



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.

GCFW Practical Assignment

Version 1.8, Revised September 10, 2002

Fortune Favors the Prepared Mind¹

Richard Turk

January 2003

¹ I've seen this quote attributed to Louis Pasteur, but have been unable to verify it.

Table of Contents

Part I	Introduction and Overview of Practical Submission
Part II	Data Collection
Part III	Assignment 1 – Security Architecture
Part IV	Assignment 2 – Security Policy and Tutorial
Part V	Assignment 3 – Verify the Firewall Policy
Part VI	Assignment 4 – Design Under Fire
	Appendix A – Defense in Depth and Host Security
	Appendix B – Stealth NIDS and Logging (Logging for the Paranoid)
	Appendix C – References

© SANS Institute 2003, Author retains full rights.

Part I – Introduction and Overview of Practical Submission

Abstract

This paper is submitted in fulfillment of the GIAC Certified Firewall Analyst Practical Assignment version 1.8. It begins with careful data collection, and then describes a security system including a Cisco router, Linux bridging Netfilter firewalls, and a Cisco VPN appliance. It includes testing procedures that use commonly available network tools combined with hands-on, packet-level log analysis, rather than canned or commercial network testing suites. The fourth assignment attempts to find weaknesses in the Practical design submitted by James Giesecke.

Goals

Certainly my first goal of this exercise is to submit an assignment that will be judged acceptable for the GCFW certification. My second goal is to develop a set of skills, methods and processes to improve my effectiveness as a network administrator, security analyst, consultant, and problem solver. Although personal skills and abilities will always be important elements of problem solving, other parts of the process can be formalized; that is, developed into recipes that improve the chance of producing an acceptable outcome with minimal wasted time.

The Approach

Any successful consultation or analysis must begin with good data collection, which I have divided into five topics; explanations are given below. For three of these I have also developed a worksheet to facilitate complete, accurate and consistent data collection.

Data Collection

- Consulting Job Overview Worksheet: This is intended to elicit information about the “Big Picture”, and should guide the investigation and presentation of solutions.
- Stakeholder Worksheets: A “stakeholder” is any entity or person with an interest in the success of the project, and it’s important to know in advance what each stakeholder expects. That’s not to say that all stakeholders can be fully satisfied, but satisfaction is more likely if we know what is expected or needed. Techies often overlook the human element, but humans are the final arbiters of a successful job.
- Current Network Diagram: An accurate diagram of the current network environment.
- Host Worksheets: Hardware, software and configuration information for each device in the current network structure.

- Security Policy: This is a formal security policy, if the client has one. If not, a “Technical Security Policy” can be developed through discussions with stakeholders. The security policy motivates the large-scale goals of the project.

Part II Data Collection

A Brief Note on Data Collection: This part of the submission is just the collected information and forms. The data collection sheets have page borders to emphasize that they would normally be filled out as standardized forms. There will be no additional text.

© SANS Institute 2003, Author retains full rights

Consulting Job Overview Worksheet

Name of client: GIAC Enterprises

Nature of client's business: Online distribution and resale of fortune cookie sayings

Executive contact: Company President

Technical contact(s): IT Manager

Who will evaluate our work? President and IT Manager.

What are their expectations? President expects reassurance that the business can be made secure with only minor efforts and financial cost. He has made it clear that the business relationship will not extend beyond the consultancy contract: in particular, that there will be no ongoing maintenance contract. The IT Manager expects our analysis to support his claims that significant additional computing and human resources are needed to provide full services and maintain them adequately.

Type of work requested: Improve network security and remote access services.

Design and planning: Re-design network to improve security and enhance remote access for partners and traveling sales staff.

Procurement: Recommendations only; no direct procurement.

Implementation / installation / deployment: Work with local IT on all post-design work.

Training and documentation: Train local IT to install and maintain all new systems, designs. Provide documentation of installation and maintenance procedures.

© SANS Institute 2003. All rights reserved.

Stakeholders Summary: Stakeholders are people or entities with an interest in the success of the project. The stakeholder's *expressed need* does not necessarily imply a job requirement.

Stakeholder	Expressed Needs
President of the company	Keep IT resources secure without spending a lot of money; improve functionality for sales force, if possible
Suppliers	Submit new sayings on-line
Partners	Acquire English sayings on-line
Mobile sales force	(1) Gain access to marketing and sales resources from anywhere; (2) submit orders on-line
Teleworkers	Perform all work functions from off-site
Customers	(1) Place orders; (2) pay for orders (3) check delivery status
On-site employees	Access to all work-related resources; unfettered access to Internet
IT Manager	A secure network implementation that satisfies user needs without imposing excessive risks or support burdens on IT staff

Number and character of business locations: A single corporate headquarters with approximately 70 staff. In addition, there is a worldwide traveling sales force of 50, who work from their homes and hotel rooms. Staff at headquarters consists of clerical, accounting, sales support and purchasing, in addition to senior management.

Size and character of internal IT staff: IT Manager and three full-time technical staff. The IT department uses mostly Microsoft products, with one Linux box in non-production use; also several of the staff use Linux at home, and would be comfortable incorporating it into their environment. Brief interviews with staff indicate that they are competent, but feel over-extended.

What assistance can we expect from internal IT staff: IT staff seem excited by the prospect of Executive support for IT enhancements. Also, the President has clearly stated that no additional outsourcing is planned. Consequently, IT employees do not appear to feel threatened, and will probably give full assistance with the project.

Number and character of client and host computers:

- About 70 desktop computers running Windows 2000 Professional.
- About 10 laptops at headquarters running Windows 2000 Professional
- About 50 laptops with the sales force running Windows 2000 Professional

- One Windows 2000 Server for file, print, MS Exchange and backup services
- One Windows 2000 Server for SQL Server, Internet Information Services, DNS and DHCP. This machine also hosts Outlook Web Server for Exchange.
- One Linux box in IT department: Currently used by IT staff as “poor geek’s VPN” using sshd, and as investigation/test environment for new service options.

List and character of current Internet-accessible services:

- **SQL Server** is open to permit access to the sales databases
- **IIS** is open for the public web site, GE’s proprietary web application, and Outlook Web interface for e-mail.

Overview of current network design:

See attached diagram of current structure.

Network Structure: A simple, flat network using valid Internet addresses for all hosts. The border router performs basic filtering, including most Microsoft services.

Internet Service: A T1 to a commercial ISP.

External Access to Data: Customers, suppliers and partners perform all transactions via GE’s proprietary web application. Traveling sales staff use the web application for some parts of their job, but also communicate directly with the SQL Server.

Management and Maintenance: IT staff spend a lot of time keeping the clients and servers patched.

Remote Access to Network: There is no modem pool or VPN: teleworkers use commercial ISP’s. IT staff tunnel in using sshd on the Linux box, for some management functions.

Financial constraints on proposals: The total cost of consultancy services plus any needed enhancements may not exceed \$50,000. The combined limitation was imposed at the insistence of the president of the company.

Internal political and discretionary concerns: The president and IT manager generally get along well, but have drastically different perspectives. The president’s background is strictly in sales and sales management: he thinks the IT manager is overly cautious about security and too conservative about deploying new services for

customers and sales staff. The IT manager, on the other hand, is often frustrated that he is unable to communicate the risks and costs to the president. It will be important to express our recommendations in a way that does not offend either.

Are there any existing policy documents?

Technical: There are no policy documents, only procedure documents.

HR: General business confidentiality agreements, nothing specific to IT.

Known security issues or vulnerabilities:

No previously known successful attacks. No known enemies: the fortune business is not known to play dirty.

Additional considerations regarding the environment and job:

GIAC Enterprises is a small, profitable, but low-margin company. Their executive structure and financial condition offer several challenges:

1. Any request for expenditures needs to be well justified, and expressed in language accessible to the president.
2. The president will expect no decrease in services to sales staff and customers; it would be better if we can enhance services; this may give us more credibility and executive support for security enhancements, and any complexities they introduce.
3. The final design must be acceptable to the IT manager. He must be assured that he and his staff can maintain the new structure, and that it will not impede the delivery of services.

© SANS Institute

Stakeholder Worksheet

Stakeholder Name or Category: President of company

Service / Access Needs:

On-site: file server, e-mail, web server, database; access to Internet

Remote: file server, e-mail, web server, database

Current method of remote access: Has personal DSL account.

Web / SQL Server: SSL to web server; does not use SQL Server

File Server: Does not use file or print services remotely

E-mail: Outlook Web Server for Exchange

Technical proficiency or comfort: Not comfortable with complex computer skills or operations. Changes to routine operations will require explicit instructions.

Narrative summary: Any significant changes may put off the president, making him less likely to accept other elements of the design.

© SANS Institute 2003, Author retains full rights.

Stakeholder Worksheet

Stakeholder Name or Category: Suppliers

Service / Access Requirements:

On-site: None

Remote: Submit new sayings

Current method of access: GE's web application

Technical proficiency or comfort: Unpredictable. Must assume that a simple interface is required.

Narrative summary: Any new procedures must not impede the suppliers.

© SANS Institute 2003, Author retains full rights.

Stakeholder Worksheet

Stakeholder Name or Category: Partners

Service / Access Needs

On-site: None.

Remote: Purchase English sayings on-line via web application.

Current method of remote access:

Web / SQL Server: SSL to web server; no remote access to SQL Server.

File Server: None

Technical proficiency or comfort: Unpredictable. Must assume that a simple interface is required.

Narrative summary: New procedures must not introduce any greater complexity. Explicit instructions will be required.

© SANS Institute 2003, Author retains full rights.

Stakeholder Worksheet

Stakeholder Name or Category: Mobile Sales Force

Service / Access Needs:

On-site: Same as on-site staff.

Remote: (1) access to marketing and sales resources from anywhere; (2) submit orders on-line

Current method of remote access: Dial-up accounts; hotel-supplied broadband

Web / SQL Server: SSL to web server; direct access to SQL Server via custom application.

File Server: No direct access to file or print services

E-mail: Outlook Web Server for Exchange

Technical proficiency or comfort: Members of mobile sales force are generally uncomfortable with computers. However, they are highly motivated to increase sales, and will learn whatever is necessary to maintain or increase their selling power. Changes to routine will require explicit instructions.

Narrative summary: Mobile sales force must be able to accomplish required job operations without unreasonable additional complexity. The company president is very sensitive to sales issues, so dissatisfaction among sales staff will have a negative effect on our credibility.

© SANS Institute 2003. All rights reserved. This document is for personal use only. No part of this document may be reproduced without written permission from the SANS Institute.

Stakeholder Worksheet

Stakeholder Name or Category: On-site employees

Service / Access Needs:

On-site: file server, e-mail, web server, database; access to Internet

Remote: file server, e-mail, web server, database

Current method of remote access: Personal accounts with commercial ISP's

Web / SQL Server: SSL to web server; no remote access to SQL Server.

File Server: No direct access to file or print services.

E-mail: Outlook Web Server for Exchange

Technical proficiency or comfort: Range of skills from enthusiastic power users to intransigent old-timers. Changes to routine may require explicit instructions.

Narrative summary: Employees must be able to accomplish required job operations without unreasonable additional complexity. Other than that, they have no direct influence on the design or implementation process.

© SANS Institute 2003, Author retains full rights.

Stakeholder Worksheet

Stakeholder Name or Category: Customers

Service / Access Needs:

On-site: None

Remote: Purchase bulk fortunes online

Current method of remote access: Custom web application over SSL.

Technical proficiency or comfort: Variable and unpredictable. Must assume that simplicity is required.

Narrative summary: Any new procedures must not introduce any greater complexity. Explicit instructions should be unnecessary.

© SANS Institute 2003, Author retains full rights.

Stakeholder Worksheet

Stakeholder Name or Category: IT Manager

Service / Access Needs:

On-site: A secure network implementation that satisfies user needs without imposing excessive support burdens on IT staff.

Remote: Full access to all resources

Current method of remote access: Has personal DSL account.

Web / SQL Server: SSL to web server; no remote access to SQL Server.

File Server: No direct access to file or print services.

E-mail: Outlook Web Server for Exchange

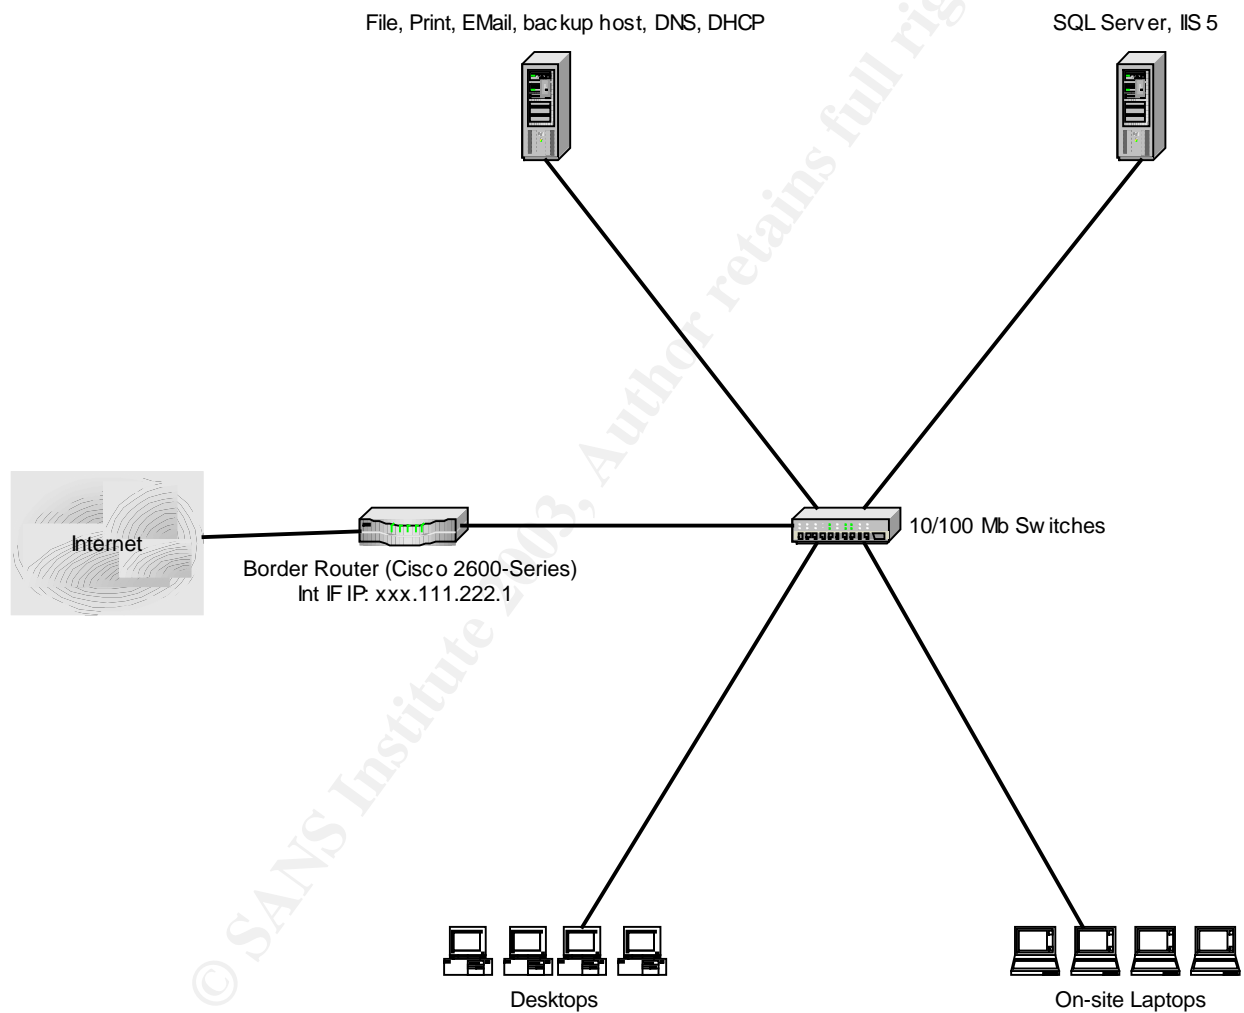
Linux Host: sshd

Technical proficiency or comfort: N/A

Narrative summary: IT Manager must be satisfied that (1) the proposal will satisfy the needs of his clients; (2) his IT team can maintain and extend the proposed solution. We assume that the IT Manager will represent the needs of his staff.

