



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Migrating Services Between Firewall Technologies

Critical considerations to be addressed when moving high availability services to new firewall technologies.

**Andrew Barratt
GSEC: GIAC Security Essentials Certification
Version 1.4b
Option 1
Submitted 17 July 2003**

Table of Contents

Abstract	3
Terminology	3
Why Replace the Technology?	3
Pre-requisite Information	4
Service Detail	4
Technical Detail	4
Contact Information	4
Get the Business Onside	4
Components within the Architecture	5
IP Strategy	5
To Re-IP or Not To Re-IP: Big Bang vs. Re-addressing	5
Big Bang	5
Re-addressing	6
User Acceptance Testing	7
Technical Capabilities of Different Firewall Technologies	7
Rule Set Implementation	8
Network Address Translation	8
Routing	8
Spoofing Rules	9
Performance of ASIC Technology	9
Session State and High Availability	10
Management Capability for Distributed Environments	10
Appendix A – Desirable NAT Features	11
Appendix B – Example of an Unexpected Complication	12
The Service Migration	12
The Problem from a Technical Perspective	12
The Solution	13
The Key Learns	14
References	15

Abstract

This paper describes the considerations that are essential to address when a corporate firewall infrastructure is replaced with new technology. The focus is on a business environment where services pass through high availability firewalls and any loss of service can immediately result in significant financial loss or worse. This paper has deliberately been written independent of specific firewall products. This is not a step-by-step guide. It is an aid to assist firewall administrators with project management issues and technical issues that could result in unplanned outages.

Fundamental to successful firewall service migrations is a sound knowledge of current security principles as they relate to the company security policy, technical understanding of the security technologies involved, understanding of the business dependency on the firewall infrastructure, and close management of all changes including potential impact.

Terminology

The following terminology is used within the context of this paper:

Firewall Appliance – The term ‘firewall appliance’ within this paper refers to a device manufactured for use as a firewall only. It has no other potential use and is designed with hardware technology such as ASIC to maximise performance of features such as VPN encryption.

Security Policy – The term ‘security policy’ within this paper is intended to incorporate a company security policy, standards, notices and guidelines. This is not in-line with the strict definitions of these documents, which therefore will be mentioned. A security policy constitutes minimum-security requirements at a high level. Standards detail metrics, products and processes for satisfying policies. Notices specify changes to existing policies/standards, or provide information on new policies/standards. Guidelines are procedures recommended for adherence to a policy.

Service – The term ‘service’ within this paper refers to a firewall policy rule, or collection of rules that define access required by an application or group of people for a single purpose. For example, a firewall service may be all rules that allow a business team to connect their PC clients to a third party’s database server farm. This is not in-line with the definition of a service many vendors may use, which relates to a single IP protocol.

Why Replace the Technology?

“There are security considerations within the TCP/IP protocol suite. To some people these considerations are serious problems, to others they are not; it depends on the user requirements.”¹ A responsible IT employee will ensure that management is aware of key threats and developments in IT security. Management must be in a position to make informed decisions about updating and enforcing security policy.

¹ Socolofsky, C. Kale, C RFC1180.

Firewall technology is used to enforce the aspects of the security policy that relate to perimeter security. A decision to replace the technology should be the result of either meeting the requirements of an updated security policy or meeting the future business requirements that adhere to the security policy and associated standards. If this is not the case, then either the security policy and/or standards require review or the deployment of new technology may be an unnecessary risk.

Pre-requisite Information

The following information should be found in documentation produced prior to migration. The documentation should be stored in more than one secure location.

Service Detail

1. Ownership of for each firewall service must be defined. The owner of a service should have a vested interest in service availability and agree responsibility for testing that the service is functioning after migration. This is key to managing the impact of changes. A firewall administrator should only own a service if the service relates to firewall infrastructure management.
2. The business impact of service outage is required for risk analysis. The firewall administrator requires this before a migration strategy can be adopted. The service owner should determine business impact.
3. Time windows within which each service can be unavailable allow for the scheduling of migrations must be clear.

Technical Detail

1. Topology diagrams for each service should be created.
2. Full firewall configuration details should be accessible offline.
3. An IP address table listing all devices the firewalls communicate with, including routers and listing the source/destination of each service connection should be current.

