



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.



SANS GCFW Practical Assignment

Version 1.9

Randy Yonekura

April 6, 2003

Abstract

This document is divided into 4 main parts. The first will define business requirements and operations, and describe the network and security architecture. The second part goes more into detail as to how the main network components are configured, including a tutorial describing the steps to follow in order to configure these components, and will discuss the security policies and rules used. The third part will describe the technical approach used to plan and conduct an audit of the primary firewall, followed by an evaluation of the results obtained. The last part will discuss a simulated attack on a rival company's network.

As recommended, I examined some of the proposed architectures described in other papers online (<http://www.giac.org/GCFW.php>), and for the most part there were two different types of implementations. The first type, from a security point of view are robust and impressive. However, from a practical and business perspective, they seem impossible to implement profitably. Multiple ISP links, eliminating every single point of failure by having redundant external and internal routers, multiple HA firewalls and IDS in every subnet, and at least a secondary system for every service, make it a solid and complex design yet somewhat impractical for the current size and budget of this company.

The designs of the second type, however, are somewhat creative by trying to get the most security and functionality with the least amount of money involved. They are simple yet efficient, and my personal favorites.

Both type of configurations seem to acquire the new equipment and implement the design from the ground up. But, what about those dreaded scenarios when we have some legacy equipment or existing infrastructure, and need to re-design or use what you already have to meet our new needs? What about when the CEO (similar to the pointy hair manager in a popular comic) insists on a particular product, or worse yet, insists on an aspect of the design be done in a certain way even if alternative options, and products, are much better suited for our needs?

For this scenario, we will assume that our CEO has indiscriminately bought equipment at a bargain price from companies that were going under, and since he was an engineer in his young days, he has some requirements of his own. Under these less than ideal circumstances, we will try to do the best we can while satisfying political requirements.

Table of Contents

Abstract.....	2
1. Assignment 1 : Security Architecture.....	7
1.1 Introduction.....	7
1.2 Business Requirements.....	7
1.3 Design Requirements.....	8
1.4 Additional Requirements.....	8
1.4.1 System configuration requirements.....	8
1.4.2 Political and other requirements.....	9
1.5 User Groups.....	10
1.6 Business Operations.....	10
1.7 Access Requirements.....	11
1.7.1 User access.....	11
Table 1: Access needed by customers.....	12
Table 2: Access needed by suppliers.....	12
Table 3: Access needed by partners.....	13
Table 4: Access needed by remote users.....	14
Table 5: Access needed by internal users.....	14
1.7.2 Additional access.....	14
Table 6: Access needed by the VPN concentrator.....	15
Table 7: Access needed by the web servers.....	15
Table 8: Access needed by the DNS servers.....	15
Table 9: Access needed by the mail servers.....	16
Table 10: Access needed by the NTP servers.....	16
Table 11: Access needed by the NIDS.....	17
Table 12: Access needed by the web proxy.....	17
Table 13: Access needed by the application server.....	17
Table 14: Access needed by the database server.....	17
Table 15: Access needed by the log server.....	17
Table 16: Access needed by the management workstations.....	18
Table 17: Access needed by the terminal concentrator.....	18
1.8 GIACE Network Architecture.....	18
1.8.1 Networks.....	19
Illustration 1: GIAC Enterprises Network Design.....	19
Table 18: IP scheme.....	20
1.8.2 Interface configuration.....	20
Table 19: Interface configuration.....	21
1.8.3 IP addresses.....	21
Table 20: IP addresses.....	22
1.8.4 Network devices.....	22
1.8.5 External service network.....	24
1.8.6 Application network.....	26
1.8.7 Management network.....	27
1.8.8 Internal service network.....	28
1.8.9 Employee network.....	30
2. Assignment 2: Security Policy and Tutorial.....	31
2.1 Security Policies for the Cisco Routers.....	31

2.1.1 ACL for the border router external interface.....	31
2.1.2 ACL for the border router internal interface.....	33
2.1.3 Alternative ACL for the border router internal interface.....	34
2.1.4 ACL for the internal router (Application interface).....	34
2.1.5 ACL for the internal router (Management interface)	35
2.1.6 ACL for the internal router (Internal service network interface).....	35
2.1.7 ACL for the internal router (Employee interface).....	36
2.2 Security Policy and Configuration for the Cisco PIX.....	37
Table 21: Interface configuration for the firewall.....	37
Table 22: Firewall policy.....	38
2.2.1 PIX configuration.....	38
2.2.2 Access Control Lists.....	42
2.2.3 ACL for the external interface (outside).....	42
2.2.4 ACL for the external service network interface (INT_service).....	43
2.2.5 ACL for the VPN network interface (INT_VPN).....	43
2.2.6 ACL for the VPN User network interface (INT_VPNUser).....	44
2.2.7 ACL for the internal network interface (inside).....	44
2.3 Security Policy for the VPN Concentrator.....	45
2.3.1 Security associations.....	45
2.3.2 IKE proposal.....	46
2.3.3 Security policies for groups.....	47
2.3.4 Filters and rules.....	48
2.3.5 Management protocols.....	48
2.4 Tutorial on Configuring Cisco Devices.....	48
2.4.1 Typographic convention.....	49
Table 23: Typographic convention.....	49
2.5 Configuring the Routers.....	49
2.5.1 Common configuration.....	50
2.6 Border Router.....	58
2.6.1 Access Control Lists.....	58
Table 24: Standard ACL declaration.....	59
Table 25: Extended ACL declaration.....	60
Table 26: IP access-list declaration.....	60
Table 27: Interface configuration for the border router.....	61
2.6.2 External interface (Serial0).....	61
2.6.3 Internal interface (FastEthernet0).....	66
2.6.4 Routing.....	67
2.6.5 Saving the configuration.....	67
2.7 Internal Router.....	68
Table 28: Interface configuration for the internal router.....	68
2.7.1 Screened network interface (FastEthernet0/0).....	68
2.7.2 Application network interface (FastEthernet0/1).....	68
2.7.3 Management network interface (FastEthernet0/2).....	70
2.7.4 Internal service network interface (FastEthernet0/3).....	70
2.7.5 Employee network interface (FastEthernet0/4).....	72
2.7.6 Routing.....	73
2.7.7 Saving the configuration.....	74
3. Assignment 3: Verify the Firewall Policy.....	75
3.1 Planning the audit.....	75

3.2 Scope.....	75
3.3 Approach.....	75
3.3.1 Tools & Equipment.....	75
3.3.2 Methodology.....	76
3.3.3 IP addresses.....	77
Table 29: Scanning combinations.....	80
3.3.4 Systems to be replaced.....	80
Table 30: Systems to be replaced.....	80
3.3.5 Syntax.....	80
3.3.6 Costs, time and risks.....	80
Table 31: Expected duration of the audit.....	81
3.3.7 Expected results.....	81
3.4 Conducting the audit.....	81
3.5 Scanning firewall interfaces.....	82
3.5.1 Scanning the firewall from eth0 (outside).....	82
Table 32: Scanning firewall interfaces from eth0.....	82
3.5.2 Scanning the firewall from eth1 (external service network).....	82
Table 33: Scanning firewall interfaces from eth1.....	82
3.5.3 Scanning the firewall from eth2 (VPN network).....	82
Table 34: Scanning firewall interfaces from eth2.....	83
3.5.4 Scanning the firewall from eth3 (VPN User network).....	83
Table 35: Scanning firewall interfaces from eth3.....	83
3.5.5 Scanning the firewall from eth4 (inside).....	83
Table 36: Scanning firewall interfaces from eth4.....	83
3.6 Scanning systems behind the firewall.....	83
3.6.1 Scanning systems behind the firewall from eth0 (outside).....	84
Table 37: Scanning systems behind the firewall from eth0.....	84
3.6.2 Scanning systems behind the firewall from eth1 (external service network).....	84
Table 38: Scanning systems behind the firewall from eth1.....	84
3.6.3 Scanning systems behind the firewall from eth2 (VPN network).....	84
Table 39: Scanning systems behind the firewall from eth2.....	84
3.6.4 Scanning systems behind the firewall from eth3 (VPN User network).....	85
Table 40: Scanning systems from the partners network.....	85
Table 41: Scanning systems from the remote sales network.....	85
Table 42: Scanning systems from the telecommuters network.....	86
Table 43: Scanning systems from the remote administrators network.....	86
3.6.5 Scanning systems behind the firewall from eth4 (inside).....	86
Table 44: Scanning systems behind the firewall from eth3.....	86
3.7 Scanning systems behind the firewall by replacing a system.....	86
3.7.1 Scanning by replacing a system from eth0 (outside).....	86
Table 45: Scanning systems as the border router.....	87
Table 46: Scanning systems as the public NTP servers and antivirus site...87	87
3.7.2 Scanning by replacing a system from eth1 (external service network).....	87
Table 47: Scanning as the external NTP server.....	88
Table 48: Scanning as the external web server.....	88
Table 49: Scanning as the external mail server.....	89
Table 50: Scanning as the external DNS server.....	89
Table 51: Scanning as the external NIDS.....	90

3.7.3 Scanning by replacing a system from eth2 (VPN network)	90
Table 52: Scanning as the VPN concentrator (public).....	90
3.7.4 Scanning by replacing a system from eth3 (VPN_user network)	90
Table 53: Scanning as the VPN concentrator (private).....	90
3.7.5 Scanning by replacing a system from eth4 (inside)	90
Table 54: Scanning as the internal DNS server.....	91
Table 55: Scanning as the internal mail server.....	91
Table 56: Scanning as the web proxy.....	91
Table 57: Scanning as the internal NTP server.....	92
3.8 Evaluate the audit.....	92
Table 58: Results from scanning systems behind the firewall.....	92
Table 59: Results from scanning by replacing a system.....	93
3.9 Improvements and recommendations.....	94
Illustration 2: Potential improvements for the external service network.....	95
Illustration 3: Potential improvements for the application network.....	96
4. Assignment 4 – Design Under Fire.....	97
Illustration 4: Vivekanand Chudgar's design.....	97
4.1 Reconnaissance.....	97
4.2 Research.....	98
4.3 Direct attacks against the firewall.....	99
4.3.1 Vulnerability research.....	99
4.3.2 Syslog attack.....	99
4.3.3 IKE aggressive mode attack.....	100
4.3.4 Flood attacks.....	101
4.4 Denial of service attack.....	102
4.4.1 Approach.....	102
4.4.2 Denial of service attack countermeasures.....	103
4.5 Compromising an internal system through the perimeter system.....	103
4.5.1 Further research.....	103
4.5.2 Microsoft IIS WebDAV remote compromise vulnerability.....	104
4.5.3 Microsoft IIS WebDAV remote compromise vulnerability countermeasures	105
4.6 Conclusion.....	105
5. References.....	106

1. Assignment 1 : Security Architecture

1.1 Introduction

GIAC Enterprises started as a web site hosted by one of the many Free-Web-Hosting sites online, where our current CEO and founder, would sell cookie sayings using a few CGI forms. Users filled out a plain HTTP form with their credit card information, which would then be emailed to him, and he would later email them the cookie saying. Simple (and risky), but somehow the number of customers kept growing. As the difficulty to come up with original and creative material all the time increased, he was forced to find suppliers that could help him with the growing demand. Suppliers would send emails with the cookie sayings and charge his credit card after he reviewed and acknowledged them.

The cookie saying mania kept rising, and our CEO decided to expand his business to multiple languages and eventually partnered with others around the world to translate and re-sell his cookie sayings. This brought a lot more customers, but with it, a lot of overhead for GIACE. There simply was not enough time to process all the customer transactions, email the cookie sayings to customers and translating partners, review and approve the data from the suppliers, and email payment information to suppliers and partners.

As a solution to his problem, our CEO hired a few programmers to develop an application that could take care of all these transactions. This turned out to be an amazing improvement to the overall process, and gave GIACE an astounding boost.

Bandwidth consumption grew beyond the quota offered by the hosting site, so our CEO rented a small office, bought a dozen of used PCs for his employees, his own domain, and an ISDN link with the local ISP. He brought along the developers so they could keep improving their application, and also hired some more people to help with sales. The infrastructure consisted of the PCs used by sales and the developers, the web server and the system running the application.

The crash of the dot.com's came, and ironically, it multiplied the current demand for online cookie sayings. Unemployed people just needed to know what the future held for them, and for many, a cookie saying was the answer. The higher exposure also brought along several attacks to our system and compromised some of the financial data we had.

Acknowledging the pressing need to upgrade our infrastructure, and improve security, our CEO took advantage of the situation, leased a T1 line and bought as much equipment as he thought was needed from these bankrupt companies, hoping to once again transform GIACE, an e-commerce company pioneering in the online fortune cookie saying world, into an even greater company.

1.2 Business Requirements

Keeping business goals and profitability in mind, we outline the following requirements:

- **Prompt deployment:** Imperative in order to maintain our market share and to start the needed positive cash flow.
- **Low cost:** Provide the highest performance and availability possible with the resources and budget available. Low cost maintenance is another must.
- **Security and confidentiality:** Maintain a high level of trust with our customers, suppliers and partners by ensuring that all transactions and information are

secure. We cannot stay in business if our clients do not feel comfortable providing sensitive information, such as their credit card number.

1.3 Design Requirements

In addition to our business requirements, there are more technical requirements that will make a set of rules that we will keep in mind when making design choices:

- **Centralized and secure management:** Without making this our Achilles' heel, our administrators and security staff should be able to monitor and configure the infrastructure from a central location. This will also help maintain a 1,000 ft. high view of the network, while also giving them the option to zoom in on key components.
- **Uniform vendors:** Based on the fact that now most of our existing equipment is Sun Microsystems hardware and Cisco devices, we might as well take advantage of the uniformity and make it a requirement. Therefore, whenever possible, we will use the same hardware vendor across the design. There are many advantages to this. From the hardware side, it helps reduce costs and decrease downtime since it is more likely that the same spare parts can be used on multiple systems. The main advantage is that our staff would only have to be familiar with one environment, and have fewer places to check for patches and updates. One important point to mention is that this might lead to the possibility of having the same vulnerability across all systems with the same platform.
- **Defense in depth:** Multiple layers of protection for the network and critical systems should be implemented based on resources available.
- **Scalability and growth:** It is important to make sure the design can accommodate future growth, without major investments, downtime or redesign.
- **Low-cost technologies:** When deemed sensible, complement our layers of security with free or low-cost well-known technologies, as long as they have been out there long enough to be proven reliable for their purpose (i.e. Snort, PGP, OpenSSH, Netfilter, etc.) and have a large industry presence.
- **Need-to-access privileges:** When defining network access, both internal and external, it is imperative to only grant access to the services and networks needed. This should be tailored to the specific type of users. (i.e. access granted to a supplier is not the same as access granted to remote sales staff, for example).
- **Explicit Access:** Any traffic that has not been explicitly defined, should be dropped.

1.4 Additional Requirements

1.4.1 System configuration requirements

Every infrastructure system needs to be configured following these requirements:

- All of our systems (with the exception of the Management workstations, which will be running Trusted Solaris) will be using a hardened and minimized Solaris 9. All the recommended and security patches from SunSolve¹ will be applied regularly. Additional patches depending on the function and the software of the system

¹ <http://sunsolve.sun.com/pub-cgi/show.pl?target=home>

might need to be installed.

- Solaris hardening tools like YASSP², Titan³, and the Solaris Toolkit⁴ will be applied.
- Configuration and hardening guides found in many of the Sun Blueprints online (links are included in the reference section) will also be implemented.
- All systems are configured to mitigate damage from DoS attacks (decreased connection timeouts, larger queues and connection tables, etc.).
- The only login access will be through the console, with the exception of those systems that require access by developers or administrators. In which case, TCPWrappers⁵ will also be installed to only allow connections from those servers that need access.
- Only those ports and daemons needed by the system to perform its duties will be available.
- Tripwire⁶ will be installed to ensure file and system integrity.
- At the time of this writing, we could not find a suitable Solaris HIDS supporting 2.9 that met our needs. However, in the meantime we will install and configure Swatch⁷ to monitor system logs. We are aware this will not provide real-time alerts, but we expect to include one of the existing Solaris HIDS as soon as they add support for Solaris 2.9
- All infrastructure servers will send their logs to the log server.
- Each server will be configured so that it can be JumpStarted and rebuilt from scratch in less than 2 hours.

1.4.2 Political and other requirements

One important requirement for the new infrastructure imposed by the CEO is to use the newly acquired capital equipment as much as possible. This legacy equipment is not necessarily the latest technology, but he hopes it will meet the current needs and give us some room for growth.

Here is the list of systems we have to work with:

- (6) Sun Ultra 60⁸ with only one 450 MHz processor
- (9) Sun Enterprise 220R⁹ with 2 450 MHz processors
- (4) Sun Netra T1 105¹⁰ with a 440 MHz processor
- (2) Sun StorEdge T3+¹¹
- (1) Sun StorEdge D1000¹²
- (1) Cisco 7204¹³ Series Router
- (1) Cisco 3640¹⁴ Router with two 32-port asynchronous modules
- (1) Cisco 3620¹⁵ Router with two 4-port Ethernet modules
- (1) Cisco PIX 515E¹⁶

2 <http://w.w.yass.org>

3 <http://www.fish.com/titan>

4 <http://www.sun.com/blueprints/tools>

5 <http://sunfreeware.com/programlistsparc9.html#tcpwrappers>

6 <http://www.tripwire.org>

7 <http://swatch.sourceforge.net>

8 http://sunsolve.sun.com/handbook_pub/Systems/U60/spec.html

9 http://sunsolve.sun.com/handbook_pub/Systems/E220R/spec.html

10 http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/spec.html

11 http://sunsolve.sun.com/handbook_pub/Systems/T3/spec.html

12 http://sunsolve.sun.com/handbook_pub/Systems/D1000/spec.html

13 <http://www.cisco.com/en/US/products/hw/routers/ps341/ps346/index.html>

14 <http://www.cisco.com/en/US/products/hw/routers/ps274/ps278/index.html>

15 <http://www.cisco.com/en/US/products/hw/routers/ps274/ps276/index.html>

16 http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet09186a0080091b15.html

- (5) Cisco FastHub 100 Series¹⁷
- Four (4) Cisco 1538 Microhub¹⁸

It is worth mentioning that with the exception of the 1538 Microhubs, the 3620 Router, the storage, and the PIX 515E, all of these systems have already reached their End-Of-Life with Cisco and Sun.

Also, our CEO has been pretty emphatic on the following aspects of the design and maintenance:

- The border router must be biggest network device: the Cisco 7204, and the Cisco 3620 must be used as the internal router.
- **No multi-functioning.** Each system must perform only one function and one function only.
- **Simplicity.** Complexity often leads to misconfigurations. Simple should not be synonymous with weak or insecure, but it should keep the technical complexity low. This should keep maintenance and hardware costs low as well.
- **Diligent maintenance.** All administrators are required to subscribe to multiple vulnerability mailing lists, actively monitor vulnerability sites, and keep the infrastructure current with current patches and vulnerability fixes. A system that is affected by a vulnerability might be grounds for job termination of the responsible administrator due to neglect of his/her responsibilities, if the vulnerability has been publicly known and fixes have been released for over a week.

1.5 User Groups

There are five sets of users that need access to our network. Each performs a different set of operations, and must have its corresponding access privileges and restrictions. Therefore, it is important to define the different types of users:

- **Customers:** Companies or individuals that purchase bulk online fortunes.
- **Suppliers:** Companies that supply GIACE with their fortune cookie sayings.
- **Partners:** International companies that translate and resell fortunes.
- **Remote Users:**
 - Remote sales staff
 - Remote administrators
 - Telecommuters
- **Internal Users:** GIACE employees located on the internal network.

1.6 Business Operations

All business transactions are performed one way or another using GIACE's own software, which is the heart of all its business operations and developed internally. For reference, we will call it Fortune Software. The system running this software, which we will refer to as the Fortune server or simply the Application server, interacts with the database server to create/edit customer profiles; perform operations with fortune cookie data: download, upload, editing, translation; store and retrieve information used by the sales staff and partners: contact numbers, contracts, etc.; and user management, which concerns user

¹⁷ <http://www.cisco.com/en/US/products/hw/hubcont/ps853/index.html>

¹⁸ <http://www.cisco.com/en/US/products/hw/hubcont/ps209/ps208/index.html>

access control within the application.

The data contained in the database server is obviously our most valuable asset. Therefore, we go through great lengths to ensure that it **only** talks to our application server. Moreover, **only** the internal and external web servers can talk to the application server.

Let us briefly describe our business process:

- **Customers** connect to our external web server to make a purchase or update/create their profile. The external web server will interact with our application server (the system running our Fortune Software), which will then handle the transactions with the database server. There is never a direct connection between the customer and our application server.
- **Suppliers** will connect to our external web server to upload their fortunes. The web server processes and relays this to the application server, which upon verification will upload it to the database server.
- **Partners** will connect to the internal web server through the VPN. The internal web server interfaces with the application server, allowing them to update fortune and sales data.
- **Local and remote sales staff** will use the internal web server, via the local network or VPN, to interface with the application server to perform all the necessary commercial transactions relating to our business and the players involved: **customers, suppliers and partners**.

The above business process give us a large view of the overall process. However, let us keep in mind that it is oversimplified for convenience's sake, leaving out technical aspects that are invisible to the users. There might be multiple instances of web, application, and database servers providing load balancing and redundancy. For example, the web server the users see and use, could in reality be multiple web servers and load balancers, and the traffic from the customer to the web server, or the traffic from the web server to the application might go through multiple proxies.

As resources increase, and assuming the needs justify it, the design can be improved to allow for some of the features mentioned above, without affecting the overall business process seen by the players involved.

1.7 Access Requirements

1.7.1 User access

It is important to define in detail the type of access that is required by our users before sketching our network design and implementation. The network and infrastructure have to provide access to our defined User Groups to accomplish our Business Operations while meeting our Business and Design Requirements.

Customers

As most e-commerce companies, our customers will be performing their transactions through our external web server on our external service network. Customers can purchase fortunes without having to create accounts, however, if

they create an account, this creates and safely stores their profile on our database. Strong password policies shall be enforced on these accounts, for ours and the customer's protection.

For general browsing, HTTP is acceptable, but for logons and other transactions, a 128-bit SSL shall be required. No client certificates are needed, but we will use Verisign¹⁹ certificates to ensure server authenticity to our customers. This provides an acceptable level of confidentiality and authenticity that boosts the confidence of many of our online customers when they are typing their credit card information.

Expected operations are:

- General browsing
- User registration and login
- Profile editing
- Purchase and download of fortune sayings

Even though we do not accept credit card information via email, customers still need to be capable of communicating via email for general questions and customer service.

Origin	Port/Transport	Destination
Internet (Customers)	53/udp	External DNS
Internet (Customers)	25/tcp	External mail relay
Internet (Customers)	80/tcp, 443/tcp	External web server

Table 1: Access needed by customers

Suppliers

Similar to our customers, our suppliers will perform their transactions through our external web server. However, stronger password policies are enforced on these accounts. All transactions from our suppliers will use a 128-bit SSL encryption. Moreover, a mutual SSL authentication will be used to ensure client authenticity as well.

Expected operations are:

- User login
- Profile editing
- Upload of fortune cookie data

All e-mail between GIACE and its suppliers will require the use of GnuPG²⁰ 1.2.1 or higher, which should be compatible²¹ with PGP²² 5.x or higher. Any correspondence that cannot be authenticated and/or was not encrypted shall be disregarded.

Origin	Port/Transport	Destination
Internet (Suppliers)	53/udp	External DNS
Internet (Suppliers)	25/tcp	External mail relay
Internet (Suppliers)	80/tcp, 443/tcp	External web server

Table 2: Access needed by suppliers

19 <http://www.verisign.com/products/onsite/ssl/faq.html> : According to Verisign, " VeriSign SSL Certificates are instantly recognized by 98% of the installed base of web browsers"

20 <http://www.gnupg.org>

21 IDEA is still patented

22 <http://www.pgp.com>

Partners

Partners connect to the internal web server through our VPN. This server will not use or require server/client certificates. However, VPN group and user passwords will be enforced by strong password policies. In addition, partners can only access the web server in our internal network.

Expected operations are:

- User login
- Profile editing
- Download and upload of fortune cookie data
- Access to sales information

All e-mail between GIACE and its partners will require the use of GnuPG 1.2.1 or higher, which should be compatible with PGP 5.x or higher. Any correspondence that cannot be authenticated and/or was not encrypted shall be disregarded.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Internet (VPN client)	500/udp	VPN concentrator
Internet (VPN client)	IP protocol 50 (ESP)	VPN concentrator
Internet (Partners)	25/tcp	External mail relay
VPN User Network (Partners)	80/tcp, 443/tcp	Internal web server

Table 3: Access needed by partners

Remote Users

We have three types of remote users:

- Remote sales staff
- Remote administrators
- Telecommuters

All three connect to our internal network through our VPN. Remote sales staff need access to the internal web server, and the internal mail server; telecommuters only need access to the internal mail server; and remote administrators only need access to the management workstations. We have decided not to allow access to the employee network. This might be inconvenient for those who want to access their workstations, but currently the needs do not justify it.

Expected operations:

- Access the internal web server (Remote sales staff)
- Access the management workstations (Remote administrators)
- Access the internal mail server (Remote sales staff & telecommuters)

Systems used by our remote users **must** have a personal firewall, and an anti-virus software installed. For this we will use Norton Internet Security 2003²³, which comes bundled with a firewall, virus protection, IDS, and privacy control among its features. Moreover, our remote users must be instructed to configure the Cisco VPN client to disable all local LAN connections when it connects to our network, which is the default.

All three connect to the VPN concentrator the same way, but based on

²³ http://www.symantec.com/sabu/nis/nis_pe/features.html

their group and user information, they will be restricted on where they can go within our network.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Internet (VPN client)	500/udp	VPN concentrator
Internet (VPN client)	IP Protocol 50 (ESP)	VPN concentrator
VPN User Network (Remote sales staff & Telecommuters)	25/tcp	Internal mail server
VPN User Network (Remote sales staff)	80/tcp, 443/tcp	Internal web server
VPN User Network (Administrators)	22/tcp	Mgmt. workstations

Table 4: Access needed by remote users

Internal Users

Internal users do not have absolute access to the internal network. Access to the application network is only given to users connected from specific subnets within the employee network.

All internal users can connect to the internal service network, and the outside world. All access to the outside world goes through an internal proxy and it is only limited to web access. As needs change, this might be revised. Internal users do not have access to the administrators network or the application network. The exception to this are the systems on a subnet primarily used by the Fortune Software developers, who need to connect to the production application and web server for application upgrades and maintenance. Local sales staff, who connect from a dedicated subnet, also have access to the web server within the application network

Expected operations:

- Access to the internal service network
- Access to the application network from selected subnets
- Controlled access to the outside world

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Employee Network	53/udp	Internal DNS server
Employee Network	25/tcp	Internal mail server
Employee Network (Sales & Developers subnet)	80/tcp, 443/tcp	Internal web server
Employee Network	80/tcp, 443/tcp	Web proxy
Employee Network (Developers subnet)	22/tcp	Application Network

Table 5: Access needed by internal users

1.7.2 Additional access

Here is the complete network access that our systems/networks need in order to perform their duties properly. Let us note that some of these entries **will not** have to be enforced or specified with the access control lists of our network devices if the systems are on the same subnet, but they are mentioned now for completeness' sake.

