

DaVinci Academy of Science and The Arts

Policy Number: 706

Policy Section: 700 – Technology

POLICY TITLE: Internal Firewall Design and Security

Revision History

Effective Date	Action Date	Revised
8 October 2008	New Policy	New Policy

Electronic Information Security Policy
Effective Date: 17 September 2008
Revision Date:

1. OVERVIEW

This document is for the DaVinci Academy and its appliance purchases specifically for firewalls and security related devices. It should help you to select a firewall product(s) suitable for the DaVinci Academy internal network. It presents the different classes of firewall available and highlights significant features. It also outlines design guidelines which will enable us to determine our own requirements and select the most appropriate product.

2. OBJECTIVES

1. Identify the features necessary in your internal firewall.
2. Classify firewall products.
3. Select the best firewall product for your internal firewall.

3. APPLIES TO

Ethernet/IP based firewall products

4. DESIGN GUIDELINES

This document considers the requirements for an internal firewall in a network, the types of devices which can meet those requirements and the options available for their deployment. Unfortunately intrusions into networks **both from external and internal users** have become a regular event, which means that the DaVinci Academy must install protection from these intrusions. A firewall costs money and creates an impediment to traffic flow. You must therefore ensure that our firewall is designed to be as cost effective and efficient as possible.

5. NETWORK ARCHITECTURE

5.1 In a network architecture there will generally be three zones:

1. **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. It feeds data through to the perimeter network via a perimeter firewall
2. **Perimeter network:** This network, often called the DMZ (demilitarized 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
3. **Internal Networks:** The internal networks link the internal servers, such as SQL Server and the internal users.

5.2 In an organization there will frequently be two different firewalls, the perimeter firewall and the internal firewall. Although the tasks of these firewalls are similar, they also have a different emphasis as the perimeter firewall focuses on providing a limitation to untrusted external users, whereas the internal firewall focuses on preventing external users accessing the internal network and limiting what internal users can do. For further information on perimeter firewall design see the networks depicted in Figure 1.




Figure 1 Network architecture

6. DESIGN INPUTS

6.1 A firewall checks incoming IP packets and blocks those it detects as intrusive. Some blocking can be done by recognizing by default that certain packets are illegal. Alternatively, you can configure the firewall to block certain packets. The TCP/IP protocol was designed many years ago without any concept of hacking or intrusion and has many weaknesses. For example, the ICMP protocol was designed as a signaling mechanism within TCP/IP but this is open to abuse and can lead to such problems as denial of service attacks. An internal firewall has more exacting requirements than a perimeter firewall. This is because internal traffic is more difficult to control as its legitimate destination may be any server in the internal network.

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

6.2a Budget

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 potential 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.

6.2b Existing Facilities

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.

6.2c **Availability**

Does our organization require the firewall to be available at all times? If you are offering a public Web server facility with constant availability then you will require almost 100% uptime. With all firewalls there is always a chance of failure so how can you mitigate against that? The availability of a firewall can be improved by two methods:

- **Redundant components:** Duplicating some of the 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 cannot have any redundant options as it is expensive to add resilience, particularly as it usually adds no more 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 the firewall connects to. However, depending on the firewalls, it may also double the throughput to compensate. In theory all firewalls from the smallest to the largest could be duplicated, but in practice it also needs a software switch over mechanism which may not be present in smaller firewalls.

6.2d **Scalability**

What are the throughput requirements of the firewalls? 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. How can we handle an increase? We must select a firewall solution that can scale up as the throughput increases. Can the firewall grow by adding more components, or could you install another firewall in parallel?

6.2e **Features Required**

Which firewall features are required? Based on risk assessments conducted against the services provided in the organization, you can determine which types of firewall features are required to protect the assets that provide the services. Will VPNs (Virtual Private Networks) be required, as this affects the design?

7. SYSTEM ATTACKS AND DEFENSE

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

7.1 External Attacks: The Internet is often used as a tool by people 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 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.

7.2 Internal Attacks: Not all attacks are Internet-based; you must also protect sensitive information from internal users of the network. Most organizations have sensitive information that should be protected from certain users (Grades, student names and other personal items) on the internal network, including employees, vendors, contractors, and customers.