Contact Information

1. Contact details for the service owner and those assigned the task of testing the service after migration should be available.
2. Contact details for third parties potentially impacted by a service migration should be current. These should include at least a technical point of contact and escalation details in case the primary contact is unreachable.
3. Contact details for vendor support including hours of operation are essential.

Appendix B provides an account of how the availability of above information is essential to risk management and can prevent service outages.

Get the Business Onside

A request to the owner of a service that he/she should coordinate testing within a prescribed change window is not likely to receive a great response. The investment in time and resource should be encouraged through

discussion of the benefits of the new firewall infrastructure in terms of security, availability, functionality and cost. Advising the key parties of their involvement in the project at an early stage will better allow them to provide input to the schedule. Ensuring that there is awareness of senior support for the project is essential to success.

Components within the Architecture

There are many functions that some firewall technologies can provide. They may include features that implement packet inspection, network address translation, VPN, encryption, authentication, authorisation, accounting, intrusion detection, proxy services, web caching, load balancing and more. A single technology may provide a good solution to one function, but a limited solution to another.

If budget allows, separating the implementation of different functions using more than one type of device may provide strong security with the required functionality. It can provide strong defence in depth. It can also allow for replacement of only specific components of the architecture when business requirements require that specific functions need to be improved. To divide the infrastructure into components can minimise the risk of changes, but is an expensive solution in terms of investment in technology and the associated support cost.

IP Strategy

There are basically two strategies for allocating IP addresses that can be used when migrating a complex set of services from one set of firewalls to another. Each has pros and cons. Combinations of the two strategies can also be used. This is an area where a bit of forethought and even creativity can reduce risk and/or effort. The best strategy depends on the existing topology, the acceptable risk, the capability of the firewalls utilised, the available IP space, and the time available to implement the migration.

To Re-IP or Not To Re-IP: Big Bang vs. Re-addressing

It must be decided if new IP address space will be used for the new firewall interfaces, the networks directly connected to the new firewalls, and for the NAT address space. Note that this decision must be made regardless of whether firewalls have IP addresses assigned to operational interfaces or are running in 'transparent mode' without displaying IP addresses to the network.

Big Bang

Basically, the big bang strategy refers to unplugging the existing firewalls and connecting the new, pre-configured firewalls using the same cables. In the utopian world, no other devices require configuration and all services work through the new firewalls just as they did the old firewalls (with the exception of any improvements). There is no requirement to change the IP addresses of

any other devices connected to the network. Due to the efficiency of using this strategy, it is the most commonly chosen option.

This strategy requires that the rule set on the new firewalls mirrors exactly the rule set on the existing firewalls, including the IP address assignments, spoofing rules and NAT particulars. Additionally, all other utilised functions of the new firewalls need to operate with the same behaviour as the original firewalls. With the exception of improvements, the 'black box' is unchanged.

The big bang strategy can be the most efficient strategy to use when replacing a firewall infrastructure. A migration of all services can be completed in one change window. However, efficiency comes at a cost. The big bang strategy can be a relatively high-risk operation for complex environments. If any service does not function correctly, and is not be fixed within the window of time allocated for the change, then the entire migration needs to be backed out. Additionally, coordinating with service owners the testing of each service may not be possible in the scheduled time frame. This depends entirely on how many services must be tested.

A worse scenario is also possible. If there are problems encountered only when the service enters production mode, such as open of business Monday morning, then the firewall administrator has a difficult situation. The options are either to correct the problem or back out the entire migration. Backing out the entire migration affects all services at a time that will cause an outage. In some business environments, the time frame required to take either of these options may not be acceptable and will cause significant financial loss to the business or worse.

A tip for those implementing this strategy is to be aware that some devices connected to the same networks as the new firewalls, may not efficiently learn the new firewall interface MAC addresses matching the reused IP addresses. A change to the ARP tables on these devices may need to be forced by sending approximately five ping packets to the device from the new firewalls. Firewalls that implement gratuitous ARP should not see this issue.

