they can be purchased for as little as $100, the cost is much higher for an enterprise firewall, since the price is often based on the number of concurrent sessions, throughput, and availability requirements.

- Complex configuration and management

  Because high-end hardware firewalls have much greater capability than low-end firewalls, they are also more complex to configure and manage.

A variety of products are available that add firewall capability to a high-end server, providing robust fast protection on standard hardware and software systems. The benefits of this approach are the use of familiar hardware or software, which provides a reduced number of inventory items, simplified training and management, reliability, and expandability. Many of the high-end hardware firewall products are implemented on an industry-standard hardware platform running an industry-standard operating system (but hidden from view) and therefore have little difference, either technically or in performance, from a server firewall. However, because the operating system is still visible, the server firewall feature can be upgraded and made more resilient by techniques such as clustering.

Because the server firewall is a server running a commonly-used operating system, additional software, features, and functionality can be added to the firewall from a variety of vendors (not just one vendor, which is the case with a hardware firewall). Familiarity with the operating system can also lead to more effective firewall protection, because some of the other classes need considerable expertise for full and correct configuration.

This class is suitable where there is a high investment in a particular hardware or software platform, because using the same platform for the firewall makes the management task simpler.

The caching capability of this class can also be very effective.

Table 5. Class 5 - High-end Server Firewall

Firewall Attribute	Value
Features supported	Most high-end server firewalls support static packet filters and NAT. They may also support stateful inspection and/or application layer filtering.
Configuration	Typically manual

Firewall Attribute	Value
Block or allow IP addresses	Yes
Block or allow protocol/port numbers	Yes
Block or allow incoming ICMP messages	Yes
Control outgoing access	Yes
Application protection	Potentially
Audible or visible alerts	Yes
Log file of attacks	Yes
Real-time alerts	Yes
VPN support	Potentially
Remote management	Yes
Manufacturer support	Good
High-availability option available	Yes

Firewall Attribute	Value
Number of concurrent sessions	>50,000 (across multiple network segments)
Modular upgradeability (hardware or software)	Yes
Other	Commonly used operating system
Price range	High

Server firewalls offer the following general advantages and disadvantages.

The advantages of server firewalls include:

- High performance

  When run on a suitably sized server, these firewalls can offer high levels of performance.

- Integration and consolidation of services

  Server firewalls can make use of various features of the operating system on which they run. For example, firewall software that runs on the Microsoft Windows Server™ 2003 operating system can take advantage of the Network Load Balancing functionality built into the operating system. Additionally, the firewall could serve as a VPN server, again utilizing functionality in the Windows Server 2003 operating system.

- Availability, resilience, and scalability

  Because this firewall runs on standard personal computer hardware, it has all the availability, resilience, and scalability features of the personal computer platform on which it runs.

The disadvantages of server firewalls include:

- Requires high-end hardware

  For high performance, most server firewall products require high-end hardware in terms of central processing unit (CPU), memory, and network interfaces.

- Susceptible to vulnerabilities

Because server firewall products run on well-known operating systems, they are susceptible to the vulnerabilities present in the operating system and other software running on the server. Although this is also the case for hardware firewalls, their operating systems are not usually as familiar to attackers as most server operating systems.

A perimeter firewall exists to serve the requirements of users outside the boundaries of the organization. User types may include:

- Trusted

  Employees of the organization, such as branch office workers, remote users, or users that work from home.

- Semi-trusted

  Business partners of the organization, for whom a higher level of trust exists than with untrusted users. However, it is often still a somewhat lower level of trust than that with the organization's employees.

- Untrusted

  For example, users of the organization's public Web site.

It is important to consider the fact that the perimeter firewall is particularly open to external attack, because it must be broken for an intruder to get further into your network. It therefore becomes an obvious goal to break.

Firewalls used in a border capacity are an organization's gateway to the outside world. In many large organizations, the firewall class implemented here is typically a high-end hardware or server firewall, although some organizations use router firewalls. When selecting the firewall class to use as a perimeter firewall there are a number of issues that should be considered. The following table highlights these issues.