7.3 Intrusion Threats: 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. However, after discovering the presence of a server the hacker might attempt a more serious attack. This means that all intrusions should be considered potentially harmful. Some of the major intrusions are:

7.3a 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 displayed by sniffer products. This means that 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 and many of the potential sniffing intruders will be your own users, inside a firewall. Free sniffer software can be readily downloaded from the Internet and your users could be running it on their PCs, examining packets as they pass. If you are running Microsoft Windows operating systems on your PCs, users would normally require administrator access to run a sniffer so that limits the number of users who may attempt to sniff. However, your administrator users, who may be numerous, would be able to run a sniffer. Apart from accessing confidential data they may see clear text passwords, as mentioned above. As many people will use the same password for every application, intruders can infer what encoded passwords will be and gain further access. There are various measures to counter sniffing. The primary measure is the use of strong encrypted passwords, but this is beyond the scope of this document.

7.3b IP Spoofing: IP spoofing occurs when the source address of an IP packet is changed to hide the identity of the sender. The routing operation in the Internet only uses the destination address to send a packet on its way and ignores the source

address. A hacker can therefore send a destructive packet to your system and disguise the source so that you do not know 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 (DoS) attacks, which are described later in this document.

7.3c Denial of Service Attacks: DoS attacks are among the hardest to prevent. DoS attacks are different from other types of attacks in that they do not cause permanent damage to your network. Instead, they try to stop the functioning of the network by bombarding a particular computer (either server or network device), or by degrading the throughput of network links to a point where performance is sufficiently abysmal to cause ill-will among customers and loss of business to the organization. A distributed DoS (DDoS) is an attack initiated from many other computers concentrating the bombardment on your system. The attacking computers have not initiated the attack themselves but have allowed themselves to be infiltrated due to their own security vulnerabilities.

7.3d 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. This means that the best defense is usually 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. It exploited a known problem in Microsoft SQL Server™ 2000 for which Microsoft had issued a fix in August 2002. This worm exploited the fact that many administrators had not applied the recommended update and did not have adequate firewalls in place (which could have dropped packets destined for the port that the worm used). 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.

7.3e 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. Network reconnaissance does no harm itself. However, discovering which addresses are in use can help someone to launch a hostile attack. 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, network basic input/output system (NetBIOS), Hypertext Transfer Protocol (HTTP), and Simple Mail Transport Protocol (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 a firewall, but turning off some of these services may restrict your network diagnostics capabilities.

8. DEVICE DEFINITION

8.1 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.

8.2 An internal firewall generally provides 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 typically does not 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.

9. FIREWALL FEATURES

9.1 Depending on the features that a firewall supports, traffic is allowed or blocked using a variety of techniques. These techniques offer different degrees of protection based on the capabilities of the firewall.

9.2 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
- Application layer filtering

In general, firewalls that provide complex features will also support simpler features. However, you should read vendor information carefully when choosing a firewall because there can be subtle differences between the implied and the actual capability of a firewall. Selection of a firewall typically involves inquiring about the features as well as testing to ensure that the product can indeed perform according to specifications.

10. NETWORK ADAPTER INPUT FILTERS

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 and port numbers for UDP and TCP, as well as the protocol of the traffic, TCP, UDP, and generic routing encapsulation (GRE). Network adapter input filtering allows a quick and efficient denial of standard incoming packets that meet the rule criteria

configured in the firewall. However, it can easily be evaded, as it only matches headers of the IP traffic and works on the basic assumption that the traffic being filtered follows IP standards and is not crafted to evade the filtering.

11. STATIC PACKET FILTERS

11.1 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 inbound as well as outbound communications to an interface. In addition, 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 if the traffic coming into the interface appears to be return traffic based on the conventions of the IP headers.

11.2 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.

12. NETWORK ADDRESS TRANSLATION

12.1 In the worldwide IP address range, certain address ranges are designated as *private addresses*. These address ranges 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.