Re-addressing

The re-addressing strategy requires use of new IP address space for the firewall interfaces, the networks directly connected to the firewalls, and for the NAT address space. It allows the new infrastructure to be deployed in parallel to the existing infrastructure. In high availability environments, the lower risk of this option may prove the most significant consideration leading to adoption of the strategy. Appendix B is an example of a service migration implemented in this way.

Each independent service can be migrated separately where topology allows. This means that the firewall migration can be staged. Each stage can involve less risk than a big bang strategy. The cost of this strategy is high in terms of resource required and timeframe for migration of all services.

When a service is migrated, devices on the directly connected networks are moved to the segments used in the new infrastructure. Therefore there may be a requirement for a third party to change IP addresses of devices connected to the same networks as the firewalls. A third party may also need to make routing changes on the third party network to reach the newly assigned IP addresses.

All devices used as components of an independent service topology that are directly connected to the re-addressed networks are physically repatched during the change window.

Firewall policy and firewall routing changes can be made in advance of service migrations. Routing changes to the business's internal routers can be done in advance if new NAT ranges are in use. The old IP address ranges can be released on completion of all migrations.

User Acceptance Testing

The most important phase of a firewall service migration is the testing phase. The owners of the service or applications that use the service should be organised to do a full service test if possible. The firewall administrators should not do this. The responsibility for testing and therefore acceptance of the new technology should lie in the hands of the service owner. Appendix B provides an account of how a strong test plan can prevent outages.

Technical Capabilities of Different Firewall Technologies

In order for the any strategy to work, the new firewalls should support a superset of all features currently utilised on the existing firewalls. If this is not the case, then it may be possible to make changes to other devices communicating via the firewalls to create an acceptable solution. If this is not an option, then network topology redesign may be a solution. For example, strategically placing an additional router with the required functionality can solve a technical problem, and also provide defence in depth. If there is no work-around, then the choice of firewall architecture should be reconsidered.

There are many different firewall technologies available, each with some unique technical characteristics. The best technology for a company can only be selected once the current and future company requirements are analysed. The technology must be capable of enforcing the parts of the firewall policy relevant to perimeter control. The technology must be flexible enough to meet current and near future business requirements that are in adherence to the security policy.

The initial and ongoing costs play a large part in this decision and are not covered in this paper. Lists of commercial firewall vendors and resellers can

be found on the Internet.² The same applies to the skills required for migration and for ongoing support that must be available to the company.

Rule Set Implementation

Although firewall rules may look similar on many firewalls, the way in which the firewalls implement the rules may be significantly or subtly different. It is critical to understand clearly what analysis is done to packets and in what order firewall policies, routing, and NAT are implemented.

Network Address Translation

Network Address Translation (NAT) is a standard feature on modern firewalls. However, there can be a world of difference between the different NAT implementations on each product. Although some vendors provide a very flexible feature set, some provide a feature set tailored to the SOHO market. These are not appropriate for complex environments. There are many implementations that are somewhere in between, but are rapidly improving.

Appendix A is a table of example NAT scenarios that can be of assistance when assessing whether a new technology can meet a company's NAT requirements. A complex environment may require many of the combinations of NAT specified in the table. Many firewalls currently do not support all the combinations.

Certainly a good NAT implementation allows connections to be initiated using any interface as the ingress interface and any other interface as the egress interface (rules permitting) without compromise to the NAT options available. Also desirable is built-in implementation of ARP responses for NAT IP addresses within address ranges of locally connected networks. This means that a firewall administrator does not need to configure a proxy-ARP solution at the operating system level.

In cases where a firewall does not support a desired NAT scenario, a redesign of architecture can often solve the problem. For example, an SSH server behind a firewall that restricts access based on source IP as a means of host security, may need to be reconfigured to allow a different source IP.

Routing

Support for the required routing protocols is advantageous. The alternative is to use separate routers supporting the required protocols and redistribute routes if necessary. This can increase security through defence in depth if implemented securely, but can also increase complexity. If the routers utilised are easily compromised then the solution is flawed.

² "Firewall Product Overview" <http://www.thegild.com/firewall/>

Common LAN protocols such as RIP v1 & v2 are not supported by some newer firewall implementations. Some older firewalls do not support WAN protocols such as BGP.

Some firewall appliances make use of virtual routers. For example, a 'trusted' and 'un-trusted' router can be configured on a single device, and specified routes can be inherited between the two. If used intelligently, this feature can be utilised to work around poor IP address allocation schemes.

