Global Information Assurance Certification Paper

Dominique Galland

SANS Monterey 2000

Practical Assignment for Skill Track 2:
Firewalls, Perimeter Protection and Virtual Private Networks

Table of contents

Notes:

1) In the examples provided, I have created a fictious company named "Everest". I have decided to use a 10. addressing scheme for all internal hosts. I have also chosen a set of external public addresses which may represent real addresses currently in use on the Internet. Readers are very welcome to use any of the examples provided in this document, but should make sure they replace these addresses beforehand with their own in order to prevent conflicts.

2) For the purposes of this paper, the term "DMZ" (De -Militarized Zone) refers to a screened subnet to which some degree of protection is provided by the firewall. It defers from the definition used in SANS' litterature where such a subnet is called a screened network, and where a DMZ is a network with no protection or only router access control.

1. Assignment 1

The security architecture which has to be designed faces several challenges:
1) Connectivity is a must -have. Since the firm's primary business is on -line sales o n the Internet, any loss of network connectivity would have a huge impact on the firm's earnings.
2) It is also clear that the firm has a deep need to communicate efficiently with its suppliers and partners.
3) One additional challenge is the recently comp leted merger/acquisition. It is essential to provide connectivity to our new unit(s) which may not be located very closely to our offices - and without compromising the security level.
However the planned earnings of the company mean that shortage of cash is not to be considered an issue in this case. All reasonable solutions are within our reach.

In order to implement the proposed requirements, my proposition is described by the diagram below.