Table 6. Perimeter Firewall Class Choice Issues

Issue	Typical Characteristics of a Firewall Implemented in This Capacity
Required firewall features, as specified by the security administrator	This is a balance between the degree of security required versus the cost of the feature and the potential degradation of performance that increased security may cause. While many organizations want the maximum security for a perimeter firewall, some are not willing to take the performance hit. For example, very high-volume Web sites not involved with e-commerce may allow lower levels of security, based on higher levels of throughput obtained by using static packet filters instead of application layer filtering.
Whether the device will be a dedicated physical device, provide other functionality,	As the gateway between the Internet and the enterprise's network, the perimeter firewall is often implemented as a dedicated device, in order to

Issue	Typical Characteristics of a Firewall Implemented in This Capacity
or be a logical firewall on a physical device	minimize the attack surface and accessibility of internal networks that would occur if the device were breached.
Manageability requirements for the device, as specified by the organization's management architecture	Some form of logging is typically used, while an event monitoring mechanism is also often required. Remote administration may not be allowed here, in order to prevent a malicious user from remotely administering the device and only local administration will be allowed.
Throughput requirements will likely be determined by the network and service administrators within the organization	These will vary for each environment, but the power of the hardware in the device or server and the firewall features being used will determine the overall network throughput available.
Availability requirements	As the gateway to the Internet in large enterprises, high levels of availability are often required, especially when a revenue-generating Web site is protected by a perimeter firewall.

In the following discussion, the term bastion host means a server located in your perimeter network that provides services to both internal and external users. Examples of bastion hosts include Web servers, and VPN servers.

Typically, your perimeter firewall will need the following rules implemented, either by default or by configuration:

- Deny all traffic unless explicitly allowed.
- Block incoming packets that claim to have an internal or perimeter network source IP address.
- Block outgoing packets that claim to have an external source IP address (traffic should only originate from bastion hosts).
- Allow for UDP-based DNS queries and answers from the DNS resolver to DNS servers on the Internet.
- Allow for UDP-based DNS queries and answers from the Internet DNS servers to the DNS advertiser.
- Allow external UDP-based clients to query the DNS advertiser and provide an answer.
- Allow TCP-based DNS queries and answers from Internet DNS servers to the DNS advertiser.
- Allow outgoing mail from the outbound SMTP bastion host to the Internet.

- Allow incoming mail from the Internet to the inbound SMTP bastion host.
- Allow proxy-originated traffic from the proxy servers to reach the Internet.
- Allow proxy-responses from the Internet to be directed to the proxy servers on the perimeter.

The hardware requirements for a perimeter firewall are different for software-based and hardware-based firewalls, as summarized below:

- Hardware-based firewall

  These devices usually run specialized code on a custom-built hardware platform. They are typically scaled (and priced) based on the number of connections they can handle and the complexity of the software that is to be run.

- Software-based firewalls

  These are also configured based on the number of concurrent connections and the complexity of the firewall software. Calculators exist that can compute the processor speed, memory size, and disk space needed for a server, based on the number of connections supported. You should take into account other software that may also be running on the firewall server, such as load balancing and VPN software. Also, consider the methods for scaling the firewall both upward and outward. These methods include increasing the power of the system by adding additional processors, memory, and network cards, and also using multiple systems and load balancing to spread the firewall task across them (see the "Scalability" section later on in this module). Some products take advantage of symmetrical multiprocessing (SMP) to boost performance. The Network Load Balancing service of Windows Server 2003 can offer fault tolerance, high availability, efficiency, and performance improvements for some software firewall products.

To increase the availability of the perimeter firewall, it can be implemented as a single firewall device with redundant components or as a redundant pair of firewalls incorporating some type of failover and/or load balancing mechanism. The advantages and disadvantages of these options are presented in the following subsections.

A single firewall without redundant components is depicted in Figure 2:


Cc700828.pfnoredundancy-0-0(en-us,TechNet.10).gif

Figure 2. Single Firewall Without Redundant Components

The use of a single firewall without redundant components offers the following advantages and disadvantages.

The advantages of a single firewall with no redundancy include:

- Low cost

  Because there is only one firewall, the hardware and licensing costs are low.

- Simplified management

Management is simplified, because there is only one firewall for the site or enterprise.

- Single logging source

  All traffic logging is central to one device.

The disadvantages of a single firewall with no redundancy include:

- Single point of failure

  There is a single point of failure for inbound and/or outbound Internet access.

- Possible traffic bottleneck

  A single firewall could be a traffic bottleneck, depending on the number of connections and the throughput required.

A single firewall tier with redundant components is depicted in Figure 3:


Cc700828.pfwithredundancy-0-0(en-us,TechNet.10).gif

Figure 3. Single Firewall with Redundant Components

Use of a single firewall with redundant components offers the following advantages and disadvantages.