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

13. STATEFUL INSPECTION

13.1 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 slightly slower than static packet filtering. However, it ensures that traffic is only allowed to pass 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.

13.2 Certain firewalls may store more information (such as IP fragments sent and received) in the state table while others store less. 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 must read the firewall documentation carefully. The stateful inspection feature typically assists in mitigating the risk posed by network reconnaissance and IP spoofing.

14. CIRCUIT-LEVEL INSPECTION

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

15. APPLICATION LAYER FILTERING

15.1 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. This processing includes inspecting, screening or blocking, redirecting, and modifying the 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). Typically, third-party tools for content screening such as virus detection, lexical analysis, and site categorization can be added to your firewall.

15.2 An application layer firewall has the ability to inspect many different protocols based on the traffic that passes through it. Unlike a proxy firewall that usually inspects Internet traffic such as HTTP, FTP download, and SSL, the application layer firewall has much greater control over the way any traffic travels through it. 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 the traffic originates internally.

15.3 The application layer firewall feature assists in mitigating the risk posed by IP spoofing, DoS, some application layer attacks, network reconnaissance, and virus/Trojan horse attacks. The drawbacks of an application layer firewall are 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 application layer firewalls is determining what the firewall is capable of doing at the application layer. Application layer filtering is widely used for protecting publicly exposed services. If your organization has an online store that collects credit card numbers and other personal information about customers, it is prudent to take the highest level of precautions in protecting this information. The application layer feature ensures that the traffic being

passed over a port is appropriate. Unlike packet filter or stateful inspection firewalls, which simply look at the port and source and destination IP address, firewalls that support the application layer filtering feature have the ability to inspect the data and the commands being passed back and forth.

15.4 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 to look for DNS-specific commands as it goes through the firewall. This additional layer of protection ensures that users or attackers cannot conceal information in allowed types of traffic.

16. FIREWALL CLASSES

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

16.2 Grouping firewalls into classes allows for the abstraction of the hardware from the requirements of the service. Service requirements can then be matched against class features. As long as a firewall fits into a specific class, it can support all of the services that the class of firewalls is required to support.

16.3 The various classes are as follows:

- Class 1 - Personal firewalls
- Class 2 - Router firewalls
- Class 3 - Low-end hardware firewalls
- Class 4 - High-end hardware firewalls
- Class 5 - High-end server firewalls

16.4 It is important to understand that some of these classes overlap. This is by design, as 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 your organization can select a model that suits its present and future requirements. Apart from the price and feature set, firewalls can be classified on the basis of performance (or throughput). However, manufacturers do not provide any figures of throughput for most classes of firewalls. Where they are provided (typically 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.

16.5 The following subsections define firewall classes in detail.

16.5a Class 1-Personal Firewall

A personal firewall is defined as a software service running in an operating system that provides 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 personal 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, a personal firewall has limited performance and will degrade 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 in general a lower level of protection is needed on a personal computer.

Personal firewalls may come free-of-cost in an operating system or at a very low cost. They are suitable for their intended purpose but should not be considered for use in an enterprise, large businesses or large educational facilities even for small satellite offices, due to their restricted performance and functionality. They are, however, particularly suitable for mobile users on laptop computers.

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

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
Block or allow incoming ICMP messages	Yes

Firewall Attribute	Value
Control outgoing access	Yes
Application protection	Possibly
Audible or visible alerts	Possibly
Log file of attacks	Possibly
Real-time alerts	Depends on the product
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)

Advantages

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 Windows XP. 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 workable out-of-the-box configurations with straightforward configuration options.

Disadvantages

The disadvantages of personal firewalls include:

- **Difficult to manage centrally**
Personal firewalls need to be configured on every client, which adds to 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 single personal computers. Using them on a personal computer that serves as a router for a small network will lead to degraded performance.

16.5b Class 2-Router Firewall

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 that 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 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

Firewall Attribute	Value
Real-time alerts	In many cases
VPN Support	Often in lower-end routers, not as common in higher-end routers. Separate dedicated devices or servers for this task are available
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

Advantages

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 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 and no retraining is required. Network cabling is simplified because no additional hardware is installed, which also simplifies network management.