© SANS Institute 2003, Author retains full rights.

GIAC Enterprises Current Network Structure



Host Data Collection Worksheet

Host description: Desktop computers

Number of host(s): 70

Purpose of host (s): Office applications, e-mail, Internet access software (e.g. web browsers)

Internally accessible services: None

Internet-accessible services: None

IP Configuration:

Address: Assigned by DHCP

Netmask: 255.255.255.0

OS, version, patch level:

Windows 2000 Professional; SP3 + security patches

How heavily used is this host?

Typical: Low to moderate

Peak times: Moderate

Applications, version, patch level:

(Note: not all desktops have all software installed: varies by department and job function)

- Microsoft Office 2000 Professional, SP1 (including MS Outlook)
- Internet Explorer 6.0 + security patches
- Small business accounting client (accounting department only)
- Vendor management client (works with SQL Server back end)
- Sales management client (works with SQL Server back end)

© SANS Institute 2003, Author retains full rights.

Host Data Collection Worksheet

Host description: Portable computers

Number of host(s): 60

Purpose of host (s): Office applications, e-mail, sales software, Internet access software (e.g. web browsers)

Internally accessible services: None

Internet-accessible services: None

IP Configuration:

Address: Assigned by DHCP

Netmask: 255.255.255.0

OS, version, patch level:

Windows 2000 Professional; SP3 + security patches

How heavily used is this host?

Typical: Low to moderate

Peak times: Moderate

Applications, version, patch level:

(Note: not all desktops have all software installed: varies by department and job function)

- Microsoft Office 2000 Professional, SP1 (including MS Outlook)
- Internet Explorer 6.0 + security patches
- Custom sales software (Visual Basic program which works standalone or while connected to the SQL Server back end)

© SANS Institute 2003. Author retains full rights.

Host Data Collection Worksheet

Host description: File Server

Number of host(s): 1

Purpose of host (s): File, Print, MS Exchange, backup host, DNS, DHCP

Internally accessible services: File, Print, MS Exchange, backup host, DNS, DHCP

Internet-accessible services: DNS, SMTP

IP Configuration:

Address: xxx.111.222.32

Netmask: 255.255.255.0

OS, version, patch level:

Windows 2000 Professional; SP3 + security patches

How heavily used is this host?

Typical: Moderate

Peak times: Moderate-high (backups)

Applications, version, patch level:

- IIS 5.0 + security patches
- MS Exchange 2000 + Service Pack 3
- Veritas Backup Exec 8.6

© SANS Institute 2003, Author retains full rights.

Host Data Collection Worksheet

Host description: Web / E-Commerce Server

Number of host(s): 1

Purpose of host (s): SQL Server, IIS, Outlook Web Server for Exchange

Internally accessible services: SQL Server, IIS, Outlook Web Server for Exchange

Internet-accessible services: SQL Server, IIS, Outlook Web Server for Exchange

IP Configuration:

Address: xxx.111.222.29

Netmask: 255.255.255.0

OS, version, patch level:

Windows 2000 Professional; SP3 + security patches

How heavily used is this host?

Typical: Moderate

Peak times: Moderate-high (backups)

Applications, version, patch level:

- IIS 5.0 + security patches
- SQL Server 2000 + security patches
- Outlook Web Server for Exchange 2000 + patches

© SANS Institute 2003, Author retains full rights.

Technical Security Policy

About this Document

This document will guide the analysis and decision processes for this consulting job. It will help us select the right hardware, software and methods to provide the level of service expected by client. Normally, we use the client's internal Security Policy to help us develop this Technical Security Policy; however the client had not yet developed a comprehensive Security Policy, so we have had to rely on interviews with senior management and the IT Manager for the necessary guidance.

Computer Security Is a Risk Management Problem

Absolute network security can never be achieved since any legitimate business operation leaves open the possibility of abuse. The goal of this Technical Security Policy is to define the business needs of the company and to prescribe a set of policy statements that will greatly reduce risks in the course of business.

Business Needs ²

Assets to be protected:

1. Data: client, financial and product data are essential to the continuation of the business. The loss of data will result in the loss of profit, and may result in failure of the enterprise.
2. Computing Resources: computing facilities are purchased and maintained at significant cost; if these resources are expropriated for unauthorized uses, they will be unavailable for company business. This may result in a loss of productivity and profit.
3. Reputation: GIAC Enterprises relies on its reputation with customers, suppliers and partners for continued operation; if their trust is lost, they may choose to do business with a competing company.

Threats to business assets:

1. Intrusion: Unauthorized access to company resources may result in compromise to all key assets.
2. Data Theft: Stolen intellectual product may be given to competitors; stolen financial information may pose a threat to the customers' finances, in addition to any threat to GE.
3. Data Alteration: Altered intellectual property must be recalled and replaced; damage to reputation may be irreparable. Altered financial records may result in legal difficulties.

² Zwicky, Elizabeth D., et al, Building Internet Firewalls, Second Edition. O'Reilly, 2000. 4-7.

Primary Risks

1. Intentional or accidental disclosure of credentials
2. Open access to host services
3. Credit card data theft

Essential Services:

1. Sell product to clients via GE's web application.
2. Accept new products from suppliers via GE's web application.
3. Provide products to partners via GE's web application.
4. Support business operations on the internal network.
5. Support mobile sales force and teleworkers.

Other Constraints (e.g. financial)

1. We have been given a strict budget for this job: the combined cost of hardware, software and our services must not exceed \$50,000.
2. We have been told not to propose more change than IT can reasonably absorb since they are small and heavily loaded already.
3. The President of GE has told us explicitly that this is strictly a design and rollout project: there will be no contract for ongoing services to support the changes.

Elements of Security:

Our consultancy divides network security into three parts: defense, mitigation and maintenance. There is significant overlap among these parts, but this division provides a useful means of verifying that a security implementation has addressed all major concerns.

Defense: Many security professionals view defense as the totality of “security”, and it is certainly essential. Defense has two components:

- Placement of barriers to unauthorized behavior without imposing significant impediments to business processes
- Monitoring the network and hosts for breaches

Clearly, the types and numbers of barriers will depend on the internal structure of the environment: the concept of *Defense In-Depth* asserts that no single barrier can prevent all potential attacks on a heterogeneous network. Instead, multiple layers of security are selectively placed at appropriate locations to address

different aspects of the security problem. Common components of *defense in-depth* are³

- 1) Router configuration
 - i) Access lists to permit only authorized traffic (ingress and egress filtering)
 - ii) Hardening of the router itself
- 2) Firewall implementation
 - i) A rule base that supports the organization's Security Policy and business operations.
 - ii) A rule base that includes robust logging of illegal or questionable activity
 - iii) Hardening of the firewall itself
- 3) Intrusion Detection Systems to detect and log inappropriate or abnormal traffic flow.
- 4) Host-based security
 - i) OS access control lists
 - ii) OS configuration to decrease vulnerability (e.g. disable unneeded services)
 - iii) OS software security updates
 - iv) Application configuration to decrease vulnerability
 - v) Application software security updates
 - vi) Host-based firewall or IDS
- 5) Human Resource policies to sanction inappropriate use of resources, and sharing of credentials among employees, partners and others users of the internal resources.

Mitigation⁴: If a security compromise occurs the organization will usually have four priorities:

1. Discover the compromise as quickly as possible
2. Stop the attack or exploitation (e.g. shut down the host, block the attacker)
3. Restore the compromised service to full functionality after appropriate adjustment of the defense strategy
4. Analyze records of the compromise to determine how the attack occurred, and who committed it.

³ SANS Institute, *Firewalls, Perimeter Protection and VPN's, Day 3*. SANS Institute. 2002

⁴ Some of these ideas are taken from ⁴ Zwicky, Elizabeth D., et al, *Building Internet Firewalls, Second Edition*. O'Reilly, 2000.

Successful implementation of these tasks requires preparation, including

- IDS installation, configuration and monitoring
- Routine log analysis
- An incident response plan
- Filesystem integrity checking
- A comprehensive data backup plan and disaster recovery plan

Maintenance: Often overlooked in security plans, maintenance is an essential component of any security plan. No computing environment is static: new exploits are discovered, new software is installed, and business processes evolve. For example, wireless access points now provide a means to infiltrate a LAN from outside the building walls; as wireless technology evolves, it will pose increasing challenges to security. It's imperative that the security policy and implementation evolve as well. Maintenance includes:

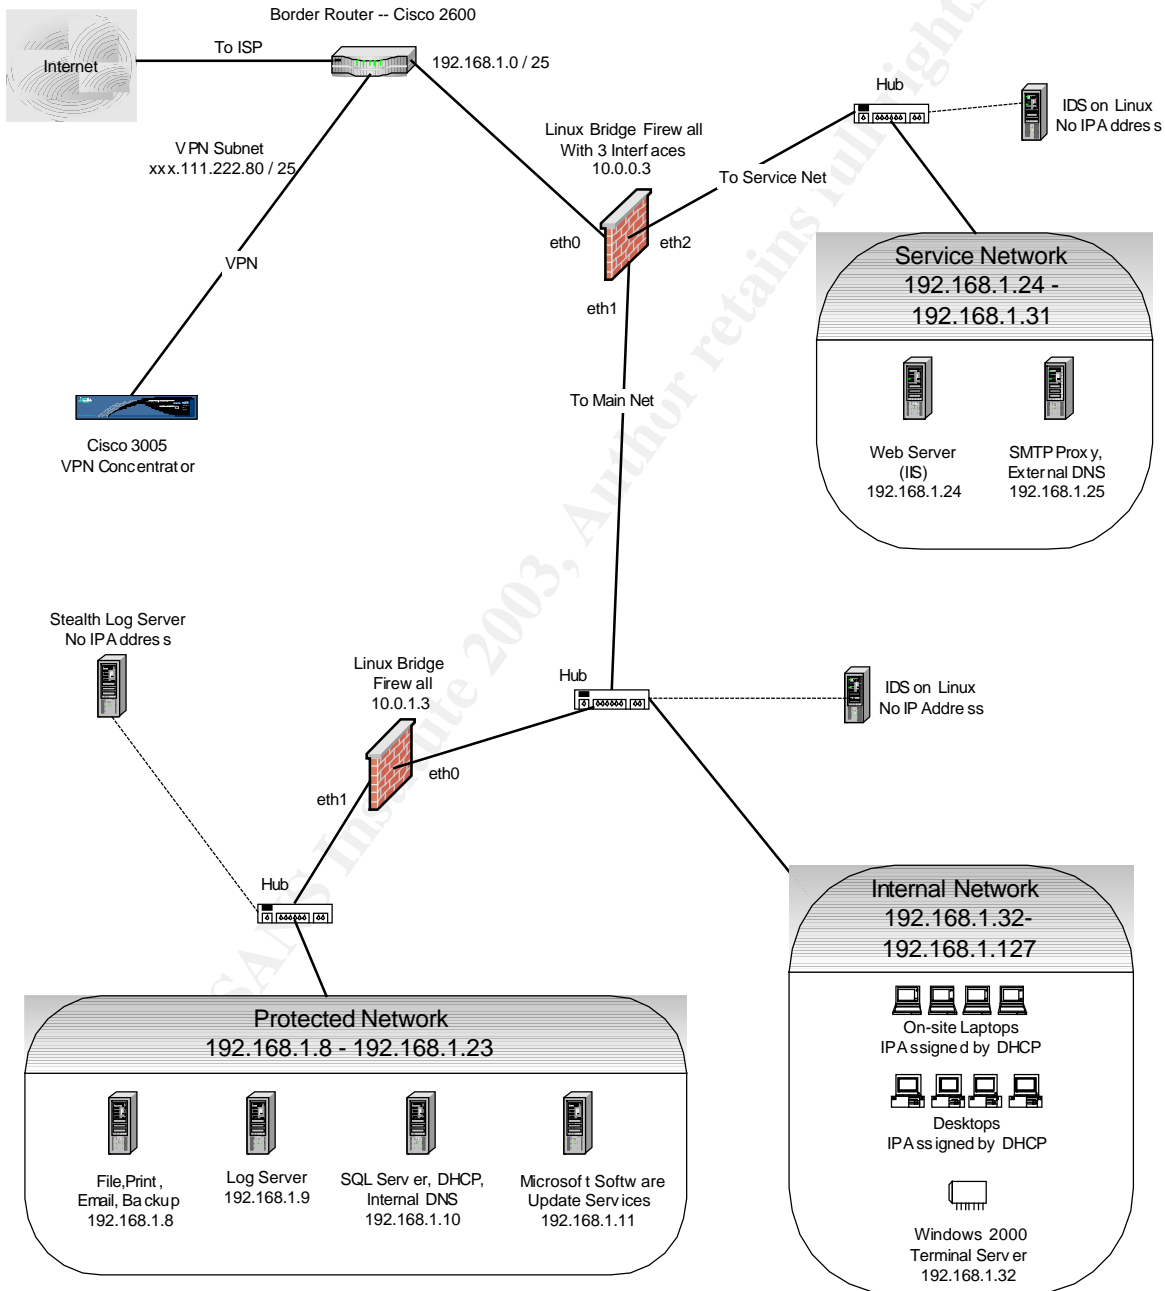
1. Reviewing firewall, IDS and host security logs for unusual activity
2. Applying application and OS software security patches
3. Periodic review of the Security Policy to assure that it addresses current business processes and security issues
4. Periodic review of the configuration of each security component, e.g. routers, firewalls, hosts, to verify that their configuration is appropriate for the current Security Policy
5. Remaining aware of current technologies and their potential as tools for intrusion

Policy Statements for Our Proposed Security and Service Enhancements

1. No proposed security enhancement may significantly impair current business practices. Changes requiring small adjustment to procedure are acceptable. Ideally, the proposal should increase service levels.
2. The new plan must be maintainable with current staffing levels: proposed changes must not impose an unreasonable new burden on the IT department.
3. To stay within budget, we will use free or open source software and commodity hardware where possible.
4. We will minimize IT burden and cost to the extent possible by re-using existing components; we prefer that each host or server continue all or part of its current role.
5. Where possible, we will design solutions based on Cisco, Microsoft Windows and Linux, as these are the products most familiar to the IT staff.
6. The security proposal must address the following areas:
 - a. Vulnerable services
 - b. Poor access control
 - c. Viruses

Part III Assignment 1 – Security Architecture

GIAC Enterprises Proposed Network Design Net xxx.111.222.0/24



Proposed Addressing Scheme

Border Router External Interface: No change: GE does not plan to request a change in IP allocation from their ISP.

Border Router Internal Interface (Ethernet 0): To provide an additional layer of safety, we will implement NAT on the internal interface. Some hosts have services that must be accessible from the Internet or the VPN network; for those, we will configure static NAT assignments. The remaining hosts (mostly clients) will share a pool of outside addresses using Port Address Translation.

Border Router VPN Interface (Ethernet 1): The Cisco 3005 can implement IPSec through a NAT using NAT Traversal (NAT-T); however, NAT-T is still an IETF Working Group Draft⁵, and we prefer to use established standards rather than drafts to ensure long-term stability of the implementation. Instead, the VPN subnet will be allocated xxx.111.222.80/25.

Security and Service Component Evaluations

Filtering Router

Current: Cisco 2600

Analysis: The 2600 is fully adequate as a simple filtering router at the border, but does not have sufficient power for a full firewall rule set. A new Cisco router with sufficient power would be prohibitively expensive, and we can satisfy the firewall problem by other means. We will review the router policy and recommend changes, including NAT.

Change: No hardware change needed

Location: There is no other place for a border router; NAT and basic filtering are best done at the border.

Firewall(s)

Current: None

Analysis: There is a clear benefit to be gained by separating the Internet-facing services from the other internal hosts using a firewall. The web server and an SMTP proxy must communicate with the T1 Internet connection and the internal network. Although normal business operations will not generate a huge amount

⁵Kivinen, T., et. al. "Negotiation of NAT-Traversal in the IKE". 24 June 2002.
[URL: http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ipsec-nat-t-ike-03.txt](http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ipsec-nat-t-ike-03.txt) (9 January 2003).

of traffic through the firewall, it will also be necessary to run backups of web and SMTP proxy hosts; unless a dedicated backup solution is placed in the Internet-facing Service Network, the backup traffic will also pass through the firewall.