VPN concentrator: The public interface of the VPN concentrator needs to send logs, and IKE/ESP traffic out.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
VPN concentrator (public)	500/udp	Internet (VPN clients)
VPN concentrator (public)	514/udp	Log server
VPN concentrator (public)	IP protocol 50 (ESP)	Internet (VPN clients)

Table 6: Access needed by the VPN concentrator

Web servers: Both are expected to reply back to internet or remote/local users and require access to our application server through TCP port 5000. The internal web server should reply to the queries from local and VPN users (remote sales staff and partners). The internal web server will also be accessed by developers for upgrades and maintenance via ssh, so replies must be allowed to leave the network. Both need to send logs and have access to NTP servers.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Internal web server	22/tcp	Developers subnet
External web server	80/tcp, 443/tcp	Internet (Customers & Suppliers)
Internal web server	80/tcp, 443/tcp	VPN User network (Remote sales staff & Partners)
Internal web server	80/tcp, 443/tcp	Employee network
External web server	123/udp	External NTP server
Internal web server	123/udp	Internal NTP server
External web server	514/udp	Log server
Internal web server	514/udp	Log server
External web server	5000/tcp	Application server
Internal web server	5000/tcp	Application server

Table 7: Access needed by the web servers

DNS servers: The external DNS server must be able to access the internet, and return recursive inquiries to the internal DNS server. The internal DNS server should be able to reply to queries from our internal users, relay DNS queries to our external DNS server and receive its replies. Both need to send logs and have access to NTP servers.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
External DNS server	53/udp	Internet
External DNS server	53/udp	Internal DNS server
Internal DNS server	53/udp	External DNS server
Internal DNS server	53/udp	Employee Network
Internal DNS server	53/udp	Internal Service Network
External DNS server	123/udp	External NTP server
Internal DNS server	123/udp	Internal NTP server
External DNS server	514/udp	Log server
Internal DNS server	514/udp	Log server

Table 8: Access needed by the DNS servers

Mail servers: The external mail relay needs access to the internet, and be able to forward mail to our internal mail server. Inversely, all outgoing mail should be forwarded to our external mail server. Moreover, the internal mail server needs to reply to local and VPN users (except partners and administrators). Our mail relay has a cron job to check daily for new virus identities (IDEs) and download them. To be consistent, we will sanitize the IP of the virus company and make it c.c.c.c. Both need to send logs and have access to NTP servers.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
External mail relay	25/tcp	Internet
External mail relay	25/tcp	Internal mail server
Internal mail server	25/tcp	External mail relay
Internal mail server	25/tcp	VPN User Network (Remote sales staff & Telecommuters)
Internal mail server	25/tcp	Employee Network
External mail relay	80/tcp	Antivirus download page
Antivirus download page	80/tcp	External relay
External mail relay	123/udp	External NTP server
Internal mail server	123/udp	Internal NTP server
External mail relay	514/udp	Log server
Internal mail server	514/udp	Log server

Table 9: Access needed by the mail servers

NTP servers: In order to keep our clocks in sync, we will use well-known public NTP servers. Traffic should be allowed in both directions. Our internal NTP server will query our external NTP server. Both NTP servers must be able to reply to queries. For this design, we have sanitized the IPs of the public NTP servers we use. The IPs are **a.a.a.a** and **b.b.b.b**. Both need to be able to send logs.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
External NTP server	123/udp	Public NTP servers
Public NTP servers	123/udp	External NTP server
Internal NTP server	123/udp	External NTP server
External NTP server	123/udp	Internal NTP server
Internal NTP server	123/udp	Internal Service Network
External NTP server	123/udp	External Service Network
External NTP server	514/udp	Log server
Internal NTP server	514/udp	Log server

Table 10: Access needed by the NTP servers

NIDS: NIDS need to be able to send logs and sync up with the NTP servers.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
External NIDS	123/udp	External NTP server
Internal NIDS	123/udp	Internal NTP server
External NIDS	514/udp	Log server
Internal NIDS	514/udp	Log server

Table 11: Access needed by the NIDS

Internal web proxy: Must be able to send web traffic, receive replies, and pass them to the local users.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Web proxy	80/tcp, 443/tcp	Internet
Internet	80/tcp, 443/tcp	Web proxy
Web proxy	80/tcp, 443/tcp	Employee Network
Web proxy	123/udp	Internal NTP server
Web proxy	514/udp	Log server

Table 12: Access needed by the web proxy

Application server: It is imperative that it can communicate with the database server, and reply to the web servers. As well as accept traffic from the developers subnet. All traffic to/from the database servers uses TCP port 5000. The application server needs to send logs and have access to the internal NTP server.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Application server	22/tcp	Developers subnet
Application server	123/udp	Internal NTP server
Application server	514/udp	Log server
Application server	5000/tcp	Database server
Application server	5000/tcp	External web server
Application server	5000/tcp	Internal web server

Table 13: Access needed by the application server

Database server: It can only be accessed by the application server through TCP port 5000.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Database server	5000/tcp	Application server

Table 14: Access needed by the database server

Log server: The log server needs to be able to get in sync with the internal NTP server.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Log server	123/udp	Internal NTP server

Table 15: Access needed by the log server

Management workstations: Must be able to reply to our VPN users (administrators) and communicate with our terminal concentrator.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Mgmt. workstations	22/tcp	VPN User Network (Administrators)
Mgmt. workstations	22/tcp	Terminal concentrator

Table 16: Access needed by the management workstations

Terminal concentrator: It can only be accessed from the management workstations. Our network infrastructure (routers, firewall, switches), and critical servers can be accessed through their console. In fact, that is the only way to administer most of them.

<i>Origin</i>	<i>Port/Transport</i>	<i>Destination</i>
Terminal concentrator	22/tcp	Mgmt. workstations
Terminal concentrator	Serial	Infrastructure

Table 17: Access needed by the terminal concentrator

1.8 GIACE Network Architecture

GIACE's proposed network has eight major networks, three considered external and five internal networks:

- External Service Network
- VPN Network (external)
- VPN User Network (external)
- Application Network (internal)
- Management Network (internal)
- Internal Service Network
- Employee Network (internal)
- Screened Network (internal)

GIAC Enterprises Network Design

x.x.x.0/28

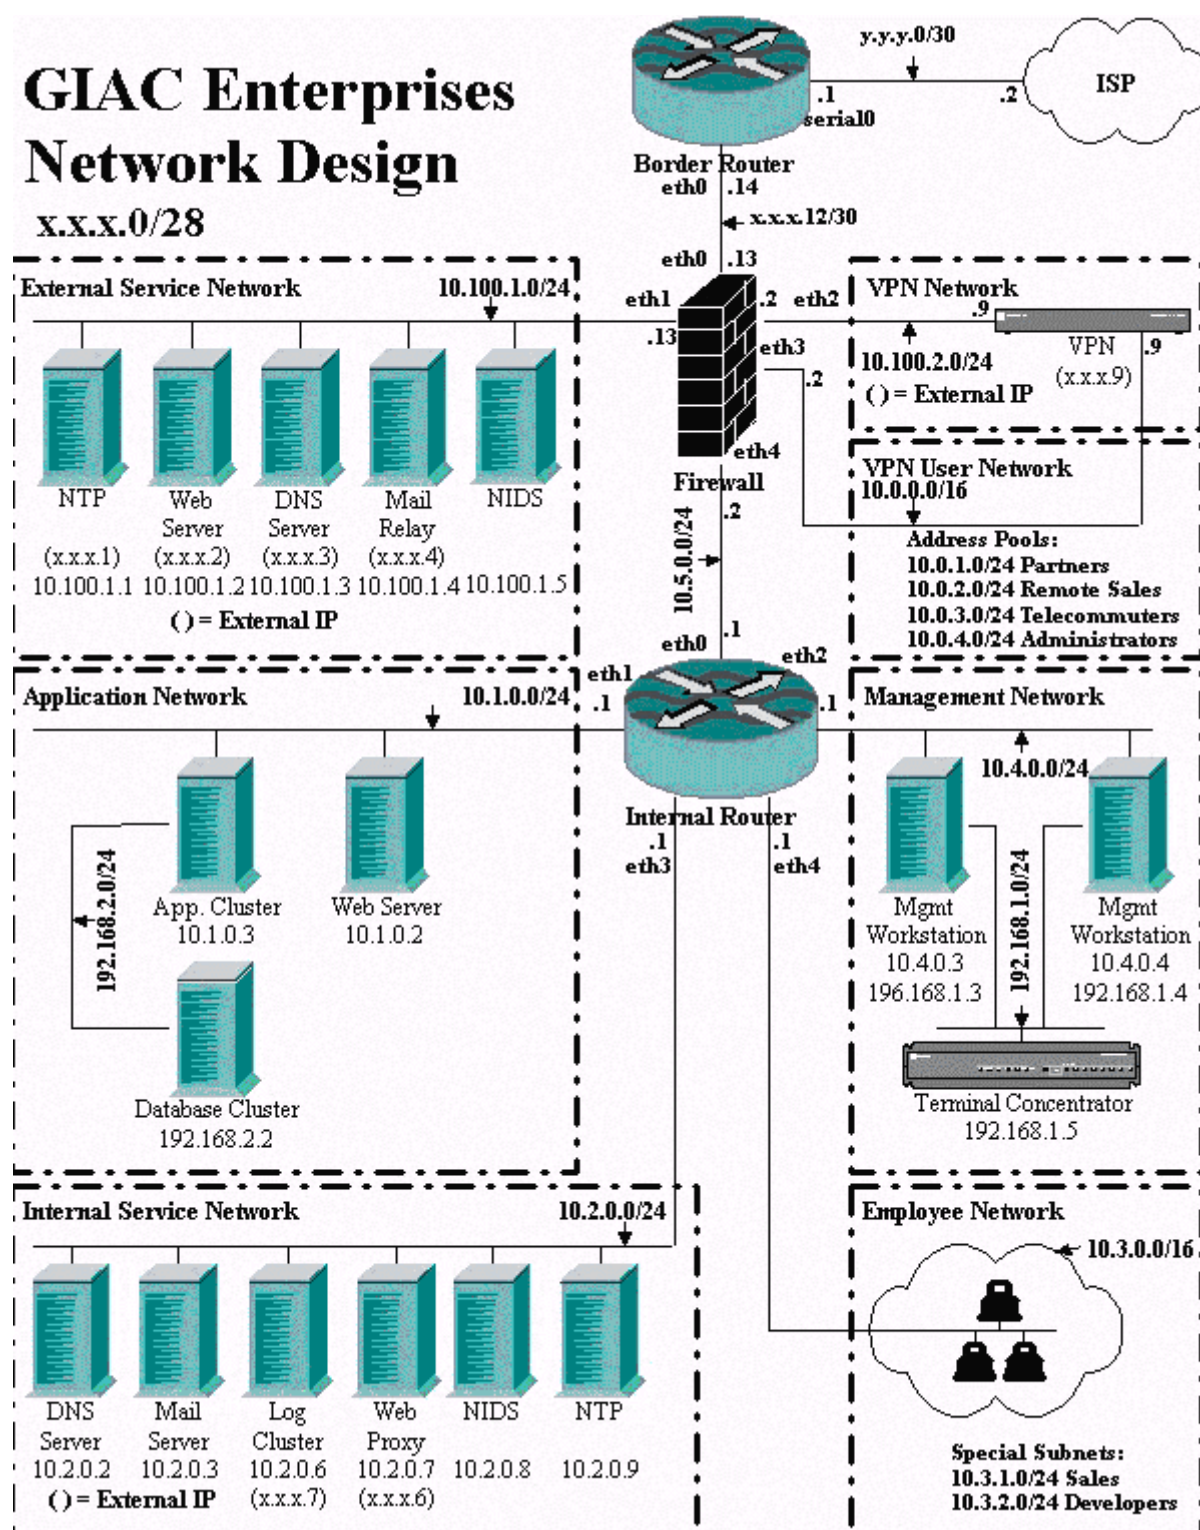


Illustration 1: GIAC Enterprises Network Design

1.8.1 Networks

This level of segmentation, based on the function offered, provides an efficient level of access control for our needs, ensuring that only traffic (both internal

and external) that has been explicitly allowed will have access to services on other networks. A breach on the network perimeter, or attack from the inside, will not be sufficient to access all critical systems.

The external IP address range allocated to our company by the ISP is: x.x.x.0/28, with a serial IP address for the border router of y.y.y.1/30. We have sanitized these IPs as a precautionary measure to ensure that no real sites within this range are unintentionally, or deliberately, attacked based on the design.

GIACE uses the RFC-1918²⁴ private addresses 10.0.0.0/8 and 192.168.0.0/16 internally. Networks on the 192.168.0.0/16 range are merely private networks between systems and are not routed, but we list them for the sake of completeness.

The systems on the external service network and the VPN concentrator use private IPs on the 10.100.1.0/24 and the 10.100.2.0/24 network, respectively. However, we use static NAT to give them an external IP. The only exception is the external NIDS, which does not have a public IP address.

Network Segment	Address
GIACE's external IPs	x.x.x.0/28
External router->ISP	y.y.y.0/30
External service network	10.100.1.0/24
VPN network	10.100.2.0/24
VPN User network	10.0.0.0/16
VPN address pool (Partners)	10.0.1.0/24
VPN address pool (Remote sales staff)	10.0.2.0/24
VPN address pool (Telecommuters)	10.0.3.0/24
VPN address pool (Administrators)	10.0.4.0/24
Application network	10.1.0.0/24
Internal service network	10.2.0.0/24
Employee network	10.3.0.0/16
Sales subnet	10.3.1.0/24
Developers subnet	10.3.2.0/24
Management network	10.4.0.0/24
Screened network	10.5.0.0/24
Private (Mgmt.->Terminal concentrator)	192.168.1.0
Private (App->Database network)	192.168.2.0

Table 18: IP scheme

1.8.2 Interface configuration

For the sake of thoroughness, we list how the interfaces of our network devices are configured.

Device	Interface	Device/Network
Border router (external)	serial0	ISP
Border router (internal)	ethernet0	Firewall

²⁴ <http://www.isi.edu/in-notes/rfc1918.txt>

Device	Interface	Device/Network
Firewall	ethernet0	Border router
Firewall	ethernet1	External service network
Firewall	ethernet2	VPN concentrator (Public)
Firewall	ethernet3	VPN concentrator (Private)
Firewall	ethernet4	Internal router (Screened network)
Internal router	ethernet0	Firewall (Screened network)
Internal router	ethernet1	Application network
Internal router	ethernet2	Management network
Internal router	ethernet3	Internal service network
Internal router	ethernet4	Employee network

Table 19: Interface configuration

1.8.3 IP addresses

Here is the complete list of IPs used in our infrastructure, and the external IPs used by our services. External IPs used for static NAT are listed in parenthesis:

System	IP
ISP serial link	y.y.y.2
Border router serial0	y.y.y.1
Border router ethernet0	x.x.x.14
Firewall eth0 (outside)	x.x.x.13
Firewall eth1 (External Service Network)	10.100.1.13
Firewall eth2 (VPN Network)	10.100.2.2
Firewall eth3 (VPN User Network)	10.0.0.2
Firewall eth4 (Screened Network)	10.5.0.2
External NTP	10.100.1.1 (x.x.x.1)
External web server	10.100.1.2 (x.x.x.2)
External DNS server	10.100.1.3 (x.x.x.3)
External Mail relay	10.100.1.4 (x.x.x.4)
External NIDS	10.100.1.5
Web proxy's external IP	(x.x.x.6)
Log server's external IP	(x.x.x.7)
VPN concentrator (Public Interface)	10.100.2.9 (x.x.x.9)
VPN concentrator (Private Interface)	10.0.0.9
Internal router eth0 (Screened Network)	10.5.0.1
Internal router eth1 (Application Network)	10.1.0.1
Internal router eth2 (Management Network)	10.4.0.1
Internal router eth3 (Internal Service Network)	10.2.0.1
Internal router eth4 (Employee Network)	10.3.0.1
Application cluster	10.1.0.3

<i>System</i>	<i>IP</i>
Application system1 eth0	10.1.0.4
Application system2 eth0	10.1.0.5
Application cluster (Private)	192.168.2.3
Application system1 eth1 (Private)	192.168.2.4
Application system2 eth1 (Private)	192.168.2.5
Database server (Private)	192.168.2.2
Database system1 (Private)	192.168.2.6
Database system2 (Private)	192.168.2.7
Internal web server	10.1.0.2
Management workstation1 eth0	10.4.0.3
Management workstation2	10.4.0.4
Management workstation1 eth1 (Private)	192.168.1.3
Management workstation2 eth1 (Private)	192.168.1.4
Terminal concentrator (Private)	192.168.1.5
Internal DNS server	10.2.0.2
Internal mail server	10.2.0.3
Log cluster	10.2.0.6
Log system1	10.2.0.4
Log system2	10.2.0.5
Web proxy's internal IP	10.2.0.7
Internal NIDS	10.2.0.8
Internal NTP	10.2.0.9
Public NTP server1	a.a.a.a
Public NTP server2	b.b.b.b
Antivirus download site	c.c.c.c

Table 20: IP addresses

1.8.4 Network devices

Border Router

Our CEO insists that we use our massive Cisco 7204 running Cisco IOS 12.2 as our border router. This system can handle a load much greater than our current needs (or the needs in the foreseeable future), and its integrated VPN can have more connections than our VPN device. It also came with many modules to support services and protocols we have no current use for (ATM, Voice/Video/Data Integration, etc.). Unfortunately, this product reached its End-of-Life on April 30, 2000²⁵.

All of the modules and features we do not need will be removed/disabled.

Besides providing connectivity to the outside world, we will be using our border router to do some basic packet filtering.

The router will be directly connected to the uplink with our ISP and our

²⁵ <http://www.cisco.com/en/US/products/hw/routers/ps341/ps346/index.html>

firewall. It will be configured to send logs to our log server, and use our external NTP server for time synchronization.

For a complete list of features:

http://www.cisco.com/en/US/products/hw/routers/ps341/products_data_sheet09186a008008872b.html

Firewall

Our second line of defense is a Cisco PIX 515E firewall running Cisco PIX Firewall 6.3. This family of firewalls is well-known and enjoys a fairly large user base, which ensures that any new found vulnerabilities are more likely to be reported promptly, together with the appropriate security fix from Cisco. They have earned Common Criteria EAL4,²⁶ ICSA Firewall²⁷ and IPsec²⁸ certification.

Even though the firewall supports a good number of features (VPN, encryption, DHCP, etc.), we will use it mainly for stateful access control.

We will be using four interfaces on the PIX 515E: one interface for the external service network, one for the internal network, and two for the VPN.

Our firewall will provide stateful filtering, ensure that services and networks are only accessed by the users who need access, and also provide some stateful packet inspection for some of the protocols passing through it.

The pix will be configured to send logs to our log server, and use our external NTP server for time synchronization.

The complete data sheet can be found at:

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet09186a0080091b15.html

VPN concentrator

The VPN device is a Cisco 3015 VPN Concentrator running Cisco VPN Concentrator Software 3.6. All the clients will be using Cisco VPN Client 3.7²⁹, which supports a great variety of operating systems.

The VPN concentrator will provide access to our remote users and partners to selected services and subnets inside the internal network.

Traffic to and from this system will go through our firewall in order to protect it, and to protect our internal network from our VPN users. All VPN traffic will use ESP to protect our sensitive information, and be in tunnel mode in order to protect our internal address scheme. Traffic flow is also determined based on what group the user belongs to. There are four VPN groups: partners, remote sales staff, telecommuters, and remote administrators.

It will also be configured to send logs to the log server. However, the VPN concentrator will be our only network device that will not use NTP. The reason we have decided not to enable NTP on it, is mainly because it does not handle authentication, keys, encryption or any of the features we would prefer. The only two settings that the concentrator allows are: IP/name of the NTP server and the frequency of the updates. We will update the time manually, and periodically monitor it to make sure it does not drift too much.

Each group has its own IP address pool:

- 10.0.1.0/24 for Partners
- 10.0.2.0/24 for Remote sales staff
- 10.0.3.0/24 for Telecommuters

²⁶ <http://commoncriteria.org/docs/EALs.html#EAL4>

²⁷ <http://www.icsalabs.com/html/communities/firewalls/index.shtml>

²⁸ <http://www.icsalabs.com/html/communities/ipsec/index.shtml>

²⁹ <http://www.cisco.com/en/US/products/sw/secursw/ps2308/ps3875/index.html>

- 10.0.4.0/24 for Remote administrators

As mentioned in the Access Requirement section, each of these user groups will only access the services they need. Any traffic going to a service they should not have access to will be dropped and logged by our firewall.

For a complete set of features and product description:

http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_data_sheet09186a0080091e4f.html

Internal Router

The internal router is a Cisco 3620 with two 4-port Ethernet modules running Cisco IOS 12.2. It connects the internal networks and one of its interfaces is attached to the firewall. It also enforces, using static filters, that internal users only have access to the services and networks they need.

This router will be configured to send logs to the log server, and use the internal NTP server for time synchronization.

The Cisco 3640 reached its End-of-Life on November 15, 2002³⁰.

1.8.5 External service network

With the exception of the NTP and NIDS, all of the servers on this network are directly accessible to the public. Only traffic directed to the services offered will enter this network, and only expected responses and logs can leave the network.

External NTP server

The external NTP is running Solaris 9 FCS on an Ultra 60. Solaris' xntpd uses Network Time Protocol (NTP) version 3 standard, as defined by RFC1305³¹. NTPV4 has not been formalized in an RFC yet. However, we will install NTP v4.1.1 from www.ntp.org to take advantage of new encryption and authentication features.

We will configure the NTP server and its clients to use encryption and authentication, and using NTP access lists, we can limit which systems can connect to it. For information on configuring NTP:

<http://www.eecis.udel.edu/~mills/ntp/html/index.html>

Based on our location, we will choose our public NTP servers from the following list:

<http://www.eecis.udel.edu/~mills/ntp/servers.html>

We will follow the rules of engagement and have an email notification and/or prior arrangement with the administrators of the public servers we select before we start connecting to them.

For this design, we have sanitized the IPs of our public NTP servers, so they are: a.a.a.a and b.b.b.b.

The external NTP server will serve the systems on the external service network, the border router, the firewall and the internal NTP server.

For more information on NTP, additional links have been included in our reference section.

The Ultra 60 reached its End-Of-Life on July 2002³²

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

³⁰ http://www.cisco.com/en/US/products/hw/routers/ps274/prod_bulletin09186a0080107b23.html

³¹ <http://www.eecis.udel.edu/~mills/database/rfc/rfc1305>

³² http://sunsolve.sun.com/handbook_pub/Systems/U60/U60.html

External web server

The external web server is running Solaris 9 FCS on a Sun Enterprise 220R. The server will be using Apache 2.0.44³³ and will have a Verisign³⁴ certificate to ensure authenticity to clients. Whenever possible, it will be configured to show little or bogus information about itself.

This server is accessed by our customers and suppliers.

Documentation can be found at:

<http://httpd.apache.org/docs-2.0>

When resources allow it, a reverse proxy server would be good addition to our configuration.

The Sun Enterprise 220R reached its End-Of-Life on August 2002.³⁵

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

External DNS server

The external DNS server is running Solaris 9 FCS on a Sun Netra T1 105 using version 9.2.2 of BIND.³⁶ The server will be the authoritative DNS server for our mail and web server. It will not have entries for the other servers on our external service network, our VPN network, or internal servers. All external queries will be answered non-recursive with the address of its root server. Inquiries from our internal DNS server and external service systems, however, will be answered recursively. DNS should be configured so that the output of all these queries does not exceed 512 bytes, and not to use TCP. DNS is configured to run as a user other than root for security purposes. No zone transfers are allowed.

The Netra T1 reached its End of Life on November, 2002.³⁷

The BIND 9 Administration and configuration manual can be found at:

<http://www.nominum.com/content/documents/bind9arm.pdf>

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

External mail relay

The external mail server is running Solaris 9 FCS on a Sun Enterprise 220R. The mail server will be using Sendmail 8.12.8.³⁸ This system will merely forward inbound messages to the internal mail server and forward outbound messages. There are some issues when installing this version of Sendmail with Solaris 9. For information regarding Sendmail and Solaris:

<http://www.sendmail.org/vendor/sun>

To complement our Sendmail server, we will be using MailScanner³⁹ to scan messages for dangerous files and trojan horses. MailScanner is a neat set of tools and scripts that can be configured with most mail servers and operating systems. MailScanner can be configured to handle mail spamming, and to scan for viruses using Sophos⁴⁰ virus identities (IDEs). As part of the daily cron job, it will

33 http://httpd.apache.org/docs-2.0/new_features_2_0.html

34 <http://www.verisign.com>

35 http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

36 <http://www.isc.org/products/BIND/bind9.html>

37 http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/Netra_t1_105.html

38 <http://www.sendmail.org/8.12.8.html>

39 <http://www.sng.ecs.soton.ac.uk/mailscanner>

40 <http://www.sophos.com/products/software/antivirus/savunix.html>

check and download any new virus IDE files.

With this configuration, no malicious or infected attachments will easily be entering our network.

Sendmail will also be configured to run as a user other than root for security purposes. For outgoing messages, the headers are stripped to protect the internal address scheme, and the internal mail server.

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁴¹

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

External NIDS

The external DNS server is running Solaris 9 FCS on a Sun Netra T1 105 with Snort 1.9.1.⁴² If any malicious pattern is detected, Snort will send information to our log server. Rules and patterns will be tweaked for our network traffic, so we can ignore certain packets and alert us if some are detected.

We have placed our NIDS behind the firewall to prevent the many false alarms and great amount of logging it would have otherwise generated. By placing it on the external service network, we can monitor that our filters are working properly, and since there is always the possibility that our exposed servers get compromised, we find it very useful to monitor all the network activity on this network.

The Netra T1 reached its End-of-Life on November, 2002.⁴³

Snort's user manual and IDS deployment guides can be found at:

<http://www.snort.org/docs>

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

1.8.6 Application network

This network hosts our application cluster, database cluster and our internal web server, and should only be accessible by sales staff (local and remote), partners, and developers.

Internal web server

The internal web server is running Solaris 9 FCS on a Sun Enterprise 220R. The server will be using Apache 2.0.44. It will be accessed by partners and remote sales staff through the VPN User network, and locally by sales staff and developers. No Verisign certificate will be installed on it, but like the external web server, it uses 128-bit SSL for its secure transactions. Whenever possible, it will be configured to show little or bogus information about itself.

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁴⁴

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Application server

The application server is really a cluster is composed of two Sun Enterprise 220Rs running Solaris 9 FCS and Sun Cluster⁴⁵ 3.0 software for High

⁴¹ http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

⁴² <http://www.snort.org>

⁴³ http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/Netra_t1_105.html

⁴⁴ http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

⁴⁵ <http://www.sun.com/software/cluster>

Availability. Without an application server, none of our transactions with our customers, suppliers or partners can be performed. Therefore it is critical that there is an instance of our application server running at all times. While the real IPs of both systems is 10.1.0.4 and 10.1.0.5, the virtual IP of the cluster, which is used by both of the web servers, is 10.1.0.3.

Both of these systems have a private interface to interact with the database server: 192.168.2.4 and 192.168.2.5, with a virtual IP for the cluster of 192.168.2.3. It is crucial that these systems are configured **not** to act as routers since they are multi-homed.

For additional information on Sun Clustering and how to make them more secure, we have included some links in the reference section.

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁴⁶

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Database server

The database server is also a cluster is composed of two Sun Enterprise 220Rs running Solaris 9 FCS and Sun Cluster⁴⁷ 3.0 software for High Availability. Without a database server, none of our transactions with our customers, suppliers or partners can be performed. Therefore it is critical that there is an instance of our database server running at all times. While the real IPs of both systems is 192.168.2.6 and 192.168.2.7, the virtual IP of the cluster, which is used by the application cluster, is 192.168.2.2. The database server is only connected to the application cluster. These units are sharing two Sun StorEdge T3+s⁴⁸ in a partner-pair configuration providing hardware RAID 5 with mirroring. Each T3 will have a RAID 5 copy of the data. This configuration will provides the maximum level of redundancy, availability, and data integrity for our critical data.

The T3s will not be on the network and can only be configured through their consoles, which are connected to the terminal concentrator.

For additional information on Sun Clustering and how to make them more secure, we have included some links in the reference section.

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁴⁹

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

1.8.7 Management network

The management network only has the management workstations, which are connected to the terminal concentrator.

Management workstations

These systems are Sun Ultra 60s running Trusted Solaris 8 4/01⁵⁰, which has earned a Common Criteria of EAL4.⁵¹ The systems are only used to monitor logs and connect to the terminal concentrator, which has access to the consoles of all critical systems in our infrastructure.

Both systems have a private interface (192.168.1.3 and 192.168.1.4,

46 http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

47 <http://www.sun.com/software/cluster>

48 <http://www.sun.com/storage/t3wg>

49 http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

50 <http://www.sun.com/software/solaris/trustedsolaris/index.html>

51 <http://www.commoncriteria.org/ccc/epl/productType/epldetail.jsp?id=110>

respectively) that allows them to connect to our terminal concentrator.

Since they are the only systems that can access the terminal concentrator, thus the console of every critical infrastructure system, it has been hardened with Trusted Solaris to make security breaches even more difficult.

It is crucial that these systems are configured **not** to act as routers since they are multi-homed.

The Ultra 60 reached its End-Of-Life on July 2002⁵²

We have included additional information on Trusted Solaris in our reference section.

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Terminal concentrator

The terminal concentrator is a Cisco 3640 running Cisco IOS 12.2 with two 32-port Asynchronous Network Modules⁵³, the consoles of all infrastructure systems, T3+s, and network devices are connected to it. This system has been configured to only accept ssh connections from the management workstations for administration and to access the consoles.

1.8.8 Internal service network

This network is composed of the internal DNS server, mail server, log cluster, web proxy, NIDS, and NTP server.

Internal DNS server

The internal DNS server is running Solaris 9 FCS on a Sun Netra T1 105 using version 9.2.2 of BIND.⁵⁴ The server will only have entries for the internal systems. All other queries will be sent to the external DNS server, which is configured recursively for queries from the internal DNS server. DNS is configured to run as a user other than root for security purposes. No zone transfers are allowed.

This system will serve the internal service network and the employee network.