Disadvantages

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 this configuration 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 performance of the router, especially when it is already under attack.

16.5c Class 3-Low-end Hardware Firewall

At the low end of the hardware firewall market are Plug and Play units requiring little or no configuration. These devices often incorporate switch and/or VPN functionality. Low-end hardware firewalls are suitable for small businesses and 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 if you require one.

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
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

Firewall Attribute	Value
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

Advantages

The advantages of low-end hardware firewalls include:

- **Low cost**
Low-end firewalls can be purchased inexpensively.
- **Simple Configuration**
Almost no configuration is required.

Disadvantages

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 email and/or a Web site.
- **Limited upgradeability**
Usually there can be no hardware upgrades, though there are often periodic firmware upgrades available.

16.5d Class 4-High-end Hardware Firewall

At the high end of the hardware firewall market there are high performance, highly resilient products suitable for the enterprise, large businesses, large educational facilities 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 that maintains the current table of connections through automatic stateful synchronization.

You should use firewalls 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 head office 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
High-availability option available	Yes
Number of concurrent sessions	> 7500-500,000
Modular upgradeability (hardware or software)	Yes
Price range	High

Advantages

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 which usually only filter at Layer 3 and perhaps Layer 4 of the OSI model, high-end hardware firewalls provide filtering at Layers 5 through 7 for well-known applications.

Disadvantages

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 and is often based on the number of concurrent sessions, throughput, and availability requirements.

Complex configuration and management

Because this class of firewalls has much greater capability than low-end firewalls, it is also more complex to configure and manage.

16.5d Class 5-High-end Server Firewall

High-end server firewalls add firewall capability to a high-end server, providing robust fast protection on standard hardware and software systems. The benefit of this approach is the use of familiar hardware or software. This 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 industry standard hardware platforms with industry standard operating systems running on

them (but hidden from view) and therefore have little difference, technically and 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 the 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, as 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
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/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

Advantages

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 features in the operating system they run on. 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 also 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.

Disadvantages

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.

17. INTERNAL FIREWALL USAGE

17.1 An internal firewall exists to control access to and from the internal network. User types may include:

- **Trusted:** Employees of the organization, which can be internal users going out to the perimeter zone or the Internet, external users such as branch office workers, remote users, or users that work from home.
- **Semi-trusted:** Business partners of the organization, for which a higher level of trust exists than with untrusted users. However, it is often a lower level of trust than that for the organization's employees.
- **Untrusted:** For example, users of the organization's public Web site.

17.2 Untrusted users from the Internet should, in theory, only be accessing your Web servers in your perimeter zone. If they need access to your internal servers, for example to check stock levels, the trusted Web server makes the enquiry on their behalf, so untrusted users should never be allowed through the internal firewall.

17.3 There are a number of issues that you should consider when you select the firewall class that is going to be used in this capacity. The following table highlights these issues.

Table 6. Internal 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 firewall serving in this capacity, some are not willing to accept the accompanying reduction in performance. For very high-volume non e-commerce Web sites for example, lower levels of security may be allowed 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, or be a logical firewall on a physical device	This depends upon the performance required, the sensitivity of the data and how frequently access is required from the perimeter zone.
Manageability requirements for the device as specified by the organization's management architecture	Some form of logging is generally used, while an event monitoring mechanism is also usually required. You may choose not to allow remote administration here to prevent a malicious user from remotely administering the device.
Throughput requirements will likely be determined by the network and service	These will vary for each environment, but the power of the hardware in the device or server and the

Issue	Typical Characteristics of a Firewall Implemented In This Capacity
administrators within the organization	firewall features being used will determine the overall network throughput available.
Availability requirements	Again this depends upon the access requirements from the Web servers. If they are primarily for handling information requests which are satisfied by providing Web pages, the flow to the internal networks will be low. However, high levels of availability will be required for an e-commerce situation.

18. INTERNAL FIREWALL RULES

18.1 Internal firewalls monitor traffic between the perimeter and internal zones of trust. The technical requirements for the internal firewalls are considerably more complex than those of the perimeter firewalls due to the complexity of the traffic types and flows between these networks.