The president of the company does not permit total lockdown of desktop and laptop computers, so the clients are susceptible to infection by viruses and other malicious code. Therefore, we will also recommend a second firewall to separate the client network from the enterprise-critical services (“Protected Network”).

Option 1: Commercial firewall appliance from Cisco, e.g. PIX506E or PIX515E.

Analysis of Option 1: The PIX506E is reasonably priced at about \$1300 + annual maintenance; however, it suffers a fatal deficiency: it has fixed configuration with two Ethernet interfaces. Since we would like to split off a separate service network, this would not serve our purposes. On the other hand, the PIX515E has possibilities: it has three Ethernet interfaces and sufficient power to handle traffic generated by host backups even with a substantial rule base installed. The cost of the PIX515E is about \$2600 + \$450 per year support.

Option 2: Linux-based firewall with subnetting at the firewall

Analysis of Option 2: The current generation of Linux firewall code, Netfilter, is highly regarded in the industry. Furthermore, commodity hardware is fast and cheap, making this a very cost effective solution, with price estimated at less than \$1,000 for a low-end Dell server.

Option 3: Linux-based firewall using the Layer2 bridge code⁶

Analysis of Option 3: This option offers the same benefits of Option 2, but includes a few additional features:

- A bridge firewall requires no subnetting or other reconfiguration of the network: it's a “drop-in” device.
- Because it works at Layer 2 rather than Layer 3, it's invisible to IP traffic; a potential attacker will not detect its presence, and so will not know to attack it.
- Since the bridge has no IP address⁷, conventional IP-based tools cannot attack it; a successful attack would require a bridge-specific attack tool.

Firewall Selected: For simplicity and cost, we choose the Linux bridge firewall.

⁶ See the Bridge tutorial.

⁷ This is not strictly true: the bridge will not have an IP address involved in routing, but it will have an IP address for management purposes only, which will be accessible only from the interior of the network. Furthermore, the bridge's IP address can be arbitrary – it need not conform to the same addressing scheme used for the remainder of the network.

Location: The first firewall should be just inside the router to provide protection to the entire network. If a second firewall is installed, it should protect the most valuable assets, the servers that contain GE's critical business data.

Cost: About \$1,000 for a Dell PowerEdge 600SC running Red Hat Linux 8.0.

VPN: Remote Access for Employees

Current: None. Remote access is not controlled.

Options for Change

Option 1: Cisco VPN Concentrator (3000-series)

Analysis of Option 1: This series of concentrators, obtained through an acquisition of Altiga, have several desirable features: (1) they are fully self-contained, providing authentication, IPSec management and policy support; (2) they have an easy-to-use web interface; (3) Configuration is compatible across the product line, making this a scalable solution. Furthermore, the VPN Concentrator comes with client software that is easy to configure for automatic VPN initiation. This combination of features should permit deployment and maintenance without significant addition burdens on the IT staff. The cost of a low-end unit will be about \$2500 + annual maintenance.

Option 2: Cisco PIX Firewall with VPN Support

Analysis of Option 2: Although the PIX firewalls do support VPN's, their support is more geared toward WAN-to-WAN VPN's, rather than client remote access. Also, the PIX's web interface, the PIX Device Manager, is not as easy to use as the VPN 3000 Concentrator products. If GIAC Enterprises already had (or planned to purchase) a PIX firewall, it might be worthwhile to expend the additional effort needed to implement client VPN access. Otherwise, this is not a good option.

Option 3: Free S/WAN (Open Source)

Analysis of Option 3: The main advantages of Free S/WAN are related to its availability as a free, Open Source product. It can be run under Linux on commodity hardware. There are two chief detractors: (1) it lacks an easy web interface for configuration, and, since GE's IT staff are not familiar with this product, additional ramp-up time and effort would be required; (2) it lacks the Cisco's easy client configuration, instead relying on the existing capabilities of the given client (in this case, Windows 2000). Cost estimate: about \$1,000 for a low-end Dell server.

Option 4: Use the Cisco 2600 router as the VPN.

Analysis of Option 4: In its current configuration, the router doesn't have enough processing power to handle the VPN duties in addition to its regular duties. An add-in module can be purchased for about \$1600; this would give it the necessary power to handle the VPN duties. Nevertheless, from an administrative perspective, it would be easier to have the router perform routing, and leave VPN to another device.

VPN Selected: Cisco VPN 3005 Concentrator. The ease of setup and configuration make up for the cost increase over the FreeS/WAN solution. The PIX is not even considered a serious contender here. We will, however, recommend to the IT Manager that he begin investigation of FreeS/WAN as an option for future VPN needs, should they arise.

Location: Since NAT makes IPSec more complicated, we suggest that the VPN be placed off a separate interface on the router. Once decrypted, the incoming traffic can be NAT'ed, and can pass through the standard filter set on the router. The diagram indicates just one of the Cisco 3005's interface in use: although not widely done, the VPN 3000 series can be operated in this manner through special configuration of the interfaces⁸. The advantage of this method is that the router can screen both inbound and outbound traffic to the VPN. Fortunately, the router has enough processing power to accommodate this.

Hosts

Host #1 (File Server)

Current: Windows 2000 Server (File, Print, E-Mail, Backup, Internal DNS, DHCP)

Change: File, print, e-mail, and backup services cannot remain exposed to the Internet. This server will be moved into the trusted zone, and DNS will be split: internal queries will continue to this host. Another host should be placed in the service network with zone database containing data only for Internet-accessible resources. One issue is the decision of which host OS will be used for external DNS. We observe two important facts (1) internal DNS already runs on Windows 2000; (2) it's very likely that at least one Windows 2000 server will be placed in the service network for other purposes (see Host #4 below). The obvious conclusion is that external DNS should run on Windows.

Location: Contains critical business information, so it should be located in the most protected segment of the network.

⁸ One Cisco support engineer called this "IPSec On a Stick".

Host #2 (Web / SQL Server)

Current: Windows 2000 (IIS, SQL Server)

Change: SQL Server cannot remain exposed to the Internet, but IIS is the web platform for GE's enterprise online partner and sales application. The obvious choice is to move IIS to a different server in the Service Network, and move the SQL Server into the protected network.

Location: Contains critical business information, so it should be located in the most protected segment of the network.

Host #3 (IIS)

Current: IIS running on Host #2

Need for Change: The large majority of GE's business passes through their web server. For security, stability and performance reasons, it makes good sense to run web services on a dedicated host. Current workload is moderate, including Outlook web services; a mid-range server will be fully adequate.

Cost Estimate: \$6,000 for Dell PowerEdge 600SC server hardware and Windows Server OS

Location: Provides Internet accessible services, so it must be located on the Service Network.

Host #4 (External DNS, SMTP Proxy)

Current: See Hosts 1 and 2 above. No SMTP proxy.

Need for Change: External DNS will be separated from internal DNS. Also, an SMTP proxy is needed to buffer the Microsoft Exchange server from direct Internet-based attack. DNS and SMTP proxy can easily coexist on a single, low-end sever. The only decision to be made is what kind of SMTP proxy to recommend.

SMTP Proxy Option #1: Simple proxy, using Microsoft's SMTP service or other simple SMTP proxy on Windows.

SMTP Proxy Option #2: Proxy with virus scanning and removal capability. We recommend Symantec's Antivirus for SMTP product because GE already uses the Symantec corporate product. It's known to be reliable, and also provides the ability to strip certain types of attachments that are likely to contain viruses,

Trojan horses or other malicious programs. We are currently attempting to find cost estimates for this product, but it is likely to be several thousand dollars. Rough estimate: \$5,000 for Symantec software.

Analysis: Option #2 is preferred due to the added virus protection, and will be included in the presentation to the client.

Cost Estimate: \$2,000 for Dell PowerEdge 600SC server hardware and Windows Server OS, plus estimated \$5,000 for Symantec software. Total: \$7,000.

Location: Provides Internet accessible services, so it must be located on the Service Network.

Host #5 (Windows 2000 Terminal Services):

Current: None

Need for Change: The client has expressed a need for better employee access to company resources from off-site. Installation of a Windows 2000 Terminal Server can provide very satisfactory access to resources from almost anywhere in the world; adequate service is available even over a decent modem connection. This is a service enhancement, and an opportunity to restrict direct access to the internal network from unsecured and untrusted home computers.

Estimated Cost: About \$6,000 for Dell PowerEdge server hardware; About \$3,500 for Windows Server, Client Access Licenses and Terminal Services Access Licenses for 10 connections.

Location: Provides services that are similar to those of a desktop computer, so it should be treated as such. It will be placed in the Internal Network, along with the other end user computers.

Host #6 (Intrusion Detection System)

Current: None

Need for Change: No security system is complete without intrusion detection (1) to verify that router and firewall policies are working as designed; (2) to catch intrusion attempts that cannot be blocked by a firewall.

Analysis: Commercial IDS systems are very expensive, and not proven to be significantly better than the premier Open Source solution, Snort⁹; therefore

⁹ Newman, David. "Crying wolf: False alarms hide attacks". NetworkWorldFusion. 24 Jun 2002.
URL: <http://www.nwfusion.com/techinsider/2002/0624security1.html> (9 January 2003)

Snort on Linux will be our recommendation. We will run Snort in passive mode, attached to a hub behind the firewalls.

Estimated Cost: \$1,200 for a Dell PowerEdge 600SC server.

Location: IDS systems should be located immediately behind the firewalls to audit all traffic before it reaches the hosts.

Host #7 (Syslog Server)

Centralized logging is imperative for monitoring and incident response. Fortunately, logging does not require powerful hardware. Ideally, there should be more than one Syslog server, so we will ask the IT Manager to find a castoff workstation to serve as the secondary. If we place both Syslog servers in the protected zone on a hub, we can run the secondary syslog sever in stealth mode, making it nearly immune to attack¹⁰.

Estimated Cost: About \$800 for a Dell PowerEdge 600 SC server running Linux 8.0.

Location: Logging is a critical element of monitoring and mitigation, so the log server should be located in the most protected area possible that still permits it to function properly. In this case, it is located in the Protected Network, with other critical servers.

Host # 8 (Windows Software Update Services)

The GE IT staff spend a fair amount of time keeping desktops and notebooks patched with the latest Microsoft security updates. To give them more time to spend on other tasks, we will recommend that they install a separate host with Windows Software Update Services¹¹.

Cost: About \$800 for the Microsoft Windows Server software, and \$800 for a Dell PowerEdge 600SC server.

Location: This machine services the desktops and notebooks, and does not need to be Internet accessible. It should be located in the Protected Network for maximum security.

Business Operations in the New Network Environment

Access to Internal Resources:

¹⁰ See Appendix C for details.

¹¹ Microsoft Corp. "Software Update Services Overview White Paper". 20 June 2002.

URL:<http://www.microsoft.com/windows2000/windowsupdate/sus/susoverview.asp>. (9 January 2003).

- Customers, suppliers and partners will continue to use GE's custom web application. They should notice no change.
- GE employees on site should notice no change.
- Teleworkers will definitely experience two significant changes:
 1. They will need to login to VPN before they can make any connections to internal services. In some cases login may be automated.
 2. They will be required login to Windows Terminal Services, and perform all tasks from there. At first this may appear a little cumbersome, but this should provide off-site workers with much better access to internal resources.
- The mobile sales force will experience changes similar to those of teleworkers:
 1. They must login to the VPN in order to connect to the SQL Server with their sales software client; in some cases this will be automated, but training will still be required.
 2. They will have access to Windows Terminal Services, which will give them many of the benefits of being on-site.

Internet Access: The most common Internet activity is Web browsing, so even with NAT enabled they may notice no change. The company's T1 line is too small to support video conferencing¹² or other high-intensity applications, so there should be no loss in current service.

¹² NAT interferes with some video conferencing standards.

Part IV Assignment 2 – Security Policy and Tutorial

Router Policy¹³

The following contains just the security portion of the configuration. I have not included setup such as machine name, or instructions to implement NAT on the router. I assume that NAT occurs at Ethernet 0, and that the mapping is 1-to-1 for Internet-accessible services, and pooled for clients. The outside addresses for NAT are in the range xxx.111.222.2 – xxx.111.222.127.

Notes:

- The router has enough computational power to permit extended access lists, but reflexive would likely consume too much CPU power.
- Some of the access rules (e.g. illegal address blocking) have been placed on both the internal net interface (Ethernet 0) and on the VPN interface (Ethernet 1) instead of inbound on the serial interface. I have chosen to do this to assure that they will function even if the VPN is compromised.
- Selective ordering of router access list rules will improve throughput significantly if common traffic is near the top of the list, so Microsoft traffic has been placed near the top of the lists. For the router configuration, the order of rules is largely unimportant because there are few rules, nearly all of which block (hopefully) unusual traffic.

General Configuration

logging 192.168.1.9	Syslog server
logging trap debug	set logging level
logging console emergencies	
no cdp	disable Cisco Discovery Protocol
no service tcp-small-servers	disable unneeded or obsolete services
no service udp-small-servers	
no service finger	
no ip http server	
no snmp	
no ip source-route	disable source routing; can be used to evade firewall policy
no ip bootp server	router will not act as bootp server
no ip domain-lookup	don't resolve addresses

¹³ Thanks to Mark Hofman; I studied his router policy before attempting to assemble mine.

#Main subnet

interface Ethernet 0	
ip address 192.168.1.1 255.255.255.0	set address and subnet mask
ip access-group flt_out in	name of filter for traffic entering router
ip access-group flt_in out	name of filter for traffic leaving router
no ip redirects	ICMP redirects and unreachable can reveal internal structure
no ip unreachable	
no ip directed-broadcast	suppress broadcast-based attacks
no ip proxy-arp	don't do proxy arp

#VPN subnet

interface Ethernet 1	
ip address xxx.111.222.129 255.255.255.128	set address and subnet mask
ip access-group flt_vpn_in out	name of filter for traffic leaving router
ip access-group flt_vpn_out in	name of filter for traffic entering router
no ip redirects	ICMP redirects and unreachable can reveal internal structure
no ip unreachable	
no ip directed-broadcast	suppress broadcast-based attacks
no ip proxy-arp	don't do proxy arp

#Serial Interface to WAN

interface serial 0	
ip address xxx.111.210.30 255.255.255.0	IP address and subnet mask
ip access-group flt_ser_in in	name of ingress filter for WAN link
no ip redirects	ICMP redirects and unreachable can reveal internal structure
no ip unreachable	
no ip directed-broadcast	suppress broadcast-based attacks
no ip proxy-arp	don't do proxy arp

#Access List for WAN

ip access-list flt_ser_in	name the access list
deny ip xxx.111.222.0 0.0.0.255 any log-input	suppress spoofing
deny tcp any host xxx.111.222.1 log-input	block TCP to router interfaces from outside
deny tcp any host xxx.111.222.129 log-input	
deny tcp any host xxx.111.210.30 log-input	
permit any any	let all other traffic through; will do primary filtering on Ethernet interfaces

#Access List for Internal Net

#Coming into main net (out from Ethernet 0)	
ip access-list extended flt_in	name the access list
deny udp any any range 135 139 log-input	block NetBIOS
deny tcp any any range 135 139 log-input	
deny tcp any any eq 445 log-input	block Microsoft SMB
deny udp any any eq 445 log-input	
deny udp any any eq tftp log-input	block tftp
permit udp host xxx.111.222.3 host xxx.111.222.9 eq syslog	permit syslog from VPN to log server
deny udp any any eq syslog log-input	block all other syslog
deny ip 0.0.0.0 0.255.255.255 any log-input	block illegal, private, non-routable or reserved addresses
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	
deny ip 192.0.2.0 0.0.0.255 any log-input	
deny ip 169.254.0.0 0.0.255.255 any log-input	
deny ip 224.0.0.0 7.255.255.255 any log-input	
deny ip host 0.0.0.0 any log-input	
permit ip any any	permit whatever remains