The Netra T1 reached its End-Of-Life on November, 2002.⁵⁵

The BIND 9 Administration and configuration manual can be found at:

<http://www.nominum.com/content/documents/bind9arm.pdf>

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Internal mail server

The internal mail server is running Solaris 9 FCS on a Sun Enterprise 220R. The mail server will be using Sendmail 8.12.8.⁵⁶ This system will handle inbound messages, and forward outbound messages to the external mail server for delivery. There are some issues when installing this version of Sendmail with Solaris 9. Sendmail will also be configured to run as a user other than root for security purposes. For information regarding Sendmail and Solaris:

<http://www.sendmail.org/vendor/sun>

52 http://sunsolve.sun.com/handbook_pub/Systems/U60/U60.html

53 http://www.cisco.com/en/US/products/hw/routers/ps274/products_data_sheet09186a0080091b8b.html

54 <http://www.isc.org/products/BIND/bind9.html>

55 http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/Netra_t1_105.html

56 <http://www.sendmail.org/8.12.8.html>

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁵⁷

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Internal log cluster

The log server is really a cluster is composed of two Sun Ultra 60s running Solaris 9 FCS and Sun Cluster⁵⁸ 3.0 software for High Availability. Since we have a central logging design, it is important to maintain a log server running at all times. While the real IPs of both systems is 10.2.0.4 and 10.2.0.5, the virtual IP of the cluster is 10.2.0.6.

Swatch⁵⁹ will be running on the log server to help monitor logs, but primarily to set up alarms. Depending on the severity of the trigger, it will email or page our administrators. Setting up paging is trivial since most paging devices in the market allow customers to receive emails. Our scripts will just send an email, using our mail server. Therefore no new holes in our firewall have to be open.

For additional information on Sun Clustering and how to make them more secure, we have included some links in the reference section.

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Second to our database server, this is one of the most critical systems in our network. We have configured our infrastructure to use a different facility, depending on the type of service performed, when sending logs. This will make it easier to segregate logs.

Since it is a central log system, we have spent a great deal of time writing scripts to manipulate the logs. Our log cluster is using a fully loaded Sun StoreEdge D1000 with software-based RAID 5. There is enough space so that even in the event we get a DoS attack and a great amount of logs are generated, we will not run out of space, but even in such event, we have configured our log monitoring scripts using heuristics and based on the traffic and the timestamps, to compensate for this and group/erase log entries if necessary, so that we never run out of space.

The Ultra 60 reached its End-Of-Life on July 2002⁶⁰

Web proxy

The web proxy is running Solaris 9 FCS on a Sun Enterprise 220R. The server will be using Apache 2.0.44 with the Proxy module installed⁶¹. The only way for internal employees to access the internet, is through the web proxy. Moreover, it is only restricted to business-related activities.

We have allocated the external IP: x.x.x.6 as the external IP for this server using Static NAT.

Certain addresses have been banned for security purposes. Addresses of sites that use port 80 to bypass the firewall, such as GoToMyPC; and addresses used by instant messengers (Yahoo, Hotmail, AOL, etc.)

The Sun Enterprise 220R reached its End-Of-Life on August 2002.⁶²

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

⁵⁷ http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

⁵⁸ <http://www.sun.com/software/cluster>

⁵⁹ <http://swatch.sourceforge.net>

⁶⁰ http://sunsolve.sun.com/handbook_pub/Systems/U60/U60.html

⁶¹ http://httpd.apache.org/docs/mod/mod_proxy.html

⁶² http://sunsolve.sun.com/handbook_pub/Systems/E220R/E220R.html

Internal NIDS

The internal DNS server is running Solaris 9 FCS on a Sun Netra T1 105 with Snort 1.9.1.⁶³ If any malicious pattern is detected, Snort will send information to our log server. Rules and patterns will be tweaked for our network traffic, so we can ignore certain packets and alert us if some are detected.

We have added a NIDS on our internal service network to monitor that our external and internal filters are functioning properly. Systems on the internal service network can get compromised by a great variety of reasons: internal users might unintentionally have downloaded malicious code; our VPN perimeter, which has access to the internal mail server, might get compromised; the external DNS/mail servers were compromised and used to compromised the internal DNS/mail server; or perhaps it was done internally. Regardless, the NIDS can help us detect any suspicious activity originating from our internal network.

The Netra T1 reached its End-of-Life on November, 2002.⁶⁴

Snort's user manual and IDS deployment guides can be found at:

<http://www.snort.org/docs>

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

Internal NTP Server

The internal NTP is running Solaris 9 FCS on an Ultra 60. Solaris' xntpd uses Network Time Protocol (NTP) version 3 standard, as defined by RFC1305⁶⁵. NTPV4 has not been formalized in an RFC yet. However, we will install NTP v4.1.1 from www.ntp.org to take advantage of new encryption and authentication features.

We will configure the NTP server and its clients to use encryption and authentication, and using NTP access lists, we can limit which systems can connect to it. For information on configuring NTP:

<http://www.eecis.udel.edu/~mills/ntp/html/index.html>

The internal NTP server will keep in sync using the external NTP server, and have systems on the internal service network as clients.

The Ultra 60 reached its End-Of-Life on July 2002⁶⁶

For more information on NTP, additional links have been included in our reference section.

Additional information on how to configure and secure the system can also be found at the end of this document in the reference section.

1.8.9 Employee network

This network consists of multiple subnets, workstations, printers, etc. An entire class B (10.3.0.0/16) has been dedicated to this network, however, only two of its subnets are of particular interest to us: The subnets used by the local sales staff (10.3.1.0/24), and the developers and testers of our Fortune Software (10.3.2.0/24). All the other subnets are of no particular interest to our design. This allows for growth within the employee network, without affecting our network infrastructure.

⁶³ <http://www.snort.org>

⁶⁴ http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/Netra_t1_105.html

⁶⁵ <http://www.eecis.udel.edu/~mills/database/rfc/rfc1305>

⁶⁶ http://sunsolve.sun.com/handbook_pub/Systems/U60/U60.html

2. Assignment 2: Security Policy and Tutorial

2.1 Security Policies for the Cisco Routers

In order to satisfy our access requirements for the different types of users and enforce the need-to-access policy outlined in the first section, we need to define the security policies to be implemented in our routers. This includes the specific sets of access control lists, services offered and configuration settings particular to the routers. The requirements were defined to take care of business needs while keeping security considerations in mind.

Access Control Lists

We have designed our network to limit all traffic, regardless of whether it is internal or external, to access only the services offered and only by users that need explicit access to those services or networks. Even if a particular system or network, internal or external, is compromised, any damage that can be caused from these systems is limited and cannot easily affect all critical systems in our infrastructure.

We have chosen to filter packets as they enter an interface and pass through the routers. This will prevent wasting resources routing a packet that might get dropped at the other end.

Access Control Lists are processed sequentially until a match is found. Therefore to maximize efficiency and performance, whenever the order allows it, rules that are used more often should be listed before rules that we expect to have fewer matches.

To take advantage of the higher filtering capabilities, we will be using extended access control lists on our routers. We are counting on our firewall to maintain state for those incoming connections we do let in.

2.1.1 ACL for the border router external interface

Give our access list a meaningful name.

ip access-list extended INGRESS

Protect ourselves and the firewall first

deny ip any host y.y.y.1 log-input

deny ip any host x.x.x.14 log-input

deny ip any host x.x.x.13 log-input

Prevent spoofed packets from coming in.

deny ip x.x.x.0 0.0.0.15 any log-input

Drop multicast (224.0.0.0-239.0.0.0) addresses, since we do not use it internally.

deny ip 224.0.0.0 31.255.255.255 any log-input

Do not let fragmented traffic in.

deny ip any any fragments log-input

Drop all ICMP traffic. It can be exploited in many ways, plus we can do without it.

deny icmp any any log-input

Drop the loopback address. It should never be seen online.
deny ip 127.0.0.0 0.255.255.255 any log-input

Private addresses (RFC 1918⁶⁷) should not be seen coming from the external interface either.

deny ip 10.0.0.0 0.255.255.255 any log-input
deny ip 172.16.0.0 0.15.255.255 any log-input
deny ip 192.168.0.0 0.0.255.255 any log-input

IANA's reserved addresses⁶⁸ should never be seen online.
deny ip 0.0.0.0 0.255.255.255 any
deny ip 1.0.0.0 0.255.255.255 any
deny ip 2.0.0.0 0.255.255.255 any
... (not all addresses are shown)
deny ip 255.0.0.0 0.255.255.255 any

Block the top 20 list of attacking networks from incidents.org.⁶⁹ These are REAL Networks, so we will check the application logs to make sure none of our existing customers, partners, or suppliers come from these addresses before we block them. We can also add the IP of usual attackers to our network in here.

Note: the incidents.org list is updated very often.

deny ip 61.38.36.0 0.0.0.255 any log-input
deny ip 203.130.130.0 0.0.0.255 any log-input
deny ip 80.38.177.0 0.0.0.255 any log-input
deny ip 219.41.0.0 0.0.0.255 any log-input
deny ip 61.79.96.0 0.0.0.255 any log-input
deny ip 210.4.143.0 0.0.0.255 any log-input
deny ip 148.245.53.0 0.0.0.255 any log-input
deny ip 218.48.8.0 0.0.0.255 any log-input
deny ip 211.255.136.0 0.0.0.255 any log-input
deny ip 210.108.158.0 0.0.0.255 any log-input
deny ip 195.159.152.0 0.0.0.255 any log-input
deny ip 211.255.136.0 0.0.0.255 any log-input
deny ip 163.25.96.0 0.0.0.255 any log-input
deny ip 24.214.48.0 0.0.0.255 any log-input
deny ip 140.123.33.0 0.0.0.255 any log-input
deny ip 68.52.187.0 0.0.0.255 any log-input
deny ip 63.149.88.0 0.0.0.255 any log-input
deny ip 193.224.167.0 0.0.0.255 any log-input
deny ip 193.154.8.0 0.0.0.255 any log-input
deny ip 24.102.205.0 0.0.0.255 any log-input

Finally we will be letting some of the traffic through to our services. We will be listing the rules in decreasing order. From the service we expect will see more traffic, to the service that we expect will receive least traffic.

External web server: Allow customers and suppliers to access our web server
permit tcp any host x.x.x.2 eq 80
permit tcp any host x.x.x.2 eq 443

External IP (using NAT) of web proxy: Allow replies to the external IP of the web proxy used by the internal users.

Note: The established check can easily be fooled, but our firewall can verify that it is a valid reply.

⁶⁷ <http://www.isi.edu/in-notes/rfc1918.txt>

⁶⁸ <http://www.iana.org/assignments/ipv4-address-space>

⁶⁹ <http://isc.incidents.org/top10.html>

permit tcp any host x.x.x.6 eq 80 established
permit tcp any host x.x.x.6 eq 443 established

VPN concentrator: Allow telecommuters, remote sales staff, partners, and remote administrators to connect to the VPN concentrator. This includes: UDP port 500 for IKE, and TCP (ESP) traffic.

permit udp any host x.x.x.9 eq 500 log-input
permit esp any host x.x.x.9 log-input

External mail relay (mail): Allow mail to access the mail relay.
permit tcp any host x.x.x.4 eq 25

External DNS server: Allow public access to the external DNS server. We will only be letting UDP traffic in.

permit udp any host x.x.x.3 eq 53

External NTP server: Allow the public NTP servers to sync our external NTP server. The sanitized IP of the public NTP servers are: a.a.a.a and b.b.b.b.

Note: The established check can easily be fooled, but our firewall can verify that it is a valid reply.

permit udp host a.a.a.a host x.x.x.1 eq 123 established
permit udp host b.b.b.b host x.x.x.1 eq 123 established

External mail relay (virus updates): The cron job only runs once a day, but the mail server needs to download virus-updates from the antivirus page. For this design, we sanitize the IP of the virus site to c.c.c.c.

Note: The established check can easily be fooled, but our firewall can verify that it is a valid reply.

permit tcp host c.c.c.c host x.x.x.4 established log-input

Explicit drop: This might generate a lot of unnecessary logs, but could be disabled if desired, or tune to reduce 'noise'

deny ip any any log-input

2.1.2 ACL for the border router internal interface

Use a meaningful name for the ACL: EGRESS

ip access-list extended EGRESS

Protect ourselves first.

deny ip any host y.y.y.1 log-input
deny ip any host x.x.x.14 log-input

Let source-quenchers out for flow optimization. Since ICMP expired TTL, unreachable, echo replies, etc. can be used to map our network, we will only let quenchers out and drop the rest.

permit icmp x.x.x.0 0.0.0.15 any source-quench

External service network, VPN, and web proxy: We are letting all traffic from the external service network (including the external IP of the web proxy), and the VPN network out. At this point, we assume the firewall has done its job.

permit tcp x.x.x.0 0.0.0.15 any
permit udp x.x.x.0 0.0.0.15 any

Explicit drop: Explicitly drop everything else. We do this to prevent spoofed packets,

packets that can be used to map our network or reveal our internal address scheme, and packets that can be used to launch attacks from leaving. We do want to log it so we can monitor and investigate later if needed.

```
deny ip any any log-input
```

2.1.3 Alternative ACL for the border router internal interface

If it does not degrade performance or causes the router to drop packets under high loads, an alternative ACL could be:

```
ip access-list extended EGRESS  
deny ip any host y.y.y.1 log-input  
deny ip any host x.x.x.14 log-input  
permit icmp x.x.x.0 0.0.0.15 any source-quench  
permit tcp host x.x.x.2 eq 80 any  
permit tcp host x.x.x.2 eq 443 any  
permit tcp host x.x.x.6 eq 80 any  
permit tcp host x.x.x.6 eq 443 any  
permit tcp host x.x.x.4 eq 25 any  
permit esp host x.x.x.9 any  
permit udp host x.x.x.9 eq 500 any  
permit udp host x.x.x.3 eq 53 any  
permit udp host x.x.x.1 eq 123 host a.a.a.a  
permit udp host x.x.x.1 eq 123 host b.b.b.b  
permit tcp host x.x.x.4 eq 80 host c.c.c.c log-input  
deny ip any any log-input
```

2.1.4 ACL for the internal router (Application interface)

Use a meaningful name for the ACL: APP_EGRESS

```
ip access-list extended APP_EGRESS
```

Protect ourselves and the firewall first.

```
deny ip any host 10.1.0.1 log-input  
deny ip any host 10.2.0.1 log-input  
deny ip any host 10.3.0.1 log-input  
deny ip any host 10.4.0.1 log-input  
deny ip any host 10.5.0.1 log-input  
deny ip any host x.x.x.13 log-input  
deny ip any host 10.100.1.13 log-input  
deny ip any host 10.100.2.2 log-input  
deny ip any host 10.0.0.2 log-input  
deny ip any host 10.5.0.2 log-input
```

Application server: Allow traffic to the external web server from the application. Our firewall can verify that it is a valid reply. Also, allow logs and NTP queries.

```
permit tcp host 10.1.0.3 eq 5000 host 10.100.1.2 established  
permit udp host 10.1.0.3 host 10.2.0.9 eq 123  
permit udp host 10.1.0.3 host 10.2.0.6 eq 514
```

Web server: Allow web traffic back to the local sales staff, the remote sales staff, partners, and developers. Our firewall can verify that it is a valid reply. Also, allow logs and NTP queries.

```
permit tcp host 10.1.0.2 eq 80 10.3.1.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 443 10.3.1.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 80 10.0.2.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 443 10.0.2.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 80 10.0.1.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 443 10.0.1.0 0.0.0.255 established  
permit tcp host 10.1.0.2 eq 80 10.3.2.0 0.0.0.255 established
```

```
permit tcp host 10.1.0.2 eq 443 10.3.2.0 0.0.0.255 established
permit udp host 10.1.0.2 host 10.2.0.9 eq 123
permit udp host 10.1.0.2 host 10.2.0.6 eq 514
```

Developers access: Allow ssh traffic to go back to the developers and testers. Our firewall can verify that it is a valid reply.

```
permit tcp any 10.3.2.0 0.0.0.255 eq 22 established
```

Explicit drop: Drop everything else and log it

```
drop ip any any log-input
```

2.1.5 ACL for the internal router (Management interface)

Use a meaningful name for the ACL: MGMT_EGRESS

```
ip access-list extended MGMT_EGRESS
```

Protect ourselves and the firewall first.

```
deny ip any host 10.1.0.1 log-input
deny ip any host 10.2.0.1 log-input
deny ip any host 10.3.0.1 log-input
deny ip any host 10.4.0.1 log-input
deny ip any host 10.5.0.1 log-input
deny ip any host x.x.x.13 log-input
deny ip any host 10.100.1.13 log-input
deny ip any host 10.100.2.2 log-input
deny ip any host 10.0.0.2 log-input
deny ip any host 10.5.0.2 log-input
```

Remote administrators: Allow SSH traffic back to the VPN address pool of remote administrators. Our firewall can verify that it is a valid reply.

```
permit tcp any eq 22 10.0.4.0 0.0.0.255 established
```

Explicit drop: Drop everything else and log it

```
drop ip any any log-input
```

2.1.6 ACL for the internal router (Internal service network interface)

Use a meaningful name for the ACL: SERVICE_EGRESS

```
ip access-list extended SERVICE_EGRESS
```

Protect ourselves and the firewall first

```
deny ip any host 10.1.0.1 log-input
deny ip any host 10.2.0.1 log-input
deny ip any host 10.3.0.1 log-input
deny ip any host 10.4.0.1 log-input
deny ip any host 10.5.0.1 log-input
deny ip any host x.x.x.13 log-input
deny ip any host 10.100.1.13 log-input
deny ip any host 10.100.2.2 log-input
deny ip any host 10.0.0.2 log-input
deny ip any host 10.5.0.2 log-input
```

Web proxy: Allow all web traffic from the web proxy out except to local sites.

```
deny tcp host 10.2.0.7 10.0.0.0 255.0.0.0 log-input.
permit tcp host 10.2.0.7 any eq 80
permit tcp host 10.2.0.7 any eq 443
```

Internal mail server: Allow mail traffic to the employee network, external mail relay, remote sales staff VPN address pool, and telecommuters VPN address pool.

```
permit tcp host 10.2.0.3 eq 25 10.3.0.0 0.0.255.255
permit tcp host 10.2.0.3 eq 25 host 10.100.1.4
permit tcp host 10.2.0.3 eq 25 10.0.2.0 0.0.0.255
permit tcp host 10.2.0.3 eq 25 10.0.3.0 0.0.0.255
```

Internal DNS server: Allow DNS traffic to the employee network and external DNS server. Our firewall can verify that it is a valid reply.

```
permit udp host 10.2.0.2 eq 53 10.3.0.0 0.0.255.255 established
permit udp host 10.2.0.2 eq 53 host 10.100.1.3
```

Internal NTP server: Allow NTP traffic to the external NTP server

```
permit udp host 10.2.0.9 eq 123 host 10.100.1.1
```

Explicit drop: Drop everything else and log it

```
drop ip any any log-input
```

2.1.7 ACL for the internal router (Employee interface)

Use a meaningful name for the ACL: SERVICE_EGRESS

```
ip access-list extended EMPLOYEE_EGRESS
```

NTP replies: Allow NTP replies to the application cluster, the internal web server and the internal router:

```
permit udp host 10.2.0.9 eq 123 host 10.1.0.3
permit udp host 10.2.0.9 eq 123 host 10.1.0.2
permit udp host 10.2.0.9 eq 123 host 10.2.0.1
```

Protect ourselves and the firewall first.

```
deny ip any host 10.1.0.1 log-input
deny ip any host 10.2.0.1 log-input
deny ip any host 10.3.0.1 log-input
deny ip any host 10.4.0.1 log-input
deny ip any host 10.5.0.1 log-input
deny ip any host x.x.x.13 log-input
deny ip any host 10.100.1.13 log-input
deny ip any host 10.100.2.2 log-input
deny ip any host 10.0.0.2 log-input
deny ip any host 10.5.0.2 log-input
```

Web proxy: Allow all web traffic to the web proxy.

```
permit tcp any host 10.2.0.7 eq 80
permit tcp any host 10.2.0.7 eq 443
```

Internal web server: Allow local sales staff access to internal web server.

```
permit tcp 10.3.1.0 0.0.0.255 host 10.2.0.7 eq 80
permit tcp 10.3.1.0 0.0.0.255 host 10.2.0.7 eq 443
```

Internal mail server: Allow access to the internal mail server.

```
permit tcp any host 10.2.0.3 eq 25
```

Internal DNS server: Allow access to the internal DNS server.

```
permit udp any host 10.2.0.2 eq 53
```

Application network: Allow developers ssh access to the application network

```
permit tcp 10.3.2.0 0.0.0.255 10.1.0.0 0.0.0.255 eq 22
```

Explicit drop: Drop everything else and log it

drop ip any any log-input

2.2 Security Policy and Configuration for the Cisco PIX

The Cisco PIX Firewalls run a different operating system than the routers we have configured. Usual commands sometimes have an slightly different syntax, and needless to say, commands from Cisco IOS might not work the same way (or even exist) on the PIX firewalls, and vice versa. We have chosen to assign the same security level to the interfaces connected to the external service network, VPN network and VPN user network to ensure no traffic flows between them.

The following table outlines the interface configuration of the firewall:

Interface	Interface Name	Security Level	IP	Device/Network
ethernet0	outside	0	x.x.x.13	Border Router
ethernet1	INT_service	50	10.100.1.13	External Service Network
ethernet2	INT_VPN	50	10.100.2.2	VPN Network
ethernet3	INT_VPNuser	50	10.0.0.2	VPN User Network
ethernet4	inside	100	10.5.0.2	Internal Router

Table 21: Interface configuration for the firewall

For convenience we will outline the traffic that needs to go through the firewall. It is important to mention that: replies are not listed; communicating between different security zones requires the use of *nat/static*, commands depending on whether we are going from a higher security level to a lower security level or vice versa; interfaces with the same security level cannot talk to each other.

Source	Interface	Destination	Interface	
Border router	outside	External NTP server	INT_service	static
Border router	outside	Log cluster	inside	static
Internet	outside	External mail relay	INT_service	static
Internet	outside	External web server	INT_service	static
Internet	outside	VPN concentrator	INT_VPN	static
Internet	outside	External DNS server	INT_service	static
External NTP server	INT_service	Internet (Public NTP servers)	outside	nat
External mail server	INT_service	Internet (mail servers)	outside	nat
External mail server	INT_service	Internet (antivirus page)	outside	nat
External DNS server	INT_service	Internet (DNS servers)	outside	nat
External mail server	INT_service	Internal mail server	inside	static
External service network	INT_service	Log cluster	inside	static
External web server	INT_service	Application cluster	inside	static
VPN concentrator	INT_VPN	Internet (VPN client)	outside	nat
VPN concentrator	INT_VPN	Log cluster	inside	static
VPN (Remote sales staff & telecommuters)	INT_VPNuser	Internal mail server	inside	static

Source	Interface	Destination	Interface	
VPN (Remote sales staff & partners)	INT_VPNUser	Internal web server	inside	static
VPN (Remote administrators)	INT_VPNUser	Management workstations	inside	static
Internal NTP server	inside	External NTP server	INT_service	nat
Internal mail server	inside	External mail server	INT_service	nat
Internal DNS server	inside	External DNS server	INT_service	nat
Web proxy	inside	Internet	outside	nat

Table 22: Firewall policy

In our access-control lists we use names for the interfaces and the hosts involved. Therefore, we will list the names first and cover some security issues.

2.2.1 PIX configuration

Names

!We will declare the IPs we will use in our configuration.

!Enable IP and name association

names

!Declare IP of the interfaces

!Declare the IP of the external interface

name x.x.x.13 FW_OUTSIDE

!Declare the IP of the external service network interface

name 10.100.1.13 FW_EXTservice

!Declare the IP of the VPN network interface

name 10.100.2.2 FW_VPN

!Declare the IP of the VPN User network interface

name 10.0.0.2 FW_VPNUser

!Declare the IP of the internal network interface

name 10.5.0.2 FW_INSIDE

!Declare IP of border router

name x.x.x.14 BORDER_ROUTER

!Declare IP of internal router

name 10.5.0.1 INTERNAL_ROUTER

!External service network

!Declare IP of the external NTP server

name 10.100.1.1 EXT_NTP

!Declare IP of the external web server

name 10.100.1.2 EXT_WEB

!Declare IP of the external DNS server

name 10.100.1.3 EXT_DNS

!Declare IP of the external mail server

name 10.100.1.4 EXT_MAIL

!Declare IP of the external NIDS

name 10.100.1.5 EXT_NIDS

!VPN network

!Declare IP of the VPN concentrator public interface

name 10.100.2.9 VPN_PUB

!VPN User Network

!Declare IP of the VPN Concentrator private interface

name 10.0.0.9 VPN_PRIV

!Application network
!Declare IP of the application cluster
name 10.1.0.3 APP_CLUSTER
!Declare IP of the internal web server
name 10.1.0.2 INT_WEB

!Internal service network
!Declare IP of the internal DNS server
name 10.2.0.2 INT_DNS
!Declare IP of the internal mail server
name 10.2.0.3 INT_MAIL
!Declare IP of the log cluster
name 10.2.0.6 LOG_CLUSTER
!Declare IP of the web proxy
name 10.2.0.7 WEB_PROXY
!Declare IP of the internal NTP
name 10.2.0.9 INT_NTP

!Management network
!Declare IP of management workstations.
name 10.4.0.3 MGMT1
name 10.4.0.4 MGMT2

!Declare IP of internet hosts we use in our configuration
!Declare IP of public NTP server1
name a.a.a.a PUB_NTP1
!Declare IP of public NTP server2
name b.b.b.b PUB_NTP2
!Declare IP of antivirus site
name c.c.c.c ANTIVIRUS_SITE

!Declare IP of external IPs used with static NAT
!Declare external IP of the external NTP server
name x.x.x.1 NAT_NTP
!Declare external IP of the external web server
name x.x.x.2 NAT_WEB
!Declare external IP of the external DNS server
name x.x.x.3 NAT_DNS
!Declare external IP of the external mail server
name x.x.x.4 NAT_MAIL
!Declare external IP of the web proxy
name x.x.x.6 NAT_PROXY
!Declare external IP of the web proxy
name x.x.x.7 NAT_LOG
!Declare external IP of the VPN Concentrator public interface
name x.x.x.9 NAT_VPN

Interfaces

!Give a name to the interfaces and assign them a security level
!External interface
nameif ethernet0 outside security0
!External Service Network Interface
nameif ethernet1 INT_service security50
!VPN network interface
nameif ethernet2 INT_VPN security50
!VPN User network interface
nameif ethernet3 INT_VPNUser security50
!Internal interface
nameif ethernet4 inside security100

!Assign an IP address and netmask to each interface
 !Assign an IP to the external interface
ip address outside FW_OUTSIDE 255.255.255.252
 !Assign an IP to the external service network interface
ip address INT_service FW_EXTservice 255.255.255.0
 !Assign an IP to the VPN network interface
ip address INT_VPN FW_VPN 255.255.255.0
 !Assign an IP to the VPN User network interface
ip address INT_VPNuser FW_VPNuser 255.255.255.0
 !Assign an IP to the internal network interface
ip address inside FW_INSIDE 255.255.255.0

Network Address Translation (NAT)

Lower security to higher security level.

Outside -> External service network/Internal service network/VPN network:

!Allow traffic flow from between the outside and the external service network
 !Allow traffic from the border router to the external NTP server
static(INT_service, outside) NAT_NTP EXT_NTP netmask 255.255.255.255
 !Allow traffic from the border router to the log cluster
static(inside, outside) NAT_LOG LOG_CLUSTER netmask 255.255.255.255
 !Allow traffic from outside to the external mail server
static(INT_service, outside) NAT_MAIL EXT_MAIL netmask 255.255.255.255
 !Allow traffic from outside to the external web server
static(INT_service, outside) NAT_WEB EXT_WEB netmask 255.255.255.255
 !Allow traffic from outside to the VPN concentrator
static(INT_VPN, outside) NAT_VPN VPN_PUB netmask 255.255.255.255
 !Allow traffic from outside to the external DNS server
static(INT_service, outside) NAT_DNS EXT_DNS netmask 255.255.255.255

External service network -> Internal service network/Application network:

!Allow traffic flow between the external service network and the internal service network
 !This allows: Logs from all systems on the external service network and mail from the mail relay
static (inside, INT_service) 10.2.0.0 10.2.0.0 netmask 255.255.255.0
 !Allow traffic flow between the external web server and the application cluster
static (inside, INT_service) 10.1.0.0 10.1.0.0 netmask 255.255.255.0

VPN network -> Internal service network:

!Allow traffic flow between the VPN Network and the internal service network
 !This allows the VPN concentrator to send logs.
static (inside, INT_VPN) 10.2.0.0 10.2.0.0 netmask 255.255.255.0

VPN User network -> Internal service network/Application network:

!Allow traffic flow between the VPN User network and the internal service network
 !This allows: access to mail for remote sales staff and telecommuters
static (inside, INT_VPNuser) 10.2.0.0 10.2.0.0 netmask 255.255.255.0
 !Allow traffic flow between VPN User network and application network
 !This allows: access to the internal web server for remote sales staff and partners
static (inside, INT_VPNuser) 10.1.0.0 10.1.0.0 netmask 255.255.255.0
 !This allows: access to the management workstations for remote administrators
static (inside, INT_VPNuser) 10.4.0.0 10.4.0.0 netmask 255.255.255.0

Higher security level to lower security level

External service network -> Outside:

!Allow traffic out of the external service network
 !Allow the external NTP server to go outside
global (outside) 1 NAT_NTP
nat (INT_service) 1 EXT_NTP 255.255.255.255

!Allow the external mail server to go outside for mail delivery and virus updates

global (outside) 2 NAT_MAIL

nat (INT_service) 2 EXT_MAIL 255.255.255.255

!Allow the external DNS server to go outside

global (outside) 3 NAT_DNS

nat (INT_service) 3 EXT_DNS 255.255.255.255

VPN network ->Outside

!Allow the VPN concentrator to send IKE/ESP traffic out

global (outside) 4 NAT_VPN

nat (INT_VPN) 4 EXT_VPN 255.255.255.255

Internal service network -> External service network/Outside:

!Allow traffic out of internal service network

!Allow the internal NTP server to communicate with the external NTP server

nat (inside) 0 INT_NTP 255.255.255.255

!Allow the internal mail server to communicate with the external mail server

nat (inside) 0 INT_MAIL 255.255.255.255

!Allow the internal DNS server to communicate with the external DNS server

nat (inside) 0 INT_DNS 255.255.255.255

!Allow the web proxy to go outside

global (outside) 5 NAT_PROXY

nat (inside) 5 WEB_PROXY 255.255.255.255

Stateful Packet Inspection

!Use Cisco's Adaptive Security Algorithm (ASA)

!HTTP traffic

fixup protocol http 80

!DNS traffic

fixup protocol domain 53

!mail traffic

fixup protocol smtp 25

Logging

!Enable logging

logging on

!Disable logging to the console

no logging console

!Use facility4 as the logging facility

logging facility 20

!Include time stamps on logs

logging timestamp

!Send informational logs to syslog

logging trap informational

!No buffering on the PIX. This could be turned on when troubleshooting.

no logging buffered

!Declare log server and interface

logging host inside LOG_CLUSTER

ICMP

By default **ALL** ICMP traffic that is directed to a PIX interface will be accepted, regardless of the ACL. This is not good. We will disable ICMP directed to all interfaces.