Spoofing Rules

Most firewalls provide some form of anti-spoofing protection. Anti-spoofing protection allows for denial or rejection of IP packets based on knowledge of whether the source IP addresses is permitted on an ingress interface. This should be considered a required feature when assessing firewall technology. "Corporate network administrators should implement filtering to ensure their corporate networks are not the source of [IP spoofing] problems."³ The principle of blocking spoofed traffic is considered a standard responsibility for internet service providers.

There are many IP spoofing attacks take advantage of send packets with bogus source IP addresses. They can result in major denial-of service scenarios. "In some cases, the system may exhaust memory, crash, or be rendered otherwise inoperative."⁴

Anti-spoofing features applied on a per-interface basis can allow a firewall to take a heavy load of traffic. The firewall typically does a spoof check prior to in-depth analysis of the incoming packet. It is recommended that all firewalls selected to be Internet facing support strong anti-spoofing features.

Performance of ASIC Technology

In some cases, firewall technology is upgraded primarily due to an increased performance requirement. Firewall appliances integrating hardware and software processing are generally faster than their generic server equivalents that have firewall applications installed to a multipurpose operating system.

In particular, the performance overhead of encryption for VPN delivery is significant. ASIC (application-specific integrated circuit) technology has dramatically changed expectations of firewall performance. It allows functions such as encryption to be efficiently performed by hardware. Use of the data encryption standard ASIC can increase VPN performance dramatically. Where performance is a consideration, assess ASIC technology found in many firewall appliances.

³ P. Ferguson, D. Senie. RFC2827.

⁴ CERT[®] Advisory CA-1996-21

Session State and High Availability

Stateful firewalls are generally considered superior to those that do not maintain state knowledge of current connections. A stateful firewall will typically maintain a table consisting of protocol state relevant to at least IP, TCP and some application layer protocols. This state table is used to increase the efficiency of a firewall due to reduced packet inspection. Most importantly, stateful firewalls implement high availability.

A good stateful firewall pair constantly passes state information through redundant connections and can fail-over within seconds. Current connections passing through a high availability firewall pair are not reset as a result of the fail-over. This is essential to providing a high availability environment. Do not implement a technology that is not stateful if high availability is required.

Management Capability for Distributed Environments

An area of firewall technology that has a major impact on efficient technical management of large environments is the enterprise management platform. Current solutions are still developing rapidly and some solutions are far better than others. It is worth spending the resource to thoroughly investigate this area if it is planned to deploy and subsequently manage a large number of firewalls. It is recommended to use the management solution when the firewalls are originally deployed, rather than adopting it at a later stage, due to the resource required migrating configurations to the management platform.

© SANS Institute 2003, P. 10

Appendix A – Desirable NAT Features

Assumptions:

Ingress interface IP is 10.0.8.1

Egress interface IP is 10.0.9.1

All network masks are /24.

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
1	10.0.1.1	10.0.2.1	TCP 1234	10.0.9.1	original	original
NAT source IP is the IP of the egress interface.						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
2	10.0.1.1	10.0.2.1	TCP 1234	10.0.9.2	original	original
NAT source IP an IP on the same network as the egress interface.						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
3	10.0.1.1	10.0.2.1	TCP 1234	10.0.7.1	original	original
NAT source IP is a dynamically allocated from a pre-defined pool of addresses (such as 10.0.7.0/24).						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
4	10.0.1.1	10.0.2.1	TCP 1234	10.0.3.1	original	original
NAT source IP is a pre-defined address on a network known to exist on the ingress side of the firewall (such as 10.0.3.0/24).						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
4	10.0.1.1	10.0.2.1	TCP 1234	10.0.4.1	original	original
NAT source IP is a pre-defined address on a network unknown to the firewall (such as 10.0.4.0/24).						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
5	10.0.1.1	10.0.2.1	TCP 1234	original	10.0.9.3	original
NAT destination IP is an IP on the same network as the egress interface.						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
6	10.0.1.1	10.0.2.1	TCP 1234	original	10.0.4.2	original
NAT destination IP is a pre-defined address on a unique network (such as 10.0.4.0/24).						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
7	10.0.1.1	10.0.2.1	TCP 1234	original	original	TCP 4321
NAT service is a predefined TCP or UDP port.						