#Egress List for Internal Net

ip access-list extended flt_out	name the filter
deny udp any any range 135 139 log-input	block NetBIOS
deny tcp any any range 135 139 log-input	
deny tcp any any eq 445 log-input	block Microsoft SMB
deny udp any any eq 445 log-input	
deny udp any any eq snmp log-input	block outgoing snmp
deny udp any any eq snmptrap log-input	block outgoing snmp traps
deny udp any any eq tftp log-input	block outgoing tftp
deny udp any any eq syslog log-input	block outgoing syslog
deny ip any 0.0.0.0 0.255.255.255 log-input	block illegal, private, non-routable or reserved addresses; Note: if these are generated inside our network then we have either a misconfigured device or an internal security risk problem.
deny ip any 10.0.0.0 0.255.255.255 log-input	
deny ip any 172.16.0.0 0.15.255.255 log-input	
deny ip any 192.168.0.0 0.0.255.255 log-input	
deny ip any 192.168.2.0 0.0.0.255 log-input	
deny ip any 169.254.0.0 0.0.255.255 log-input	
deny ip any 224.0.0.0 15.255.255.255 log-input	
deny ip any 127.0.0.0 0.255.255.255 log-input	
deny ip any 255.0.0.0 0.255.255.255 log-input	
deny ip any 224.0.0.0 7.255.255.255 log-input	
permit ip any any	
	permit everything else

#Access List for VPN Net (out from ethernet 1)

ip access-list extended flt_vpn_in	name the access list
permit ip xxx.111.222.0 0.0.0.255 any	permit anything from the main subnet
permit udp any eq 500 any eq 500 log-input	ISAKMP
permit udp any eq 10000 any eq 10000 log-input	VPN 3000 uses this for IPSec over UDP to pass through NAT's; not in use in this environment, but may be required anyhow
permit ah any any	permit all AH packets for IPSec
permit esp any any	permit all ESP
deny ip any any log-input	drop everything else

#Egress List for VPN Net (in to ethernet 1)

ip access-list extended flt_vpn_out	name the access list
permit ip any xxx.111.222.0 0.0.0.255 any	permit any traffic headed to internal net
permit udp any eq 500 any eq 500 log-input	permit any ISAKMP traffic
permit udp any eq 10000 any eq 10000 log-input	VPN 3000 uses this for IPSec over UDP to pass through NAT's; not in use in this environment, but may be required anyhow
permit ah any any	permit all AH traffic
permit esp any any	permit all ESP traffic
deny ip any any log-input	drop everything else

#VTY Access List

ip access-list vty_in	name the list
permit tcp host 192.168.1.9 eq telnet log-input	permit the log server to telnet in
deny ip any any log-input	block everything else

#VTY Configuration

line vty 0	
access-class vty_in in	Note: since there's no access-class configuration for the "out" direction, all traffic is permitted.
exec-timeout 5 0	5 minute inactivity timeout on the connection
vty 1-4 as for vty 0	not repeated here
banner motd ^CWARNING: Authorized use only. All activity is monitored and reviewed. Violators may be prosecuted.^C	

VPN Policy

Address of VPN Concentrator: xxx.111.222.82

Client Address Management: Allocated by VPN device

Client Address Range: xxx.111.222..96 to xxx.111.222.127

Syslog Server: xxx.111.222.9 (statically NAT'ed to 192.168.1.9)

SMTP Server for Alerts: xxx.111.222.25 (statically NAT'd to 192.168.1.25)

Servers:

FTP disabled

HTTP disabled

HTTPS enabled for management

TFTP disabled

Telnet disabled

SNMP disabled

SNMP disabled

SSL enabled for management

SSH disabled

XML enabled (configuration files are managed as XML documents)

Authentication: Internal server

Tunnel Policies

PPTP: Disabled

L2TP: Disabled

IPSec: Enabled

IKE Proposals:

CiscoVPNClient-3DES-MD5
IKE-3DES-MD5
IKE-3DES-MD5-DH1
IKE-DES-MD5
IKE-3DES-MD5-DH7
IKE-3DES-MD5-RSA
CiscoVPNClient-3DES-MD5-DH5
CiscoVPNClient-AES128-SHA
IKE-AES128-SHA

Group Configuration: All users will go through a standard group, as shown above; however, some mobile sales staff may work in countries that have restrictive encryption laws. This would require a separate group with weaker encryption.

IPSec SA: ESP-3DES-MD5

Firewall Policy (Main Firewall)

As discussed earlier, the firewall is Netfilter operating in conjunction with the Linux Ethernet bridge: it's not a conventional routing firewall. Below are the shell scripts used to create the bridge and firewall rule set. The Linux bridge code is included in more recent distributions of Red Hat Linux (versions 7.3 and above), or it can be installed from source obtained from <http://bridge.sourceforge.net>. For our purposes, we will assume the system is running Red Hat 8.0.

# Shell script to create the Linux Ethernet bridge	
brctl addbr br0	create the bridge
brctl addif br0 eth0	add Ethernet interfaces to the bridge
brctl addif br0 eth1	
brctl addif br0 eth2	
ifconfig eth0 0.0.0.0 promisc	remove IP addresses from IF's; put them into promiscuous mode
ifconfig eth1 0.0.0.0 promisc	
ifconfig eth2 0.0.0.0 promisc	
ifconfig br0 10.0.0.3	assign an arbitrary IP address to the bridge virtual interface
echo 1 >/proc/sys/net/ipv4/ip_forward	enable IP forwarding among interfaces
echo 0 >/proc/sys/net/ipv4/conf/all/accept_source_route	disable source-routing on all IF's
echo 1>/proc/sys/net/ipv4/icmp_echo_ignore_broadcasts	block ICMP broadcasts
#!/bin/bash	It's a BASH script

#This script creates the firewall rule set.	
#Set shell variables to simplify the script	
IPTABLES='/sbin/iptables'	shortcut to iptables command
EXT_IFACE='eth0'	descriptive names for Ethernet IF's
INT_IFACE='eth1'	
SVC_IFACE='eth2'	
BAD_GUYS='217.23.165.75'	List of banned IP addresses: www.nic.im (arbitrary example: Isle of Man)
INT_SUBNET='192.168.1.0/24'	hosts behind the firewall
VPN_NET='xxx.111.222.80/25'	Descriptive name for the VPN subnet
SVC_NET='192.168.1.24/28'	Service network
LOCAL_BCAST='192.168.1.255'	Local broadcast address
ROUTER='192.168.1.1'	router
TRUSTED_HOST='192.168.1.9'	management station
LOG_SERVER='192.168.1.9'	
INT_FW_ADDR='10.0.0.3' # 10.0.0.3 is not a typo: see Linux Bridge Tutorial for details.	The firewall's internal management address
WEB_SERVER='192.168.1.24'	Web server SMTP proxy External DNS Internal mail server SQL Server Windows Terminal Server
SMTP_PROXY='192.168.1.25'	
EXT_DNS='192.168.1.25'	
INT_MAIL='192.168.1.8'	
SQL_SVR='192.168.1.10'	
TRM_SVR='192.168.1.32'	
# Clear out any existing firewall configuration	
\$IPTABLES -F	flush all chains
\$IPTABLES -X	delete all user chains
# set policies to drop by default	
\$IPTABLES -P INPUT DROP	INPUT chain
\$IPTABLES -P OUTPUT DROP	OUTPUT chain
\$IPTABLES -P FORWARD DROP	FORWARD chain
modprobe ip_conntrack_ftp	load the state tracking module for ftp
modprobe ip_conntrack	load the general state tracking module
#	

# For easier maintenance, create user chains to handle logging.	
#	
\$IPTABLES -N LOGDROPSPOOF	For each of the cases below 1. create a new chain, 2. log with descriptive info 3. drop the packet
\$IPTABLES -A LOGDROPSPOOF -j LOG --log-level info --log-prefix " SPOOF "	
\$IPTABLES -A LOGDROPSPOOF -j DROP	
\$IPTABLES -N LOGDROPADDR	
\$IPTABLES -A LOGDROPADDR -j LOG --log-level info --log-prefix " BAD DEST ADDR "	
\$IPTABLES -A LOGDROPADDR -j DROP	
\$IPTABLES -N LOGDROPBADGUY	
\$IPTABLES -A LOGDROPBADGUY -j LOG --log-level info --log-prefix " BADGUY "	
\$IPTABLES -A LOGDROPBADGUY -j DROP	
# Port not permitted (i.e. well known ports for certain protocols)	
\$IPTABLES -N LOGDROPPORT	
\$IPTABLES -A LOGDROPPORT -j LOG --log-level info --log-prefix " PORT "	
\$IPTABLES -A LOGDROPPORT -j DROP	
#Attempt to connect to FW	
\$IPTABLES -N LOGDROPFWCONN	
\$IPTABLES -A LOGDROPFWCONN -j LOG --log-level info --log-prefix " FWCONN "	
\$IPTABLES -A LOGDROPFWCONN -j DROP	
#Misc new traffic without SYN set	
\$IPTABLES -N NEWNOTSYN	
\$IPTABLES -A NEWNOTSYN -j LOG --log-level info --log-prefix " NEWNOTSYN "	
\$IPTABLES -A NEWNOTSYN -j DROP	
#Possible service network compromise	
\$IPTABLES -N LOGDROPSVCNET	
\$IPTABLES -A LOGDROPSVCNET -j LOG --log-level info --log-prefix " SVCNET "	
\$IPTABLES -A LOGDROPSVCNET -j DROP	
#Log & Accept	Log and accept certain types of traffic
\$IPTABLES -N LOGACCEPT	
\$IPTABLES -A LOGACCEPT -j LOG --log-level info	

\$IPTABLES -A LOGACCEPT -j ACCEPT	
#Log & Accept test traffic; used in Assignment 3	
\$IPTABLES -N LOGACCEPTTEST	
\$IPTABLES -A LOGACCEPTTEST -j LOG --log-level info --log-prefix " TESTOK "	
\$IPTABLES -A LOGACCEPTTEST -j ACCEPT	
#	
# Start of Rule Set	
#	
#Block Known Attackers	
#	
for i in \$BAD_GUYS;	loop through list of banned addresses
do	
\$IPTABLES -A FORWARD -s \${i} -j LOGDROPBADGUY	create a block rule
done	
#	
# Inappropriate SYN flag	
#	
\$IPTABLES -A FORWARD -p tcp ! --syn -m state --state NEW -j NEWNOTSYN	No state info (NEW session), but no SYN flag: that smells odd
#	
#Detect spoof attempts	
#	
\$IPTABLES -A FORWARD -i \$EXT_IFACE -s \$ROUTER -j ACCEPT	let router traffic through
\$IPTABLES -A FORWARD -i \$EXT_IFACE -s \$INT_SUBNET -j LOGDROPSPOOF	block spoofed traffic
\$IPTABLES -A FORWARD -i \$EXT_IFACE -d ! \$INT_SUBNET -j LOGDROPPADDR	drop traffic not headed to normal addresses
\$IPTABLES -A FORWARD -o \$EXT_IFACE -s ! \$INT_SUBNET -j LOGDROPSPOOF	block internally-generated spoofs
#	
#Allow ESTABLISHED,RELATED traffic	
#	
\$IPTABLES -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT	match traffic against state table & permit appropriate packets

#	
# kill noise: NBT, SMB, broadcast	
#	
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p tcp -- dport 135:139 -j DROP	NetBIOS
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p udp -- dport 135:139 -j DROP	NetBIOS
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p tcp -- dport 445 -j DROP	SMB
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p udp -- dport 445 -j DROP	SMB
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p udp -- dport 69 -j DROP	tftp
\$IPTABLES -A FORWARD -i \$EXT_IFACE -p udp -- dport 161:162 -j DROP	snmp
\$IPTABLES -A FORWARD -d \$LOCAL_BCAST -j DROP	internal broadcast
\$IPTABLES -A FORWARD -d 255.255.255.255 -j DROP	broadcast
#	
# drop & log some outbound protocols	
#	
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p tcp -- dport 135:139 -j LOGDROPPORT	NetBIOS
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 135:139 -j LOGDROPPORT	NetBIOS
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 69 -j LOGDROPPORT	tftp
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 161:162 -j LOGDROPPORT	snmp
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 514 -j LOGDROPPORT	syslog
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p tcp -- dport 143 -j LOGDROPPORT	IMAP
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 143 -j LOGDROPPORT	IMAP
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p tcp -- dport 110 -j LOGDROPPORT	POP3
\$IPTABLES -A FORWARD -o \$EXT_IFACE -p udp -- dport 110 -j LOGDROPPORT	POP3
#	
# Access control for the firewall itself	
#	
\$IPTABLES -A INPUT -s 127.0.0.1 -i lo -j ACCEPT	permit all traffic to/from

\$IPTABLES -A OUTPUT -s 127.0.0.1 -o lo -j ACCEPT	loopback IF Notice that these are the INPUT and OUTPUT chains, not the FORWARD chain.
\$IPTABLES -A INPUT -i \$INT_IFACE -p tcp --dport 22 -s \$TRUSTED_HOST -d \$INT_FW_ADDR -j ACCEPT	Permit SSH to FW from trusted host (INPUT chain)
\$IPTABLES -A OUTPUT -o \$INT_IFACE -p tcp --sport 22 -d \$TRUSTED_HOST -s \$INT_FW_ADDR -j ACCEPT	Permit SSH from FW to trusted host (OUTPUT chain)
\$IPTABLES -A INPUT -i ! \$INT_IFACE -j LOGDROPFWCONN	Drop and log inappropriate traffic headed to the FW's IP (INPUT chain)
\$IPTABLES -A INPUT -s ! \$TRUSTED_HOST -j LOGDROPFWCONN	
\$IPTABLES -A FORWARD -d \$INT_FW_ADDR -j LOGDROPFWCONN	Should never happen
\$IPTABLES -A OUTPUT -o eth1 -p udp --dport 514 -d \$LOG_SERVER -j ACCEPT	permit syslog outbound from the firewall (OUTPUT chain)
# # Detect suspicious outbound traffic, such as netbus & BO. # It's not practical to detect all Trojan activity at the firewall -- that's better # left to an IDS. This is just to illustrate that a few checks can be placed here # when a new Trojan variant is released. # # See the following page for a list of common Trojans ports # http://www.sans.org/newlook/resources/IDFAQ/oddports.htm #	Netbus
\$IPTABLES -A FORWARD -i ! \$EXT_IFACE -p tcp --sport 12345 -j LOG --log-level info --log-prefix " TROJAN "	
\$IPTABLES -A FORWARD -i ! \$EXT_IFACE -p tcp --sport 31337 -j LOG --log-level info --log-prefix " TROJAN "	
\$IPTABLES -A FORWARD -i ! \$EXT_IFACE -p tcp --sport 8787 -j LOG --log-level info --log-prefix " TROJAN "	BO
#	
#Allow access to/from GE's external services	
#	
\$IPTABLES -A FORWARD -o \$SVC_IFACE -p tcp --dport 80 -d \$WEB_SERVER -j LOGACCEPT	permit public access to web server

\$IPTABLES -A FORWARD -o \$SVC_IFACE -p tcp --dport 25 -d \$SMTP_PROXY -j LOGACCEPT	permit public access to SMTP proxy
\$IPTABLES -A FORWARD -i \$SVC_IFACE -p tcp --dport 25 -s \$SMTP_PROXY -d \$INT_MAIL -j LOGACCEPT	permit transfer from SMTP proxy to internal mail
\$IPTABLES -A FORWARD -i \$SVC_IFACE -p tcp --sport 25 -s \$SMTP_PROXY -d ! \$INT_SUBNET -j LOGACCEPT	permit connections from SMTP proxy to outside
\$IPTABLES -A FORWARD -o \$SVC_IFACE -p udp --dport 53 -d \$EXT_DNS -j ACCEPT	permit public access to external DNS server & queries to outside DNS servers
\$IPTABLES -A FORWARD -i \$SVC_IFACE -p udp --sport 53 -s \$EXT_DNS -j ACCEPT	
\$IPTABLES -A FORWARD -i \$SVC_IFACE -p tcp --dport 1433 -s \$WEB_SERVER -d \$SQL_SVR -j ACCEPT	permit web server access to SQL Server
\$IPTABLES -A FORWARD -i \$EXT_IFACE -o \$INT_IFACE -s \$VPN_NET -d \$TRM_SVR -j LOGACCEPT	permit VPN clients access to Terminal Server
#	
# log & drop unexpected outbound traffic	
#	
\$IPTABLES -A FORWARD -i \$SVC_IFACE -o \$EXT_IFACE -j LOGDROPSVCNET	outbound connections from service network
#	
# allow traffic from internal hosts that was not previously blocked	
#	
\$IPTABLES -A FORWARD -i \$INT_IFACE -s \$INT_SUBNET -j LOGACCEPT	

Comments on the Order of Firewall Rules

The general order of rules for traffic on the FORWARD chain is

1. Block obviously undesirable traffic (known attackers, spoofs, etc)
2. Permit existing established connections and traffic related to previous outbound packets
3. Block noise
4. Permit inbound access to specific services
5. Block improper outbound traffic
6. Permit remaining outbound traffic

This firewall has at least a Pentium 4 1.7GHz processor. That's a lot of processing power for such a small rule set, so I have not made much effort to optimize the rules; rather, I prefer to have them well organized into service-related groups. Since this is a stateful firewall implementation, only the first packet of a TCP connection will traverse the bulk of the rules; subsequent packets will match the RELATED,ESTABLISHED rule near the top of the rule set, and will be accepted quickly.