!Block all ICMP traffic directed to all interfaces

icmp deny any outside

icmp deny any INT_service

icmp deny any INT_VPN

icmp deny any INT_VPNuser

icmp deny any inside

Flood protection and fragments

We will enable Cisco's Flood Defender against flood attacks and block all fragmented traffic on all interfaces.

!Enable flood protection

floodguard enable

!Block packets if they are not whole.

fragment chain 1 outside

fragment chain 1 INT_service

fragment chain 1 INT_VPN

fragment chain 1 INT_VPNuser

fragment chain 1 inside

IP audit

While we will not be dropping packets, or sending resets automatically if an IDS signature is detected (this could be used as a DoS), we would like to generate alerts if attack or information patterns are seen. The IDS on the PIX leaves a lot to be desired, but it does not hurt to have it enabled.

ip audit attack action alarm

ip audit info action alarm

!Create a named ip audit for attacks and for information signatures and apply them to all interfaces

ip audit name attack_alert attack action alarm

ip audit name info_alert info action alarm

!Apply ip audit actions

ip audit interface outside attack_alert

ip audit interface outside info_alert

ip audit interface INT_service attack_alert

ip audit interface INT_service info_alert

ip audit interface INT_VPN attack_alert

ip audit interface INT_VPN info_alert

ip audit interface INT_VPNuser attack_alert

ip audit interface INT_VPNuser info_alert

ip audit interface inside attack_alert

ip audit interface inside info_alert

2.2.2 Access Control Lists

We have chosen to filter packets as they enter an interface. This will prevent wasting resources routing a packet that might get dropped at the other end. An important aspect to mention is that stateful replies are allowed back out. Therefore, we will not list them in our access-control lists.

2.2.3 ACL for the external interface (outside)

Drop spoofed: Just to make sure spoofed traffic does not enter our network.

access-list from_outside deny ip 10.0.0.0 0.255.255.255 any

External web server: Allow access to the external web server

access-list from_outside permit tcp any host NAT_WEB eq www

access-list from_outside permit tcp any host NAT_WEB eq https

VPN concentrator: Allow access to the VPN concentrator

access-list from_outside permit udp any host NAT_VPN eq isakmp

access-list from_outside permit esp any host NAT_VPN

External mail server: Allow access to the external mail relay

access-list from_outside permit tcp any host NAT_MAIL eq smtp

External DNS server: Allow access to the external DNS server

access-list from_outside permit udp any host NAT_DNS eq domain

Logs: Allows logs from the router to access the log cluster.

access-list from_outside permit udp host BORDER_ROUTER host NAT_LOG eq syslog

External NTP server: Allow NTP requests from the border router to the external NTP server.

access-list from_outside permit udp host BORDER_ROUTER host NAT_NTP eq ntp

Explicit drop: Explicitly drop everything else.

access-list from_outside deny any any

Apply ACL to interface :

access-group from_outside in interface outside

2.2.4 ACL for the external service network interface (INT_service)

External web server: Allow external web server to send traffic to the application cluster.

access-list from_service permit tcp host EXT_WEB host APP_CLUSTER eq 5000

External mail relay (mail): Allow the external mail relay to deliver outbound mail and forward mail to the internal mail server.

access-list from_service permit tcp host EXT_MAIL any eq smtp

External DNS server: Allow the external DNS server to perform recursive queries for internal users.

access-list from_service permit udp host EXT_DNS any eq domain

Logs: Allow logs from systems on the external service network to the log cluster.

access-list from_service permit udp 10.100.1.0 0.0.0.255 host LOG_CLUSTER eq syslog

External NTP server: Allow the external NTP server to query the public NTP servers.

access-list from_service permit udp host EXT_NTP host PUB_NTP1 eq ntp

access-list from_service permit udp host EXT_NTP host PUB_NTP2 eq ntp

External mail relay (virus update): Allow the external mail relay to update its virus definitions from the antivirus download page.

access-list from_service permit tcp host EXT_MAIL host ANTIVIRUS_SITE eq www

Explicit drop: Explicitly drop everything else.

access-list from_service deny any any

Apply ACL to interface :

access-group from_service in interface INT_service

2.2.5 ACL for the VPN network interface (INT_VPN)

VPN concentrator: Permit ESP and IKE traffic out

access-list from_VPN permit esp host VPN_PUB any

access-list from_VPN permit udp host VPN_PUB any eq isakmp

Logs: Allow logs from the VPN concentrator to the log cluster.
access-list from_VPN permit udp host VPN_PUB host NAT_LOG eq syslog

Explicit drop: Explicitly drop everything else.
access-list from_VPN deny any any

Apply ACL to interface :
access-group from_VPN in interface INT_VPN

2.2.6 ACL for the VPN User network interface (INT_VPNUser)

Internal web server: Allow partners and remote sales staff to access the internal web server.

!Partners
access-list from_VPNUser permit tcp 10.0.1.0 0.0.0.255 host INT_WEB eq www
access-list from_VPNUser permit tcp 10.0.1.0 0.0.0.255 host INT_WEB eq https
!Remote sales staff
access-list from_VPNUser permit tcp 10.0.2.0 0.0.0.255 host INT_WEB eq www
access-list from_VPNUser permit tcp 10.0.2.0 0.0.0.255 host INT_WEB eq https

Internal mail server: Allow remote sales staff and telecommuters access to the internal mail server.

!Remote sales staff
access-list from_VPNUser permit tcp 10.0.2.0 0.0.0.255 host INT_MAIL eq smtp
!Telecommuters
access-list from_VPNUser permit tcp 10.0.3.0 0.0.0.255 host INT_MAIL eq smtp

Management workstations: Allow remote administrators to access their workstations.

access-list from_VPNUser permit tcp 10.0.4.0 0.0.0.255 host MGMT1 eq ssh
access-list from_VPNUser permit tcp 10.0.4.0 0.0.0.255 host MGMT2 eq ssh

Explicit drop: Explicitly drop everything else.
access-list from_VPNUser deny any any

Apply ACL to interface :
access-group from_VPNUser in interface INT_VPNUser

2.2.7 ACL for the internal network interface (inside)

Web proxy: Allow internal users to access the outside world through the web proxy. Block traffic going to any internal sites.

access-list from_inside deny tcp host WEB_PROXY 10.0.0.0 0.0.0.255
access-list from_inside permit tcp host WEB_PROXY any eq www
access-list from_inside permit tcp host WEB_PROXY any eq https

Internal DNS server: Allow the internal DNS server to send queries to the external DNS server.

access-list from_inside permit udp host INT_DNS host EXT_DNS eq domain

Internal mail server: Allow the internal mail relay to forward mail to the external mail relay.

access-list from_inside permit tcp host INT_MAIL host EXT_MAIL eq smtp

Internal NTP server: Allow the internal NTP server to query the external NTP server.

access-list from_inside permit udp host INT_NTP host EXT_NTP eq ntp

Explicit drop: Explicitly drop everything else.
access-list from_inside deny any any

Apply ACL to interface :
access-group from_inside in interface inside

2.3 Security Policy for the VPN Concentrator

We will be using a Cisco 3015 VPN Concentrator running Cisco VPN Concentrator Software 3.6 to provide remote access to our partners, and remote users (remote sales staff, telecommuters, and remote administrators).

At this particular time, all of our users will connect using Cisco's VPN client and there will be no LAN-to-LAN connections.

As part of the security policy, we require that all of our remote users (partners included) are instructed on the concept that from our point of view, the remote system they use to connect to our VPN network is one of the weakest points of our perimeter. This also includes general guides and assistance on safe internet practices. It will require some training and guidance from our part, but we believe it is well worth the effort. Therefore, we will repeatedly stress the importance of having the product we have chosen: Norton Internet Security 2003, to be enabled and properly configured at all times. No split tunneling is allowed, and the VPN client should be configured to disable access to the local LAN. The IDS, firewall, system and mail antivirus software included in that bundle, will help us mitigate some of the risks.

To ensure consistency and compatibility with our settings, and make it easier to instruct our remote users to establish a VPN connection, we will insist that they only use the Cisco VPN client.

2.3.1 Security associations Tunneling Protocols

Although PPTP and L2TP are popular tunneling protocols with Microsoft systems, we have decided to use IPsec as our only tunneling protocol based on security aspects. PPTP and L2TP will be disabled.

Encryption, Authentication and Encapsulation Mode

Also for security reasons, we have chosen ESP over AH. While ESP will not provide the same level of authentication that AH provides. We sacrifice authenticating non-mutable headers for ESP's encryption capabilities. Considering the nature of the data that will be sent, we prefer to have this information as protected as we can. Not to mention, AH will not work properly with NAT.

Tunnel encapsulation mode will also be chosen over transport encapsulation mode. By providing this extra layer of protection, we can encrypt and authenticate the entire original packet, and with it, our internal IP address scheme.

Authentication Algorithm

The VPN concentrator offers us three choices:

- None
- ESP/MD5/HMAC-128(default)
- ESP/SHA/HMAC-160

Despite the higher overhead, we have chosen the SHA-1 hash function with the 160-bit key over the default. Considering that our remote users will only be using a web browser to the application cluster, a mail client, or ssh client. The overhead is acceptable for this non network-intensive traffic.

Encryption Algorithm

For encryption algorithm, we have three choices:

- None
- DES-56
- 3DES-168 (default)

On this particular case, we will use the default encryption mode, 3DES-168, which provides the highest level of encryption.

Perfect Forward Secrecy

We have four options on perfect forward secrecy:

- Disabled (default)
- Group 1 (768-bits)
- Group 2 (1024-bits)
- Group 7 (ECC)

We do not want keys on phase 2 to be based on phase 1 keys, therefore we will enable perfect forward secrecy, and choose Group 2 (Diffie-Hellman Group 2) to generate sessions keys with generator and prime numbers of 1024 bits. We consider that the higher overhead is acceptable.

Lifetime Measurement

Lifetime measurement presents four choices:

- Time
- Data
- Both
- None

This will determine how long before an SA expires and has to be renegotiated. We have chosen to check “Both”, and decrease the default time and data values to: 5000KB for data, and 14400 seconds (4hrs) for time.

Negotiation Mode

With security in mind, we choose Main mode over Aggressive mode. The extra exchanges are worth the protection of the identities of the communicating parties.

2.3.2 IKE proposal

We will create a new proposal with the following values:

- Authentication Mode: Preshared Keys (XAUTH)
- Authentication Algorithm: SHA/HMAC-160
- Encryption Algorithm: 3DES-168
- Diffie-Hellman Group: Group 2 (1024-bits)
- Lifetime Measurement: Both

- Data Lifetime: 5000
- Time Lifetime: 14400

We acknowledge that a digital certificate would be preferable to pre-shared keys, regardless of XAUTH. Therefore, we hope that once we have trained and instructed our non technical remote users and partners on how to install and manage them using the Certificate Manager that comes with the VPN client, we will change the authentication mode. However, the prompt deployment and business need will have to take precedence over more complex configurations on the client side. We expect to have difficulties as it is, instructing our remote users on configuring the IDS/Firewall and the VPN client.

2.3.3 Security policies for groups

We will define our base group and make sure any and all new groups inherit these security-related values:

- **Access Hours:** Considering we have partners on multiple timezones, and some of our users are night owls, we will not have any restrictions.
- **Simultaneous Logins:** We will change this to 1. (Default: 3)
- **Minimum Password Length:** This time, we will keep the default: 8.
- **Allow Alphabetic-Only Passwords:** This should be disabled, in order to prevent weak passwords and offer some protection against dictionary attacks.
- **Idle Timeout:** The default is 30 minutes, which is reasonable.
- **IPSec SA:** We will use the SA we created.
- **IKE Keepalives:** This will be enabled to detect and remove unresponsive connections.
- **Group Lock:** this will be enabled to ensure users are not authenticated without regard of their assigned group. By default, this is disabled. Therefore, a user could authenticate using a different group and group password than the one s/he belongs to.
- **Authentication:** This will be done internally.
- **Reauthentication on Rekey:** This will be enabled to authenticate the user on every IKE (phase 1) rekey.
- **Banner:** Always a good habit to include warning messages to inform the user that all traffic is being monitored, and only authorized individuals should be accessing these resources.
- **Allow Password Storage on Client:** Definitely not. Anybody with physical access to the laptop or remote computer could run the Cisco VPN dialer and connect to our VPN network.
- **Split Tunnel Policy:** All traffic should go through the tunnel, therefore we will select 'Tunnel everything'.

Once we have created our four sets of groups: remote sales staff, partners, telecommuters, and remote administrators, we will assign each group its own address pool:

- 10.0.1.1 - 10.0.1.254 for partners
- 10.0.2.1 - 10.0.2.254 for remote sales staff
- 10.0.3.1 - 10.0.3.254 for telecommuters
- 10.0.4.1 - 10.0.4.254 for remote administrators.

2.3.4 Filters and rules

By default, the filter applied to the public interface has the following rules:

- GRE In (forward/in)
- IPSEC-ESP In (forward/in)
- IKE In (forward/in)
- PPTP In (forward/in)
- L2TP In (forward/in)
- ICMP In (forward/in)
- VRRP In (forward/in)
- GRE Out (forward/out)
- IKE Out (forward/out)
- PPTP Out (forward/out)
- L2TP Out (forward/out)
- ICMP Out (forward/out)
- VRRP Out (forward/out)

We do not need most of those rules, however, we do need to create one more rule to send logs. We will create the following rule:

- **Log Out:**

Direction: Outbound
Action: Forward
Protocol: UDP
Source: 10.100.2.9/0.0.0.0
Destination: 10.2.0.2/0.0.0.0
Source Port: SYSLOG (514)
Destination Port: SYSLOG(514)

The new order of the rules will be:

- IPSEC-ESP In (forward/in)
- IKE In (forward/in)
- IKE Out (forward/out)
- Log Out (forward/out)

In addition, we will ensure that both filters (the ones applied to the public and private interface):

- Have a default action of: Drop
- Do not allow source routing
- Do not allow fragments

2.3.5 Management protocols

Since we will be doing all the management through the console, we will disable all the management protocols available: ftp, http, https, tftp, telnet, SNMP, SSL, ssh, and XML.

2.4 Tutorial on Configuring Cisco Devices

When configuring network devices like our routers, PIX, or VPN:

- It is important that we work offline and make sure that there are no interfaces connected to the external network, or any of the internal networks. The configuration will be done through the console. So the only physical cable

connected to the device must be connected to the console port and be coming from our terminal server.

- All cables must be labeled. This prevents plugging them on the wrong ports. Not only it is a good security habit, but it might save hours of troubleshooting.
- We will have the planned configuration ready on the management workstations, so that all we need to do is copy-paste the appropriate commands through the console. This can prevent typos due to human error, and makes it easier to reconfigure the device for scratch anytime we need to.
- When adding or making changes to the configuration, we will use comments whenever possible. Comments on Cisco are lines that start with the “!” character. Adding remarks on ACL is also a good habit to have, especially when the set of rules is large.
- Time will be scheduled to audit the security policies implemented by the device and ensure the configuration works before we go live.

2.4.1 Typographic convention

We will use the following convention for our device configuration section.

<i>Typeface or Symbol</i>	<i>Meaning</i>	<i>Example</i>
AaBbCc123	Commands that need to be typed.	<i>Router# show config</i> (On this example, only “ show config ” is typed)
AaBbCc123	On-screen output and does not need to be typed.	<i>Router(config)# hostname ER</i> (On this example, “Router(config)#” represents the prompt and is not typed)
<AaBbCc123>	This is used to indicate that it must be replaced with a real name or value.	<i>Router# enable secret 0 <password></i> (On this example, “<password>” is replaced with the the password to be used)

Table 23: Typographic convention

The only exception to this convention are Cisco comments, which are not shown on bold for readability purposes. They should still be typed and included in our configuration files.

Also, the first time we show a command in the tutorial section, we will include a footnote with a link to Cisco documentation online that describes the command in detail. Obviously, the footnote number (in superscript) should not be typed.

2.5 Configuring the Routers

When we are ready to change the system configuration, we need to be in enabled/privileged mode, which is usually indicated with a “#” prompt. During the rest of this section, we might indicate to make changes from the Global Configuration Mode, Interface Configuration mode, or ACL configuration mode, which enable us to make global changes, changes that apply to the interface only, or a specific ACL respectively.

The command “**configure terminal**”⁷⁰ (or “**config t**” for short) allows us to access the global configuration mode from privileged mode. From here we can make

70 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft202.htm#1029179>

global changes to the configuration. Global Configuration mode has a “Router(config)#” prompt.

The Interface Configuration mode applies to a specific interface (whether it is virtual or physical), and as the name indicates, it only affects that interface. This mode can be accessed from the global configuration mode using the interface command. (e.g. the command “Router# **interface**⁷¹ **gigabitethernet0/2**” will access the interface configuration mode for that particular interface). The Interface Configuration mode, will display a “Router(config-if)#” prompt.

The “**exit**”⁷² command, or Ctrl-Z can be used to exit global or interface configuration mode.

Once changes have been made and we have made sure that they work, we can save them with: the command:

```
Router# copy system:running-config nvram:startup-config73
```

or

```
Router# write memory74
```

A trivial detail to mention, but sometimes catches administrators off-guard, is the distinction between the startup-config and the running-config.

Router# **show config**⁷⁵ or Router# **show startup-config**⁷⁶ will show non-volatile configurations, which might not be the same as the running configuration.

Router# **show running-config**⁷⁷ will show what the current running configuration is. This is NOT saved if the device is rebooted.

2.5.1 Common configuration

The following are common configurations for all of our Cisco routers. These commands need to be done in global configuration mode.

Bogus Names

We will give all of our devices ambiguous names. No need to make it obvious what the function of the device is by just looking at the name.

!The following command, applies only to the border router

!Name the border router.

```
Router# hostname78 hf84n-24
```

!The following command, applies only to the border router

!Name the internal router

```
Router# hostname hnfk47z
```

We will continue the rest of the common configuration using the prompt of the border router “**hf84n-24**”, but the prompt will obviously be different for other devices.

Encrypted Password

We definitely want to use the password encryption service, even if this encryption can easily be broken. It is much better than plain text.

!Enable the password-encryption service

71 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tir/irftdce.htm#1019390>

72 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122txr/xftscmd1.htm#1028591>

73 http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_r/ffrprt2/frf006.htm#1024612

74 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft204.htm#1018899>

75 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft201.htm#1018291>

76 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft202.htm#24087>

77 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft202.htm#1022094>

78 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1018257>

hf84n-24(config)# **service password-encryption**⁷⁹

!We want to disable the enable password

hf84n-24(config)# **no enable password**⁸⁰

!Use secret password

hf84n-24(config)# **enable secret 0** <cleartext-password-goes-here>⁸¹

!or use "enable secret 5 <encrypted-password-goes-here>" if copy-pasting from another

!router or configuration file. Either way, the line will be displayed later as

!enable secret 5 \$1\$h45L9brv/ram5sagPWQXh1", for example.

Services

Cisco routers, by default, have a great variety of services available. Since we do not need them for our design, we will be disabling some of them and/or making sure they are disabled, since they could be potential vulnerabilities, yet some services will be enabled for security and performance purposes. Some of these services are enabled/disabled globally, and some per interface.

- **SNMP:** SNMP monitoring will not be used. It will be disabled.
- **Small services (echo, discard, daytime, chargen, etc.):** These ports (for TCP and UDP) will not be used anywhere in our design. They will be disabled.
- **Finger:** Sensitive information can be leaked out and used for further attacks or social engineering attacks. It will be disabled.
- **Http server:** No management will be done through the network. It will be disabled.
- **BOOTP:** Another service we do not need and could be used as a DoS. It will be disabled.
- **DHCP:** Our systems all have statically configured IP addresses. It will be disabled.
- **CDP:** Our routers will not be using Cisco Discovery Protocol. It will be disabled.
- **DNS:** No name resolution on network devices. We prefer real IP addresses on our logs, plus we can avoid the risks of cache poisoning. It will be disabled.
- **PAD:** No need for the packet assembler/disassembler option. It will be disabled.
- **Keepalives:** We will enable TCP keepalives to detect and get rid of dead sessions.
- **NTP (interface):** For security reasons, NTP will be disabled on all interfaces other than the one connected to the NTP server.
- **MOP (interface):** No need for the Maintenance Operations Protocol to be enabled. It will be disabled.
- **Proxy-Arp (interface):** No need for this to be enabled in our design. It will be disabled.

!Disable SNMP since we will not be using it to monitor our systems

hf84n-24(config)# **no snmp server**⁸²

!Disable small services (i.e. echo, discard, daytime, chargen, etc.

79 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/fothercr/sftpass.htm#1017841>

80 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/fothercr/sftpass.htm#1017394>

81 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/fothercr/sftpass.htm#1017497>

82 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/fft303.htm#1183584>

!They are usually disabled by default, it does not hurt to make sure they are disabled.

```
hf84n-24(config)# no service tcp-small-servers83  
hf84n-24(config)# no service udp-small-servers84
```

!Definitely disable finger since it can provide dangerous information and can be a vulnerability

```
hf84n-24(config)# no service finger85
```

!Disable web administration

```
hf84n-24(config)# no ip http server86
```

!Disable bootp

```
hf84n-24(config)# no ip bootp server87
```

!Disable CDP.

!CDP (Cisco Discovery Protocol) is a Proprietary Layer 2 protocol to exchange information between routers. We do not need it.

```
hf84n-24(config)# no cdp run88  
hf84n-24(config)# no cdp advertise-v289
```

!We do not want to resolve names

```
hf84n-24(config)# no ip domain-lookup90
```

!Disable DHCP

```
hf84n-24(config)# no service dhcp91
```

!We do not need PAD enabled.

!PAD is a packet assembler/disassembler

```
hf84n-24(config)# no service pad92
```

Routing

There are some configuration settings that involve the routing of packets that need to be configured for security and functionality purposes.

- **Source-Routing:** There is no legitimate reason to allow source routing packets in, and they would most likely be malicious in nature. It will be disabled.
- **ICMP Redirects:** Another feature that would most likely be used as a DOS. It will be disabled.
- **Zero and Classless subnets:** Classless subnets will be enabled. However, subnet-zero will not.

!Definitely do not let this in unless it's really needed for legitimate reasons

```
hf84n-24(config)# no ip source-route93
```

!No redirects

```
hf84n-24(config)# no icmp redirect94
```

83 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019632>

84 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019697>

85 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019486>

86 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft105.htm#1020224>

87 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1077239>

88 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft304.htm#1017573>

89 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft304.htm#1032226>

90 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1018214>

91 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftidhcp.htm#1066359>

92 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122twr/x25cmds/wftx251.htm#1045320>

93 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip2.htm#1019433>

94 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip2.htm#1019389>

!Disable zero subnet and allow classless
hf84n-24(config)# **no ip subnet-zero**⁹⁵
hf84n-24(config)# **ip classless**⁹⁶

Network Booting

Booting from the network can be dangerous. So we will be disabling network booting and the service config feature.

!Booting from the network can be dangerous.
!Disable network booting.
hf84n-24(config)# **no boot network**⁹⁷
hf84n-24(config)# **no service config**⁹⁸

Warnings

It is a good habit to have deterring messages on all infrastructure devices, even if our network is designed to only allow administrators to ever see them. MOTD and banner logins stating that access is only allowed by authorized personnel will be configured.

!Set up the MOTD
!The character "!" is the character we use to start and end the input
hf84n-24(config)# **banner motd**⁹⁹ !
This system belongs to GIAC Enterprises.
If you do not have EXPLICIT access, leave immediately.
Unauthorized access is subject to civil and criminal prosecution.
!

!Set up a login warning
hf84n-24(config)# **banner login**¹⁰⁰ ! **Authorized Personnel Only !**

Console Only

Since all management to our network devices will be done through the console. We will make sure that no access will be accepted through VTYS or aux ports. For this particular section, the prompt will change as we accessed the particular devices.

!Disable access to VTY and aux port.
!Timeout the console after 60 seconds of inactivity.
hf84n-24(config)# **line**¹⁰¹ **vtty 0 4**
hf84n-24(config-line)# **no login**¹⁰²
hf84n-24(config-line)# **transport input**¹⁰³ **none**
hf84n-24(config-line)# **transport output**¹⁰⁴ **none**
hf84n-24(config-line)# **exit**

hf84n-24(config)# **line aux 0**
hf84n-24(config-line)# **no login**
hf84n-24(config-line)# **transport input none**
hf84n-24(config-line)# **transport output none**

95 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1020462>
96 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1018033>
97 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft202.htm#1017676>
98 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft202.htm#1017913>
99 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft104.htm#1017571>
100 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft104.htm#1017507>
101 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tdr/drflcmo.htm#1064507>
102 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/tftlosho.htm#998262>
103 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/tfttrx.htm#1083564>
104 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/tfttrx.htm#1083654>


```
hf84n-24(config-line)# exit
```

Time-outs

For security reasons, we configure the console to log the user out if there is no activity after 60 seconds. For this particular section, the prompt will change as we accessed the particular devices.

!Timeout the console after 60 seconds

```
hf84n-24(config)# line con 0
```

```
hf84n-24(config-line)# session-timeout105 1
```

```
hf84n-24(config-line)# exec-timeout106 0 60
```

```
hf84n-24(config-line)# exit
```

Logging

Logging is crucial in our design. Our routers will be configured to enable logging. Each device will be using a different logging facility to facility the segregation of logs in our log server. For security purposes, the IP address of the log server and the interface used to send logs to will be specified. Logs need to have accurate time information for better log analysis.

As a precautionary feature in case communication with the log server is affected and to facility troubleshooting, we will also send some logs to the console and have a small buffer for the logs.