The advantages of a single firewall with redundant components include:

- Low cost

  Because there is only one firewall, the hardware and licensing costs are low. The cost of the redundant components, such as a power supply, is not high.

- Simplified management

  Management is simplified because there is only one firewall for the site or enterprise.

- Single logging source

  All traffic logging is central to one device.

The disadvantages of a single firewall with redundant components include:

- Single point of failure

  Depending on the number of redundant components, there may still be a single point of failure for inbound and/or outbound Internet access.

- Cost

  The cost is higher than a firewall without redundancy, and may also require a higher class of firewall to be able to incorporate redundancy.

- Possible traffic bottleneck

  A single firewall could be a traffic bottleneck, depending on the number of connections and the throughput required.

A fault tolerant firewall set would include a mechanism to duplicate each of the firewalls, as shown in Figure 4.

 Figure 4. Fault Tolerant Firewalls

Figure 4. Fault Tolerant Firewalls

Use of a fault tolerant firewall set offers the following advantages and disadvantages.

The advantages of a fault tolerant firewall set include:

- Fault tolerance

  Using pairs of servers or devices can help provide the required level of fault tolerance.

- Central logging

  All traffic logging is central to a pair of devices with good connectivity between them.

- State sharing possible

  Depending on the device vendor, firewalls in this tier may be able to share the state of sessions between them.

The disadvantages of a fault tolerant firewall set include:

- Increased complexity

  The setup and support of this type of solution is more complex due to the multi-path nature of the network traffic.

- Complex configuration

  The separate sets of firewall rules can lead to security holes and support issues if not correctly configured.

In the preceding scenarios, the firewall could be hardware- or software-based. In the previous figures, the firewall is serving as the gateway between the organization and the Internet but the border router is placed outside the firewall. This router is extremely vulnerable to intrusion, and so it also must have certain firewall features configured. Limited firewall capabilities could be implemented without a full firewall feature set, relying on the firewall device to prevent total intrusion. Alternatively, the firewall could be consolidated within the router with no additional stand-alone firewall device.

When implementing a fault tolerant firewall set (often referred to as a cluster), there are two primary approaches, as described in the following sections.

In an active/passive fault tolerant firewall set, one device handles all the traffic while the other device does nothing. There is typically a convention through which both devices are communicating either the availability and/or the state of the connection to partner nodes. This communication is often called a heartbeat, which each system signals to the other, several times a second, to ensure connections are being handled by the partner node. When the passive node does not receive a heartbeat from the active node at a specific user-defined interval, it then assumes the active role.

An active/passive fault tolerant firewall set is depicted in the Figure 5:


Figure 5. Active/Passive Fault-Tolerant Firewall Set

Figure 5. Active/Passive Fault-Tolerant Firewall Set

The use of an active/passive fault tolerant firewall set has the following advantages and disadvantages.

The advantages of the active/passive fault tolerant firewall set include:

- Simple configuration

  This configuration is simple to set up and troubleshoot, because only a single network path is active at any one time.

- Predictable failover load

  Because the whole traffic load switches to the passive node at failover, it is easy to plan for the traffic that the passive node is expected to manage.

The disadvantages of the active/passive fault tolerant firewall set include:

- Inefficient configuration

  The active/passive fault tolerant firewall set is inefficient, because the passive node provides no useful function to the network during normal operation.

In an active/active fault tolerant firewall set, two or more nodes are actively listening to all of the requests sent to a virtual IP address that every node shares. The load is distributed between the nodes through algorithms unique to the fault tolerance mechanism in use, or through static user-based configuration, so that each node is actively filtering different traffic at the same time. In the event that one node fails, the surviving nodes distribute the processing of the load that the failed node had previously assumed.

An active/active fault tolerant firewall set is depicted in Figure 6:


Figure 6. Active/Active Fault Tolerant Firewall Set

Figure 6. Active/Active Fault Tolerant Firewall Set

Use of an active/active fault tolerant firewall set offers the following advantages and disadvantages.

The advantages of the active/active fault tolerant firewall set include:

- Greater efficiency

  Because both firewalls are providing a service to the network, this configuration is more efficient than an active/passive fault tolerant firewall set.

- Higher throughput

  During normal operation, this configuration can handle higher levels of traffic compared with the active/passive configuration, because both firewalls can provide service to the network simultaneously.