Other Approaches to Rule Set Structure

This rule set is fairly small and the firewall's CPU is powerful, so I placed a high value on ease of maintenance over raw performance: rules which apply to a specific service are kept close together. With a substantially larger rule set, I might have chosen to divide the rules differently for better performance. One of the most common methods¹⁴ is to create a separate chain to test traffic between each pair of interfaces, for example

```
$IPTABLES -A FORWARD -i eth0 -o eth1 -j FROM-ETH0-TO-ETH1
```

The chain FROM-ETH0-TO-ETH1 would include all tests for traffic passing from eth0 to eth1. A substantial amount of time can be saved since the tests for incoming and outgoing interfaces are performed just once each.

Testing for Pathological Packets

Additional rules could be added to test for scan signatures such as SYN/FIN and ACK, and Christmas Tree flag patterns. I omitted these only for brevity. The IDS system ought to catch them.

¹⁴ Vestergaard, Peter. "FIREWALLS, PERIMETER PROTECTION AND VPNS PRACTICAL ASSIGNMENT". 26 Oct 2001. URL: http://www.giac.org/practical/Peter_Vestergaard_GCFW.zip. (9 January 2003)

Tutorial: The Linux Bridge Firewall & Netfilter

Section 1: The Linux Bridge

Just as the name states, a *bridge firewall* is a firewall that runs on a bridging host, rather than on a routing host. A *bridge* is very much like an Ethernet switch: in fact, an Ethernet switch is one specific kind of bridge: bridges also exist for other types of physical networking technologies. An Ethernet bridge works at Layer 2 of the OSI model to connect different segments of a physical network, passing frames selectively among the segments. Since we're working with an Ethernet bridge, let's use that as our example.

1. When a frame comes in to the bridge (think switch here), the bridge examines its source and destination MAC addresses.
2. The bridge determines which of its ports received the frame, and saves the port designator and the source MAC address together in a table.
3. The bridge searches the table of known MAC addresses; if the destination MAC address is found, the bridge sends the frame just to the one port associated with the destination MAC address.
4. If the destination MAC address is not in the table, the switch must *flood* the packet: that is, it sends it to all ports, except the sending port.
5. The port-MAC table has a timeout to permit stale associations to be removed. For the Linux bridge the timeout is 5 minutes.

In a normal Ethernet environment, the destination MAC address is nearly always found in the MAC table, so the frame will not be flooded. In particular, for TCP, flooding can only happen during the three-way handshake phase, or if there has been a long silence in the communication. In an active, two-way communication, only a few frames will flood during the course of the connection. Also, it's important to understand that the bridge cannot make use of IP information when it decides where to send a frame: its decision is made entirely using its internal MAC table.

How This Relates to Firewalls

The current bridge code in the Linux kernel is a recent re-write done by Lennert Buytenhek of the Netherlands. In rewriting the code, he made sure that it could interoperate with the Linux firewall tools, ipchains and Netfilter, so it's possible to filter all packets passing through the bridge. For example, suppose you have a bridge firewall with three interfaces, and a frame comes in carrying an IP packet. The following process occurs:

1. The frame enters the source address carrying the IP packet.
2. The bridge determines whether to unicast or flood the packet, depending on the state of its port-MAC table. It makes a copy of the Ethernet frame for each destination port.

3. Each copy of the original frame is passed through the firewall code to determine whether it should be accepted or dropped.

Why You Might Use a Bridge Firewall Instead of a Routing Firewall

There are several important reasons:

1. A bridge firewall is a “drop-in” device, requiring no change to IP network structure. With a routing firewall, it’s usually necessary to create subnets off of the internal interfaces¹⁵. Not so with a bridge: since it works at Layer 2 it’s like dropping in an unmanaged Ethernet switch: no re-configuration of your subnet structure is required. This is particularly helpful if you have no authority over your network structure – you can install a full-featured bridge firewall without arousing the ire, or even the notice, of the router administrators¹⁶.
2. The bridge can be configured without an IP address. This means:
 - a. It can’t be probed using standard IP-based tools; essentially, it’s invisible to IP traffic, and therefore...
 - b. It can’t be attacked using standard IP-based tools.This is not to say that a bridge firewall is immune to attack, just that it’s more difficult than for a routing firewall.
3. Because a bridge is not specific to a given location in your network structure, you can keep a spare on the shelf and just drop it in if your current firewall fails. This is particularly true if you’re using an inline firewall, i.e., a firewall with just two interfaces.
4. Even if you have an existing router-based firewall, the bridge firewall can make an excellent second line of defense because it adds only minimal maintenance cost.
5. Since it works at Layer 2, it will not interfere with non-IP protocols, e.g. IPX.

Note: most people choose to assign a non-routable IP address to the bridge for management purposes. However, because the IP address is unnecessary to the function of the bridge, it should be unrelated to your standard IP address assignment scheme. This will ensure that your bridge is invisible and inaccessible except from specific, trusted hosts that have been configured with additional routing information.

Disadvantages of the Bridge Firewall

In some cases a bridge firewall may not be appropriate.

1. You may prefer to have a routing firewall to break up a larger network. Of course, you could also do this with a standalone router.
2. If your environment requires nearly 100% perfect security, you may consider the occasional packet floods to be a security risk. Bear in mind that flooding doesn’t pass any packets unless the firewall accepts them; however, sometimes a packet

¹⁵ Proxy ARP can help work around this, but it’s usually not worth the hassle.

¹⁶ The bridge code supports Spanning Tree Protocol for fail over; if this feature is enabled it would be apparent to a router admin who monitors logs regularly.

may be transmitted to several interfaces instead of just one. It's possible to mitigate this problem, as described in the Netfilter tutorial given later.

3. The flooding issue requires that you exercise a little more care in the construction of your Netfilter rule base.¹⁷

Bridge Firewall Installation and Configuration

The bridge code is included in Redhat Linux 7.3 and 8.0, so I'll use this for my example. If you're using another Linux distribution, or if you want to build your own kernel, you can find appropriate information at the URL's provided at the end of the tutorial.

The Ethernet bridge consists of two pieces of software:

- The kernel component that implements the bridge. This is already compiled into the standard Redhat kernels.
- The bridge control utility, *brctl*, which allows you to create and configure bridges using your existing Ethernet interfaces.

Installation

1. Start with a standard RedHat 7.3 or 8.0 distribution. You should install at least two Ethernet interfaces, but you may install as many as you like. I will assume that you plan to use eth0, eth1 and eth2 for your bridge.
2. Create a bridge configuration script. The script will configure the Ethernet interfaces and collect them into the bridge structure. Here's an example:

Command	Explanation
ifconfig eth0 0.0.0.0 promisc up	Remove IP configuration, put interface in promiscuous mode, and bring it up.
ifconfig eth1 0.0.0.0 promisc up	
ifconfig eth2 0.0.0.0 promisc up	
brctl addbr br0	Create a bridge named "br0"
brctl addif br0 eth0	Add the three interfaces to the bridge.
brctl addif br0 eth1	
brctl addif br0 eth2	
echo 1 >/proc/sys/net/ipv4/ip_forward	enable ip forwarding
ifconfig br0 10.0.0.3 netmask 255.0.0.0 promisc up	Assign an arbitrary, non-routable IP address to the bridge. This address should be unrelated to your normal addressing scheme.
route add -host 192.168.1.34 br0	Tell Linux how to talk to your trusted host.

¹⁷ Buytenhek, Lennert. "Bridging and Firewalling". 1 Oct 2002.
[URL: http://bridge.sourceforge.net/docs/bridge-firewall.html](http://bridge.sourceforge.net/docs/bridge-firewall.html). (9 January 2003)

Run the following command on the trusted host:	You must tell your trusted host how to talk to the bridge.
<pre>route add -net 10.0.0.0 netmask 255.0.0.0 eth0</pre>	

3. Once these configuration commands have been issued, the bridge is up and passing packets. Of course, now you need to build a firewall rule base using Netfilter.

Section 2: Netfilter¹⁸

To construct a Netfilter rule set you use the Linux command *iptables* with the appropriate command-line options. A typical command looks like this:

(1) `iptables -A FORWARD -s <source addr> -d <destination addr> -j DROP`

Generally, a Netfilter ruleset is created by collecting a list of these commands in a shell script.

Netfilter divides IP traffic into three categories:

INPUT	The Netfilter host is the ultimate destination of the packet
OUTPUT	The Netfilter host is the source of the packet
FORWARD	The Netfilter host is forwarding the packet from an external source to an external destination

There is a separate *chain*, or list of rules, for each of these categories. On a dedicated firewall most of the rules are in the FORWARD chain. Look again at Example 1 above: it creates a rule that drops any packet that has the given source and destination addresses. The option “-A FORWARD” indicates that this rule is added to the FORWARD chain, and so this rule applies only to packets that are to be forwarded from one external host to another. The option “-j DROP” means to drop the packet. The -j option can be read more generally as “jump to” the target specified. The more common targets are:

DROP	Drop the matched packet
ACCEPT	Accept the matched packet
LOG	Create a log entry about the packet.

(2) `iptables -A FORWARD -s 192.168.1.4 -d 10.3.4.67 -j ACCEPT`

¹⁸Russell, Rusty. “Linux 2.4 Packet Filtering HOWTO”. 24 January 2002.URL:
<http://www.netfilter.org/documentation/HOWTO//packet-filtering-HOWTO.html> (9 January 2003)

Example 2 shows a rule in the *FORWARD* chain to accept traffic with the given source and destination addresses.

A Netfilter rule can also select packets based on the protocol and port.

```
(3) iptables -A FORWARD -p tcp --dport 80 -d 10.3.4.90 -j ACCEPT
```

Example 3 is a rule to accept TCP packets with destination port 80 and destination IP 10.3.4.90.

Netfilter can even select based on the incoming and outgoing interfaces.

```
(4) iptables -A FORWARD -i eth0 -d 192.168.1.3 -j DROP
```

A packet *inbound from* eth0 directed to this host will be dropped. Similarly,

```
(5) iptables -A FORWARD -o ! eth1 -d 192.168.1.3 -j DROP
```

This rule will drop packets for this host if they are *not (!)* to be forwarded to interface eth1. This can be very important when using Netfilter in conjunction with the Linux bridge; this will be explained further near the end of the tutorial.

You can also test packets for certain protocols.

```
(6) iptables -A FORWARD -p udp --dport 19 --j DROP
```

This drops all UDP traffic for the chargen service. You can test traffic for these protocols: tcp, udp, icmp.

Netfilter is an extensible firewall, with several plug-in modules available. Two of the most popular modules are the connection tracking modules for IP and FTP. Using the connection modules makes Netfilter a stateful firewall for these protocols¹⁹. To include the connection tracking modules in your Netfilter configuration, add the following lines to your Netfilter creation script.

```
(7) modprobe ip_conntrack_ftp  
modprobe ip_conntrack
```

Once the modules are loaded, the connection (i.e. state) tracking is performed automatically. However, state testing for a particular packet must be done manually.

```
(8) iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT
```

¹⁹ The adjective “stateful” is relative: only TCP has a true connection state that can be tracked. However, it is possible for a firewall to be smart about some other traffic. For example, if Netfilter sees an ICMP echo request go out, it can expect to see a reply.

In example 8, the expression “-m state” tells Netfilter that it will need to work with the state module (ip_conntrack). In Netfilter, the double-dash “--” expresses a sub-option, sometimes an option for a module. In this case, “--state ESTABLISHED, RELATED” is an option for the state module; it tests whether the current packet meets a least one of the following conditions: it’s part of an established TCP connection or related to an existing connection. If at least one of these is true, then the packet matches the rule, and is accepted. The supported connection states are

INVALID	The packet is not associated with any known connection
ESTABLISHED	Part of an established TCP connection
NEW	The packet is starting a new connection
RELATED	The packet is starting a new connection, but is associated with an existing connection, such as FTP data transfer.

Logging in Netfilter is done through the *LOG* target. A typical logging rule looks like this.

```
(9) iptables -A FORWARD -p udp --dport 19 -j LOG --log-prefix  
“CHARGEN”
```

The log entry will include the text prefix “CHARGEN”, which makes it easy to *grep* for similar entries. Notice that the rule in example 9 does not also drop the packet, so a second rule with “-j DROP” is needed for that. This is rather inefficient; a better method is shown later.

Netfilter allows the user to create new chains to supplement the standard INPUT, OUTPUT and FORWARD chains. User-defined chains are often employed to break a complex rule set into more manageable parts. Consider this example

```
(10) iptables -N TCP_PACKETS  
iptables -A TCP_PACKETS -j LOG --log-prefix “A TCP Packet”  
iptables -A FORWARD -p tcp -j TCP_PACKETS
```


The first command creates a new chain, and names it “TCP_PACKETS”. The second command adds a new rule to the TCP_PACKETS chain that simply creates a log entry labeling the packet as TCP. The third command adds a command to the *FORWARD* chain to redirect all TCP packets to the new chain. By creating separate rules for TCP, UDP, ICMP and other protocols, a large rule set can be broken down into more manageable parts. This can also improve throughput since selection based on protocol happens just once, rather than multiple times throughout the rule set.

In the case of a bridge firewall, however, there is another way to improve efficiency and also to overcome most of the risks of frame *flooding* that were described in the bridge tutorial. Rather than employ user chains to handle different protocol types, I prefer to have them divide traffic by source and destination interface. For example,

```
(11) iptables -N FromEth0ToEth1
iptables -N FromEth0ToEth2
iptables -N FromEth1ToEth0
iptables -N FromEth1ToEth2
iptables -N FromEth2ToEth0
iptables -N FromEth2ToEth1
```

There are several advantages to this approach:

1. Security policies often govern how traffic should flow from one segment to another: one might be a highly protected segment with valuable assets, and another might be the service network. Collecting rules according to policy can make it easier to enforce the policy and avoid errors.
2. In the case of a bridge firewall, this approach can eliminate the problems of flooding: a few rules at the start of each user chain can block stray traffic escaping to inappropriate segments. Furthermore, if network addresses are cleverly assigned, this flood-protection can be accomplished with a single firewall rule in each user chain. Consider this network with three segments.



Each of the networks has been assigned a range of addresses on even binary boundaries. This permits us to take advantage of Netfilter's ability to work with address masks. For example,

```
(12) iptables -N FromEth0ToEth1
iptables -A FromEth0ToEth1 -d ! 192.168.1.16/28 -j DROP
iptables -A FORWARD -i eth0 -o eth1 -j FromEth0ToEth1
```

If the destination of the packet is not on the local segment, it's simply dropped, and the possibility of flooding is eliminated.

Finally, consider logging in example 13, below.

```
(13) iptables -N LOGDROPCHARGEN
iptables -A LOGDROPCHARGEN -j LOG --log-prefix " CHARGEN "
iptables -A LOGDROPCHARGEN -j DROP
iptables -A FORWARD -p udp --dport 19 -j LOGDROPCHARGEN
```

Creating user chains to handle logging has two advantages:

1. It eliminates the need to write two identical rules: one for -j LOG and the second for -j DROP.
2. It puts the log-prefix in a single location, so it can be changed more easily, if the need arises.