!enable logging

```
hf84n-24(config)# logging on107
```

!Only send informational (severity level 5) logs to the console

```
hf84n-24(config)# logging console108 informational
```

!Set the buffer size

```
hf84n-24(config)# logging buffered109 12000
```

!Trap informational (severity level 5)

```
hf84n-24(config)# logging trap110 informational
```

!Show substantial timestamps in the logs

```
hf84n-24(config)#service timestamps111 log datetime msec show-timezone localtime
```

```
hf84n-24(config)#service timestamps debug datetime msec show-timezone localtime
```

!Select Timezone

```
hf84n-24(config)# clock timezone112 PST -8
```

The following lines are for the border router only

!Specify the external IP (NAT) of the log cluster

```
hf84n-24(config)# logging x.x.x.7
```

!Declare the Interface used to send logs.

```
hf84n-24(config)# logging source-interface113 FastEthernet0
```

!Specify facility

```
hf84n-24(config)# logging facility114 local5
```

The following lines are for the internal router only

!Specify the internal IP of the log cluster

¹⁰⁵<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/tftlosho.htm#999375>

¹⁰⁶<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft104.htm#1017909>

¹⁰⁷<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1018658>

¹⁰⁸<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1018034>

¹⁰⁹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1072328>

¹¹⁰<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1018863>

¹¹¹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1019290>

¹¹²<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft301.htm#1018092>

¹¹³<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1090413>

¹¹⁴<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122ttr/fft302.htm#1018173>

```
hnfk47z(config)# logging 10.2.0.6
!Declare the Interface used to send logs.
hnfk47z(config)# logging source-interface FastEthernet0/3
!Specify facility
hnfk47z(config)# logging facility115 local6
```

NTP

In conjunction with logging. Keeping track of time is critical for log analysis and incident handling. For security issues, we will use NTP authentication and trusted keys. Moreover, our routers will only use our own NTP servers for time synchronization. Specifying the interface to receive NTP packets from, and applying an ACL to only accept NTP packets from our server is another security issue. This will prevent an exploitation of a known NTP vulnerability on Cisco IOS.¹¹⁶

The following lines apply to the border router only.

```
!Create an access list to ensure NTP access to our external NTP server
hf84n-24(config)# access-list117 20 remark Permit access to external NTP server
hf84n-24(config)# access-list 20 permit x.x.x.1
hf84n-24(config)# access list 20 deny any
!Use NTP Authentication
hf84n-24(config)# ntp authenticate118
!Enter Authentication MD5 Key
hf84n-24(config)# ntp authentication-key119 1 md5 Tick!Tock
!Authenticate Using Key
hf84n-24(config)# ntp trusted-key120 1
!Enforce access using access listed
hf84n-24(config)# ntp access-group121 peer 20
!Specify NTP server: External NTP server and the interface
hf84n-24(config)# ntp server122 x.x.x.1 key 1 source FastEthernet0
!Use NTP to update-calendar
hf84n-24(config)# ntp update-calendar123
```

The following lines apply to the internal router only.

```
!Create an access list to ensure NTP access to our internal NTP server
hnfk47z(config)# access-list 20 remark Permit access to internal NTP server
hnfk47z(config)# access-list 20 permit 10.2.0.9
hnfk47z(config)# access list 20 deny any
!Use NTP Authentication
hnfk47z(config)# ntp authenticate
!Enter Authentication MD5 Key
hnfk47z(config)# ntp authentication-key 1 md5 Tick!Tock
!Authenticate Using Key
hnfk47z(config)# ntp trusted-key 1
!Enforce access using access listed
hnfk47z(config)# ntp access-group peer 20
!Specify NTP server: External NTP server and the interface
hnfk47z(config)# ntp server 10.2.0.9 key 1 source FastEthernet0/3
```

115 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft302.htm#1018173>

116 <http://www.cisco.com/warp/public/707/NTP-pub.shtml>

117 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftp1.htm#1017823>

118 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1018522>

119 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1018574>

120 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019137>

121 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1077485>

122 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1048540>

123 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019193>

Flood and DoS Protection

Our routers will be configured to protect themselves from packet flooding and denial of service attacks. This can be accomplished by allocating the minimum amount of time for process-level tasks, and using the TCP-Intercept feature to take handle SYN attacks.

In addition, no packets directed to broadcast addresses will be allowed to prevent flooding.

The following section applies to the border router only.

!Protect ourselves from a packet floods

!Do not spend all the time handling interrupts from packets.

!Process-level tasks are to be handled no less frequently than every 500 milliseconds

hf84n-24(config)# scheduler interval¹²⁴ 500

!Protect against SYN attacks

hf84n-24(config)# access-list 110¹²⁵ remark TCP Intercept SYN Attacks

hf84n-24(config)# access-list 110 permit tcp any x.x.x.0 0.0.0.15

!Enable TCP Intercept to protect against SYN flooding.

hf84n-24(config)# ip tcp intercept list¹²⁶ 110

!Make sure TCP Intercept is in intercept mode, instead of watch mode

hf84n-24(config)# tcp intercept mode¹²⁷ intercept

!Change how long a connection will be managed after no activity.

!The default is 86400 seconds (24 hr.). We want 60 seconds.

hf84n-24(config)# ip tcp intercept connection-timeout¹²⁸ 60

!Change how long to wait to reach established state before sending a reset

!The default is 30 seconds. We want 10 seconds

hf84n-24(config)# ip tcp intercept watch-timeout¹²⁹ 10

!When dropping incomplete connections, start with the oldest

hf84n-24(config)# ip tcp intercept drop-mode¹³⁰ oldest

!Change how long after a reset or FIN-exchange to stop managing the connection

!The default is 5 seconds, make it 3 seconds.

hf84n-24(config)# ip tcp intercept finrst-timeout¹³¹ 3

!Change the maximum number of incomplete connections before we start to get

!aggressive. Default is 1100. Make it 2000.

hf84n-24(config)# ip tcp intercept max-incomplete high¹³² 2000

!When to stop being aggressive. Default is 900. Make it 1500

hf84n-24(config)# ip tcp intercept max-incomplete low¹³³ 1500

!Change how many connections requests we can get in one minute before we get
!aggressive.

hf84n-24(config)# ip tcp intercept one-minute high¹³⁴ 2000

!When to stop being aggressive

hf84n-24(config)# ip tcp intercept one-minute low¹³⁵ 1500

!NOTE: These values should be tuned for the amount and flow of the traffic we get. It would also be a good idea to examine the router behavior under a flood.

Keepalives

Detect and get rid of “Dead” sessions.

!Enable the tcp-keepalives-in service

¹²⁴<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft301.htm#1019389>

¹²⁵<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1017448>

¹²⁶<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017525>

¹²⁷<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017769>

¹²⁸<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017392>

¹²⁹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017995>

¹³⁰<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017421>

¹³¹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017493>

¹³²<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017594>

¹³³<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017680>

¹³⁴<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017821>

¹³⁵<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tsr/ttrafwlr/sftenl.htm#1017907>

hf84n-24(config)# **service tcp-keepalives-in**¹³⁶

Interface-specific

We have stripped our border router to only have two interfaces: FastEthernet0 and Serial0.

The internal router has eight interfaces: FastEthernet0/0 to FastEthernet0/7.

The terminal concentrator only has one interface: FastEthernet3/0.

Packets that can be used to determine information about network will not be accepted/generated: echo-replies, unreachablees, mask-requests, TTL-expiration, etc.

The following are settings to be configured on **each** interface of the system. Therefore the prompt will be different since we will be in Interface configuration mode.

Steps for the border router

```
hf84n-24(config)# int serial137 0
!Disable MOP (Maintenance Operations Protocol)
hf84n-24(config-if)# no mop enabled138
!No proxy-arping
hf84n-24(config-if)# no proxy-arp139
!Do not send unreachablees
hf84n-24(config-if)# no ip unreachablees140
!Watch out for Smurf attacks
hf84n-24(config-if)# no ip directed-broadcast141
!Nobody should be asking for our network mask
hf84n-24(config-if)# no ip mask-reply142
!Make sure CDP is disabled
hf84n-24(config-if)# no cdp enable143
!Enable the interface
hf84n-24(config-if)#no shutdown144
```

```
hf84n-24(config)# int FastEthernet0
!Disable MOP (Maintenance Operations Protocol)
hf84n-24(config-if)# no mop enabled
!No proxy-arping
hf84n-24(config-if)# no proxy-arp
!Do not send unreachablees
hf84n-24(config-if)# no ip unreachablees
!Watch out for Smurf attacks
hf84n-24(config-if)# no ip directed-broadcast
!Nobody should be asking for our network mask
hf84n-24(config-if)# no ip mask-reply
!Make sure CDP is disabled
hf84n-24(config-if)# no cdp enable
!Enable the interface
hf84n-24(config-if)#no shutdown
```

!Exit Interface configuration mode

¹³⁶<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft302.htm#1019216>
¹³⁷<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tdr/drfeip.htm#1017823>
¹³⁸<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tir/irftlop.htm#1017692>
¹³⁹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1020409>
¹⁴⁰<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip2.htm#1019840>
¹⁴¹<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1018109>
¹⁴²<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1067984>
¹⁴³<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tfr/fft304.htm#1017473>
¹⁴⁴<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tir/irftsip.htm#1018004>

```
hf84n-24(config-if)# exit
```

The previous steps must be done for the each interface of the internal router in the same manner.

Unused Interfaces

Any and all unused interfaces and ports will be administratively shutdown. This will be done on switches as well. While this can generate a higher administration overhead when connecting new systems or devices, it is worth implementing, especially on switches/hubs used by critical systems.

The following section applies only to the internal router since it is the only device with unused interfaces.

The following section only applies to interfaces that are not being used: FastEthernet0/5, FastEthernet0/6, and FastEthernet0/7

```
!Make sure any interface not being used is shutdown
hnfk47z(config)# int <Interface-name-goes-here>
hnfk47z(config-if)# shutdown
```

The previous steps must be repeated for each interface.

```
!Exit Interface configuration mode
hnfk47z(config-if)# exit
```

2.6 Border Router

The border router is exposed to all traffic from the Internet. Therefore it is best to filter all the traffic that should not be allowed in under any circumstances, before it hits our firewall.

2.6.1 Access Control Lists

Access Control Lists (ACL) are the rules that determine what traffic to let in or block. By default, routers do not have ACL, therefore all the packets are routed.

An ACL can perform one of two actions on a packet: "permit", which will route it; or "deny", which will drop the packet.

As the name indicates, an ACL is a sequential list of permit and deny conditions that will be applied to an interface. When a packet is received on an interface, the router compares the fields in the packet against the ACL applied on the interface to verify that the packet meets the criteria specified. This is done sequentially until it finds a matching criteria, at which point the packet will be dropped or forwarded depending on the action declared on the matching ACL. Therefore, the order of these conditions is crucial and must be well defined. If no matches are found, the router will drop the packet.

Cisco routers have different types of access control lists, however, we will be concerned with the two most basic static types: Standard and Extended.

Since most access control lists are usually referred to by number, each type has its own range and can be easily identified. Standard ACL have a range of 1-99, and extended ACL have a range of 100-199. However, extended access control list can also be referred to with a name.

The difference between them is the number of fields that are examined. Standard access control lists filter packets based only on the source address, but extended access control lists can check a greater variety of options (source address, destination, protocol options, etc.). For this reason, processing standard ACL can be processed much faster than extended ACL.

Access Control Lists are processed as packets enter an interface or as they are leaving. There are advantages and reasons for doing inbound or outbound filtering, however, we have chosen to filter packets as they enter an interface and pass through the router. This will prevent wasting resources routing a packet that might get dropped at the other end.

A very important restriction to keep in mind is that only one type of ACL can be applied in one interface per direction. So, we reiterate what was previously stated: Access control lists must be carefully planned and well defined.

We have chosen to use extended access control lists on our router, and let the firewall keep maintain state with the incoming connections.

When writing an ACL, whenever the order allows it, most used rules must be used before rules that we expect to have fewer matches. This is critical for performance. Also, it is always a good habit to use the remark¹⁴⁵ command when writing an ACL.

The declaration of an ACL depends on its type. For an standard ACL:
access-list <ACL number> <action> <source> <wildcard> [log]

Field	Explanation
ACL Number	Number of the access list. This is an integer number from 1 to 99 or from 1300 to 1999.
action	Only the actions deny and permit are allowed. deny : To block a packet if it matches. permit : To let the packet pass if it matches.
source	There are two ways to declare the source. The source can be declared by using the IP address of the host, or in the case of the network, by specifying the network in conjunction with the wildcard. The any keyword, which matches everything, is the equivalent of using a source of 0.0.0.0 with a wildcard 255.255.255.255.
wildcard	(optional) Used with the source to specify a network. The wildcard can be think of as the inverse of the netmask for the particular network.
log	(optional) Log if a match is found.

Table 24: Standard ACL declaration

Example, to create a standard ACL numbered "10" that denies all traffic from 10.0.0.0/8 and log it:

Router(config)#access-list 10 deny 10.0.0.0 0.255.255.255 log

Declaring an extended ACL is more complex since extended access control lists check more fields than an standard ACL.

access-list <ACL number> <action> <protocol> <source> <wildcard> <src port> <destination> <wildcard> <dst port> <options>

Field	Explanation
ACL Number	Number of the access list. This is an integer number from 100 to 199 or from 2000 to 2699.
action	Only the actions deny and permit are allowed. deny : To block a packet if it matches. permit : To let the packet pass if it matches.
protocol	Name of the internet protocol (e.g. eigrp, gre, icmp, igmp, igrp, ipinip, nos, ospf, pim, tcp, udp) or an integer number 0-255 representing the protocol. The ip keyword will match any protocol.

¹⁴⁵<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip2.htm#1030586>

Field	Explanation
source	There are three ways to declare the source. An individual host: Using the host keyword followed by its IP. A network: Specifying the network in conjunction with the wildcard. The any keyword, which matches everything, is the equivalent of using a source of 0.0.0.0 with a wildcard 255.255.255.255.
wildcard	(optional) Used with the source to specify a network. The wildcard can be think of as the inverse of the netmask for the particular network.
src port	Source TCP/UDP port.
destination	There are three ways to declare the destination. An individual host: Using the host keyword followed by its IP. A network: Specifying the network in conjunction with the wildcard. The any keyword, which matches everything, is the equivalent of using a destination of 0.0.0.0 with a wildcard 255.255.255.255.
wildcard	(optional) Used with the destination to specify a network. The wildcard can be think of as the inverse of the netmask for the particular network.
dst port	Destination TCP/UDP port.
options	There are many options available, that we will not discuss. However, we will use the following options: log-input : To log if a match is found. It will include the interface. established : Check if the TCP datagram has the ACK, FIN, PSH, RST, SYN or URG control bits set. fragments : If the packets is a fragment.

Table 25: Extended ACL declaration

For example, to creates an extended ACL number “110” that allows all TCP traffic to port 80 on the host 12.12.12.2 from the network 12.34.30.0/24:

Router(config)#access-list 110 permit tcp 12.32.30.0 0.0.0.255 host 12.12.12.2 eq 80

For more detailed information on declaring extended access control lists:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1017448>

Access control lists can also be declared using the “**ip access-control list**” command. The main difference is that this command will allow any ACL to be named, and it will take us into access-list configuration mode, where we can define the rest of the ACL.

ip access-list <type> <name>

type	Extended or standard
name	Number or name of the ACL. Names must start with an alphabetic character.

Table 26: IP access-list declaration

For example, to declare an standard ACL named INGRESS:

Router(config)# ip access-list standard INGRESS

Router(config-std-nacl)#

After this, the prompt will change to indicate we are in ACL declaration

mode. The rest of the ACL can be declared as defined previously for standard and extended ACLs, with the exception that the “*access-list <ACL number>*” prefix is omitted.

For example:

```
Router(config-std-nacl)#deny 10.0.0.0 0.255.255.255 log
```

Three more trivial details to mention, which are sometimes overlooked, are:

- Enabling access control lists will drop everything that does not match by default, so it is important not to forget to explicitly permit the traffic we need.
- Cisco uses an inverse netmask (or wildcard bits) in their access control lists. Therefore, a familiar 255.255.255.0 becomes 0.0.0.255 when used in an ACL.
- Both standard and extended access lists have an extra range of numbers for additional use: 1300-1999 for standard, and 2000-2699 for extended. Therefore, just because the number used to refer to the ACL does not match the first range, does not necessarily mean the ACL is neither standard or extended.

The following table outlines the physical configuration of the border router:

<i>Interface</i>	<i>IP</i>	<i>Device/Network</i>
serial0	y.y.y.1	ISP
ethernet0	x.x.x.14	Firewall

Table 27: Interface configuration for the border router

2.6.2 External interface (Serial0)

As part of our ingress filtering, we will be filtering out:

- Traffic directed to the interfaces of the router or the firewall
- Traffic from the loopback address
- Traffic from a multicast address
- Spoofed traffic using our external addresses
- Fragmented traffic
- ICMP Traffic
- Traffic from a private address
- Traffic from a reserved address
- Traffic from known attackers
- Traffic not destined to a matching service and port.
- Traffic not from an established connection to our web proxy or mail server

In decreasing order, we expect to see more traffic going to:

The web server, web proxy, VPN concentrator, mail server (mail traffic), DNS server, NTP server, mail server (virus updates). As traffic flow changes, we can change the order, if possible, for better performance.

Here is the ACL that our external interface will use. They must be entered from the Global Configuration mode, but once we start adding entries, the prompt will change to indicate we are configuring an extended ACL.

!Enable ACL Turbo Acceleration¹⁴⁶

¹⁴⁶http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a00800873c8.html#1061578

!This allows us to process ACLs a lot faster.
hf84n-24(config)# **access-list compiled**¹⁴⁷

We will be logging most of the denied ACL at first. If the amount of logging is no longer useful, we can always change it.

!Make sure there is no ACL with the same name.
hf84n-24(config)# **no access-list extended INGRESS**
hf84n-24(config)# **ip access-list extended INGRESS**
hf84n-24(config-ext-nacl)# **remark ACL FOR EXTERNAL INTERFACE**

!Protect ourselves first and the firewall
hf84n-24(config-ext-nacl)# **deny ip any host y.y.y.1 log-input**
hf84n-24(config-ext-nacl)# **deny ip any host x.x.x.14 log-input**
hf84n-24(config-ext-nacl)# **deny ip any host x.x.x.13 log-input**
!Drop Packets that come with our IP addresses
hf84n-24(config-ext-nacl)# **remark Drop Spoofed Addresses**
hf84n-24(config-ext-nacl)# **deny ip x.x.x.0 0.0.0.15 any log-input**

!Drop multicast 224.0.0.0 to 239.0.0.0
hf84n-24(config-ext-nacl)# **remark Drop Multicast**
hf84n-24(config-ext-nacl)# **deny ip 224.0.0.0 31.255.255.255 any log-input**

!Drop all fragmented traffic
hf84n-24(config-ext-nacl)# **remark Drop Fragmented**
hf84n-24(config-ext-nacl)# **deny ip any any fragments log-input**

!Drop all ICMP traffic
hf84n-24(config-ext-nacl)# **remark Drop ICMP**
hf84n-24(config-ext-nacl)# **deny**¹⁴⁸ **icmp any any log-input**

!Drop loopback addresses
hf84n-24(config-ext-nacl)# **remark Drop Loopback**
hf84n-24(config-ext-nacl)# **deny ip 127.0.0.0 0.255.255.255 any log-input**

!Drop private addresses (RFC 1918)¹⁴⁹
hf84n-24(config-ext-nacl)# **remark Drop Private Addresses**
hf84n-24(config-ext-nacl)# **deny ip 10.0.0.0 0.255.255.255 any log-input**
hf84n-24(config-ext-nacl)# **deny ip 172.16.0.0 0.15.255.255 any log-input**
hf84n-24(config-ext-nacl)# **deny ip 192.168.0.0 0.0.255.255 any log-input**

!Drop IANA¹⁵⁰'s reserved addresses
hf84n-24(config-ext-nacl)# **remark Drop Reserved Addresses**
hf84n-24(config-ext-nacl)# **deny ip 0.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 1.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 2.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 5.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 7.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 23.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 27.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 31.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 36.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 37.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 39.0.0.0 0.255.255.255 any**
hf84n-24(config-ext-nacl)# **deny ip 41.0.0.0 0.255.255.255 any**

147 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1061578>

148 <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1018146>

149 <http://www.isi.edu/in-notes/rfc1918.txt>

150 <http://www.iana.org/assignments/ipv4-address-space>

y
y
y
y
y
y
y
y
y
y
y
y
y
v


```

hf84n-24(config-ext-nacl)# deny ip 126.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 197.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 201.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 240.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 241.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 242.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 243.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 244.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 245.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 246.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 247.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 248.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 249.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 250.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 251.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 252.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 253.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 254.0.0.0 0.255.255.255 any
hf84n-24(config-ext-nacl)# deny ip 255.0.0.0 0.255.255.255 any

```

!Block top 20 list of attacking networks from incidents.org¹⁵¹

!This are REAL Networks, so check the Application logs to make sure
!none of our existing customers, partners, or suppliers come from these
!addresses before we block them.

!We can also add the IP of usual attackers to our network in here.

```

hf84n-24(config-ext-nacl)# remark Drop Top 20 attackers
hf84n-24(config-ext-nacl)# deny ip 61.38.36.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 203.130.130.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 80.38.177.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 219.41.0.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 61.79.96.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 210.4.143.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 148.245.53.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 218.48.8.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 211.255.136.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 210.108.158.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 195.159.152.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 211.255.136.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 163.25.96.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 24.214.48.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 140.123.33.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 68.52.187.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 63.149.88.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 193.224.167.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 193.154.8.0 0.0.0.255 any log-input
hf84n-24(config-ext-nacl)# deny ip 24.102.205.0 0.0.0.255 any log-input

```

!Finally we will be letting some of the traffic through to our services

!We will be listing them in the decreasing order. From the service that
!we expect to see more traffic, to the service that we expect least traffic.

!External web server

!Allow customers and suppliers to access our web server

```

hf84n-24(config-ext-nacl)# remark Permit Web Access
hf84n-24(config-ext-nacl)# permit152 tcp any host x.x.x.2 eq 80
hf84n-24(config-ext-nacl)# permit tcp any host x.x.x.2 eq 443

```

¹⁵¹ <http://isc.incidents.org/top10.html>

¹⁵² <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip2.htm#1019870>

!Replies to the virtual IP of the web proxy
!Allow internal users access to receive web replies.
!Note: The establish check can be easily fooled, but our firewall can verify that this is a valid reply.
hf84n-24(config-ext-nacl)# remark Permit Replies to Web Proxy
hf84n-24(config-ext-nacl)# permit tcp any host x.x.x.6 eq 80 established
hf84n-24(config-ext-nacl)# permit tcp any host x.x.x.6 eq 443 established

!VPN concentrator
!Allow telecommuters, remote sales staff, partners, and remote administrators
!to connect to the VPN concentrator
!UDP Port 500 is needed to allow IKE
!and ESP Traffic must be let in. Since it is encrypted, we cannot do much about it.
!Therefore, the firewall, will handle this traffic once it is unencrypted.
hf84n-24(config-ext-nacl)# remark Permit VPN Traffic
hf84n-24(config-ext-nacl)# permit udp any host x.x.x.9 eq 500 log-input
hf84n-24(config-ext-nacl)# permit esp any host x.x.x.9 log-input

!External mail relay
!Allow mail to access the mail relay
hf84n-24(config-ext-nacl)# remark Permit Mail Access
hf84n-24(config-ext-nacl)# permit tcp any host x.x.x.4 eq 25

!External DNS server
!We will only be letting UDP traffic in.
!Allow public to resolve our IPs
hf84n-24(config-ext-nacl)# remark Permit DNS Access
hf84n-24(config-ext-nacl)# permit udp any host x.x.x.3 eq 53

!NTP servers
!Allow the public NTP servers to sync our external NTP server
!The Public NTP servers we use are: a.a.a.a and b.b.b.b
!Note: The establish check can be easily fooled, but our firewall can verify that this is a valid reply.
hf84n-24(config-ext-nacl)# remark Permit NTP Traffic
hf84n-24(config-ext-nacl)# permit udp host a.a.a.a host x.x.x.1 eq 123 established
hf84n-24(config-ext-nacl)# permit udp host b.b.b.b host x.x.x.1 eq 123 established

!Virus-Updates
!The cron job only runs once a day, but the mail server needs to download
!virus-updates from the antivirus page.
!For this design, we sanitize the IP of the virus site to c.c.c.c
!Note: The establish check can be easily fooled, but our firewall can verify that this is a valid reply.
hf84n-24(config-ext-nacl)# remark Permit Reply for Daily Virus Update
hf84n-24(config-ext-nacl)# permit tcp host c.c.c.c host x.x.x.4 established log-input

!Drop everything else
!Logging it might generate a lot of unnecessary logs, but could be enabled if
!desired
hf84n-24(config-ext-nacl)# remark Drop everything else
hf84n-24(config-ext-nacl)# deny ip any any

!Leave ACL configuration mode.
hf84n-24(config-ext-nacl)# exit

Having defined our Ingress filter, we can apply them to the interface and give it an IP, which we have not done yet:

!Enter Interface configuration mode for the external interface
hf84n-24(config)# int serial0
!Assign our allocated serial IP to the interface

```

hf84n-24(config-if)# ip address153 y.y.y.1 255.255.255.252
!Apply our INGRESS ACL to inbound traffic
hf84n-24(config-if)# ip access-group154 INGRESS in
!Leave Interface Configuration Mode
hf84n-24(config-if)# exit

```

2.6.3 Internal interface (FastEthernet0)

As part of our egress filtering, we will be blocking:

- Traffic directed to the border router interfaces and the firewall
- Traffic that can be used to map our network
- Traffic that does not come from our internal network.

We will be logging most of the denied ACL at first. If the amount of logging is no longer useful, we can always change it.

```

!Use a meaningful name for the ACL: EGRESS
!Make sure there is no ACL with the same name.
hf84n-24(config)# no access-list extended EGRESS
hf84n-24(config)# ip access-list extended EGRESS
hf84n-24(config-ext-nacl)# remark ACL FOR INTERNAL INTERFACE

```

```

!Protect ourselves first
hf84n-24(config-ext-nacl)# remark Block traffic directed to our interfaces
hf84n-24(config-ext-nacl)# deny ip any host y.y.y.1 log-input
hf84n-24(config-ext-nacl)# deny ip any host x.x.x.14 log-input

```

```

!Let these out for flow optimization
!Since ICMP regarding TTL, unreachable, echo replies, etc. can be used to
!map out network, we will only let this type of ICMP out and drop the rest
hf84n-24(config-ext-nacl)# remark Allow source-quenchers out
hf84n-24(config-ext-nacl)# permit icmp x.x.x.0 0.0.0.15 any source-quench

```

```

!We are letting all traffic from the external service network
!(including the virtual IP of the web proxy), and the VPN out.
!At this point, we assume the firewall has done its job.
hf84n-24(config-ext-nacl)# remark Permit DMZ traffic out
hf84n-24(config-ext-nacl)# permit tcp x.x.x.0 0.0.0.15 any
hf84n-24(config-ext-nacl)# permit udp x.x.x.0 0.0.0.15 any

```

```

!Drop packets that do not match.
hf84n-24(config-ext-nacl)# remark Drop everything else
!Be a good neighbor and drop everything else.
hf84n-24(config-ext-nacl)# deny ip any any log-input

```

```

!Assuming it does not degrade performance or
!causes the router to drop packets, an alternative ACL could be:
!hf84n-24(config)# no access-list extended EGRESS
!hf84n-24(config)# ip access-list extended EGRESS
!hf84n-24(config-ext-nacl)# remark ACL FOR INTERNAL INTERFACE
!hf84n-24(config-ext-nacl)# remark Block traffic directed to our interfaces
!hf84n-24(config-ext-nacl)# deny ip any host y.y.y.1 log-input
!hf84n-24(config-ext-nacl)# deny ip any host x.x.x.14 log-input
!hf84n-24(config-ext-nacl)# remark Allow source-quenchers out
!hf84n-24(config-ext-nacl)# permit icmp x.x.x.0 0.0.0.15 any source-quench
!hf84n-24(config-ext-nacl)# remark Permit Web traffic out.

```

¹⁵³<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftipad.htm#1017863>

¹⁵⁴<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tip1r/p1ftip1.htm#1018640>

```

!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.2 eq 80 any
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.2 eq 443 any
!hf84n-24(config-ext-nacl)# remark Permit Web proxy traffic out
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.6 eq 80 any
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.6 eq 443 any
!hf84n-24(config-ext-nacl)# remark Permit Mail traffic out
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.4 eq 25 any
!hf84n-24(config-ext-nacl)# remark Permit VPN traffic out
!hf84n-24(config-ext-nacl)# permit esp host x.x.x.9 any
!hf84n-24(config-ext-nacl)# permit udp host x.x.x.9 eq 500 any
!hf84n-24(config-ext-nacl)# remark Permit DNS traffic out
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.3 eq 53 any
!hf84n-24(config-ext-nacl)# remark Permit NTP traffic out.
!hf84n-24(config-ext-nacl)# permit udp host x.x.x.1 eq 123 host a.a.a.a
!hf84n-24(config-ext-nacl)# permit udp host x.x.x.1 eq 123 host b.b.b.b
!hf84n-24(config-ext-nacl)# remark Permit Virus Update Traffic out
!hf84n-24(config-ext-nacl)# permit tcp host x.x.x.4 eq 80 host c.c.c.c log-input
!hf84n-24(config-ext-nacl)# remark Drop everything else
!hf84n-24(config-ext-nacl)# deny ip any any log-input

```

```

!Leave ACL configuration mode.
hf84n-24(config-ext-nacl)# exit

```

Having defined our Egress filter, we can apply them to the interface and give it an IP, which we have not done yet:

```

!Enter Interface configuration mode for the external interface
hf84n-24(config)# int FastEthernet0
!Assign our allocated serial IP to the interface
hf84n-24(config-if)# ip address x.x.x.14 255.255.255.240
!Apply our EGRESS ACL to inbound traffic
hf84n-24(config-if)# ip access-group EGRESS in
!Leave Interface Configuration Mode
hf84n-24(config-if)# exit

```

2.6.4 Routing

We will not be using any routing protocol. Instead we will define our two default routes:

```

!Make sure IP routing is enabled
hf84n-24(config)# ip routing
!Route to internal network: Firewall
hf84n-24(config)# ip route155 x.x.x.0 255.255.255.240 FastEthernet0 x.x.x.13 permanent
!Default route to the Internet: The ISP router
hf84n-24(config)# ip route 0.0.0.0 0.0.0.0 serial 0 y.y.y.2 permanent

```

```

!Leave Global Configuration Mode
hf84n-24(config-if)# exit

```

2.6.5 Saving the configuration

From the enabled/privileged mode:

```

!Save our running configuration
hf84n-24# copy running-config startup-config156
!Reboot the router to make sure there are no problems with the configuration
hf84n-24# reload157

```

¹⁵⁵<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122tcr/122tdr/drfeip.htm#1064329>

¹⁵⁶http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_r/ffrprt2/frf006.htm#1024612

¹⁵⁷http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_c/ffcp22/fcf010.htm#1008386

Proceed with reload? [confirm] (Press Return)

If the router comes back up without problems, we check that it is running the configuration we gave it:

hf84n-24# show running-config

Assuming everything is as it should be, we are ready to test the configuration.