18.2 In this section reference is made to "bastion hosts." Bastion hosts are servers located in your perimeter network that provide services to both internal and external users. Examples of bastion hosts include Web servers and VPN servers. Typically your internal firewall will need the following rules implemented, either by default or by configuration:

- Block all packets by default.
- On the perimeter interface, block incoming packets that appear to have originated from an internal IP address to prevent spoofing.
- On the internal interface, block outgoing packets that appear to have originated from an external IP address to restrict an internal attack.
- Allow UDP-based queries and response from the internal DNS servers to the DNS resolver bastion host.
- Allow UDP-based queries and responses from the DNS resolver bastion host to the internal DNS servers.
- Allow TCP-based queries from the internal DNS servers to the DNS resolver bastion host, including responses to those queries.
- Allow TCP-based queries from the DNS resolver bastion host to the internal DNS

servers, including responses to those queries.

- Allow zone transfers between the DNS advertiser bastion host and the internal DNS server hosts.
- Allow outgoing mail from the internal SMTP mail server to the outbound SMTP bastion host.
- Allow incoming mail from the inbound SMTP bastion host to the internal SMTP mail server.
- Allow traffic originating from the back end on the VPN servers to reach internal hosts and the responses to return back to the VPN servers.
- Allow for authentication traffic to the RADUIS servers on the internal network and the responses to return back to the VPN servers.
- All outbound Web access from internal clients passes through a proxy server and the responses are returned to them.
- Support Microsoft Windows 2000/2003 domain authentication traffic between network segments for both the perimeter domain and the internal domain.
- Support at least five network segments.
- Perform stateful inspection of packets between all network segments that they join (Circuit Layer Firewall - Layer 3 and Layer 4).
- Support high availability features like stateful failover.
- Route traffic between all connected network segments without using Network Address Translation.

19. HARDWARE REQUIREMENTS

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

- **Hardware-based firewall:** These devices typically run specialized code on a custom-built hardware platform. These firewalls 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 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. 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.

20. AVAILABILITY

20.1 To increase the availability of the firewall, it can be implemented as a single firewall device with or without 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.

20.1a SINGLE FIREWALL WITHOUT REDUNDANT COMPONENTS

A single firewall without redundant components is depicted in the following figure:

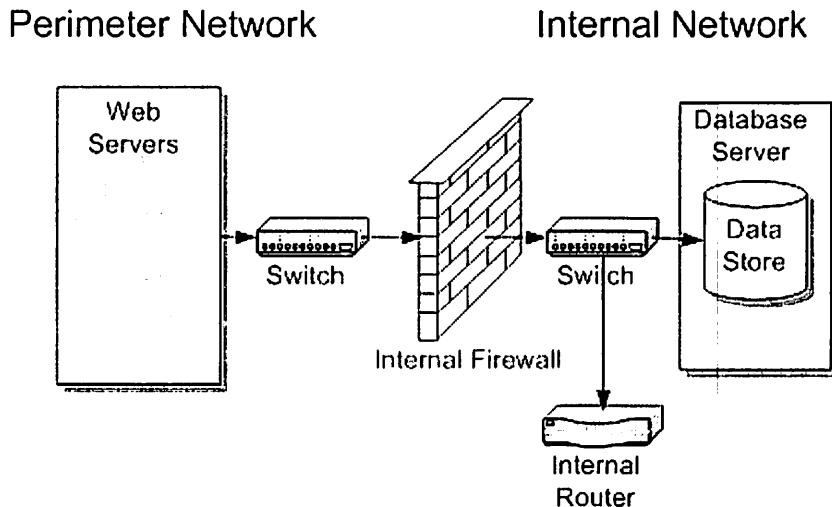


Figure 2 Single firewall without redundant components

Advantages

The advantages of a single firewall 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.