Section 3: Implementing the Policy

Implementation is now a simple sequence of steps.

1. Obtain host hardware, and install the appropriate number of NIC's. It's not necessary to preinstall additional NIC's to accommodate future needs as this can be done fairly easily when the time comes.
2. Install Linux with support for your hardware. Make sure to install a recent version of the kernel to protect against known bugs.
3. Identify and physically label your NIC's to avoid confusion. A bridge firewall can be partially functional even if two Ethernet cables are reversed; however, it won't implement the security policy correctly.
4. Create the shell script to create the bridge and include the NIC's.
5. Create the script that invokes *iptables* to build the rule base.
6. Connect the Ethernet cables to the appropriate segments of your network.

Documentation and Links

<http://bridge.sourceforge.net/>

<http://www.netfilter.org/>

Conclusion

The Linux bridge firewall configuration is an excellent way to create a powerful yet inexpensive firewall that is invisible to IP traffic. It is particularly well suited to these circumstances:

- Small, non-subnetted networks
- Networks where the bridge-firewall administrator can't change the subnet structure
- Situations where an inline firewall is needed, e.g. to protect a small group of high-value hosts
- As a second line of defense, in partnership with a routing firewall

Part V Assignment 3 – Verify Firewall Policy

Audit Plan

The purpose of this section is to audit the firewall policy, not to assess the vulnerabilities of hosts. Before beginning the test setup, we interview technical staff to see if there are any non-functional services, which might indicate that some traffic is inappropriately blocked. Then, to test the firewall rule set, we will use `nmap` to send TCP and UDP packets from an attack host to a target host, and the results will be captured on the target host using `tcpdump`. The logs will show which packets have traversed the firewall, and therefore whether the firewall rules function as intended.

Technical Approach: Two stages

1. Bench testing: For this test, we assemble a test firewall, an attack computer and a target computer in an isolated environment. The configuration and rule set for the live firewall will be copied to the test firewall. The attack computer uses `nmap` to perform scans against the target computer, and the target computer runs `tcpdump` continuously to monitor incoming traffic. The firewall software is Netfilter, and the rule set has been configured to permit *RELATED* and *ESTABLISHED* traffic; therefore, new traffic is of primary interest, so we will send SYN packets and empty UDP packets to the target host. The firewall will see these packets as new traffic, giving us a reliable indication of packet traversal. We will not perform explicit tests for permitted traffic in the bench test phase, but will leave this for the live test phase.
2. Live testing: We will coordinate with GE's management to select a date and time to perform testing on the firewall (this will likely occur on a weekend). We then install a scanning computer and listeners on different segments of the internal network (service network and internal network). As in Stage (1), the listeners will run `tcpdump` continuously, and the attack computer uses `nmap` to attempt penetration of the network. The scan procedure is performed from each segment of the network. In addition to the `nmap` scan, we also work with local IT staff to verify that all permitted services are functioning properly.

Practical Considerations: The bench test is performed in an isolated environment, and need not take place on GIAC Enterprises' premises. The live test should take place on a weekend or holiday, preferably at a time that does not interfere with international sales staff.

Cost and Effort: Bench testing will take 4 to 6 hours for setup, testing and analysis. Live testing will take 6-8 hours, including time to verify that permitted services function correctly. The cost will be billed at the Normal Business Hours rate of \$ per hour for bench testing, and \$\$ per hour for live testing on the weekend.

Risks and considerations: Bench testing offers no risks. Live testing using `nmap` should be a low-risk exercise since `nmap` is merely a scanner, not a vulnerability


testing tool. During live testing, the most probable risk to individual hosts would be caused by a large number of half-open TCP connections; this may cause memory allocation errors, and unpredictable host behavior. Although there is no reason to believe that live testing will pose a threat to data, we insist that GE's IT staff perform full backups of critical systems before live testing occurs.

Conduct the Audit

Stage 1: Bench Testing Procedure

Live testing of a firewall interrupts business and can pose risks to a functional environment. The purpose of the bench test is to find procedural and configuration problems before the real firewall is tested in the production environment.

1. In an isolated network environment, make the physical network connections shown in the diagram below. The test firewall should be functionally equivalent to the production firewall, though it need not be precisely identical.



2. Create a shell script to build and configure the bridge; this will be the same as the bridge script used on the production firewall²⁰. Copy the firewall creation script from the production computer.
3. Notice that the firewall rule base gives special treatment to certain IP addresses and groups of addresses; compile a list of all such special addresses, and include a representative from each special group. For example,

192.168.1.9 is the log server
192.168.1.32 is the Windows Terminal server
192.168.1.35 represents a generic client computer
63.100.47.46 is a generic external host (www.sans.org)

Use this list to create the following shell script to assign all of these special case addresses to the target computer's Ethernet card. The target computer will then accept packets for all of these IP addresses.

```
# This script assigns multiple IP addresses to the single physical
# Ethernet interface to permit bench testing of the firewall rules.

ifconfig eth0:8 192.168.1.8
ifconfig eth0:9 192.168.1.9
ifconfig eth0:10 192.168.1.10
ifconfig eth0:11 192.168.1.11
ifconfig eth0:24 192.168.1.24
ifconfig eth0:25 192.168.1.25
ifconfig eth0:32 192.168.1.32
ifconfig eth0:35 192.168.1.35
ifconfig eth0:99 63.100.47.46
```

4. On the attack computer, compose the following shell script to add static routes for all of the target addresses. Compose another shell script to create static ARP entries for each of the target addresses. The static routes and static ARP entries are required because there is no intermediary router to forward packets: the static configurations assure that the attack computer will be able to send the packet to the target.

```
#Script to add static routes for attack computer.
#This is necessary because the attack computer will spoof
#several source addresses

route add -net 192.168.1.0 netmask 255.255.255.0 eth0
route add -net 10.0.0.0 netmask 255.0.0.0 eth0
route add -host 63.100.47.46 eth0
```

²⁰ Not exactly the same; see step 9 below.

```
#Script to add static ARP entries on attack computer. There is  
#no intermediate router and the FW may block ARP. This  
#assures that Nmap will always send the scan packet, regardless  
#of whether the target host is reachable.
```

```
arp -s 192.168.1.8 00A0CC5C1111  
arp -s 192.168.1.9 00A0CC5C1111  
arp -s 192.168.1.10 00A0CC5C1111  
arp -s 192.168.1.11 00A0CC5C1111  
arp -s 192.168.1.24 00A0CC5C1111  
arp -s 192.168.1.25 00A0CC5C1111  
arp -s 192.168.1.32 00A0CC5C1111  
arp -s 192.168.1.35 00A0CC5C1111  
arp -s 63.100.47.46 00A0CC5C1111  
arp -s 10.0.0.3 00065b822222
```

5. Run all shell scripts to prepare the firewall, attack host and target host.
6. In preparation for Step 7, run *tcpdump* on the target computer to capture packets that have been permitted to pass the firewall. The command below allows us to monitor and capture the traffic simultaneously.

```
tcpdump | tee dump.txt
```

7. Use nmap, as described below, to send TCP SYN packets and UDP packets to the host through the firewall. Nmap will be invoked several times, using different spoofed source addresses, to simulate different kinds of traffic.

© SANS Institute 2003, Author retains full rights.

```
#This is the attack script; it invokes Nmap multiple times
#to spoof traffic from various sources.
#
# Notes:
# -sS is a SYN scan
# -sU is a UDP scan
# -P0 means don't ping the host first; just generate requested packets
# -e eth0 is required to spoof source addresses
# -initial_rtt_timeout sets the listen timeout to 1 millisecond
# for maximum performance. This is OK since
# we don't expect replies.
# -S is the spoofed source address
# -p1-1500 tells Nmap to scan ports 1 through 1500. There
# is nothing special in the FW rule set above 1500,
# so this will improve performance without loss of
# completeness.
# The remainder of each line consists of target addresses.

nmap -sS -P0 -e eth0 --initial_rtt_timeout 1 -S 192.168.1.33
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500
nmap -sU -P0 -e eth0 --initial_rtt_timeout 1 -S 192.168.1.33
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500

nmap -sS -P0 -e eth0 --initial_rtt_timeout 1 -S 63.100.47.46
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500
nmap -sU -P0 -e eth0 --initial_rtt_timeout 1 -S 63.100.47.46
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500

nmap -sS -P0 -e eth0 --initial_rtt_timeout 1 -S 217.23.165.75
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500
nmap -sU -P0 -e eth0 --initial_rtt_timeout 1 -S 217.23.165.75
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500

nmap -sS -P0 -e eth0 --initial_rtt_timeout 1 -S 192.168.1.1
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500
nmap -sU -P0 -e eth0 --initial_rtt_timeout 1 -S 192.168.1.1
192.168.1.8,9,10,11,24,25,32,35 63.100.47.46 10.0.0.3 -p1-1500
```

8. Examine `/var/log/messages` file on the firewall, and `dump.txt` on the target computer to assess the function of the firewall, and fill-in the appropriate summary information in the Firewall Assessment Worksheet.
9. Fine tuning: It's possible that misconfigured networking components could also prevent packets from reaching the target computer; if this were to happen, the firewall test would be invalid. To verify that permitted packets *would* pass from attacker to target, we add two temporary firewall rules to permit passage of all

TCP and UDP traffic with destination port of 1499. In all tests, we should see these packets arrive at the target computer. (port 1499 was chosen simply because it has no special significance to the real rule set) Another way to test network configuration is to perform an initial test with an empty rule set and a policy of ACCEPT; once network configuration is verified, we would switch to the real firewall rule set. However, I prefer the method used here because it leaves less room for procedural error.

10. Repeat Steps 6-8 for each possible configuration of attack and target computers, as listed below:

Attack	Target
Eth0	Eth1
Eth0	Eth2
Eth1	Eth0
Eth1	Eth2
Eth2	Eth0
Eth2	Eth1

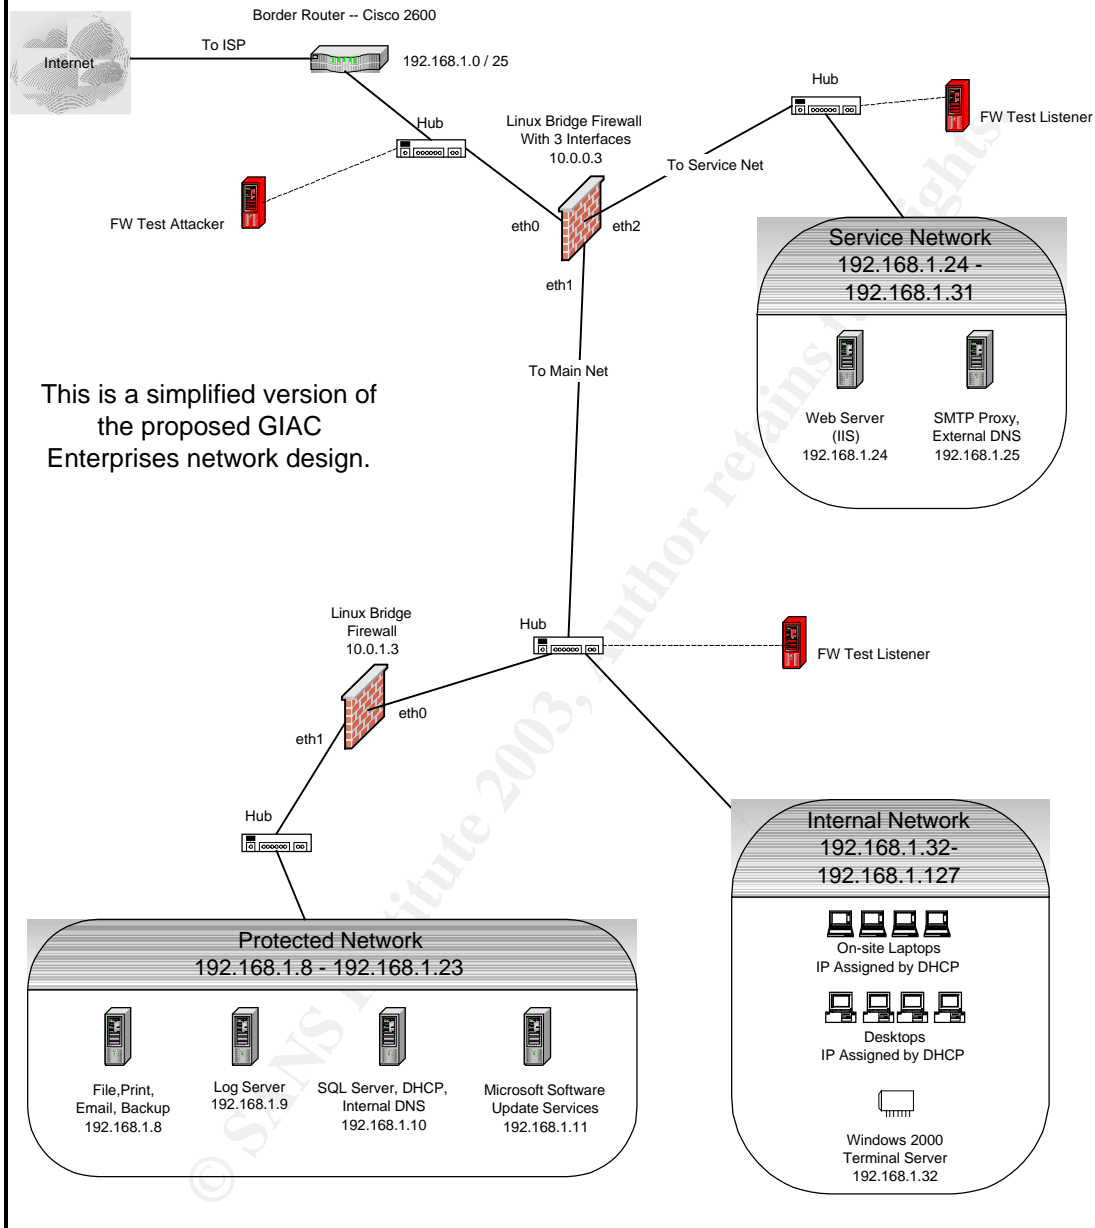
Stage 2: Live Testing Procedure

The live testing procedure is similar to the bench test procedure, except that we will not need to create proxy IP addresses on the target host. Instead, we will install simple passive listeners behind two of the interfaces of the firewall; we could even use existing IDS systems as listeners.

1. Build the attack computer. This can be the same one used in the bench tests.
2. Build two listener computers: these are generic Linux computers with one Ethernet card and capacity to run *tcpdump*; these are shown in red in the diagram.
3. Install the attack computer and listeners as shown in the diagram below.
4. Launch *tcpdump* on the listeners and redirect their output to a file.
5. Run the same *nmap*²¹ commands used in the bench test.
6. Collect the *tcpdump* files from the listeners and analyze as for the bench test.
7. Perform functional testing of each major service component, e.g. SMTP, DNS, Exchange e-mail, SQL Server.

²¹ Be sure to disable IDS systems before launching the attack.

GIAC Enterprises Firewall Test Network Diagram



Evaluate the Audit

Analysis of Test Results

First, we examine the log files from the bench test. These files are much too large to reproduce here, so excerpts are given below, with commentary.

Excerpts from /var/log/messages with commentary below.