2.7 Internal Router

The internal router will be connecting all of the internal networks and the firewall; and enforcing our security and access policy for the internal network. On this router we have decided to apply our access control lists as the interfaces receive outbound packets. We expect the packets dropped by this to be low (noise), but in case of an internal attack (by an employee or malicious code that somehow made its way in), they will show in our logs.

The following table outlines the physical configuration of the internal router:

<i>Interface</i>	<i>IP</i>	<i>Device/Network</i>
ethernet0	10.5.0.1	Firewall
ethernet1	10.1.0.1	Application Network
ethernet2	10.4.0.1	Management Network
ethernet3	10.2.0.1	Internal Service Network
ethernet4	10.3.0.1	Employee Network

Table 28: Interface configuration for the internal router

2.7.1 Screened network interface (FastEthernet0/0)

A direct link to the internal interface of the firewall. However, during audits, we can add additional systems on this network to observe the traffic leaving the internal network and the traffic that goes through the firewall.

We will not be doing any ingress filtering on this network since we do not need to duplicate the efforts of the firewall on this subnet.

!Enter Interface configuration mode for the screened network interface

hnfk47z(config)# int FastEthernet0/0

!Assign an IP to the interface

hnfk47z(config-if)# ip address 10.5.0.1 255.255.255.0

!Leave Interface Configuration Mode

hnfk47z(config-if)# exit

2.7.2 Application network interface (FastEthernet0/1)

This network hosts our valuable application cluster, and the database. Therefore we need to be careful about the traffic leaving this network.

We will be filtering the traffic leaving the network. In decreasing order, we expect most of the traffic to go from:

- The application server to the external web server
- The internal web server to the local sales subnet
- The internal web server to the remote sales VPN Address pool
- The internal web server to the Partners VPN Address pool to the web

server

- The application network to the development network

!Use a meaningful name for the ACL: APP_EGRESS

!Make sure there is no ACL with the same name.

```
hnfk47z(config)# no access-list extended APP_EGRESS
```

```
hnfk47z(config)# ip access-list extended APP_EGRESS
```

```
hnfk47z(config-ext-nacl)# remark ACL FOR APPLICATION INTERFACE
```

!Protect ourselves first and the firewall

```
hnfk47z(config-ext-nacl)# deny ip any host 10.1.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.2.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.3.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.4.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host x.x.x.13 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.100.1.13 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.100.2.2 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.0.0.2 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.2 log-input
```

!Allow application traffic to the external web server

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.

```
hnfk47z(config-ext-nacl)# remark Permit Traffic out to external web server
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.3 eq 5000 host 10.100.1.2 established
```

!Allow web traffic back to the local sales staff

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.

```
hnfk47z(config-ext-nacl)# remark Permit Traffic out to Local sales staff
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 80 10.3.1.0 0.0.0.255 established
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 443 10.3.1.0 0.0.0.255 established
```

!Allow web traffic back to the remote sales staff

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.

```
hnfk47z(config-ext-nacl)# remark Permit Traffic out to Remote sales staff
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 80 10.0.2.0 0.0.0.255 established
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 443 10.0.2.0 0.0.0.255 established
```

!Allow web traffic back to the partners

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.

```
hnfk47z(config-ext-nacl)# remark Permit Traffic out to Partners
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 80 10.0.1.0 0.0.0.255 established
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.1.0.2 eq 443 10.0.1.0 0.0.0.255 established
```

!Allow ssh traffic back to the developers

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.

```
hnfk47z(config-ext-nacl)# remark Permit All Traffic to Developers
```

```
hnfk47z(config-ext-nacl)# permit tcp any eq 22 10.3.2.0 0.0.0.255 established
```

!Drop everything else and log it

```
hnfk47z(config-ext-nacl)# remark Drop everything else
```

```
hnfk47z(config-ext-nacl)# drop ip any any log-input
```

!Leave ACL Configuration Mode

```
hnfk47z(config-ext-nacl)# exit
```

Having defined our Egress filter, we can apply them to the interface and give it an IP, which we have not done yet:

!Enter Interface configuration mode for the external interface
hnfk47z(config)# int FastEthernet0/0
!Assign our allocated serial IP to the interface
hnfk47z(config-if)# ip address 10.1.0.1 255.255.255.0
!Apply our EGRESS ACL to inbound traffic
hnfk47z(config-if)# ip access-group APP_EGRESS in
!Leave Interface Configuration Mode
hnfk47z(config-if)# exit

2.7.3 Management network interface (FastEthernet0/2)

Only SSH traffic should leave this network, and it should only go to the VPN address pool for remote administrators.

!Use a meaningful name for the ACL: MGMT_EGRESS
!Make sure there is no ACL with the same name.
hnfk47z(config)# no access-list extended MGMT_EGRESS
hnfk47z(config)# ip access-list extended MGMT_EGRESS
hnfk47z(config-ext-nacl)# remark ACL FOR MANAGEMENT INTERFACE
!Protect ourselves first and the firewall
hnfk47z(config-ext-nacl)# deny ip any host 10.1.0.1 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.2.0.1 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.3.0.1 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.4.0.1 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.1 log-input
hnfk47z(config-ext-nacl)# deny ip any host x.x.x.13 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.100.1.13 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.100.2.2 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.0.0.2 log-input
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.2 log-input

!Allow SSH Traffic back to the VPN address pool for remote administrators
!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.
hnfk47z(config-ext-nacl)# remark Permit all SSH Traffic out to remote administrators
hnfk47z(config-ext-nacl)# permit tcp any eq 22 10.0.4.0 0.0.0.255 established
!Drop everything else and log it
hnfk47z(config-ext-nacl)# remark Drop everything else
hnfk47z(config-ext-nacl)# deny ip any any log-input
!Leave ACL Configuration Mode
hnfk47z(config-ext-nacl)# exit

Having defined our Egress filter, we can apply them to the interface and give it an IP, which we have not done yet:

!Enter Interface configuration mode for the external interface
hnfk47z(config)# int FastEthernet0/2
!Assign our allocated serial IP to the interface
hnfk47z(config-if)# ip address 10.4.0.1 255.255.255.0
!Apply our EGRESS ACL to inbound traffic
hnfk47z(config-if)# ip access-group MGMT_EGRESS in
!Leave Interface Configuration Mode
hnfk47z(config-if)# exit

2.7.4 Internal service network interface (FastEthernet0/3)

Our internal network hosts some services that are indirectly exposed to factors from outside of our network: email attachments, DNS replies, web content through the web proxy, etc. In addition, since these services are available to all users on the employee network, an internal attack is also a possibility. Despite all the

mechanisms we have available, these systems could be compromised. Therefore, we will enforce our global security policy of only allowing traffic as it needed and outlined by our access requirements. By limiting the traffic that can leave this network we can help localize any damage that can be done from this network.

We will be performing some Egress filtering on the traffic leaving this network. In decreasing order, we expect most of the traffic to go from:

- The web proxy to the employee network
- The web proxy to the Internet
- The mail server to the employee network
- The mail server to the external relay
- The mail server to the Remote sales VPN address pool
- The mail server to the Remote Users VPN address pool
- The internal DNS to the employee network
- The internal DNS to the external DNS.
- The internal NTP server to the external NTP server

!Use a meaningful name for the ACL: SERVICE_EGRESS

!Make sure there is no ACL with the same name.

```
hnfk47z(config)# no access-list extended SERVICE_EGRESS
```

```
hnfk47z(config)# ip access-list extended SERVICE_EGRESS
```

```
hnfk47z(config-ext-nacl)# remark ACL FOR SERVICE INTERFACE
```

!Protect ourselves first and the firewall

```
hnfk47z(config-ext-nacl)# deny ip any host 10.1.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.2.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.3.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.4.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.1 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host x.x.x.13 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.100.1.13 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.100.2.2 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.0.0.2 log-input
```

```
hnfk47z(config-ext-nacl)# deny ip any host 10.5.0.2 log-input
```

!Allow all web traffic from the web proxy

```
hnfk47z(config-ext-nacl)# remark Permit Web traffic to employee network
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.7 any eq 80
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.7 any eq 443
```

!Allow mail traffic to the employee network

```
hnfk47z(config-ext-nacl)# remark Permit Mail Traffic to Employee Network
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.3 eq 25 10.3.0.0 0.0.255.255
```

!Allow mail traffic to the external mail relay

```
hnfk47z(config-ext-nacl)# remark Permit Mail Traffic to external mail relay
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.3 eq 25 host 10.100.1.4
```

!Allow mail traffic to the remote sales staff VPN address pool

```
hnfk47z(config-ext-nacl)# remark Permit Mail Traffic to Remote sales Address Pool
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.3 eq 25 10.0.2.0 0.0.0.255
```

!Allow mail traffic to the remote users VPN address pool

```
hnfk47z(config-ext-nacl)# remark Permit Mail Traffic to Remote sales Address Pool
```

```
hnfk47z(config-ext-nacl)# permit tcp host 10.2.0.3 eq 25 10.0.3.0 0.0.0.255
```

!Allow DNS traffic to the employee network

!Note: The established check can easily be fooled, but our firewall can verify this is a valid reply.
hnfk47z(config-ext-nacl)# **remark Permit DNS Traffic to Employee Network**
hnfk47z(config-ext-nacl)# **permit udp host 10.2.0.2 eq 53 10.3.0.0 0.0.255.255 established**

!Allow DNS traffic to the external DNS server
hnfk47z(config-ext-nacl)# **remark Permit Mail Traffic to external DNS server**
hnfk47z(config-ext-nacl)# **permit udp host 10.2.0.2 eq 53 host 10.100.1.3**

!Allow nTP traffic to the external NTP server
hnfk47z(config-ext-nacl)# **remark Permit Mail Traffic to external NTP server**
hnfk47z(config-ext-nacl)# **permit udp host 10.2.0.9 eq 123 host 10.100.1.1**

!Drop everything else and log it
hnfk47z(config-ext-nacl)# **remark Drop everything else**
hnfk47z(config-ext-nacl)# **deny ip any any log-input**

!Leave ACL Configuration Mode
hnfk47z(config-ext-nacl)# **exit**

Having defined our Egress filter, we can apply them to the interface and give it an IP, which we have not done yet:

!Enter Interface configuration mode for the external interface
hnfk47z(config)# **int FastEthernet0/3**
!Assign our allocated serial IP to the interface
hnfk47z(config-if)# **ip address 10.2.0.1 255.255.255.0**
!Apply our EGRESS ACL to inbound traffic
hnfk47z(config-if)# **ip access-group SERVICE_EGRESS in**
!Leave Interface Configuration Mode
hnfk47z(config-if)# **exit**

2.7.5 Employee network interface (FastEthernet0/4)

We will be performing some Egress filtering on the traffic leaving this network. In decreasing order, we expect most of the traffic to go from:

- Any subnet on the network to the web proxy
- The sales subnet to the internal web server
- Any subnet on the network to the mail server
- Any subnet on the network to the DNS server
- The development subnet to the application network

!Use a meaningful name for the ACL: SERVICE_EGRESS
!Make sure there is no ACL with the same name.
hnfk47z(config)# **no access-list extended EMPLOYEE_EGRESS**
hnfk47z(config)# **ip access-list extended EMPLOYEE_EGRESS**
hnfk47z(config-ext-nacl)# **remark ACL FOR EMPLOYEE INTERFACE**

!Protect ourselves first and the firewall
hnfk47z(config-ext-nacl)# **deny ip any host 10.1.0.1 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.2.0.1 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.3.0.1 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.4.0.1 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.5.0.1 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host x.x.x.13 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.100.1.13 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.100.2.2 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.0.0.2 log-input**
hnfk47z(config-ext-nacl)# **deny ip any host 10.5.0.2 log-input**

```
!Allow all web traffic to the web proxy
hnfk47z(config-ext-nacl)# remark Permit Web traffic to proxy
hnfk47z(config-ext-nacl)# permit tcp any host 10.2.0.7 eq 80
hnfk47z(config-ext-nacl)# permit tcp any host 10.2.0.7 eq 443
```

```
!Allow sales access to internal web server
hnfk47z(config-ext-nacl)# remark Permit Local sales to access internal web server
hnfk47z(config-ext-nacl)# permit tcp 10.3.1.0 0.0.0.255 host 10.2.0.7 eq 80
hnfk47z(config-ext-nacl)# permit tcp 10.3.1.0 0.0.0.255 host 10.2.0.7 eq 443
```

```
!Allow access to mail server
hnfk47z(config-ext-nacl)# remark Permit access to mail server
hnfk47z(config-ext-nacl)# permit tcp any host 10.2.0.3 eq 25
```

```
!Allow access to DNS server
hnfk47z(config-ext-nacl)# remark Permit access to DNS server
hnfk47z(config-ext-nacl)# permit udp any host 10.2.0.2 eq 53
```

```
!Allow developers ssh to the application network
hnfk47z(config-ext-nacl)# remark Permit Traffic from developers
hnfk47z(config-ext-nacl)# permit tcp 10.3.2.0 0.0.0.255 10.1.0.0 0.0.0.255 eq 22
```

```
!Drop everything else and log it
hnfk47z(config-ext-nacl)# remark Drop everything else
hnfk47z(config-ext-nacl)# deny ip any any log-input
```

```
!Leave ACL Configuration Mode
hnfk47z(config-ext-nacl)# exit
```

Having defined our Egress filter, we can apply them to the interface and give it an IP, which we have not done yet:

```
!Enter Interface configuration mode for the external interface
hnfk47z(config)# int FastEthernet0/4
!Assign our allocated serial IP to the interface
hnfk47z(config-if)# ip address 10.3.0.1 255.255.255.0
!Apply our EGRESS ACL to inbound traffic
hnfk47z(config-if)# ip access-group EMPLOYEE_EGRESS in
!Leave Interface Configuration Mode
hnfk47z(config-if)# exit
```

2.7.6 Routing

We will not be using any routing protocol. Instead we will define a default route to the firewall

```
!Internal Route to the screened network
hnfk47z(config)# ip route 10.5.0.0 255.255.255.0 FastEthernet0/0 permanent
!Internal Route to the application network
hnfk47z(config)# ip route 10.2.0.0 255.255.255.0 FastEthernet0/1 permanent
!Internal Route to the management network
hnfk47z(config)# ip route 10.4.0.0 255.255.255.0 FastEthernet0/2 permanent
!Internal Route to the Internal service network
hnfk47z(config)# ip route 10.2.0.0 255.255.255.0 FastEthernet0/3 permanent
!Internal Route to the Internal employee network
hnfk47z(config)# ip route 10.2.0.0 255.255.255.0 FastEthernet0/4 permanent
!Default route to the Internet: The firewall
hnfk47z(config)# ip route 0.0.0.0 0.0.0.0 10.5.0.2 permanent
```

```
!Leave Global Configuration Mode
```

hf84n-24(config-if)# **exit**

2.7.7 Saving the configuration

From the enabled/privileged mode:

!Save our running configuration

hnfk47z# **copy running-config startup-config**

!Reboot the router to make sure there are no problems with the configuration

hnfk47z# **reload**

Proceed with reload? [confirm] (Press Return)

If the router comes back up without problems, we check that it is running the configuration we gave it:

hnfk47z# **show running-config**

Assuming everything is as it should be, we are ready to test the configuration.

© SANS Institute 2003, Author retains full rights

3. Assignment 3: Verify the Firewall Policy

3.1 Planning the audit

The first step when auditing the policy of a firewall is the planning phase. This includes :

- The scope of the audit and the objectives we hope to accomplish.
- A detailed description of the technical approach, methodology and tools to be used.
- An estimate of the costs, risks, time and technical involvement.
- Expected results.

3.2 Scope

The main objective of the audit is to verify the security policy of the firewall. This implies that we will ensure our rules and configuration do what we intended when we defined them. We accomplish this by verifying that all traffic we intended to be blocked is being blocked, all legitimate traffic we intended to let through is passing through, and the firewall behaves to certain stimuli the way we expected. This is not a vulnerability assessment of the firewall or the rest of the infrastructure.

3.3 Approach

We will break our approach into several sections that will be repeated for each interface:

Scanning firewall interfaces: First, we check for basic connectivity and scan the firewall interfaces themselves to make sure no services are available. Several types of scans will be employed.

Scanning systems behind the firewall: We will attempt to scan hosts and services behind the other interfaces to determine what traffic generates a response, what traffic is passed through, and what traffic is silenced.

Scanning by replacing a system: Third, we will replace servers sitting on the current interface and perform a test similar to the second step to determine what traffic can and cannot pass through the firewall when it originates from that server. This will help us to estimate how much damage can be done when a system is compromised or it is spoofed.

Scans and Traffic that will be used

- **Basic connectivity:** We will try to test connectivity with simple ICMP echo-requests.
- **Connect scan:** Basic TCP connect() scan.
- **SYN scan:** Sends SYN packets to open a TCP connection.
- **FIN scan:** Sends SYN/FIN packets.
- **Xmas tree scan:** Uses packets with the SYN, FIN, URG, and PUSH bit set.
- **Null scan:** Sends packets without any TCP flags.
- **ACK scan:** Send packets with the ACK flag on.
- **UDP scan:** Uses UDP packets.

3.3.1 Tools & Equipment

- **Ping(1M)**¹⁵⁸: Standard command on every Solaris system (/usr/sbin/ping). It sends an ICMP echo-request datagram and expects an ICMP echo-reply. It can be configured to send UDP packets to several ports expecting ICMP port unreachables. It can also be configured to use source routing and different TTL values. We will use this tool to test for basic connectivity.
- **Tcpdump**¹⁵⁹ & **Snoop(1M)**¹⁶⁰: Both tcpdump and snoop can capture and analyze network packets. Snoop is a native command on every Solaris system (/usr/sbin/snoop), while tcpdump is not. While tcpdump has a few more features and it is more flexible when writing filtering expressions (from the command line or from a file), we find snoop convenient and easy to use since it can be found on any system. For this particular audit, either one can be used. We will use these tools to monitor what traffic passes through the firewall.
- **Nmap**¹⁶¹: Nmap is a great tool to explore and scan ports. This will be the main tool we will use to determine what traffic is blocked, what services are available, and even OS fingerprinting if we wanted to.
- **Network hubs**: Almost all of the interfaces of the firewall are directly connected to a device: the border router, the VPN concentrator, and the internal router. This makes the monitoring of the traffic on those network difficult. Therefore, when the test requires it, we will connect these devices and the firewall with a network hub (and attach a laptop to the hub) so that traffic can be monitored.
- **Laptops**: Two generic laptops running Solaris will be used to conduct the audit. One will be used to generate the traffic, and the other, will be used to monitor the traffic that leaves other interfaces on the firewall. Currently, the laptops belong to the staff, and we hope that GIACE acquires systems that can be used for auditing soon.

3.3.2 Methodology

Prior to performing the audit:

- Partners, suppliers, remote and local users will be notified days in advance.
- We will disconnect the router from the outside world for security reasons.
- We make sure that our laptops have nmap installed and working properly.
- We verify that there is no unique data on any of the critical infrastructure systems. All servers that provide a service have been configured so that they can be JumpStarted¹⁶² and re-installed from scratch anytime. This is done as a precautionary method since some of the tests might crash the systems.
- We will connect a hub to eth4 (inside interface) and connect the internal router to it. The reason for this is that we want to allow the

¹⁵⁸Solaris man page

¹⁵⁹http://www.tcpdump.org/tcpdump_man.html

¹⁶⁰Solaris man page

¹⁶¹<http://www.insecure.org/nmap>

¹⁶²<http://www.sun.com/software/solaris/8/ds/ds-solwebstart/#6>

firewall to send logs to the log cluster, even when we are generating traffic from this interface.

The audit will be performed by two capable members of our staff. This will ensure that details are not overseen by a single pair of eyes, and that the audit is followed strictly.

Scanning firewall interfaces: When performing this part of our approach, we will have a laptop directly connected to the interface we are testing. The only exception will be eth4 (inside interface), where we will connect the laptop to the hub instead. On all interfaces, the laptop will have an IP that matches the network it is connected to, and will generate all the scanning traffic. Results will be observed by the tool used to generate the scanning traffic. The logs on the log cluster will also be checked to verify that the firewall generated the entries it is supposed to. Since the firewall has 5 interfaces, this test will be performed 5 times. Total duration of this test: (35 minutes per interface) = 2 hours and 55 minutes.

Scanning systems behind the firewall: This step requires the use of the other laptop, and might require a second hub in some combinations. The traffic generating laptop will be directly connected to the interface where we want the traffic to originate from (with the exception of the inside interface, where we will connect to that hub instead). The second laptop will be connected to the interface where the traffic is going to. If the traffic is going to an interface other than eth1 (external service network), then another hub will be required to connect the firewall interface, the network device, and the monitoring interface. The reason for this is to observe the traffic that makes it across the firewall that might not generate a reply that will be received by the traffic generating laptop. Results will be observed by the tool used to generate the traffic and the monitoring laptop. The logs on the log cluster will also be checked to verify that the firewall generated the entries it is supposed to. Considering that traffic from any network or system is only allowed to go to selected networks and systems, we do not expect some of these combinations to be too interesting. Each scan takes about 35 minutes. Total duration of this test: (48 scans * 35 minutes) = 28 hours.

Scanning by replacing a system: This is similar to the previous test, with the difference that the generating laptop will be assigned the IPs of some of the systems behind the interface where the traffic is being originated. Results will be observed by the tool used to generate the traffic and the monitoring laptop. The logs on the log cluster will also be checked to verify that the firewall generated the entries it is supposed to. Needless to say, it will not be productive to use the IP of every system we have. Therefore, we will use IPs of systems that have inbound access for that interface. Considering that traffic from any host or network is only allowed to go to a few networks and systems, we believe that most of the combinations on this test will not be any more interesting than in the previous test. Each scan takes about 35 minutes. Total duration of this test: (121 scans * 35 minutes) = 70 hours and 35 minutes.

3.3.3 IP addresses

The IP addresses of the firewall interfaces we will be scanning are:

- Eth0 interface (x.x.x.13)
- Eth1 interface (10.100.1.13)
- Eth2 interface (10.100.2.2)

- Eth3 interface (10.0.0.2)
- Eth4 interface (10.5.0.2)

It would be a waste of time to try to scan the ports of every possible IP address in our class B network (10.0.0.0/16) since most of them are not used. Also, since most of the firewall interfaces have the same security level, it would be pointless to try to scan systems between those interfaces, due to the fact the firewall cannot route packets between them. Moreover, it would be fruitless to scan those networks that have no interaction whatsoever with the firewall (i.e. the employee network), but for the sake of completeness we will scan an IP address from that subnet.

Therefore, for efficiency we will scan the IPs used by systems or interfaces with different security levels.

When we are scanning systems behind the outside interface, obviously we will not scan every network on the internet. It would take a very long time, and since the router has been unplugged, it would be a waste of time. However, to make sure packets can go out, we will use an arbitrarily chosen internet IP (d.d.d.d), the IP address of one of the public NTP servers and the antivirus site. We will see the packets leaving the outside interface, but needless to say, we will not see responses.

Here is a table with the scans we will be performing:

Scan from	Network destination	System
Outside (Eth0)	Public network (x.x.x.0/28)	External NTP server (x.x.x.1)
Outside (Eth0)	Public network (x.x.x.0/28)	External web server (x.x.x.2)
Outside (Eth0)	Public network (x.x.x.0/28)	External DNS server (x.x.x.3)
Outside (Eth0)	Public network (x.x.x.0/28)	External mail server (x.x.x.4)
Outside (Eth0)	Public network (x.x.x.0/28)	Web proxy (x.x.x.6)
Outside (Eth0)	Public network (x.x.x.0/28)	Log cluster (x.x.x.7)
Outside (Eth0)	Public network (x.x.x.0/28)	VPN concentrator (x.x.x.9)
External service network (eth1)	Internet	Public NTP server (a.a.a.a)
External service network (eth1)	Internet	Antivirus site (c.c.c.c)
External service network (eth1)	Internet	Random Internet IP (d.d.d.d)
External service network (eth1)	Application network (10.1.0.0/24)	Internal web server (10.1.0.2)
External service network (eth1)	Application network (10.1.0.0/24)	Application cluster (10.1.0.3)
External service network (eth1)	Management network (10.4.0.0/24)	Management workstation (10.4.0.3)
External service network (eth1)	Internal service network (10.2.0.0/24)	Internal DNS server (10.2.0.2)
External service network (eth1)	Internal service network (10.2.0.0/24)	Internal mail server (10.2.0.3)
External service network (eth1)	Internal service network (10.2.0.0/24)	Log cluster (10.2.0.6)
External service network (eth1)	Internal service network (10.2.0.0/24)	Web proxy (10.2.0.7)

Scan from	Network destination	System
External service network (eth1)	Internal service network (10.2.0.0/24)	Internal NTP server (10.2.0.9)
External service network (eth1)	Employee network (10.3.0.0/16)	Random Employee network IP (10.3.1.27)
VPN network (eth2)	Internet	Public NTP server (a.a.a.a)
VPN network (eth2)	Internet	Antivirus site (c.c.c.c)
VPN network (eth2)	Internet	Random Internet IP (d.d.d.d)
VPN network (eth2)	Application network (10.1.1.0/24)	Internal web server (10.1.0.2)
VPN network (eth2)	Application network (10.1.1.0/24)	Application cluster (10.1.0.3)
VPN network (eth2)	Management network (10.4.0.0/24)	Management workstation (10.4.0.3)
VPN network (eth2)	Internal service network (10.2.0.0/24)	Internal DNS server (10.2.0.2)
VPN network (eth2)	Internal service network (10.2.0.0/24)	Internal mail server (10.2.0.3)
VPN network (eth2)	Internal service network (10.2.0.0/24)	Log cluster (10.2.0.6)
VPN network (eth2)	Internal service network (10.2.0.0/24)	Web proxy (10.2.0.7)
VPN network (eth2)	Internal service network (10.2.0.0/24)	Internal NTP server (10.2.0.9)
VPN network (eth2)	Employee network (10.3.0.0/16)	Random Employee network IP (10.3.1.27)
VPN user network (eth3)	Internet	Random Internet IP (d.d.d.d)
VPN user network (eth3)	Application network (10.1.0.0/24)	Internal web server (10.1.0.2)
VPN user network (eth3)	Application network (10.1.0.0/24)	Application cluster (10.1.0.3)
VPN user network (eth3)	Management network (10.4.0.0/24)	Management workstation (10.4.0.3)
VPN user network (eth3)	Internal service network (10.2.0.0/24)	Internal DNS server (10.2.0.2)
VPN user network (eth3)	Internal service network (10.2.0.0/24)	Internal mail server (10.2.0.3)
VPN user network (eth3)	Internal service network (10.2.0.0/24)	Log cluster (10.2.0.6)
VPN user network (eth3)	Internal service network (10.2.0.0/24)	Web proxy (10.2.0.7)
VPN user network (eth3)	Internal service network (10.2.0.0/24)	Internal NTP server (10.2.0.9)
VPN user network (eth3)	Employee network (10.3.1.0/16)	Random Employee network IP (10.3.1.27)
Inside (eth4)	Internet	Random Internet IP (d.d.d.d)

<i>Scan from</i>	<i>Network destination</i>	<i>System</i>
Inside (eth4)	External service network (10.100.1.0/24)	External NTP server (10.100.1.1)
Inside (eth4)	External service network (10.100.1.0/24)	External web server (10.100.1.2)
Inside (eth4)	External service network (10.100.1.0/24)	External DNS server (10.100.1.3)
Inside (eth4)	External service network (10.100.1.0/24)	External mail server (10.100.1.4)
Inside (eth4)	VPN network (10.100.2.0/24)	VPN concentrator public (10.100.2.9)
Inside (eth4)	VPN User network (10.0.0.0/24)	VPN concentrator private (10.0.0.9)

Table 29: Scanning combinations

3.3.4 Systems to be replaced

We will scan from the following systems:

<i>Scan from</i>	<i>as</i>
Outside (eth0)	Public NTP server1 (a.a.a.a)
Outside (eth0)	Antivirus site (c.c.c.c)
Outside (eth0)	Border router (x.x.x.14)
External service network (eth1)	External NTP server (10.100.1.1)
External service network (eth1)	External web server (10.100.1.2)
External service network (eth1)	External DNS server (10.100.1.3)
External service network (eth1)	External mail server (10.100.1.4)
External service network (eth1)	External NIDS (10.100.1.5)
VPN network (eth2)	VPN concentrator (public) (10.100.2.9)
VPN User network (eth3)	VPN concentrator (private) (10.0.0.9)
Inside (eth4)	Internal DNS server (10.2.0.2)
Inside (eth4)	Internal mail server (10.2.0.3)
Inside (eth4)	Web proxy (10.2.0.7)
Inside (eth4)	Internal NTP server (10.2.0.9)

Table 30: Systems to be replaced

3.3.5 Syntax

The syntax used by ping is: **ping** <IP>

The syntax used by nmap is:

- **Connect scan:** **nmap -sT -P0 -n -O -vv -T 4 <IP>**
- **Syn scan:** **nmap -sS -P0 -n -O -vv -T 4 <IP>**
- **Fin scan:** **nmap -sF -P0 -n -O -vv -T 4 <IP>**
- **Xmas tree scan:** **nmap -sX -P0 -n -O -vv -T 4 <IP>**
- **Null scan:** **nmap -sN -P0 -n -O -vv -T 4 <IP>**
- **Ack scan:** **nmap -sA -P0 -n -O -vv -T 4 <IP>**
- **UDP scan:** **nmap -sU -P0 -n -O -vv -T 4 <IP>**

3.3.6 Costs, time and risks

Since the audit will be performed by our own staff we will not be paying any consulting fees to external individuals. Moreover, since all of the software tools are free, we already own the network equipment (hubs), and the laptops are gracefully borrowed from staff members, we have no expenses in terms of tools and equipment. However, the cost will be the overtime paid to our two engineers, and most importantly the opportunity cost of the potential sales that will not be done since customers will not have access to the site.