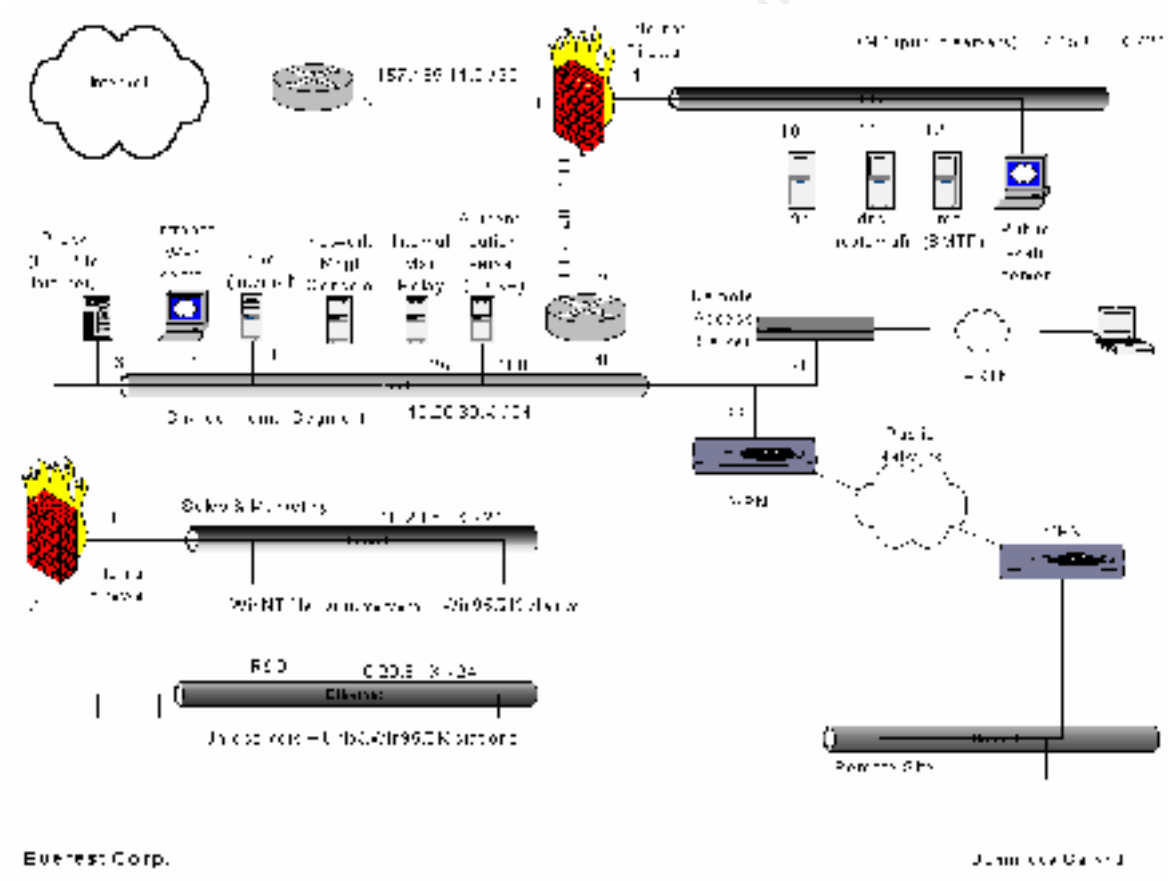

1.1. High-level overview

The internal network uses a private addressing scheme as defined by RFC 1918. All internal addresses belong to the 10.x.x.x/8 segment (variable subnetting is in use within that range).

Internally, the network is divided into 3 main segments: R&D, Sales&Marketing and the shared "public" segment. These 3 segments are separated using a low-end firewall (for example IP Chains on a Linux box) that will provide basic packet filtering based on a simple set of rules.

The Sales&Marketing segment (10.20.32/24) contains a mix of Windows machines (NT, Win95/2K).
The R&D segment (10.20.31/24) contains a mix of Windows (95/2K) and Unix machines.
The shared segment contains servers which are available for all staff.

External connectivity is achieved via 3 solutions:
- an Internet Firewall, to filter all incoming packets to and from the Internet. It also provides filtering for public servers located on the DMZ.
- users can also connect to the company's network using a dial-up remote access server.
- basic connectivity is provided to and from the network of the newly acquired company via a VPN connection. Basic protocols are available through this solution.

1.2. Detailed description

1.2.1. Architecture

1.2.1.1. Internal architecture

The Internal firewall provides packet filtering in order to prevent any unauthorized access to the data and applications used on the R&D as well as Sales&Marketing segments. It will only allow the following services:
- from these 2 segments:
  - http&ssl (tcp 80 and 443) as well as ftp (tcp 20 &21) to the proxy (10.20.30.160) and the intranet web server (.80)
  - dns queries (53 udp & tcp) to the Internal DNS server
  - SMTP (25) access to the Internal mail server (.25)
- to these 2 segments:
  - SNMP and ICMP access from the Network Management platform

1.2.1.2. External connectivity architecture

Connections to the Internet via the Internet Firewall are only allowed from the proxy (all proxied services) and the network management console. The internal dns server will use the external one as a forwarder and the internal mail relay will communicate only with the DMZ mail relay.
Network Address Translation (NAT) will be applied by the firewall to all network traffic originating from the internal network. That means that internal 10. addresses will be replaced with the public address of the firewall (157.159.11.1) as the packets' source addresses, as explained later on.

Users will get connectivity to the public segment from remote locations via a remote access server over a dial up connection on the PSTN. In order to be able to connect, users will have to submit credentials which will be matched against a LDAP database stored on the authentication server. The same credentials can also be used by users from the newly connected subsidiary, which provides connectivity via an encrypted VPN connection over a public network (frame relay or X25).

Limited connectivity is provided from the Internet to specific services on the DMZ: http/ssl traffic is allowed to the public web server, ftp traffic to the ftp server, the same being applicable to the dns and mail servers. The only server on the DMZ allowed to initiate connections to the internal network is the mail relay, for delivery of incoming messages from the Internet.

Network Address Translation (NAT) is a very important part of the architecture. NAT serves two main purposes:
1) it is used to "hide" the addresses of my internal network, so that they can not be seen from the outside.
2) when the internal addresses are not routable on the Internet (which is the case for Everest's network), NAT is used to provide connectivity to th e Internal for the internal hosts.

It works in the following way:
- Whenever the firewall inspects an outgoing packet, before sending the packet out to the Internet, it will replace the source address of the packet (private) with the address of its own publ ic interface (public). It will store this information in its state table.
- When the reply comes back, the firewall will use the packet's source address (which was the destination as it went out) as well as the tcp sequence numbers for a lookup in its state table to find out what the real origin address was on the private internal network.

1.2.2. Applications

All Unix servers on the DMZ and the internal shared network will be hardened during the installation of the software, before they are connected to the netwo rk. This means that during the installation process:
- once the standard "default" configuration is finished, all services/deamons which are not necessary for the basic operation of the server will be removed
- passwords for all login accounts will be chan ged to non-default values in accordance with the company's password policy.