#	Orig. Src. IP	Orig. Dest. IP	Orig. Service	NAT Src. IP	NAT. Dest. IP	NAT. Service
8	10.0.1.1	10.0.2.1	TCP 1234	10.0.3.1	original	TCP 4321
Any combination of the above, such as: NAT source IP is the IP of the egress interface. NAT destination IP is a pre-defined address on a unique network. NAT service is a predefined TCP or UDP port.						

Appendix B – Example of an Unexpected Complication

This appendix describes one example of a complication that the author encountered during a firewall service migration that led to implementing a back out strategy. This is just one example of many unexpected issues that can be encountered and is included for the purpose of demonstrating the importance of testing and back out strategies.

The Service Migration

A project involving the migration of a complex firewall topology required that many service migrations were scheduled over a number of weeks. New high availability firewalls had been deployed in parallel to the existing infrastructure. This took advantage of the reduced risk of staging migrations using the IP re-addressing strategy.

Midway through the project when many services had already been migrated, the technical understanding of the new technology was considered to be good. The next service migration was considered low risk from an implementation perspective, but high risk from a business impact perspective. Therefore it had been scheduled for a change window on a weekend.

The migration involved a third party changing IP addresses on their routers and servers (connected to the same network as the firewalls). It involved some routing changes on the third party network. It involved repatching of affected red zone hosts and routers. It required that internal client applications on the green zone were reconfigured to connect to new NAT IP addresses assigned to the third party's hosts. It also required testing by end users. Many parties were involved which is typical of service migrations.

The migration went smoothly to schedule on a Saturday. Testing conducted by the end users demonstrated that the applications in the green and red zones were communicating correctly through the new firewalls and that the third party was routing correctly. There were no unexpected logs or traffic to be seen on the new firewalls. The migration was considered a success.

It had been agreed in advance with the third party that a second phase of testing at only the IP level (as opposed to the application level) would also be conducted on the Sunday to confirm that no problems developed. The third party did indeed find a problem when this testing was done. Analysis uncovered that the new firewalls were behaving differently than the old firewalls in an unexpected way.

The Problem from a Technical Perspective

After a successful TCP 3-way handshake, packets from the source were being dropped by the firewalls. This only occurred when the first data packet was not sent by the source within a 60 second window of time. The new firewalls had a feature protecting against denial of service attacks that open many TCP connections without utilising them. To prevent this sort of attack filling the firewall state table, the firewalls were dropping any connections

where the first data packet after the TCP 3-way handshake was not sent within 60 seconds.

In addition to this timeout, the firewalls were also dropping connections that were previously used, but remained quiescent for 60 minutes. This firewall behaviour is also designed to prevent unused connections from filling the firewall state table or even getting hijacked by a savvy hacker.

After the successful testing that had been conducted on the Saturday, the client-to-server connections had remained open. This is quite normal for many applications. After 60 minutes of inactivity the firewall timeout removed knowledge of the connections from its state table and dropped any further data packets for that connection. When the hosts did resume sending data packets, they detected the lack of response (as no corresponding TCP packet with ACK flag set was returned).

The source host then did something interesting. It retried sending data packets, but only after waiting for some back-out timers to expire. These back-out timers were found to be hard-coded in the application. These back-out timers took slightly more than 60 seconds to expire. After sending repeat data packets without receiving a response, the host attempted to open a new connection.

Normally, opening a new connection in this scenario would solve the problem. The firewalls allowed the TCP 3-way handshake, and waited for the first data packet. However, the first data packet was sent only after the application back-out timers expired. The first data packet was sent shortly after 60 seconds. By this stage, the firewalls had dropped the newly created state table entry - at exactly 60 seconds after the third packet of the 3-way handshake.

“The introduction of a firewall and any associated tunneling or access negotiation facilities MUST NOT cause unintended failures of legitimate and standards-compliant usage that would work were the firewall not present.”⁵
The IP behaviour exhibited by the host was in accordance with TCP specification, however the firewall had a protective feature that actually prevented the service working. This was not detected in the first test window solely because the connections between that source and destination hosts had not been open in a quiescent state.