There is the possibility that our scans might crash some of the systems, but considering all of our systems are configured to be easily re-installed from scratch, the risks are very low. Regardless of whether the system crashes or not, we will check Tripwire¹⁶³ to make sure nothing was changed or corrupted from the test. If so, we will simply JumpStart¹⁶⁴ the system. JumpStarting and rebuilding one of our systems from scratch should not take more than two hours before it is back up. In the extreme event that the crash actually creates hardware malfunctions, we have a few spare parts, but we do not expect this to be a likely possibility.

The expected duration of the audit is:

<i>Action</i>	<i>Total amount of time</i>
Scanning firewall interfaces	3 hours
Scanning systems behind the firewall	28 hours
Scanning by replacing a system	71 hours
Compiling and analyzing results	4 hours
Possible changes to the configuration	(12 hours)
Time allocated in case systems have to be rebuilt	(8 hours)
Reconnect, bring site back up, and verify connectivity	2 hour

Table 31: Expected duration of the audit

Total amount of time: 108 hours (possibly 128 hours)

This means the site will not be accessible to customers for about 6 days. However, we feel the audit is necessary before we go live in order to ensure there are no overseen security flaws that might cost us more than the opportunity cost of having our site unavailable to customers for 6 days.

3.3.7 Expected results

We expect all connectivity and scans done on the firewall interfaces not to show any services or ports available. In fact, we expect nmap to believe the firewall is not even up.

On scans performed on systems behind the firewall, we expect it will only find those services available to the originating network/host.

When replacing a system with our scanning laptop, we do not expect it will find any services or ports in addition to the ones the server is supposed to have access to.

¹⁶³<http://www.tripwire.org>

¹⁶⁴<http://www.sun.com/software/solaris/8/ds/ds-solwebstart/#6>

3.4 Conducting the audit

With the laptops ready, we begin our audit.

3.5 Scanning firewall interfaces

We scan the firewall from each interface.

3.5.1 Scanning the firewall from eth0 (outside)

<i>Scan</i>	<i>Results</i>
ping x.x.x.13 ping 10.100.1.13 ping 10.5.0.2 ping 10.100.2.2	Did not generate an echo-reply. Logs were generated indicating the ICMP packets were dropped. Another log entry was generated by the Cisco IDS to indicate an echo request was received.
Connect scan Syn scan	Could not connect to any ports. Each test took about 300 seconds.
Fin scan Xmas tree scan Null scan	No TCP responses received. Nmap assumes all ports are filtered. Each test took about 300 seconds. The Cisco IDS reported some of the scans.
Ack scan	No responses. Test took about 300 seconds.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. Test took about 300 seconds.

Table 32: Scanning firewall interfaces from eth0

3.5.2 Scanning the firewall from eth1 (external service network)

<i>Scan</i>	<i>Results</i>
ping x.x.x.13 ping 10.100.1.13 ping 10.5.0.2 ping 10.100.2.2	Did not generate an echo-reply. Logs were generated indicating the ICMP packets were dropped. Another log entry was generated by the Cisco IDS to indicate an echo request was received.
Connect scan Syn scan	Could not connect to any ports. Each test took about 300 seconds.
Fin scan Xmas tree scan Null scan	No TCP responses received. Nmap assumes all ports are filtered. Each test took about 300 seconds. The Cisco IDS reported some of the scans.
Ack scan	No responses. Test took about 300 seconds.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. Test took about 300 seconds.

Table 33: Scanning firewall interfaces from eth1

3.5.3 Scanning the firewall from eth2 (VPN network)

<i>Scan</i>	<i>Results</i>
ping x.x.x.13 ping 10.100.1.13 ping 10.5.0.2 ping 10.100.2.2	Did not generate an echo-reply. Logs were generated indicating the ICMP packets were dropped. Another log entry was generated by the Cisco IDS to indicate an echo request was received.
Connect scan Syn scan	Could not connect to any ports. Each test took about 300 seconds.
Fin scan Xmas tree scan Null scan	No TCP responses received. Nmap assumes all ports are filtered. Each test took about 300 seconds. The Cisco IDS reported some of the scans.
Ack scan	No responses. Test took about 300 seconds.

<i>Scan</i>	<i>Results</i>
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. Test took about 300 seconds.

Table 34: Scanning firewall interfaces from eth2

3.5.4 Scanning the firewall from eth3 (VPN User network)

<i>Scan</i>	<i>Results</i>
ping x.x.x.13 ping 10.100.1.13 ping 10.5.0.2 ping 10.100.2.2	Did not generate an echo-reply. Logs were generated indicating the ICMP packets were dropped. Another log entry was generated by the Cisco IDS to indicate an echo request was received.
Connect scan Syn scan	Could not connect to any ports. Each test took about 300 seconds.
Fin scan Xmas tree scan Null scan	No TCP responses received. Nmap assumes all ports are filtered. Each test took about 300 seconds. The Cisco IDS reported some of the scans.
Ack scan	No responses. Test took about 300 seconds.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. Test took about 300 seconds.

Table 35: Scanning firewall interfaces from eth3

3.5.5 Scanning the firewall from eth4 (inside)

<i>Scan</i>	<i>Results</i>
ping x.x.x.13 ping 10.100.1.13 ping 10.5.0.2 ping 10.100.2.2	Did not generate an echo-reply. Logs were generated indicating the ICMP packets were dropped. Another log entry was generated by the Cisco IDS to indicate an echo request was received.
Connect scan Syn scan	Could not connect to any ports. Each test took about 300 seconds.
Fin scan Xmas tree scan Null scan	No TCP responses received. Nmap assumes all ports are filtered. Each test took about 300 seconds. The Cisco IDS reported some of the scans.
Ack scan	No responses. Test took about 300 seconds.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. Test took about 300 seconds.

Table 36: Scanning firewall interfaces from eth4

3.6 Scanning systems behind the firewall

We scan our networks from each interface and monitor the traffic from the intended interface.

3.6.1 Scanning systems behind the firewall from eth0 (outside)

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: <div> <div>25/tcp</div> <div>open</div> <div>x.x.x.4 (leaving interface: INT_service)</div> </div> <div> <div>80/tcp</div> <div>open</div> <div>x.x.x.2 (leaving interface: INT_service)</div> </div> <div> <div>443/tcp</div> <div>open</div> <div>x.x.x.2 (leaving interface: INT_service)</div> </div> <p>The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed.</p>

<i>Scan</i>	<i>Results</i>
Fin scan Xmas tree scan Null scan Ack scan	No TCP responses received. Nmap assumes all ports are filtered. No traffic leaving the firewall is seen by the monitoring laptop. The Cisco IDS reported some of the scans. No traffic leaving the firewall was seen by the monitoring laptop.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 53/udp x.x.x.3 (leaving interface: INT_service) 500/udp x.x.x.9 (leaving interface: INT_VPN) No replies were observed by the monitoring laptop.

Table 37: Scanning systems behind the firewall from eth0

3.6.2 Scanning systems behind the firewall from eth1 (external service network)

Laptop was assigned the internal IP 10.100.1.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 514/udp 10.2.0.6 (leaving interface: inside) No replies were observed by the monitoring laptop.

Table 38: Scanning systems behind the firewall from eth1

3.6.3 Scanning systems behind the firewall from eth2 (VPN network)

Laptop was assigned the internal IP 10.100.2.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No replies were observed by the monitoring laptop leaving any interface.

Table 39: Scanning systems behind the firewall from eth2

3.6.4 Scanning systems behind the firewall from eth3 (VPN User network)

Partners network:

Laptop was assigned the internal IP 10.0.1.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 80/tcp open 10.1.0.2 (leaving interface: inside) 443/tcp open 10.1.0.2 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. The were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface.

Table 40: Scanning systems from the partners network

Remote sales network:

Laptop was assigned the internal IP 10.0.2.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open 10.2.0.3 (leaving interface: inside) 80/tcp open 10.1.0.2 (leaving interface: inside) 443/tcp open 10.1.0.2 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. The were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface.

Table 41: Scanning systems from the remote sales network

Telecommuters network:

Laptop was assigned the internal IP 10.0.3.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open 10.2.0.3 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. The were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed by the monitoring laptop leaving any interface.

Table 42: Scanning systems from the telecommuters network

Remote administrators network:

Laptop was assigned the internal IP 10.0.4.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 22/tcp open 10.4.0.3 (leaving interface: inside) 22/tcp open 10.4.0.4 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed by the monitoring laptop leaving any interface.

Table 43: Scanning systems from the remote administrators network

3.6.5 Scanning systems behind the firewall from eth4 (inside)

Laptop was assigned the internal IP 10.5.0.10.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface.

Table 44: Scanning systems behind the firewall from eth3

3.7 Scanning systems behind the firewall by replacing a system

We scan our networks from each interface and monitor the traffic leaving the firewall. The traffic will be originating from the system we replace. This system is taken out of the network.

3.7.1 Scanning by replacing a system from eth0 (outside)

As the border router: x.x.x.14

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open x.x.x.4 (leaving interface: INT_service) 80/tcp open x.x.x.2 (leaving interface: INT_service) 443/tcp open x.x.x.2 (leaving interface: INT_service) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	No TCP responses received. Nmap assumes all ports are filtered. No traffic leaving the firewall is seen by the monitoring laptop. The Cisco IDS detected some of the scans. No traffic leaving the firewall is seen by the monitoring laptop.

<i>Scan</i>	<i>Results</i>
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 53/udp x.x.x.3 (leaving interface: INT_service) 123/udp x.x.x.1 (leaving interface: INT_service) 500/udp x.x.x.9 (leaving interface: INT_VPN) 514/udp x.x.x.7 (leaving interface: inside) No replies were observed by the monitoring laptop.

Table 45: Scanning systems as the border router

As the public NTP server1 and the antivirus site: a.a.a.a and c.c.c.c.

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open x.x.x.4 (leaving interface: INT_service) 80/tcp open x.x.x.2 (leaving interface: INT_service) 443/tcp open x.x.x.2 (leaving interface: INT_service) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed.
Fin scan Xmas tree scan Null scan Ack scan	No TCP responses received. Nmap assumes all ports are filtered. No traffic leaving the firewall is seen by the monitoring laptop. The Cisco IDS reported some of the scans. No traffic is observed leaving the firewall.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 53/udp x.x.x.3 (leaving interface: INT_service) 500/udp x.x.x.9 (leaving interface: INT_VPN) No replies were observed by the monitoring laptop.

Table 46: Scanning systems as the public NTP servers and antivirus site

3.7.2 Scanning by replacing a system from eth1 (external service network)

As the external NTP server: 10.100.1.1

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 123/udp a.a.a.a (leaving interface: outside) 514/udp 10.2.0.6 (leaving interface: inside) No replies were observed by the monitoring laptop. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 47: Scanning as the external NTP server

As the external web server: 10.100.1.2

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 5000/tcp open 10.1.0.2 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
Fin scan Xmas tree scan Null scan Ack scan	No TCP responses received. Nmap assumes all ports are filtered. No traffic leaving the firewall is seen by the monitoring laptop. An entry was generated indicating an ACK packet was received but no connection existed. No traffic leaving the firewall is seen by the monitoring laptop. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 514/udp 10.2.0.6 (leaving interface: inside) No replies were observed by the monitoring laptop.

Table 48: Scanning as the external web server

As the external mail server: 10.100.1.4

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open 10.2.0.3 (leaving interface: inside) The monitoring laptop observed the traffic above going to the intended systems and the replies, and all packets going to the 10.2.0.0/24 network passed the firewall and left through the inside interface. Only replies from 10.2.0.3 were seen. The following traffic was also observed leaving the firewall, but no replies were received. 25/tcp Internet (leaving interface:outside) 80/tcp c.c.c.c (leaving interface: outside) The were also logs indicating packets being blocked or allowed. They logs also show that 25/tcp packets going to some interfaces were not sent because the interfaces had the same security level. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 514/udp 10.2.0.6 (leaving interface: inside) No replies were observed by the monitoring laptop. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 49: Scanning as the external mail server

As the external DNS server: 10.100.1.3

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.

<i>Scan</i>	<i>Results</i>
UDP scan	<p>No UDP responses received. Nmap assumes all ports are filtered.</p> <p>The monitoring laptop observed that the following traffic did make it through:</p> <p>53/udp Internet (leaving interface: outside)</p> <p>53/udp 10.2.0.0/24 (leaving interface: inside)</p> <p>514/udp 10.2.0.6 (leaving interface: inside)</p> <p>No replies were observed by the monitoring laptop. There were also logs indicating packets being blocked or allowed. They logs also show that 53/udp packets going to some interfaces were not sent because the interfaces had the same security level. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.</p>

Table 50: Scanning as the external DNS server

As the external NIDS: 10.100.1.5

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	<p>Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.</p>
UDP scan	<p>No UDP responses received. Nmap assumes all ports are filtered.</p> <p>The monitoring laptop observed that the following traffic did make it through:</p> <p>514/udp 10.2.0.6 (leaving interface: inside)</p> <p>No replies were observed by the monitoring laptop.</p> <p>In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.</p>

Table 51: Scanning as the external NIDS

3.7.3 Scanning by replacing a system from eth2 (VPN network)

As the VPN concentrator (public interface): 10.100.2.9

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	<p>Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, the Cisco IDS reported some of the scans.</p>
UDP scan	<p>No UDP responses received. Nmap assumes all ports are filtered.</p> <p>The monitoring laptop observed that the following traffic did make it through:</p> <p>500/udp Internet (leaving interface: outside)</p> <p>514/udp 10.2.0.6 (leaving interface: inside)</p> <p>No replies were observed by the monitoring laptop. There were also logs indicating packets being blocked or allowed. They logs also show that 500/udp packets going to other interfaces were not routed because the interfaces had the same security level or because they could not be translated.</p>

Table 52: Scanning as the VPN concentrator (public)

3.7.4 Scanning by replacing a system from eth3 (VPN_user network)

As the VPN concentrator (private interface): 10.0.0.9

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, and the Cisco IDS reported some of the scans.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface.

Table 53: Scanning as the VPN concentrator (private)

3.7.5 Scanning by replacing a system from eth4 (inside)

As the internal DNS server : 10.2.0.2

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, and the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 53/udp 10.100.1.3 (leaving interface: INT_service) No replies were observed. The NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 54: Scanning as the internal DNS server

As the internal mail server : 10.2.0.3

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Both tests found: 25/tcp open 10.100.1.4 (leaving interface: INT_service) The monitoring laptop observed the traffic above going to the intended systems and the replies. No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed. The NIDS on the network generated alerts for the unusual traffic trying to leave the network.
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, and the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface. The NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 55: Scanning as the internal mail server

As the web proxy : 10.2.0.7

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan	Nmap received no responses. However, the following traffic did make it through: 80/tcp Internet (leaving interface:outside) 443/tcp Internet (leaving interface:outside) No other traffic was observed leaving the firewall. There were also logs indicating packets being blocked or allowed. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.

<i>Scan</i>	<i>Results</i>
Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, and the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. No packets were observed leaving any interface. The NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 56: Scanning as the web proxy

As the internal NTP server : 10.2.0.9

<i>Scan</i>	<i>Results</i>
Connect scan Syn scan Fin scan Xmas tree scan Null scan Ack scan	Nmap received no responses and no traffic was observed leaving any interface by the monitoring laptop. Logs were generated for packets being dropped, and the Cisco IDS reported some of the scans. In addition the NIDS on the network generated alerts for the unusual traffic trying to leave the network.
UDP scan	No UDP responses received. Nmap assumes all ports are filtered. The monitoring laptop observed that the following traffic did make it through: 123/udp 10.100.1.1 (leaving interface: INT_service) No packets were observed leaving any interface. The NIDS on the network generated alerts for the unusual traffic trying to leave the network.

Table 57: Scanning as the internal NTP server

3.8 Evaluate the audit

Scanning firewall interfaces:

Based on the results, there was no response by our firewall on any of its interfaces, which is what we expected.

Scanning systems behind the firewall:

Here are the results in a tabulated format of the traffic that did manage to go through the firewall:

<i>Interface</i>	<i>Port/Protocol</i>	<i>Destination</i>	<i>Received reply?</i>	<i>Expected?</i>
outside	25/tcp	x.x.x.4	Yes	Yes
outside	80/tcp, 443/tcp	x.x.x.2	Yes	Yes
outside	53/udp	x.x.x.3	No	Yes
outside	500/udp	x.x.x.9	No	Yes
INT_service	514/udp	10.2.0.6	No	Yes
INT_VPNUser (remote administrators subnet)	22/tcp	10.4.0.3 10.4.0.4	Yes	Yes
INT_VPNUser (sales and telecommuters subnet)	25/tcp	10.2.0.3	Yes	Yes
INT_VPNUser (sales and partners subnet)	80/tcp, 443/tcp	10.1.0.2	Yes	Yes

Table 58: Results from scanning systems behind the firewall

We are glad that the results only show small subset of what we intended with our security policy. Based on external scans, it can only be determined that we

have a system with TCP port 25 open and a system with TCP port 80 and 443 open. Scans from other interfaces originating from systems other than our infrastructure server do not reflect much either. The results also show that a compromised remote user (partner, telecommuter, remote sales staff, and remote administrators) has very limited systems to attack. Any scan to find these systems would be easily detected and logged. The firewall will offer some protection by inspecting the packets for port 25 and port 80 (fixup protocol setting), and SSH on the systems running Trusted Solaris should offer an adequate amount of protection.

Scanning by replacing a system:

Here are the results in a tabulated format of the traffic that did manage to go through the firewall:

<i>System replaced</i>	<i>Port/Protocol</i>	<i>Destination</i>	<i>Received reply?</i>	<i>Expected?</i>
Border router	25/tcp	x.x.x.4	Yes	Yes
Border router	80/tcp, 443/tcp	x.x.x.2	Yes	Yes
Border router	53/udp	x.x.x.3	No	Yes
Border router	123/udp	x.x.x.1	No	Yes
Border router	500/udp	x.x.x.9	No	Yes
Border router	514/udp	x.x.x.7	No	Yes
Public NTP servers and antivirus site	25/tcp	x.x.x.4	Yes	Yes
Public NTP servers and antivirus site	80/tcp, 443/tcp	x.x.x.2	Yes	Yes
Public NTP servers and antivirus site	53/udp	x.x.x.3	No	Yes
Public NTP servers and antivirus site	500/udp	x.x.x.9	No	Yes
External NTP server	123/udp	a.a.a.a	No	No
External NTP server	514/udp	10.2.0.6	No	Yes
External web server	5000/tcp	10.1.0.2	Yes	Yes
External web server	514/udp	10.2.0.6	No	Yes
External DNS server	53/udp	Internet	No	Yes
External DNS server	53/udp	10.2.0.0/24 (*)	No	NO
External DNS server	514/udp	10.2.0.6	No	Yes
External mail server	25/tcp	Internet	No	Yes
External mail server	25/tcp	10.2.0.3	Yes	Yes
External mail server	25/tcp	10.2.0.0/24 (*)	Only from 10.2.0.3	NO
External mail server	80/tcp	c.c.c.c	No	Yes
External mail server	514/udp	10.2.0.6	No	Yes
External NIDS	514/udp	10.2.0.6	No	Yes
VPN Concentrator(public)	500/udp	Internet	No	Yes
VPN Concentrator(public)	500/udp	Interfaces other than outside(*)	No	NO
VPN Concentrator(public)	514/udp	10.2.0.6	No	Yes

System replaced	Port/Protocol	Destination	Received reply?	Expected?
Internal DNS server	53/udp	10.100.1.3	No	No
Internal mail server	25/tcp	10.100.1.4	Yes	Yes
Web proxy	80/tcp, 443/tcp	Internet web servers	Yes	Yes
Internal NTP server	123/udp	10.100.1.1	No	Yes

Table 59: Results from scanning by replacing a system

The results show a subset of what is intended with the security policy of the firewall. It is worth mentioning that traffic going to the random external network we chose on the internet never left our network since we were disconnected from the outside world during the audit. Therefore, it would have been impossible for replies to be received.

(*) However, we encountered an unexpected behavior from three of our systems: the external mail server, external DNS server, and the public interface of the VPN concentrator. It seems the ACLs for the external service network interface and the ACL for the VPN network interface do not implement correctly what was intended in our security policy.

The original access control lists in question for both interfaces are:

```
access-list from_service permit tcp host EXT_MAIL any eq smtp
access-list from_service permit tcp host EXT_DNS any eq domain
access-list from_VPN permit esp host VPN_PUB any
access-list from_VPN permit udp host VPN_PUB any eq isakmp
```

The problem here is the “any” keyword. It was meant to be “any system outside of our network”, however, it will unintentionally try to send traffic to internal systems on other interfaces as well. When we conducted the audit, this traffic did not leave some interfaces only because the firewall could not translate it properly. Moreover, we did not see the same behavior with outbound IPsec since Nmap did not generate IP traffic with protocol 50.

The external mail server should be able to initiate a connection to send mail to the internal mail server and deliver mail to systems outside of our network. The external DNS server should reply to queries from the internal DNS server and query DNS servers outside of our network. The VPN concentrator should never send IPsec or IKE traffic to internal systems, only systems outside of our network.

We will change those ACL rules to be:

```
access-list from_service permit tcp host EXT_MAIL host INT_MAIL eq smtp
access-list from_service deny tcp host EXT_MAIL 10.0.0.0 0.255.255.255 eq smtp
access-list from_service permit tcp host EXT_MAIL any eq smtp

access-list from_service deny tcp host EXT_DNS 10.0.0.0 0.255.255.255 eq domain
access-list from_service permit tcp host EXT_DNS any eq domain

access-list from_VPN deny esp host VPN_PUB 10.0.0.0 0.255.255.255
access-list from_VPN permit esp host VPN_PUB any
access-list from_VPN deny udp host VPN_PUB 10.0.0.0 0.255.255.255 eq isakmp
access-list from_VPN permit udp host VPN_PUB any eq isakmp
```

These changes will correctly implement our security policy, and resolve the unexpected behavior reported in our audit.

The results show that even if one of our servers is compromised, the

access from that system is very limited. We hope that the internal logs from the systems (Tripwire¹⁶⁵, etc.), the NIDS on both our internal and external service networks, and the logs generated by the router and firewall can serve as an early warning system from internal and external attacks.

3.9 Improvements and recommendations

As resources allow it, we would like to implement some modest improvements and recommendations to the current external service network. We list them from most critical to least critical:

- **Reverse web proxies:** We feel that having a reverse proxy on our exposed external web server would be a great addition to the configuration. Secondly, adding a reverse proxy to our internal web server, which is exposed to some remote and local users, would be very beneficial as well. Our business depends on web access in order to operate properly. Protecting them as much as we can is essential.
- **Additional web servers and a load balancers:** By having multiple web servers performing the same service we can increase the traffic load we can handle.
- **Additional NTP servers:** We prefer to have them in pairs for redundancy. Two on each service network instead of only one on each.




Illustration 2: Potential improvements for the external service network

¹⁶⁵<http://www.tripwire.org>

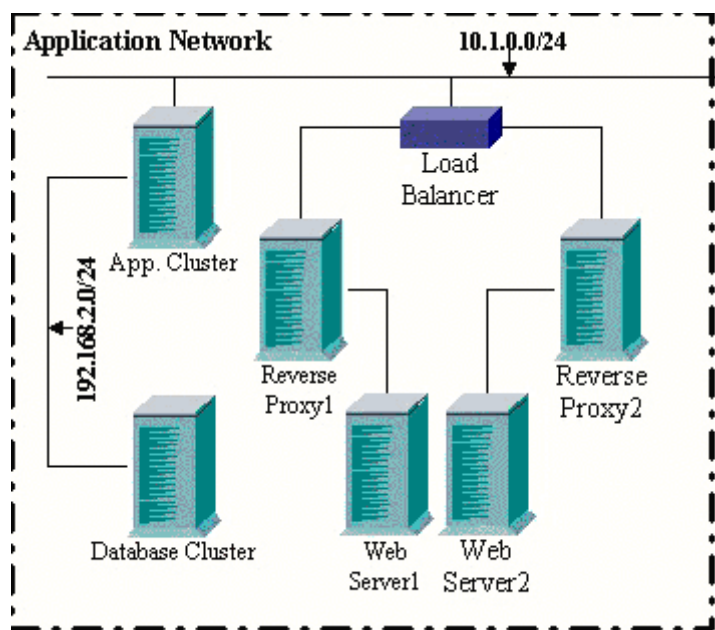


Illustration 3: Potential improvements for the application network

A not so modest improvement would be a firewall cluster for high availability, or an internal firewall.

© SANS Institute 2003, Author retains full rights.

4. Assignment 4 – Design Under Fire

The main difference between the typical “script kiddie” running a script or tool that he downloaded and managed to compile, and a potentially successful or more covert attack is definitely having a good planning and reconnaissance phase followed by methodical research and possible execution, if the required elements are met and all factors are favorable. Another main difference is commitment.

GIACE has two main competitors. After shedding all the unnecessary ethical and moral hindrances that would otherwise prevent us from performing this, we will attack one of our competitors and try to point all the attention to the other competitor.

For our attack, we have chosen Vivekanand Chudgar's design¹⁶⁶(GCFW #0381). The main reason for it, is that it is one of the five most recent designs listed, so in theory, everything probably has the most recent patches installed, making our attack more difficult and challenging.

Now that we have a target, we have three goals in mind:

- Perform a direct attack against their firewall
- Execute a denial of service attack
- Compromise an internal system through the perimeter system




Illustration 4: Vivekanand Chudgar's design

4.1 Reconnaissance

First, we must gather all the information we might need. The tools used for this are:

¹⁶⁶http://www.giac.org/practical/GCFW/Vivekanand_Chudgar.pdf

whois¹⁶⁷, **nslookup**¹⁶⁸: There are many sites online that allow people to perform whois/nslookup inquiries anonymously. This will help us find out what IP addresses they own, names, addresses, phone numbers, DNS and mail servers, etc.

traceroute(1M)¹⁶⁹: It helps us determine the route followed by packets on the way to/from the destination. There are many sites online¹⁷⁰ that will perform traceroutes for any user online.

Nmap¹⁷¹, **nessus**¹⁷²: Great tools for network and port mapping, fingerprinting, and to check for some vulnerabilities. Fingerprinting is crucial for reconnaissance. Knowing what we are dealing with, helps us narrow our search for vulnerabilities and to customize our attack.

Web browser: Names, phone numbers and additional information can be gathered by browsing public pages of the target and looking at the source code. It can also help us determine the type of web server and the platform.

Social Engineering: Despite common sense, when gathering information it is easy to get information out of IT people by simply asking. Granted, the questions are not asked directly. An all time favorite of our staff is to call somebody on behalf of a certain IT magazine. From the information we have gathered so far, we call the people we believe maintain their network and we offer a free subscription to "our" well-known magazine. In the IT industry we get these calls, all the time. However, most people do not ask for a callback number, or ask the caller for some information to verify it is a legitimate call. They might be already subscribed, in which case after some typing noises, we simply tell them we would want to renew their free subscription after answering a few questions. For the most part, we will use the same questions asked by the magazine people (type of company, number of employees, whether they have purchasing/evaluation roles, etc.) The information we need comes when we start questions like: "Does your company currently use or plan to use any of the following products? Cisco PIX, Check Point FW-1, other (please specify), etc." In this manner we can gather information about firewalls, IDS, routers, web servers, platforms, etc.