The disadvantages of the active/active fault tolerant firewall set include:

- Subject to potential overload

  If one node fails, the hardware resources on the remaining node(s) may be insufficient to handle the total throughput requirement. It is important to plan for this accordingly, understanding that performance degradation is likely to occur as the surviving nodes take on the additional workload when a node fails.

- Increased complexity

  Because the network traffic can pass through two routes, troubleshooting becomes more complex.

Security of firewall products is of paramount importance. Although there are no industry standards for firewall security, the vendor-independent International Computer Security Association (ICSA) runs a certification program aimed at testing the security of commercially available firewall products. The ICSA tests a significant number of products available in the market today (for further information, refer to www.icsalabs.com).

You must take care to ensure that a firewall achieves the requisite security standards; one way of doing this is to achieve ICSA certification. In addition, check whether your chosen firewall has an existing track record. A number of security vulnerability databases are available on the Internet; you should scan these to see how many vulnerabilities the product has been susceptible to in the past and their significance.

Unfortunately, all products (hardware- and software-based) have bugs. In addition to determining the number and severity of bugs that have affected the product you are thinking of buying, you should also assess the responsiveness of the vendor to the exposed vulnerabilities.

This section addresses the scalability requirement of a firewall solution. Scalability of firewalls is largely determined by the performance characteristics of the device, and it is wise to select a firewall that will scale to meet the scenarios it will face in practice. There are two basic ways to achieve scalability. They are:

- Vertical Scaling (Scaling Up)

  Whether the firewall is a hardware device or a software solution running on a server, varying degrees of scalability can be achieved by increasing the amount of memory, CPU processing power, and throughput of network interfaces. However, each device or server has a finite cap in terms of how far it can be vertically scaled. For example, while you may purchase a server that has sockets for four CPUs and you start with two, you will only ever be able to add two more CPUs.

- Horizontal Scaling (Scaling Out)

  Once a server has been vertically scaled to its limit, horizontal scaling becomes important. Most firewalls (hardware- and software-based) have the ability to scale out through the use of some form of load balancing. In such a scenario, multiple servers are arranged into a cluster, which is seen by the clients on the network as just one server. This scenario is essentially the same as the active/active cluster described in the "Firewall Availability" section earlier in the module. The technology used to provide this functionality may or may not be the same as that described earlier, and will be dependent on the vendor.

Scaling up hardware firewalls can be difficult. However, some hardware firewall manufacturers offer scale out solutions because their devices can be stacked to operate as a single, load balanced unit.

Some software-based firewalls are designed to scale up through the use of multiple processors. The firewall itself does not usually address multiprocessing, which is controlled by the underlying operating system. However, the firewall needs to be able to address the hardware to be able to fully use this capability. This approach allows scaling on single or redundant devices, as opposed to hardware-based firewalls, which are usually set to whatever hardware limitations are built into the device at the time of manufacture. Most firewalls are classified by the number of concurrent connections that a device can handle. Hardware devices often need to be replaced if connection requirements exceed what is available to the fixed-scale model of the device.

As discussed earlier, fault tolerance may be built into the operating system of a firewall server. In the case of a hardware firewall, fault tolerance is likely to be an extra cost.

A number of technologies are available to enhance the performance of a firewall, including:

- Gigabit Ethernet/Fiber Support
- Proxy, Reverse Web Proxy, and Caching
- SSL Off-loading Interfaces
- IPSec Off-loading Interfaces

For software-based firewalls, each one of these technologies is commercially available from multiple vendors, which keeps the costs low. While there may be similar third party solutions available for hardware devices, often they can only be obtained from the manufacturer of the hardware firewall itself.

The following sections discuss each of these performance-enhancing technologies.

Many switches, routers, and firewalls can handle Ethernet gigabit speed interfaces, and the reduction in cost of these interfaces has increased their popularity. This capability greatly reduces the likelihood of interfaces becoming bottlenecks in firewall deployments.

Typically, the caching ability is only available on software-based firewalls, because it requires the use of a disk to cache traffic or data.

SSL accelerator cards can improve the performance of publicly exposed Web sites that use SSL-based encryption by offloading the encryption processing from the CPU of the firewall. When SSL is terminated at the firewall, these devices offer significant benefits.

IPSec accelerator cards can improve the performance of publicly exposed services that use IPSec-based encryption, such as VPN. These devices offload the encryption processing from the CPU of the firewall. IPSec off-loading can be used for traffic that communicates between the internal interface of the firewall and a published service, thus ensuring that the traffic traversing the perimeter network is encrypted between the perimeter network hosts.