The Solution

The short-term solution was to execute the back out plan. This affected only the one service due to the IP re-addressing strategy adopted. This involved all parties that had participated in the migration the day before. This could be comfortably done before there was a business requirement for this service to be operating on the Monday. All parties had been made aware that the second test was going to take place. All required parties were prepared to assist if a problem was encountered.

⁵ Freed, N. RFC2979.

One long-term solution would have been to reconfigure the firewall by extending the timers discussed. This was discussed with the firewall vendor and a 'hack' was suggested that provided a partially acceptable way of doing this.

However, there were a number of parties connected via this firewall infrastructure and these timers did protect against very real potential threats. "Firewalls are a fact of life that application protocols must face. As such, application protocols SHOULD be designed to facilitate operation across firewalls, as long as such design choices don't adversely impact the application in other ways".⁶

Discussion with the third party resulted in the third party seeing the value in the shared firewall infrastructure protecting against the threats it was designed to. The security of one organisation is a valid concern of another organisation when their networks are connected.

The third party had a code change made to the application to change the behaviour of the back-out timers. This allowed the service to work with the TCP timeout restrictions enforced by the new firewalls. It was the most secure solution and was only possible due to the third party's appreciation of security principles. As a result the third party was held in high regard and the value of the relationship with the third party was strengthened.

The service migration was rescheduled and implemented with success.

The Key Learns

The following aspects of planning prevented an outage:

1. The risk analysis had identified high business impact and led to the service migration being scheduled for a weekend.
2. The test plan was comprehensive.
3. The back-out strategy was well defined and all required parties were made aware of the possibility of their involvement.
4. The relationship with the third party was strong and the technical contact at the third party was willing to work to the agreed plan.

Another key learn is that a good security and technical knowledge is required to make the right decisions when things go wrong.

⁶ Freed, N. RFC2979.

References

CERT Coordination Center. "CERT® Advisory CA-1996-21." TCP SYN Flooding and IP Spoofing Attacks. 29 Nov. 2000. URL: <http://www.cert.org/advisories/CA-1996-21.html> (23 Mar. 2003).

CERT Coordination Center. "Denial of Service Attacks" 4 Jun. 4 2001 URL: http://www.cert.org/tech_tips/denial_of_service.html (23 Mar. 2003).

Curtin, Matt. Ranum, Marcus. "Internet Firewalls: Frequently asked questions." 1 Dec 2001. URL: <http://www.interhack.net/pubs/fwfaq/> (18 Feb. 2003).

Ferguson, P Senie, D. "RFC2827." Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. May 2000. URL: <http://www.ietf.org/rfc/rfc2827.txt> (23 Mar. 2003).

"Firewall Product Overview" 26 Apr. 2003. URL: <http://www.thegild.com/firewall/> (17 May 2003).

Freed, N. "RFC2979." Oct. 2000. Behavior of and Requirements for Internet Firewalls. URL: <http://www.ietf.org/rfc/rfc2979.txt> (18 Feb. 2003).

Internet Software Consortium. "Firewalls Mailing List." 2003. URL: <http://www.isc.org/services/public/lists/firewalls.html> (18 Feb. 2003 - 17 Jul. 2003).

Newman, D. "RFC2647." Aug 1999. Benchmarking Terminology for Firewall Performance. D. Newman. URL: <http://www.ietf.org/rfc/rfc2979.txt> (6 Jul. 2003).

Socolofsky, C. Kale, C. "RFC1180." A TCP/IP Tutorial. 1 Jan 1991. URL: <http://www.ietf.org/rfc/rfc1180.txt> (22 Mar. 2003).

Smith, John. "ASICs... the website." 29 Feb. 2000. <http://www-ee.eng.hawaii.edu/~msmith/ASICs/HTML/ASICs.htm> (6 Jul. 2003).

Tanase, Matthew. IP Spoofing: An Introduction. 11 Mar 2003. URL: <http://www.securityfocus.com/infocus/1674> (23 Mar. 2003).

Welch-Abernathy, Dameon. "PhoneBoy's FireWall-1 FAQs." 15 May 2003. URL: www.phoneboy.com/ (17 May 2003).