These results were generated with the attack machine on Eth0 (from border router) and the Target machine on Eth1 (Internal Network). These examples represent only a very small part of the traffic that would be examined to verify the firewall rule set.

```
Dec 14 18:34:22 localhost kernel: SPOOF IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth2 SRC=192.168.1.33 DST=192.168.1.8 LEN=40 TOS=0x00  
PREC=0x00 TTL=39 ID=60832 PROTO=TCP SPT=57001 DPT=146 WINDOW=4096  
RES=0x00 SYN URGP=0
```

```
Dec 14 18:34:22 localhost kernel: SPOOF IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth1 SRC=192.168.1.33 DST=192.168.1.8 LEN=40 TOS=0x00  
PREC=0x00 TTL=39 ID=60832 PROTO=TCP SPT=57001 DPT=146 WINDOW=4096  
RES=0x00 SYN URGP=0
```

^^ Note that the two above entries are nearly identical: in fact, they refer
^^ to the same packet, and differ only in the PHYSOUT value. This is
^^ an indication that the bridge firewall has flooded the outbound frame
^^ since it has never received any frames from the destination.
^^ The IP packet is recognized as SPOOFed because it has source address
^^ of the internal network, but was received as external traffic.

```
Dec 14 18:34:25 localhost kernel: TESTOK IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth2 SRC=192.168.1.33 DST=192.168.1.8 LEN=40 TOS=0x00  
PREC=0x00 TTL=39 ID=792 PROTO=TCP SPT=57001 DPT=1499 WINDOW=4096  
RES=0x00 SYN URGP=0
```

^^ The above packet was accepted and labeled “ TESTOK “: for the purposes
^^ of bench testing, all traffic to port 1499 is accepted. We expect to see all such
^^ packets arrive successfully at the target machine, which will demonstrate that
^^ the test equipment is properly set up and configured. See step 9 in the Bench
^^ Testing section above for additional explanation.

```
Dec 14 18:36:33 localhost kernel: SPOOF IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth2 SRC=192.168.1.33 DST=192.168.1.8 LEN=28 TOS=0x00  
PREC=0x00 TTL=48 ID=55075 PROTO=UDP SPT=52072 DPT=468 LEN=8
```

^^ Spoofed UDP traffic


```
Dec 14 18:36:20 localhost kernel: FWCONN IN=br0 PHYSIN=eth0 OUT=  
MAC=00:06:5b:82:22:22:00:c0:4f:55:33:33:08:00 SRC=192.168.1.33 DST=10.0.0.3  
LEN=40 TOS=0x00 PREC=0x00 TTL=39 ID=20783 PROTO=TCP SPT=57001  
DPT=146 WINDOW=4096 RES=0x00 SYN URGP=0
```

^^ Attempt to connect to the firewall's IP address (10.0.0.3) from an
^^ unauthorized segment or IP address. In this case, this packet should
^^ fail on both criteria.

```
Dec 14 18:38:55 localhost kernel: BAD DEST ADDR IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth2 SRC=63.100.47.46 DST=63.100.47.46 LEN=40 TOS=0x00  
PREC=0x00 TTL=52 ID=31554 PROTO=TCP SPT=43076 DPT=580 WINDOW=1024  
RES=0x00 SYN URGP=0
```

^^ The destination address is not in our valid address space; packet is dropped.

```
Dec 14 18:39:59 localhost kernel: BADGUY IN=br0 PHYSIN=eth0 OUT=br0  
PHYSOUT=eth2 SRC=217.23.165.75 DST=192.168.1.8 LEN=40 TOS=0x00  
PREC=0x00 TTL=59 ID=32213 PROTO=TCP SPT=56703 DPT=1200 WINDOW=4096  
RES=0x00 SYN URGP=0
```

^^ The source address is on our "BADGUY" list; packet is dropped.

```
Dec 14 18:45:32 localhost kernel: FWCONN IN=br0 PHYSIN=eth0 OUT=  
MAC=00:06:5b:82:22:22:00:c0:4f:55:33:33:08:00 SRC=192.168.1.1 DST=10.0.0.3  
LEN=28 TOS=0x00 PREC=0x00 TTL=43 ID=47582 PROTO=UDP SPT=59663  
DPT=1494 LEN=8
```

>> Attempt to send UDP traffic to the firewall from an unauthorized segment
>> or IP address.

Excerpts from dump.txt. This file contains the output of Tcpdump on the target machine which was generated during the test run. If the firewall rule set is correct, all received packets should be permissible traffic. The vast majority of dump.txt consists of traffic from the router, which is permitted explicitly.

```
18:47:14.139852 192.168.1.33.57001 > 192.168.1.8.1499: S  
3761776015:3761776015(0) win 4096
```

^^ Traffic for port 1499, permitted for test purposes.

```
18:47:14.139905 arp who-has 192.168.1.33 tell 192.168.1.8
```

^^ The target machine ARPs for the attackers address to respond.

18:48:19.866526 192.168.1.33.57002 > 192.168.1.25.1499: S
2121515502:2121515502(0) win 4096

^^ Another port 1499 packet.

18:49:46.946107 192.168.1.33.52072 > 192.168.1.32.1499: udp 0

^^ UDP on port 1499.

18:50:53.787899 63.100.47.46.43076 > 192.168.1.24.http: S 896193237:896193237(0)
win 1024

^^ Packet to port 80 on the web server is permitted.

18:51:04.891912 63.100.47.46.43076 > 192.168.1.25.smtp: S
4212159808:4212159808(0) win 1024

^^ Packet to port 25 on the SMTP proxy is permitted.

18:52:31.020972 63.100.47.46.39309 > 192.168.1.25.domain: 0 [0q] (0)

^^ DNS traffic to external DNS server is permitted.

Full examination of dump.txt shows that:

- All traffic for port 1499 was permitted through
- Appropriate traffic for external services was permitted
- No other traffic was permitted

For brevity, I have included only the results from the testing of Eth0 against Eth1.
Testing for other scan/target segment pairs would be performed in the same manner,
and all results documented in the table below.

Scanner Location	Target Location	Source Address	Protocol (TCP/UDP)	Result
Eth0	Eth1	217.23.165.75 (bad guy)	TCP	Correct FW behavior
Eth0	Eth1	217.23.165.75	UDP	Correct FW behavior
Eth0	Eth1	192.168.1.1 (router)	TCP	Correct FW behavior
Eth0	Eth1	192.168.1.1	UDP	Correct FW behavior
Eth0	Eth1	192.168.1.33 (generic internal host)	TCP	Correct FW behavior

GCFW Practical
Part V – Assignment 3 – Verify Firewall Policy

Eth0	Eth1	192.168.1.33	UDP	Correct FW behavior
Eth0	Eth1	63.100.47.46 (generic external host)	TCP	Correct FW behavior
Eth0	Eth1	63.100.47.46	UDP	Correct FW behavior

Eth0	Eth2	217.23.165.75	TCP	
Eth0	Eth2	217.23.165.75	UDP	
Eth0	Eth2	192.168.1.1	TCP	
Eth0	Eth2	192.168.1.1	UDP	
Eth0	Eth2	192.168.1.33	TCP	
Eth0	Eth2	192.168.1.33	UDP	
Eth0	Eth2	63.100.47.46	TCP	
Eth0	Eth2	63.100.47.46	UDP	

Eth1	Eth0	217.23.165.75	TCP	
Eth1	Eth0	217.23.165.75	UDP	
Eth1	Eth0	192.168.1.1	TCP	
Eth1	Eth0	192.168.1.1	UDP	
Eth1	Eth0	192.168.1.33	TCP	
Eth1	Eth0	192.168.1.33	UDP	
Eth1	Eth0	63.100.47.46	TCP	
Eth1	Eth0	63.100.47.46	UDP	

Eth1	Eth2	217.23.165.75	TCP	
Eth1	Eth2	217.23.165.75	UDP	
Eth1	Eth2	192.168.1.1	TCP	
Eth1	Eth2	192.168.1.1	UDP	
Eth1	Eth2	192.168.1.33	TCP	
Eth1	Eth2	192.168.1.33	UDP	
Eth1	Eth2	63.100.47.46	TCP	
Eth1	Eth2	63.100.47.46	UDP	

Eth2	Eth0	217.23.165.75	TCP	
Eth2	Eth0	217.23.165.75	UDP	
Eth2	Eth0	192.168.1.1	TCP	
Eth2	Eth0	192.168.1.1	UDP	
Eth2	Eth0	192.168.1.33	TCP	
Eth2	Eth0	192.168.1.33	UDP	
Eth2	Eth0	63.100.47.46	TCP	
Eth2	Eth0	63.100.47.46	UDP	

Eth2	Eth1	217.23.165.75	TCP	
Eth2	Eth1	217.23.165.75	UDP	
Eth2	Eth1	192.168.1.1	TCP	

Eth2	Eth1	192.168.1.1	UDP	
Eth2	Eth1	192.168.1.33	TCP	
Eth2	Eth1	192.168.1.33	UDP	
Eth2	Eth1	63.100.47.46	TCP	
Eth2	Eth1	63.100.47.46	UDP	

Recommended changes:

1. The current rule set accepts all traffic from the router, as seen in the Tcpcdump file on the target machine. If the router were compromised, or if an attacker were able to spoof the router's IP address, the entire network would be open. Therefore, we should investigate which router traffic is needed and adjust the firewall rule set to block all other traffic.
2. The Tcpcdump file on the target machine shows that packets for GIAC Enterprises' external services (SMTP, DNS, HTTP) were accepted. However, the target machine was located on the *internal* network, not on the *service* network. To be more secure, we should adjust the firewall rules for external services to include the Netfilter "`-o eth2`" option to prevent flooded packets from reaching the internal segment. This issue is covered in more detail in the Bridge Firewall Tutorial in Assignment 2.

© SANS Institute 2003, Author retains full rights.

Part VI Assignment 4 – Design Under Fire

Common motivations for attack²²

1. Use internal facilities, perhaps for another attack
2. Denial of Service
3. Information theft
4. Amusement – mostly script kiddies
5. Practice


I will attack James Giesecke's design²³. His network diagram is on the following page.

²²Zwicky, Elizabeth D., et al, Building Internet Firewalls, Second Edition. O'Reilly, 2000 pp 7-11.

²³ Giesecke, James. "GIAC Fortune Saying Enterprises". 28 Oct 2002.

URL:http://www.giac.org/practical/James_Giesecke_GCFW.doc. (9 January 9, 2003)

GCFW Practical
Part V – Assignment 4 – Design Under Fire



The purpose of Assignment 4 *“is for the student to clearly demonstrate that they understand that firewall and perimeter systems are not magic ‘silver bullets’ immune to all attacks.”*

Attack #1: Attack on the Firewall

The goal is to subvert the firewall, leaving the network open to further attack. I will assume that this intruder is attempting to steal information, rather than just to inflict a denial of service. A flaw must be found to cause a firewall to fail. I have checked the following resources for known bugs in the Cisco PIX line.

- www.securityfocus.com
- www.sans.org
- www.cisco.com
- cve.mitre.org

There are no known flaws in the last two years that will permit a remote attacker to gain control of a reasonably configured PIX, so the question remains: without uncovering a new vulnerability, how can an attacker force the firewall out of service to permit access to the inside the network? A potential answer is Kevin Mitnick’s principle of *social engineering*. Since GIAC Enterprises is a sales-oriented company, and is heavily dependent on their Internet presence for revenue, the sales force and president of the company will want to keep the pipes open at all times. We will use this against them.

The Plan (detailed procedure follows)

1. Find a DoS vulnerability in the PIX.
2. Employ the vulnerability to make the PIX look unstable.
3. If all goes well, the sales force will complain about lost commission income, and the President of the company will insist that business continuity is more important than a “small” risk of attack. The IT manager will be instructed to take the PIX off-line until it can be “fixed”. The entire networking group will scramble to fix the problem while keeping the Internet connection open.
4. Perform very low-level scans, disguised as innocuous-looking traffic. When the firewall is disabled it will become possible to communicate with internal hosts.

Note: I have examined the approaches taken in several other practicals; in most cases, the exploit was a DoS against the FW. It’s my understanding that the point of this attack is to subvert rather than DoS the FW. I do employ a DoS attack as *part* of the overall attack plan; however, most DoS’s require a large volume of attack traffic, which would be easily detectable in the logs. The DoS used here requires just one malformed connection to cause the PIX to restart, making it much harder to identify the PIX restart as a DoS rather than a hardware problem.

The Procedure

1. James' design permits SSH traffic from partners and suppliers via SSH. Also, I was unable to find any specific rule that prevents SSH access to the PIX from the VPN. Cisco has recently announced a Denial of Service attack using SSH that should fit the need. The details are given at <http://www.cisco.com/warp/public/707/ssh-packet-suite-vuln.shtml>. This attack will cause the PIX to reboot, making it inactive for several minutes. This behavior could look very much like a faulty piece of equipment, as required by the plan.

Important note: James' design permits SSH *only* from suppliers and partners. Therefore, this exploit *assumes* that the attacker has control of a host that belongs to one of these two business associates. A technique to accomplish this is given in the third part of this assignment.

2. I have not been able to find a tool to implement this attack, nor have I found any specific documentation describing the specific TCP structure required to force the restart: this is not surprising for such a new vulnerability. The description of the vulnerability indicates that it should be fairly simple to implement, so a real attacker would likely monitor the hacker sites for the emergence of a tool or documentation.

Assuming that an attack tool had been found, the attacker would *not* flood the firewall with the exploit connections; instead, he would send them at intervals just slightly longer than the time required for a PIX restart. He might also try to hide the exploit packets among other, seemingly innocuous traffic. The goals are:

- Keep the PIX administrator(s) busy trying to stabilize the firewall so that they neglect to check their logs and notice the SSH traffic directed to the PIX.
 - Interrupt e-commerce so the President of the company starts leaning on the IT manager to "do something, fast!"
3. If the operations in Step 2 are successful, the IT manager will feel a great deal of pressure to get the IP packets (and money) flowing again. At this point, several different things might happen:
 - a. If IT has a spare PIX on hand, they might drop it in as a replacement; the attack would continue against the new PIX, but the risk of discovery would be higher (what are the chances of two *broken* PIX's?)
 - b. If there is no spare PIX, the IT manager may have his network team install a simple router to replace the routing functions of the PIX. It would not, however, replace its full stateful firewalling capability.
 - c. If no spare routers are available, the border router, a Cisco 3640, might have enough unused interfaces to handle the routing duties.


It's important to notice that the attacker need not understand the internal subnet structure of the target since the goal is quite simple: just make the router look unstable, then hope it's disabled, leaving a clear path to the inside.

- Options (b) and (c) are precisely what the attacker wants to see. Recall that the suppliers and partners have legitimate SSH access to some of GIAC Enterprises' hosts; with the PIX in place, the attacker would be limited to port 22 on these machines, so if he can access ports other than 22 he will know that the PIX has been removed. He is then free to attempt further attacks inside GE.

Analysis of Attack #1

Countermeasures: This attack exploits three flaws: (1) the SSH bug in the PIX; (2) the fact that the firewall has not been specifically configured to limit SSH to the firewall only from a trusted management host; (3) presumed weakness in the judgment of IT staff, who permitted the removal of the firewall. All of these flaws can be eliminated.

- Cisco does not make it particularly easy to download software updates: even using a valid CCO login I received the following error message when attempting to download *free* updates for the PIX:



- Nevertheless, it's essential to keep up-to-date on software bug fixes and configuration adjustments necessary to avoid well-known exploits.
- Before anything else, the firewall itself must be properly secured. Even an excellent rule base can be subverted if the firewall is open to attack.
 - Every IT group should develop an Incident Response Plan, which would provides for business continuity in case of catastrophic failure of important security components. This plan should have the informed support of top management. In short, GE's IT team should never have to scramble to decide what to do about their crashing PIX firewall: they should have a written procedure to follow.

Is this attack realistic? Certainly the DoS attack against the firewall is realistic. The real question is whether an IT manager would be persuaded to remove the firewall. In a security-conscious organization the idea of removing the firewall would not even receive serious consideration; however, I have seen organizations remove or neglect security

for reasons much less important than business continuity. In short, this attack is realistic if the target company doesn't fully appreciate the importance of security. A firewall is only a tool: it has limitations, and it must be used skillfully for maximum benefit.

Attack #2: Denial of Service from 50 cable modem/DSL zombies

Although James' practical does not explicitly say so, it appears that his GIAC Enterprises has a T3 connection to the Internet. Even 50 cable modem/DSL zombies can't fill that pipe with random traffic, so I'll have to find a less mundane way to tie up their services.

Choice of Target

Since this attacker can't simply overwhelm a T3 connection, a more specific target must be found: GIAC Enterprises does its business through its web servers, so it seems obvious that a DDoS against GE's web servers would inflict the greatest disruption to business.

The Vulnerability

In June of 2002, a flaw was discovered in the way Apache httpd handled certain types of requests (chunked encoding). The Apache.org bulletin can be found at http://httpd.apache.org/info/security_bulletin_20020620.txt. On *nix machines, this flaw can be exploited to produce a denial of service on the web server, or, in some cases, the ability to run arbitrary code on the Apache host "with the permissions of the web server child process". I found a portable Perl script at Packetstorm which exploits this flaw as a DoS; it can be found at <http://packetstormsecurity.org/0206-exploits/apache-dos.pl>. There are ports of Perl for the Windows OS, so there should be no problem adapting this as the DDoS attack tool. In fact, only a handful of zombies would be sufficient to disable all of GIAC Enterprises' Apache servers.