Inside information: The right incentive, even revenge, can motivate and/or convince an existing or former employee to provide information that can be useful, or even be instructed to perform some internal scans.

4.2 Research

Good places to check for any recent or outstanding vulnerabilities for our particular target are: Bugtraq¹⁷³, CERT¹⁷⁴, CVE¹⁷⁵, ISS's X-Force¹⁷⁶, SecuriTeam¹⁷⁷, SANS' top 20¹⁷⁸, NIPC¹⁷⁹ and the site of the vendor of the device we are planning to attack.

¹⁶⁷<http://www.whois.net> <http://www-whois.internic.net/cgi/whois>

¹⁶⁸<http://gday.bloke.com/cgi-bin/nslookup>

¹⁶⁹Solaris man page

¹⁷⁰<http://cache2.online.ee:81/cgi-bin/nph-traceroute>

¹⁷¹<http://www.insecure.org/nmap>

¹⁷²<http://www.nessus.org>

¹⁷³<http://online.securityfocus.com>

¹⁷⁴<http://www.cert.org>

¹⁷⁵<http://www.cve.mitre.org>

¹⁷⁶<http://xforce.iss.net/index.php>

¹⁷⁷<http://www.securiteam.com>

¹⁷⁸<http://www.sans.org/top20>

¹⁷⁹<http://www.nipc.gov>

4.3 Direct attacks against the firewall

4.3.1 Vulnerability research

After gathering information, we know our target is using a FW-1 firewall and since their site just went up recently (a few days ago), we can probably assume it is relatively up to date. We were happy to find very recent vulnerabilities that seem promising.

4.3.2 Syslog attack

Check Point VPN-1/Firewall-1 Remote Syslog Data Resource Consumption Vulnerability¹⁸⁰:

Bugtraq ID: 7159

From the SecurityFocus site:

It has been reported that some versions of Firewall-1 and VPN-1 may experience performance problems when allowing remote syslog traffic. An attacker could exploit this issue to deny service to legitimate users of the network serviced by the software.

Check Point FW-1 Syslog Daemon Unfiltered Escape Sequence Vulnerability¹⁸¹:

Bugtraq ID: 7161

From the SecurityFocus site:

An issue has been discovered in Check Point FW-1 syslog daemon when attempting to process a malicious, remotely supplied, syslog message. Specifically, some messages containing escape sequences are not properly filtered out. This may result in unpredictable behavior by the Check Point syslog daemon.

Both of these attacks are syslog related, and we even have code taken from bugtraq to generate malicious syslog messages:

```
[attacker]# echo -e "<189>19: 00:01:04:
Test!a!033[2J!033[2;5m!033[1;31mHACKER~
ATTACK!033[2;25m!033[22;30m!033[3q" | nc -u firewall 514
```

While, the vulnerabilities have some potential, we do not expect them to work. It is very unlikely for a properly configured firewall to accept external syslog traffic. The only possible exception might be allowing logs from the border router to pass through. Even then, sending syslog traffic by spoofing the IP of the border router would not work with a properly configured router, since it should block spoofed packets with its own IP.

Moreover, even if the traffic is allowed in, we cannot be for certain that we are successful unless we can see a drop of performance or the firewall crashes/hangs. It is still worth sending spoofed UDP packets containing our malformed logs from of our compromised systems.

A quick look at Chudgar's border router configuration tells us this attack will not work:

!These two lines will block spoofed and syslog traffic from coming in.

```
access-list 101 deny ip 202.54.1.32 0.0.0.15 any log
access-list 101 deny udp any any eq syslog
```

Syslog attack countermeasures:

To ensure an attack like that cannot succeed we can:

¹⁸⁰<http://www.securityfocus.com/bid/7159>

¹⁸¹<http://www.securityfocus.com/bid/7161>

- Block all external syslog traffic at the border router.
- Block all spoofed traffic at the border router.
- When allowing the border router to send logs, configure the firewall to only allow syslogs from the border router on the outside interface.
- Upgrade the firewall, or install Check Point's fix¹⁸². This will protect the firewall from internal attacks.

4.3.3 IKE aggressive mode attack

Check Point VPN-1 IKE Aggressive Mode Forcing Vulnerability¹⁸³

Bugtraq ID: 5920

From the SecuryFocus site:

Under some circumstances, VPN-1 can be forced into negotiating sessions in aggressive mode. If the system has been configured to a mode other than aggressive, and a user attempts to establish a session using aggressive mode, VPN-1 will negotiate the session.

Considering this issue had already been resolved before the network was configured, it is not likely that Chudgar's firewall does not have the fix for this. The approach is simple, we construct an IKE packet for the phase-1 of the negotiation, we send it and wait for the response. Responses, as well as detailed description of the standard can be found in RFC 2408¹⁸⁴. Once usernames have been found, a brute force dictionary approach can be used to guess the password.

From the bugtraq archives¹⁸⁵, we have details on the exploit, and an example of code exploiting the vulnerability:

Issue Details:

If we send an IKE Phase-1 aggressive mode packet with the following payloads:

- ISAKMP Header
- SA - Containing one proposal with four transforms
- Key Exchange - DH Group 2
- Nonce
- Identification - Type ID_USER_FQDN, Value is SecuRemote username

Example 1: This example which shows the username guessing program being run against a

Firewall-1 v4.1 SP6 system:

```
Script started on Thu Aug 22 15:15:30 2002
rsh@radon [499]% fw1-ike-userguess --file=testusers.txt --sport=0 172.16.2.2
testuser User testuser unknown.
test-ike-3des USER EXISTS
testing123 User testing123 unknown.
test-ike-des USER EXISTS
guest User guest unknown.
test-fwz-des User cannot use IKE
test-ike-cast40 USER EXISTS
test-ike-ah USER EXISTS
test-ike-hybrid IKE is not properly defined for user.
test-expired Login expired on 1-jan-2002.
```

In this example, the users "test-ike-3des", "test-ike-des",

¹⁸²<http://www.checkpoint.com/techsupport/alerts/syslog.html>

¹⁸³<http://www.securityfocus.com/bid/5920>

¹⁸⁴<http://www.ietf.org/rfc/rfc2408.txt>

¹⁸⁵<http://archives.neohapsis.com/archives/bugtraq/2002-09/0017.html>

"test-ike-cast40" and "test-ike-ah" exist and have valid IKE configurations with shared secret auth; the users "testuser", "testing123" and "guest" do not exist; and the users "test-fwz-des", "test-ike-hybrid" and "test-expired" exist but cannot use IKE for various reasons which are explained in the Firewall message.

We will not bother trying this attack. Chudgar's firewall is not vulnerable to it. Moreover, even if the attack is spread over a long period of time, the logs will probably alert the administrators long before we can find a successful match.

IKE aggressive mode countermeasures:

To ensure this attack cannot succeed we can:

- Upgrade our firewall and/or apply Check Point's patch¹⁸⁶
- Use digital certificates
- Implement a strong username and password policy not vulnerable to dictionary attacks
- Use hybrid mode for authentication

4.3.4 Flood attacks

Realistically, attacking a firewall directly should not be possible, since in a proper configuration the border router should always block all traffic directed to its own or the firewall's interfaces. However, this is not always possible in cases where the firewall is used for other services: VPN, for example. With our compromised systems we can flood the firewall with random spoofed UDP traffic going to the port listening for IKE packets, and random spoofed TCP packets going to port 264, which is waiting for SecuRemote client connections. For our own amusement, we will use IPs from other competitor when spoofing some of the packets. We will not use those IPs too much, otherwise the flood can easily be filtered.

Chudgar has SYNDefender enabled, but we noticed that the timeout was increased from 10 to 45 seconds. We believe increasing the timeout like this would only benefit two kinds of people: users connecting from a busy server on the other side of the world on a dial-up connection and going through a satellite link at some point in between, since their connections will timeout less often; and people who want to create a DoS against this network. Most traffic online under normal circumstances should not take more than 10 seconds. Moreover, SYNDefender does not protect against UDP floods.

We believe this attack would provide a decent level of success depending on how many compromised systems we use to flood the firewall.

Flood attacks countermeasures:

To minimize the effect of this kind of attack against the firewall we could:

- Only allow remote users to connect from specific IP addresses or networks, any traffic directed to the firewall that does not originate from there should be dropped by the border router. This will not protect us against traffic spoofing those addresses
- Adjust the SYNDefender timeout
- Use some of the countermeasures against DoS attacks that we will mention on the following section

¹⁸⁶<http://www.checkpoint.com/techsupport/alerts/ike.html>

4.4 Denial of service attack

Considering we already “0VV|\|” 50 systems with a cable/DSL connection, conducting a distributed DoS attack should be fairly simple and effective. There are many tools available, but we believe TFN2K¹⁸⁷ would be most effective.

TFN2K works by having a daemon on the compromised systems listening for instructions from the system acting as the master.

The syntax for it from the master is:

```
#tfn -f <file of compromised hosts>
-i <IP of system to attack>
-p <Specify a port. Otherwise, it is random>
-c <Command to be sent to the compromised hosts>
```

There are more options, but we will not mention them since we will not be using them. Also, we will use the following arguments to the command flag:

- 0: Tells the systems to stop all attacks
- 5: Flood with TCP/SYN packets
- 8: Flood with a mix of TCP/UDP/ICMP packets

For more information on TFKN and its syntax, an good article can be found at:

<http://www.securiteam.com/securitynews/5YP0G000FS.html>

There are different ways in which we can cause a denial of service: we can overwhelm their network with traffic, thus preventing legitimate traffic from entering or leaving their network; or we could flood a particular system, thus wasting most of its resources (or causing it to crash) and preventing it from serving legitimate clients.

Our target is an E-commerce company, therefore we have decided to attack its web server. Moreover, their T1 connection can easily be filled by attacking with less than half of our compromised systems.

4.4.1 Approach

We will divide our compromised systems into three attacking groups. Some will attack port 80 of the web server; some systems will have a random mix of attacks to random ports on the mail server (they will probably be blocked by the firewall); and some will remain on reserve.

Why do we have random attacks on random ports on the mail server, which will probably be filtered by the firewall? To make it harder on the administrators. Besides, the traffic will still help to fill their T1 link.

We will set up a cron(1M)¹⁸⁸ job on the master system that will run every 20 minutes. For maximum effect and efficiency, we will start the attack when most of the online transactions are performed.

From the file listing every compromised system we own, this cron job will randomly divide it and generate 3 lists of sizes 30,15 and 5 systems, and execute the following commands:

```
// SYN attack on port 80 of the web server
#tfn -f list1.txt -i <IP of web server> -p 80 -c 5
// TCP/UDP/ICMP attacks on random ports of the mail server
#tfn -f list2.txt -i <IP of mail server> -c 8
// Stop all attacks
```

¹⁸⁷It can be downloaded from: <http://packetstormsecurity.nl/groups/mixter/tfn2k.tgz> and <http://www.brain-pro.de/Seiten/ddos/tfn2k.tgz>

¹⁸⁸Solaris man page

#tfn -f list3.txt -c 0

We hope that by attacking using spoofed addresses, shuffling the attacking roles of our systems, mixing different kinds of attacks, and making some of them stop at times will make it a lot harder for the administrators, and eventually their ISP to protect themselves against the attack.

This attack should be fairly successful despite some of the settings already in place on the router and the firewall to handle flooding.

4.4.2 Denial of service attack countermeasures

It is very hard to protect a network against a DoS attack when services are offered to the public. To mitigate the effect of denial of service attacks, we could:

- Make sure our network devices use any mechanism they have available to handle flooding and prevent system starvation (e.g. TCP intercept, SYNDefender, rate-limiting settings, aggressive behaviors, etc.)
- Tune the servers to be able to handle flooding more efficiently (e.g. lower connection timeout, aggressive behaviors, bigger queues and connection tables, etc.)
- Add load balancers to the design to protect some systems from attacks and add multiple servers for the same service
- Add a secondary/backup link to the outside world. Perhaps with a different ISP. This is useful if we want to generate alerts for our administrators. Dial-out only modems for IDS or log servers are also a good addition
- Enlist the aid of our ISP when we realize we are having a DoS attack
- Upgrade the connection (e.g. T3 uplink)
- Use tools¹⁸⁹ and online scanners¹⁹⁰ to scan the systems for DoS tools

4.5 Compromising an internal system through the perimeter system

Considering the information we gathered in our reconnaissance phase, we have chosen the web server as a good candidate for this attack for a variety of reasons: It seems Microsoft IIS web servers are a mecca of vulnerabilities. In addition, we feel we can benefit greatly from compromising this system. Realistically, we do not expect to have access to internal systems if we compromise the web server. However, we might be able to corrupt their database from there, or at least deface the web site during extremely busy hours. This kind of publicity would probably affect the confidence of current users, who might not be comfortable providing their information, and scare potential customers away if the target develops a reputation for not being secure. The real objective behind it is the potential increase of our market share, and the downfall of our rival companies.

4.5.1 Further research

We compliment any information we acquired through nmap and nessus on the web server, using tools such as whisker¹⁹¹ and CIS¹⁹² to scan for vulnerabilities.

¹⁸⁹<http://www.nipc.gov/warnings/alerts/1999/trinoo.htm>

¹⁹⁰<http://housecall.antivirus.com>

¹⁹¹<http://www.wiretrip.net/rfp>

¹⁹²<http://www.cerberus-infosec.co.uk/cis.shtml>

4.5.2 Microsoft IIS WebDAV remote compromise vulnerability¹⁹³

One of the most recent vulnerabilities is the Microsoft IIS WebDAV Remote Compromise Vulnerability (CAN-2003-0109). It is possible the administrators have been very diligent and have been keeping up with patches and fixes as soon as they are released or implemented workarounds, but even at the time of this writing, it is still being exploited by many. Even IIS lockdown might not properly protect the system against it.

From the X-Force alert:

Internet Security Systems Security Alert

March 17, 2003

Microsoft IIS WebDAV Remote Compromise Vulnerability

Synopsis:

A serious vulnerability exists within the WebDAV component of Microsoft Internet Information Services (IIS) Web server. WebDAV stands for "Web-based Distributed Authoring and Versioning". WebDAV extensions are used by administrators to manage and edit Web content remotely. WebDAV is enabled by default on IIS 5 installations, and no authentication or special privileges are required for remote exploitation.

Impact:

It is possible for remote attackers to run arbitrary code on vulnerable web servers. This vulnerability is currently being exploited in the wild, and X-Force has verified the existence of a functional exploit tool. This vulnerability is in itself very serious, but the existence of robust exploits in the wild dictates that fixes or temporary workarounds should be applied immediately.

It is possible for remote attackers to run arbitrary code on vulnerable web servers. This vulnerability is currently being exploited in the wild, and X-Force has verified the existence of a functional exploit tool. This vulnerability is in itself very serious, but the existence of robust exploits in the wild dictates that fixes or temporary workarounds should be applied immediately.

There are several implementations and "proofs of concept" for our chosen vulnerability. We have chosen one that we know works¹⁹⁴, although if we needed to, we could use a perl script¹⁹⁵ that also takes advantage of the exploit.

After compiling, we simply run the command:

`#!/rs_iis_roman`

If successful, we could run any command we want, which can be very convenient. We do not expect to be able to infiltrate systems on the internal network, since the firewall has probably been configured not to allow the web server any additional access. Moreover, since there is probably a NIDS on the external service network, we will be easily detected by the unexpected traffic.

Whether we deface the web server (e.g. changing fortune prices to be \$1 million, or the typical "/0|_| |-|4V3 |333|\\ 0VV|\\|3|)" message), or do something

¹⁹³<http://www.iss.net/issEn/delivery/xforce/alertdetail.jsp?oid=22029>

¹⁹⁴http://www.rs-labs.com/exploitsntools/rs_iis.c

¹⁹⁵<http://www.security.nnov.ru/files/webdav.pl>

more devious like try to intercept credit card information, corrupt the database, install more backdoors, crash the system, or simply try to overcharge customers will depend on the mood. However, installing or downloading our tools into the web server will probably not be possible or might require us to use port 80 or 443 for this. In addition, if the system is running a tool like Tripwire, any files we change might be easily detected (unless we try to take care of Tripwire first).

Another possible outcome of running the rs_iis program is that it makes IIS crash. We would have to wait for the administrators to restart IIS before we can try again, but crashing IIS is still a good thing for us. If customers cannot access the web server, then they might have to go somewhere else.

The chances that this attack will work really depend on how diligent the administrators have been. Since it is a fairly recent vulnerability, it might succeed.

4.5.3 Microsoft IIS WebDAV remote compromise vulnerability countermeasures

In order to prevent this sort of attack and many attacks on IIS we could:

- Install Microsoft's patch¹⁹⁶
- Configure IIS lockdown to disable WebDav¹⁹⁷
- Restrict the size of the buffer IIS utilizes to process requests by using Microsoft's URL Buffer Size Registry Tool¹⁹⁸
- Install a host-based IDS and file integrity tools on all systems
- Add a reverse proxy
- Install IIS protecting tools like eEye's secure IIS¹⁹⁹
- Upgrade to a real web server on a real operating system

4.6 Conclusion

We would like to end the section by listing a few axioms on security:

No matter how safe or robust a design is, if it is not maintained properly and diligently, it becomes vulnerable over time. This should include, but it is not limited to, the network layer, operating system, and applications or services offered.

- Subscribing to multiple security and vulnerability mailing lists is definitely worth the "spam".
- Monitoring of logs should be done religiously. This can serve as an early warning system and helps us avert an attack.
- No matter how hard we try to keep attackers out, one day somebody will eventually break in. The design should keep this in mind and isolate systems and networks as much as possible to minimize damage.
- Since a system can eventually be compromised, it is important we can monitor network traffic with NIDS, system activities with host-based IDS, and file integrity tools. This can alert us when that day occurs. A contingency plan for such an occasion is a must.

¹⁹⁶<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/ms03-007.asp>

¹⁹⁷<http://support.microsoft.com/default.aspx?scid=kb:en-us:241520>

¹⁹⁸<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/Default.asp>

¹⁹⁹<http://www.eeye.com/html/Products/SecureIIS/index.html>

5. References

Cisco Router Security

http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080120f48.shtml

<http://www.cisco.com/warp/public/707/21.html>

<http://www.sans.org/rr/netdevices/disabling.php>

<http://www.cymru.com/Documents/secure-ios-template.html>

Equipment

Cisco 7204 Series Router

<http://www.cisco.com/en/US/products/hw/routers/ps341/ps346/index.html>

http://www.cisco.com/en/US/products/hw/routers/ps341/products_data_sheet09186a008008872b.html

Cisco 3640 Router

<http://www.cisco.com/en/US/products/hw/routers/ps274/ps278/index.html>

Cisco 3620 Router

<http://www.cisco.com/en/US/products/hw/routers/ps274/ps276/index.html>

Cisco PIX 515E

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet09186a0080091b15.html

Cisco VPN Concentrator 3015

<http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/ps2291/index.html>

http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_data_sheet09186a0080091e4f.html

Cisco FastHub 100 Series

<http://www.cisco.com/en/US/products/hw/hubcont/ps853/index.html>

Cisco 1538 Microhub

<http://www.cisco.com/en/US/products/hw/hubcont/ps209/ps208/index.html>

Sun Microsystems Ultra 60

http://sunsolve.sun.com/handbook_pub/Systems/U60/spec.html

Sun Microsystems Enterprise 220R

http://sunsolve.sun.com/handbook_pub/Systems/E220R/spec.html

Sun Microsystems Netra T1 105

http://sunsolve.sun.com/handbook_pub/Systems/Netra_t1_105/spec.html

Sun Microsystems StorEdge T3+

http://sunsolve.sun.com/handbook_pub/Systems/T3/spec.html

Sun Microsystems StorEdge D1000

http://sunsolve.sun.com/handbook_pub/Systems/D1000/spec.html

Firewalls

Cisco PIX Firewalls

<http://www.cisco.com/go/pix>

Firewall Auditing

<http://www.spitzner.net/audit.html>

Firewall Hardening

http://www.windowsecurity.com/whitepapers/The_Firewall_Hardening_Guide_v01__Checkpoint_Firewall1_Specific_Requirements__Miscellaneous.html

Links

Bugtraq

<http://online.securityfocus.com>

CERT

<http://www.cert.org>

Cisco IOS 12.2 Reference

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122mindx/index.htm>

Cisco IOS NTP vulnerability

<http://www.cisco.com/warp/public/707/NTP-pub.shtml>

Cisco PIX Firewall Command Reference 6.3

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_sw/v_63/cmdref/index.htm

Cisco VPN 3000 Series Concentrator Documentation 3.6

http://www.cisco.com/univercd/cc/td/doc/product/vpn/vpn3000/3_6/index.htm

Check Point FW-1 Syslog Daemon Unfiltered Escape Sequence Vulnerability

<http://www.securityfocus.com/bid/7161>

<http://www.securityfocus.com/bid/7161/exploit>

Check Point IKE alert

<http://www.checkpoint.com/techsupport/alerts/ike.html>

Check Point Syslog alert

<http://www.checkpoint.com/techsupport/alerts/syslog.html>

Check Point VPN-1 IKE Aggressive Mode Forcing Vulnerability

<http://www.securityfocus.com/bid/5920>

<http://archives.neohapsis.com/archives/bugtraq/2002-09/0017.html>

Check Point VPN-1/Firewall-1 Remote Syslog Data Resource Consumption Vulnerability

<http://www.securityfocus.com/bid/7159>

Common Criteria EAL4

<http://commoncriteria.org/docs/EALs.html#EAL4>

CVE

<http://www.cve.mitre.org>

Emerald

<http://www.sdl.sri.com/projects/emerald>

GCFW Practicals

Vivekanand Chudgar's practical (#381)

http://www.giac.org/practical/GCFW/Vivekanand_Chudgar.pdf

Mark Dubinski's practical (#367)

http://www.giac.org/practical/GCFW/Mark_Dubinsky_GCFW.pdf

Geoff Poer's practical (#369)

http://www.giac.org/practical/GCFW/Geoff_Poer_GCFW.pdf

IANA's Reserved Addresses

<http://www.iana.org/assignments/ipv4-address-space>

ICSA Firewall certification

<http://www.icsalabs.com/html/communities/firewalls/index.shtml>

incidents.org top attackers

<http://isc.incidents.org/top10.html>

IPsec certification

<http://www.icsalabs.com/html/communities/ipsec/index.shtml>

ISS's X-Force

<http://xforce.iss.net/index.php>

Microsoft IIS WebDAV Remote Compromise Vulnerability

<http://www.iss.net/issEn/delivery/xforce/alertdetail.jsp?oid=22029>

http://www.rs-labs.com/exploitsntools/rs_iis.c

<http://www.security.nnov.ru/files/webdav.pl>

<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/ms03-007.asp>

<http://support.microsoft.com/default.aspx?scid=kb;en-us:241520>

<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/Default.asp>

NIPC

<http://www.nipc.gov>

RealSecure

http://www.iss.net/products_services/enterprise_protection/rserver/index.php

Public NTP servers

<http://www.eecis.udel.edu/~mills/ntp/servers.html>

SANS Hardening Solaris

http://www.sans.org/projects/hard_solaris.htm

SANS Security Policies

<http://www.sans.org/resources/policies>

SANS top 20 vulnerabilities

<http://www.sans.org/top20>

SecuriTeam

<http://www.securiteam.com>

Sophos

<http://www.sophos.com/products/software/antivirus/savunix.html>

Sun Blueprints

“Customizing the JumpStart Boot Image for Recovery”

<http://www.sun.com/solutions/blueprints/0301/CustomBoot.pdf>

“Modeling Sun Cluster Availability”

<http://www.sun.com/solutions/blueprints/1202/817-0905.pdf>

“Rapid Recovery Techniques for the Solaris Operating Environment”

<http://www.sun.com/solutions/blueprints/0100/soe.pdf>

“Securing the Sun Cluster 3.x Software”

<http://www.sun.com/solutions/blueprints/0203/817-1079.pdf>

“The Solaris Security Toolkit - Installation, Configuration, and Usage Guide: Updated for Toolkit version 0.3”

http://www.sun.com/solutions/blueprints/0601/jass_conf_install-v03.pdf

“Solaris Operating Environment Network Settings for Security: Updated for Solaris 8 Operating Environment”

<http://www.sun.com/solutions/blueprints/1200/network-updt1.pdf>

“Solaris Operating Environment Security: Updated for Solaris 9 Operating Environment”

<http://www.sun.com/solutions/blueprints/1202/816-5242.pdf>

“Sun StorEdge T3 Array: Installation, Configuration and Monitoring Best Practices”

<http://www.sun.com/solutions/blueprints/1001/t3bp.pdf>

“Using NTP to Control and Synchronize System Clocks – Part I: Introduction to NTP”

<http://www.sun.com/solutions/blueprints/0701/NTP.pdf>

“Using NTP to Control and Synchronize System Clocks – Part 2: Basic NTP Administration and Architecture”

<http://www.sun.com/solutions/blueprints/0801/NTPpt2.pdf>

“Using NTP to Control and Synchronize System Clocks – Part 3”

<http://www.sun.com/solutions/blueprints/0901/NTPpt3.pdf>

Sun Freeware

<http://sunfreeware.com>

SunScreen Lite

<http://www.sun.com/software/securenet/lite/index.html>

Sun System Handbook

http://sunsolve.sun.com/handbook_pub/Systems

Verisign

<http://www.verisign.com/products/onsite/ssl/faq.html>

Tools & Software

Apache 2.0.44

http://httpd.apache.org/docs-2.0/new_features_2_0.html

at

Solaris man page (1)

BIND

<http://www.isc.org/products/BIND/bind9.html>

<http://www.nominum.com/content/documents/bind9arm.pdf>

CIS

<http://www.cerberus-infosec.co.uk/cis.shtml>

Cisco VPN Client v3.7

<http://www.cisco.com/en/US/products/sw/secursw/ps2308/ps3875/index.html>

cron

Solaris man page (1M)

<http://www.linuxbasis.com/security.html?/security1.html>

eEye's SecureIIS

<http://www.eeye.com/html/Products/SecureIIS/index.html>

GnuPG

<http://www.gnupg.org>

MailScanner

<http://www.sng.ecs.soton.ac.uk/mailscanner>

Nessus

<http://www.nessus.org>

NIPC DoS Scanner

<http://www.nipc.gov/warnings/alerts/1999/trinoo.htm>

Nmap

<http://www.insecure.org/nmap>

Norton Internet Security 2003

http://www.symantec.com/sabu/nis/nis_pe/features.html

Nslookup

<http://cc-www.uia.ac.be/ds/nslookup.html>

<http://gday.bloke.com/cgi-bin/nslookup>

NTP

<http://www.ntp.org>

<http://www.ntp.org/documentation.html>

<http://www.eecis.udel.edu/~mills/ntp/html/index.html>

PGP

<http://www.pgp.com>

ping

Solaris man page (1M)

Sendmail 8.12.8

<http://www.sendmail.org/8.12.8.html>

<http://www.sendmail.org/vendor/sun>

snoop

Solaris man page (1M)

Snort 1.9.1

<http://www.snort.org>

<http://www.snort.org/docs>

Sun Cluster 3.0

<http://www.sun.com/software/cluster>

Sun Security Toolkit

<http://www.sun.com/blueprints/tools>

Sun Solaris JumpStart

<http://www.sun.com/software/solaris/8/ds/ds-solwebstart/#6>

Sun Trusted Solaris 8 4/01

<http://www.sun.com/software/solaris/trustedsolaris/index.html>

<http://www.commoncriteria.org/ccc/epl/productType/epldetail.jsp?id=110>

Swatch

<http://swatch.sourceforge.net>

TcpWrappers

<http://sunfreeware.com/programlistsparc9.html#tcpwrappers>

Tcpdump

http://www.tcpdump.org/tcpdump_man.html

TFN2K

<http://www.securiteam.com/securitynews/5YP0G000FS.html>

<http://packetstormsecurity.nl/groups/mixter/tfn2k.tgz>

<http://www.brain-pro.de/Seiten/ddos/tfn2k.tgz>

Titan

<http://www.fish.com/titan>

traceroute

Solaris man page (1M)

<http://cache2.online.ee:81/cgi-bin/nph-traceroute>

Trend Micro's HouseCall

<http://housecall.trendmicro.com>

Tripwire

<http://www.tripwire.org>

whisker

<http://www.wiretrip.net/rfp>

Whois

<http://www.whois.net>

<http://www-whois.internic.net/cgi/whois>

xntpd

Solaris man page (1M)

YASSP

<http://www.yassp.org>

RFC References

RFC 1305 (NTP)

<http://www.ietf.org/rfc/rfc1305.txt>

<http://www.eecis.udel.edu/~mills/database/rfc/rfc1305>

RFC 1918 (Private Addresses)

<http://www.ietf.org/rfc/rfc1918.txt>

<http://www.isi.edu/in-notes/rfc1918.txt>

RFC 2406 (ESP)

<http://www.ietf.org/rfc/rfc2406.txt>

RFC 2408 (ISAKMP)

<http://www.ietf.org/rfc/rfc2408.txt>