Only the following applications will have connectivity to the Internet:
- Http and likes (TCP ports 80, 443, 8080, 8081 etc): all Internet connectivity will be established via the Internal web proxy, with a limited set of allowed ports. Users will not be able to connect directly to any external resources.
- DNS: the Internal DNS server will act as the primary server and serve all queries. It will forward requests about external hos ts to the public DNS server on the DMZ. However the external DMZ server will answer queries from the Internet concerning only the servers on the DMZ, it will not be configured to query the internal one.
- Mail: incoming mail will be delivered to the extern al (DMZ) server and queued there. The external server will then deliver mail to the internal one. Outgoing mail will be sent internally to the internal mail server, then pushed to the external one and then to the Internet.
- The network management station will have full ICMP and SNMP access to the Internet as well as to the DMZ in order to allow network management and troubleshooting.

1.2.3. Processes

1.2.3.1. User processes

Updates for the corporate anti -virus software will be distributed to all users on a weekly basis (or on a as-needed basis, depending on the threat level of recent alerts). Additionally, an anti -virus software will be installed on both the internal and external mail servers.
Passwords: they will be required to be at least 8 characters long, with 2 cha racters being non numbers or digits. Password changes for all login user accounts will be required at regular intervals, at least once every 6 weeks.
Training: to maintain a good security awareness, all users will receive regular training about good securi ty practices (once a year, 1 or 2 days training)

Users have the possibility to request additional access to the Internet for specific applications (streaming media, specific web applications etc...). If the business need is real, access will be granted eit her via the proxy or (if proxy access technically not possible) to a workstation installed on the internal shared segment.

1.2.3.2. Network admin processes

The network security administrator must subscribe to different mailing lists (ex: Security Alert Consensus -- Network Computing and the SANS Institute) describing the latest security warnings. More important, he must make sure that he allocates enough time to read all advisories received, so that prompt action will be taken in case some discovered vulnerabilitie s can be patched immediately.
A list of possible mailing lists include:

- lists from Cisco Systems (see http://www.cisco.com)
- lists from advisors (for example Network Computing and the SANS Institute's Security Alert Consensus, see http://www.sans.org)
- lists from anti-virus companies (see http://www.symantec.com)
- lists form software providers (http://www.microsof t.com, http://www.sun.com)

2. Assignment 2

2.1. Blocking of spoofed addresses

As a rule, all private IP addresses should not be allowed to be received from the Internet. These addresses are:

10.0.0.0 to 10.255.255.255 (10/8 prefix)
172.16.0.0 to 172.31.255.255 (172.16/12 prefix)
192.168.0.0 to 192.168.255.255 (192.168/16 prefix)

If such packets were allowed in the internal network, it could easily create havoc because the source addresses would then be considered as internal addresses, and too many such packets could be used for a denial-of-service attack.

This is going to be implemented via a standard ACL on the Internet router:

```
access-list 10 deny 10.0.0.0 0.255.255.255
access-list 10 deny 172.16.0.0 0.15.255.255
access-list 10 deny 192.168.0.0. 0.0.255.255
```

Now, the desired traffic is blocked but if we leave the access list like this, all traffic will be denied because as soon as an access list is implemented on the interface of a router, all traffic is denied by default, and only explicitly allowed packets will go through.
So in the end the ACL will look like:

```
access-list 10 deny 10.0.0.0 0.255.255.255
access-list 10 deny 172.16.0.0 0.15.255.255
access-list 10 deny 192.168.0.0. 0.0.255.255
access-list 10 permit any any
```

Then the access list can be enabled for inbound packets ("in" statement) during the configuration of the public (the one facing the ISP) interface via the command:

```
ip access-group 10 in
```

2.2. Blocking of source-routed packets

This can also be done at the router level. The command to configure it is "no ip source-route". Once this command has been applied to a specific interface of the router, any packet going through this interface (either incoming or outgoing) will be discarded if its route-source field is set.

However, it is not necessary to implement this at the router level because Everest's firewall, CheckPoint FireWall-1, drops by default all IP packets which have IP options set.

The other services will be blocked by the firewall. A ruleset has to be created. Before designing the first rule, it is necessary to define two basic objects which will represent Everest's internal network, and the DMZ, as defined in assignment 1.