The Procedure

1. Gather IP addresses of GE's web servers
2. Distribute the attack tool to all the zombies
3. Instruct the zombies to launch the attack against the web servers

There is only one catch, however: James has chosen to put Squid servers in front of his Apache servers. It's possible that Squid would alter the http request enough to thwart the attack. The attacker can easily overcome this problem after noticing that the Apache servers are used for e-commerce; therefore, the servers must support SSL. James' Squid servers can't proxy SSL traffic, but must simply pass it through; he can essentially bypass the Squid servers.

Analysis of Attack #2: Countermeasures

The attack exploits a flaw in Apache. The only sure way to defend against it is to keep such critical software up to date. Furthermore, IDS systems will not be able to detect an attack wrapped inside SSL; the encryption effectively hides the attack. Perhaps it would be possible to determine which IP addresses were involved in the attack; they could be banned at the firewall. This might be difficult if the server is very active. Also, Squid versions 2.5 and above can terminate SSL connections, which means that it's no longer necessary to perform pass-through on SSL connections: the Squid server can decrypt the https request, and perform its usual services.

Attack #3: Compromise an Internal System

The vulnerability from Attack #2 could also be used to compromise an internal *nix web server, but I'll choose a new exploit to make things a little more interesting.

The Vulnerability

On December 16 2002, eEye Digital Security announced vulnerability in Macromedia's Flash software. The details of this vulnerability can be found at <http://www.macromedia.com/v1/handlers/index.cfm?ID=23569>. A specially crafted Shockwave Flash format file can cause a buffer overflow in Flash, potentially leading to control of the host system. According to SANS, http://www.sans.org/newsletters/cva/cva1_22.php, no exploits are known to exist, but the discoverers of the vulnerability (researchers from eEye Digital Security) provide limited technical details and assert that the flaw is easy to exploit". It's not clear whether the malicious code would run at the security level of the logged-in user, or at Administrator level; in the case of Windows 95/98/ME these are equivalent. Also, many organizations grant local Administrator rights to users on Windows NT, 2000 and XP.

This is a particularly insidious vulnerability for several reasons:

- Flash is very widely deployed
- Defects in Flash do not receive the same attention as do flaws in Microsoft products, so IT staff and end-users are less likely to install security updates
- Malicious SWF files can be delivered from a web site or as an e-mail attachment

Choice of Target

We have two options here:

- Typical internal Windows 2000 clients; exploit code would be delivered through a web site or by e-mail
- Laptop used by remote sales staff; delivery could be accomplished as above, or by personal contact with the machine

For the purposes of this assignment, I'll assume that it's delivered to an internal client via malicious web site.

The Procedure

1. Find the name and e-mail address of one or members of the sales force at GIAC Enterprises; this should not be difficult.
2. Send enticing unsolicited e-mail to addresses obtained above; include a link to a "GREAT NEW PRODUCT TO BOOST YOUR SALES POTENTIAL!!!!", located on a web site under the attacker's control.
3. Since GE's network is not configured to perform outbound http proxy, the internal employee will be able to download the malicious code without impediment or alarm.
4. When the SWF file executes, it can launch the malicious code and take over the user's machine.

Once the buffer overflow and exploit has been accomplished, the malicious code could do many things, including:

- Install a DDoS zombie
- Install a remote control tool such as BO2K
- Install a keystroke logger to collect passwords (e.g. SilentLog: <http://packetstormsecurity.org/Win/SilentLog.zip>)
- Install an IP proxy or tunnel to bypass the firewall, such as netcat, zebedee or stunnel

Even if the user does not have local Admin rights to the Windows box, it should still be possible to install an IP tunnel/proxy, which will at least permit the attacker to bypass the firewall for further attacks. Furthermore, it might also be possible to exploit a flaw in Microsoft's WM_Timer message handler if there is an application running that does not trap WM_Timer messages. If so, this would permit the attacker to gain System level privilege. The details can be found here: <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS02-071.asp>.

Analysis of Attack #3:

The attack exploits a flaw in Macromedia's Flash player. This demonstrates that even seemingly innocuous software components can expose serious vulnerabilities. Obviously, it's essential that to keep up to date on security patches; but it's also important to consider which components should be installed on the client computers: if it doesn't serve the business need, perhaps it should not be installed.

This exploit could benefit an attacker in several ways:

1. Bypass the firewall for further attack through an IP proxy or remote control (e.g. BO2K)
2. Sniff passwords for attack against internal systems
3. Serve as a platform for attack on other outside systems

Countermeasures and Mitigation

- Install a proxy for outbound http. This would permit blocking of some undesirable downloads. If the proxy permits, virus/trojan scanning could also be performed before the downloaded file is given to the client.
- Install a virus scanner for incoming e-mail
- Don't install unnecessary software
- Monitor client computers for inappropriate accounts or privilege levels
- Re-image client machines regularly

Lessons Learned from Design Under Fire

1. Defense in Depth: no single security component or procedure can thwart all possible attacks. A layered defense is essential: security must be imposed on the perimeter, internal subnets and the hosts.
2. Security is not static: software must be updated to address new vulnerabilities, and IT must conduct regular audits to verify that the security structure properly implements policy.
3. A dedicated attacker can probably get in, so every security policy should include appropriate incident-handling procedures.

© SANS Institute 2003. All rights reserved. Author retains full rights.

Appendix A – Defense in Depth and Host Security

As demonstrated in Assignment 4, building a solid firewall does not provide perfect protection against attack. One of the primary lessons of the SANS Track 2 course is to preach *Defense in Depth*. In keeping with that principle, I'd like to include a little discussion on the elements that were not included in the responses to the assignments.

1. Router Configuration

Routers perform two functions: transporting packets between networks and packet filtering. Their essential functions also make them prime targets for exploit, so they must be properly secured.

- Disable unneeded services
- Use ACL's to disable remote (external) management of the router
- Use ACL's to limit management access to a few (or one) management station
- Keep router software up to date
- Monitor security and vendor notification sites for known vulnerabilities
- Monitor router logs for attempted exploits

2. Host Security

Hosts are the repositories of valuable data. If compromised, they can be used to launch further attacks internally or externally. The goal of host security is to prevent successful attacks that have not been blocked by the routers and firewalls. However, if an attack is successful, we'd like to know about it immediately. In the worst case, a host's data may be destroyed or corrupted beyond use.

- Keep host-based software updated to prevent exploit of known attacks.
- Consider installation of host-based firewalls or similar access-control tools, such as TCP Wrappers, to restrict access to sensitive services
- Install personal firewalls for remote clients
- Consider installation of host-based intrusion detection systems
- Install filesystem integrity checkers, such as TripWire or AIDE
- Perform regular, hands-on security audits
- Perform regular backups, and store backup media safely

3. Network Intrusion Detection

If an attack gets through the firewall, it's better to know about it sooner rather than later. Additional information about network intrusion detection is given in the next appendix.

- Install one or more network intrusion detection devices; make sure they're powerful enough to handle the traffic flow
- Keep their attack signature databases current
- Configure alerting for critical events
- Review logs regularly

4. Centralized Logging

Centralized logging provides two major benefits: first, if a device is compromised, the attacker can destroy the local security log, but not the central log. Second, it collects all security events in one location for easier correlation, analysis and reporting. Ideally, there should be at least two central log servers. The next appendix covers an interesting approach to centralized logging. One other essential element of centralized logging is centralized time synchronization. I have not indicated a time source on the GIAC Enterprises network, but it would be a very helpful addition.

5. Log Analysis

Security logs are useless unless someone looks at them. Routine review of "clean" logs also improves the network administrator's understanding of traffic flow in the network. This may lead to improved security policies and more appropriate configuration of security components.

© SANS Institute 2003. All rights reserved.

Appendix B – Stealth NIDS and Logging (Logging for the Paranoid)

My network design for GIAC Enterprises includes components labeled “Stealth Log Server” and “IDS on Linux No IP Address”. The purpose of this appendix is to describe these components, and why they’re helpful to the design.

Stealth Logging

A centralized log server can be a tremendous help to an enterprise with many Syslog clients. A well-placed and defended Syslog server is unlikely to be compromised by a script kiddie, but it’s a prime candidate for attack by a serious intruder. Some network administrators use “back-channel logging” to serve high-value resources. Typically, a back-channel is either a separate, isolated network (using a second NIC) or a serial interface. Both of these methods are effective, but have some shortcomings:

- Both methods require a second wiring scheme, which requires a lot of extra work if the clients are not near each other
- If a Syslog client were compromised, it’s possible that the intruder could gain access to the secondary network.
- For a serial back-channel, a network administrator would need a separate serial interface for each Syslog client. This could be a problem if there are many clients.

In many cases these problems are not too difficult to overcome. However, there’s another way to build an isolated Syslog server without building a back-channel.

The basic theory of a stealth log server is that the clients would be configured to send Syslog UDP packets to a bogus IP address. A static ARP entry would be needed to ensure that the packet would be sent. The Syslog server is located on the appropriate Ethernet segment; it would have no IP address, but it’s Ethernet interface would operate in promiscuous mode. The stealth Syslog server would sniff all traffic on the segment, and capture Syslog packets. The capture process could use one of several popular tools, for example:

- Snort running in promiscuous mode: the Snort server would have no IP address, and a rule would be created to transcribe Syslog packet payloads. I believe the simplest command to do this is

```
snort -d udp dst port 514
```

This will dump the application layer, which contains the raw Syslog message.

- TCPdump²⁴ generally records just the packet header; however, it can be persuaded to dump the entire packet instead.

```
tcpdump -i eth0 -s 1032 -w '-' udp dst port 514 >dumpfile.out
```

The expression `-i eth0` is necessary if there is no IP address assigned to the host.

For Snort, the output file will contain the raw Syslog packet. For TCPdump, the file will contain a UDP header as well. To make use of the log data, the sysadmin would need a Perl script (or something similar) to parse out the distinct Syslog fields. The technical information on the Syslog format can be found in RFC 3164. The link is given in the next appendix.

Other Variations

If there were a router between the Syslog client and server, it would be necessary to add static ARP entries on the intervening router instead of the Syslog client. One way to avoid this problem is to install a standard Syslog server, including a correct IP address. Then a stealth Syslog server is connected via a hub (*not* a switch!) to the primary Syslog server. The stealth server would be configured to log packets sent to the primary server. This is the configuration used in my network design.

²⁴ The TCPdump method occurred to me shortly after I attended SANS training. I later read about the Snort method in Bauer, Mick. "Stealth Logservers". Linux Journal. December 2001.
URL: <http://www.linuxjournal.com/modules.php?op=modload&name=NS-lj-issues/issue92&file=5476s2>. (9 January 9, 2003)

Appendix C – References

From the Practical

Zwicky, Elizabeth D., et al, Building Internet Firewalls, Second Edition. O'Reilly 2000.

SANS Institute. Firewalls, Perimeter Protection and VPN's. SANS Institute. 2002

Kivinen, T., et. al. "Negotiation of NAT-Traversal in the IKE". 24 June 2002.
URL: <http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ipsec-nat-t-ike-03.txt> (9 January 2003).

Newman, David. "Crying wolf: False alarms hide attacks". NetworkWorldFusion. 24 Jun 2002. URL: <http://www.nwfusion.com/techinsider/2002/0624security1.html> (9 January 2003)

Microsoft Corp. "Software Update Services Overview White Paper". 20 June 2002.
URL: <http://www.microsoft.com/windows2000/windowsupdate/sus/susoverview.asp>. (9 January 2003).

Vestergaard, Peter. "FIREWALLS, PERIMETER PROTECTION AND VPNS PRACTICAL ASSIGNMENT". 26 Oct 2001. URL:
http://www.giac.org/practical/Peter_Vestergaard_GCFW.zip. (9 January 2003)

Hofman, Mark. "Your Fortunes". August 2001. URL:
http://www.giac.org/practical/Mark_Hofman_GCFW.zip (9 January 2003)

Russell, Rusty. "Linux 2.4 Packet Filtering HOWTO". 24 January 2002. URL:
<http://www.netfilter.org/documentation/HOWTO/packet-filtering-HOWTO.html> (9 January 2003)

Giesecke, James. "GIAC Fortune Saying Enterprises". 28 Oct 2002.
URL: http://www.giac.org/practical/James_Giesecke_GCFW.doc. (9 January 9, 2003)

Bauer, Mick. "Stealth Logservers". Linux Journal. December 2001.
URL: <http://www.linuxjournal.com/modules.php?op=modload&name=NS-lj-issues/issue92&file=5476s2>. (9 January 9, 2003)

Sedayao, Jeff. Cisco IOS Access Lists (O'Reilly, 2001)

von Braun, Joakim. "What port numbers do well-known trojan horses use?"
SANS Intrusion Detection FAQ. URL:
<http://www.sans.org/resources/idaq/oddports.php> (9 January 2003)

Cisco, Inc. "Cisco Security Advisory: SSH Malformed Packet Vulnerabilities". 20 December 2002. URL: <http://www.cisco.com/warp/public/707/ssh-packet-suite-vuln.shtml> (9 January 2003)

Apache Team. "Updated Advisory". 20 June 2002.
http://httpd.apache.org/info/security_bulletin_20020620.txt (9 January 2003)

Wong, Luis. "apache-dos.pl". Packetstorm Security. 24 June 2002. URL:
<http://packetstormsecurity.org/0206-exploits/apache-dos.pl> (9, January 2003)

Macromedia, Inc. "MPSB02-15 - Macromedia Flash Malformed Header Vulnerability Issue". 12 December 2002.
URL: <http://www.macromedia.com/v1/handlers/index.cfm?ID=23569> (9 January 2003)

SANS Institute. "Macromedia Flash Player versions less than 6.0.65.0". SANS Critical Vulnerability Analysis. 22 December 2002. URL:
http://www.sans.org/newsletters/cva/cva1_22.php (9 January 2003)

Anonymous. "Silentlog". Packetstorm Security. 11 Feb 2002. URL:
<http://packetstormsecurity.org/Win/SilentLog.zip>. (9 January 2003)

Microsoft, Inc. "Microsoft Security Bulletin MS02-071". 11 December 2002.
URL:
<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS02-071.asp> (9 January 2003)

Other Interesting Links

Linux Bridge Home Page: <http://bridge.sourceforge.net/>

Netfilter Home Page: <http://www.netfilter.org/>

SANS Top 20 Vulnerabilities: <http://www.sans.org/top20/>

National Security Agency
Security Recommendation Guides for Windows 2000, XP, Cisco Routers, E-Mail:
<http://www.nsa.gov/snac/>

National Security Agency, Security Enhanced Linux:
<http://www.nsa.gov/selinux/index.html>

NIST Systems Administration Guidance for Windows 2000 Professional:
http://csrc.nist.gov/itsec/guidance_W2Kpro.html

Microsoft Lockdown Guides:
<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/lockdown.asp>

Security and Intrusion Tools

Security Testing Tools: <http://www.sans.org/top20/tools.pdf>

Nmap.Org: <http://www.nmap.org/>

Nessus Vulnerability Scanner: <http://www.nessus.org/>

Packetstorm Security: <http://www.packetstormsecurity.org/>

Netcat: <http://netcat.sourceforge.net/>

Zebedee: <http://www.winton.org.uk/zebedee/>

BO2K: <http://bo2k.sourceforge.net/>

Stunnel: <http://bo2k.sourceforge.net/>

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

Syslog: <http://www.ietf.org/rfc/rfc3164.txt>

Security and Internet References

Well-Known Port Numbers: <http://www.iana.org/assignments/port-numbers>

Foundstone: <http://www.foundstone.com/>

Neohapsis: <http://www.iana.org/assignments/port-numbers>

SecurityFocus: <http://www.securityfocus.org/>

NTBugTraq: <http://www.ntbugtraq.com/>

Incidents.Org: <http://isc.incidents.org/>

Cisco System: www.cisco.com

Common Vulnerabilities and Exposures: <http://cve.mitre.org>