Disadvantages

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

20.1b SINGLE FIREWALL WITH REDUNDANT COMPONENTS

A single firewall with redundant components is depicted in the following figure:

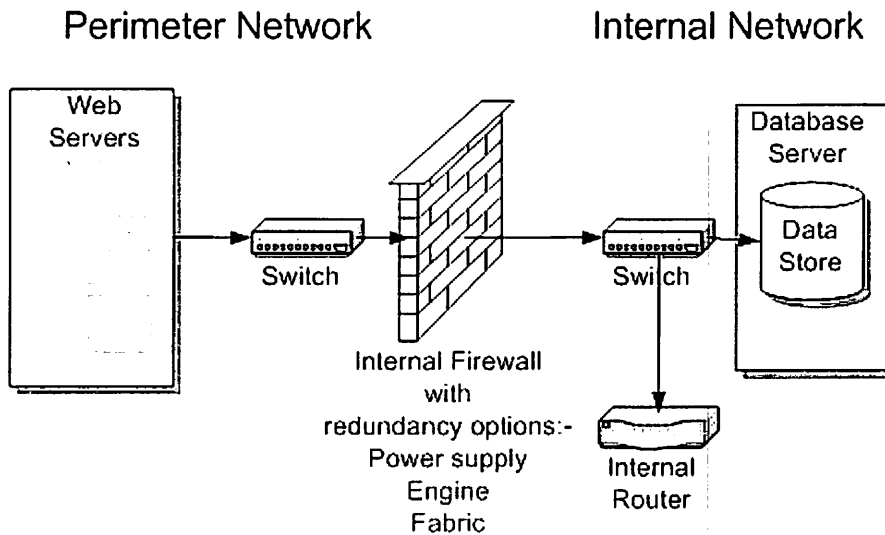


Figure 3 Single firewall with redundant components

Advantages

The advantages of a single firewall 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.

Disadvantages

The disadvantages of a single firewall 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 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 throughput required.

20.1c FAULT TOLERANT FIREWALLS

A fault tolerant firewall set would include a mechanism to duplex each of the firewalls as in the following figure.




Figure 4 Fault tolerant firewalls

Advantages

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 traffic logging
Traffic logging is more reliable as either or both firewalls may be logging activity to the other partner or a separate server.

State sharing possible

Depending on the product, firewalls in this set may be able to share the state of sessions.

Disadvantages

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 multipath 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.

Increased cost

As at least two firewalls are required the cost increases over a single firewall set.

20.2 **Fault Tolerant Firewall Configurations:** When implementing a fault tolerant firewall set (often referred to as a cluster), there are two primary approaches, as described in the following sections.

20.3 **Active/Passive Fault Tolerant Firewall Set:** In an active/passive fault tolerant firewall set, one device (also referred to as an active node) handles all the traffic while the other device (the passive node) forwards no traffic nor performs filtering but remains active monitoring the state of the active node. Typically, each node communicates its availability and/or the state of its connection to its partner node. 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. If the passive node does not receive a heartbeat from the active node for longer than a specific, user-defined interval, indicating the active node has failed, then the passive node will assume the active role. An active/passive fault tolerant firewall set is depicted in the following figure.




Figure 5 Active/passive fault-tolerant firewall set

Advantages

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

- Simple configuration
This configuration is simpler to set up and troubleshoot than the next option, active/active, 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, the traffic that the passive node is expected to manage can easily be planned for.

Disadvantages

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

- Inefficient utilization
The active/passive fault tolerant firewall set is inefficient because the passive node provides no useful function to the network during normal operation and does not increase throughput.

20.4 Active/Active Fault Tolerant Firewall Set: In an active/active fault tolerant firewall set, two or more nodes actively listen to all the requests sent to a virtual IP address that every node shares. Load is distributed between the nodes through algorithms unique to the fault tolerance mechanism in use, or through static user-based configuration. Whatever the method, the result is that each node actively filters different traffic. 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 the following figure:




Figure 6 Active/active fault tolerant firewall set

Advantages

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

- Greater efficiency
Because all firewalls are providing a service to the network, their usage is more efficient.
- Higher throughput
During normal operation, this configuration can handle higher levels of traffic compared to the active/passive configuration, because all firewalls can provide service to the network simultaneously.

Disadvantages

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 network traffic can pass through multiple routes, troubleshooting becomes

more complex.