Consolidation means either incorporating the firewall service in another device, or incorporating other services in the firewall. Consolidation benefits include:

- Lower purchase price

  By incorporating the firewall service in another service, for example in a router, you can save the cost of a hardware device, although you must still purchase the firewall software. Similarly, by incorporating other services in the firewall, you can save the cost of additional hardware.

- Reduced inventory and management costs

  By reducing the number of hardware devices, you can reduce operating costs, since fewer hardware upgrades are required, cabling is simplified, and management is simpler.

- Higher performance

  Depending upon what consolidation is achieved, you can improve performance. For example, by incorporating Web server caching in the firewall, you may cut out additional devices, allowing the services talk to each other at high speed rather than over an Ethernet cable.

Examples of consolidation include:

- Adding firewall services to the border router

  Most routers can have a firewall service available in them. The capabilities of this firewall service may be very simple in low cost routers, but high-end routers will usually have a very capable firewall service. In practice, although you may have a separate perimeter firewall, the border router should always have its firewall service active, to protect both the router itself and the border switches.

- Adding firewall services to the border switch

  Depending upon the border switch selected, it may be possible to add in the perimeter firewall as a blade, reducing costs, and improving performance.

- Adding proxy cache to the perimeter firewall

  Proxy caching stores frequently-accessed Web pages, so that the next requestor is delivered a page from the cache rather than having to re-access the Web server, which improves response times and reduces the Web

server load. Generally, this can only be incorporated in a server firewall, as it requires a local hard disk to hold the cache.

When considering consolidating other services onto the same server or device that provides the firewall service, you should take care to ensure that the use of a given service does not compromise the availability, security, manageability, or performance of the firewall. Performance considerations are also important, as the load generated by additional services will degrade the performance of the firewall service.

An alternative approach to consolidating services onto the same device, or server hosting the firewall service, is to consolidate a firewall hardware device as a blade in a switch. This approach usually costs less than a standalone firewall of any type, and can take advantage of the availability features of the switch, such as dual power supplies. Such a configuration is also easier to manage, because it is not a separate device. In addition, it usually runs faster, because it uses the switch's bus, which is much faster than external cabling.

Most Internet protocols that use version 4 of the Internet Protocol (IPv4) can be protected by a firewall, including lower-level protocols such as TCP and UDP, and higher-level protocols such as HTTP, SMTP, and FTP. Any firewall product under consideration should be reviewed to ensure that it supports the required type of traffic. Some firewalls can also interpret GRE, which is the encapsulation protocol for the point-to-point tunneling protocol (PPTP) used in some VPN implementations.

Some firewalls have built-in application layer filters for protocols such as HTTP, SSL, DNS, FTP, SOCKS v4, RPC, SMTP, H. 323, and post office protocol (POP).

This module has provided a practical process for the successful selection of firewall products. This process covers all aspects of firewall design, including the various evaluation and classification processes required to reach a solution.

No firewall is 100% safe: the only way to ensure that your network cannot be attacked electronically from the outside is to implement an air gap between it and all other systems and networks. The result would be a secure network that is virtually unusable. Firewalls enable you to implement an appropriate level of security protection when connecting your network to an external network, or when joining two internal networks.

The firewall strategies and design processes outlined in this module should be considered only as part of an overall security strategy, because a strong firewall is of limited value if there are weaknesses in other parts of the environment. Security must be applied to every component of the network, and a security policy that addresses the risks inherent in the environment must be defined for every component.

You can find further information about design and deployment of firewall services from the following URLs.

- For an overview of firewalls:

  www.microsoft.com/technet/security/guidance/networksecurity/firewall.mspx

- For detailed security information on Microsoft Windows Server 2003, refer to the "Windows Server 2003 Security Center" document:

  https://www.microsoft.com/technet/security/prodtech/windowsserver2003.mspx

- For information on Microsoft Internet Security & Acceleration Server firewall and Web proxy product, refer to:

  https://www.microsoft.com/isaserver/

- For a free e-mail notification service that Microsoft uses to send information about the security of Microsoft products to subscribers, visit the Microsoft Security Notification Service Web site:

  www.microsoft.com/technet/security/bulletin/notify.mspx

- The SANS (SysAdmin, Audit, Network, and Security) Institute security resources are available from:

  https://www.sans.org

- The Computer Emergency Response Team (CERT) organization records and publishes security alerts and a center for security expertise at:

  https://www.cert.org

Download the Complete Solution

Windows Server System Reference Architecture