21. SECURITY

21.1 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 the firewall products available in the market today. For further information, refer to the following URL: www.icsalabs.com

21.2 Care must be taken to ensure that a firewall achieves the requisite security standards, and one way of doing this is to choose a firewall that achieves ICSA certification. In addition, a track record should exist for the firewall that is chosen. A number of security vulnerability databases are available on the Internet. You should check these for information concerning the vulnerabilities of the product you are thinking of purchasing. 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, it is also important to assess the responsiveness of the vendor to the exposed vulnerabilities.

22 SCALABILITY

22.1 This section addresses the scalability requirement of a firewall solution. Scalability of firewalls is largely determined by the performance characteristics of the devices used. It is wise to select a type of firewall that will scale to meet the scenarios it will face in practice.

22.2 There are two basic ways to achieve scalability:

- 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, if you 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 and seen as one by the clients on the network. This scenario is essentially the same as the active/active cluster described in the "Availability" section in this document. The technology used to provide this functionality may or may not be the same as that which was described earlier, and will be dependent on the vendor.

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

22.4 Some software-based firewalls are designed to scale up through the use of multiple processors. Multi-processing is controlled by the underlying operating system and the firewall software does not need to be aware of additional processors; however full benefits of the multiple processors may not be achieved unless the firewall software can operate in a multi-tasking fashion. This approach allows scaling on single or redundant devices as opposed to hardware-based or device-type firewalls, which must typically adhere to whatever hardware limitations are built into them at the time of manufacture. Most device-type firewalls are classified by the number of concurrent connections that the devices can handle. Hardware devices often need to be replaced if connection requirements exceed what is available to the fixed scale model of the device.

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

23. CONSOLIDATION

23.1 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, the cost of a hardware device is saved, although the firewall software must still be purchased. Similarly, if other services can be incorporated in the firewall, the cost of additional hardware is saved.
- Reduced inventory and management costs: Reduction in the number of hardware devices leads to reduced operating costs. Since fewer hardware upgrades are required, cabling is simplified and management is simpler.
- Higher performance: Depending upon what consolidation is achieved, performance may be improved. For example, incorporating the Web server caching in the firewall may cut out additional devices, and the services talk to each other at high speed rather than over an Ethernet cable.

23.2 Examples of consolidation include:

- Adding firewall services to a router: Most routers can have a firewall service incorporated 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. You will probably have at least one router linking Ethernet segments together in your internal network. By incorporating the firewall in it you will save costs. Even if you implement specific firewall devices, implementing some firewall features in your routers may help limit internal intrusions.

- Adding firewall services to the internal switch: Depending upon the internal switch you use, it may be possible to add in the internal firewall as a blade, reducing costs, and improving performance.

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

23.4 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 does not involve a separate device. In addition, this solution usually runs faster because it uses the bus in the switch, which is much faster than using external cabling.

24. STANDARDS AND GUIDELINES

24.1 Most Internet protocols that use version 4 of the Internet protocol (IPv4) can be protected by a firewall. This includes both 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.

24.2 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). We should also consider the future of the TCP/IP protocol and IPv6 and whether this should be a mandatory requirement for any firewall, even if you are using IPv4 currently.

25. SUMMARY

25.1 This document 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.

25.2 No firewall is 100 percent safe. The only way to ensure that our 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 us to implement an appropriate level of security protection when connecting your network to an external network, or when joining two internal networks.

25.3 The firewall strategies and design processes outlined in this document should be considered only as part of an overall security strategy. A strong firewall is of limited value if there are weaknesses in other parts of the network. 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.

26. REFERENCES

Further information about design and deployment of firewall services may be found at the following URLs:

For information on Microsoft Internet Security & Acceleration Server firewall and Web cache product, refer to the following URL:

www.microsoft.com/isaserver/

For a free email notification service that Microsoft uses to send to subscribers information about the security of Microsoft products, visit the Microsoft Security Notification Service Web site at the following URL:

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

The SANS (SysAdmin, Audit, Network, and Security) Institute security resources are available at the following Web site:

www.sans.org

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

www.cert.org