Here are 2 screenshots showing what the definition of these objects would look like:

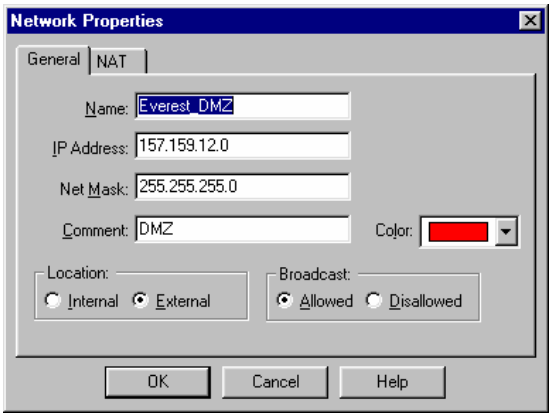

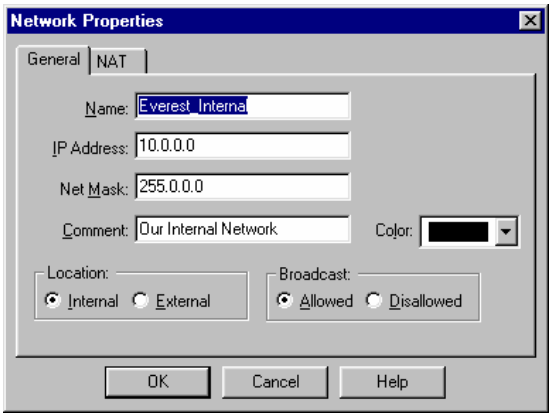

Then, using these objects (and other objects that will be created along the way), the rules can be designed.

2.3. Blocking of login services & NTP

Login services such as telnet, ftp, rlogin etc pose an obvious security risk. All of them, apart from ssh, transmit all passwords in clear text. In case of an attack, they also allow very easy access to the target hosts. Additionally, there is no reason why such services should be allowed to/from the Internet.
The NTP (Network Time Protocol) is used by servers to keep time synchronization accurate. It is not a mandatory service, and an attacker could make his footsteps more difficult to trace if he was able to modify time settings on internal servers.

This rule says that it will drop ("Action") all packets with any source address sent to any destination, for all services listed, whenever these packets are sent ("Time" equals "Any"). The “Track” option specifies that a detailed logging of these actions is required. The "Install On" options is mostly useful when a single station is used to manage different Firewalls, but as this is specific to CheckPoint FW-1, this will not be discussed in detail here.

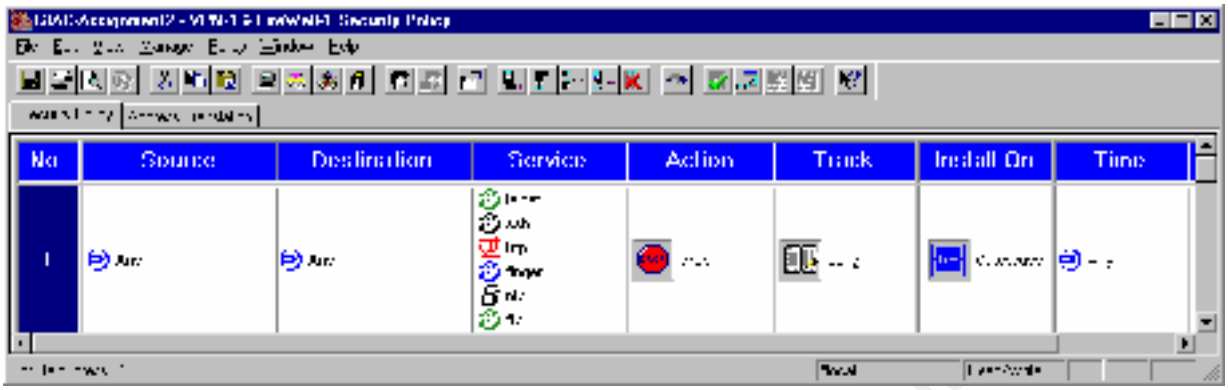

The problem with this rule is that it will prevent the administrators of all systems placed on the DMZ to remotely manage their servers ! To remedy that, we add a new rule BEFORE the deny rule in order to allow ssh to the DMZ, but only from our Internal network.
This is shown in this new screenshot:

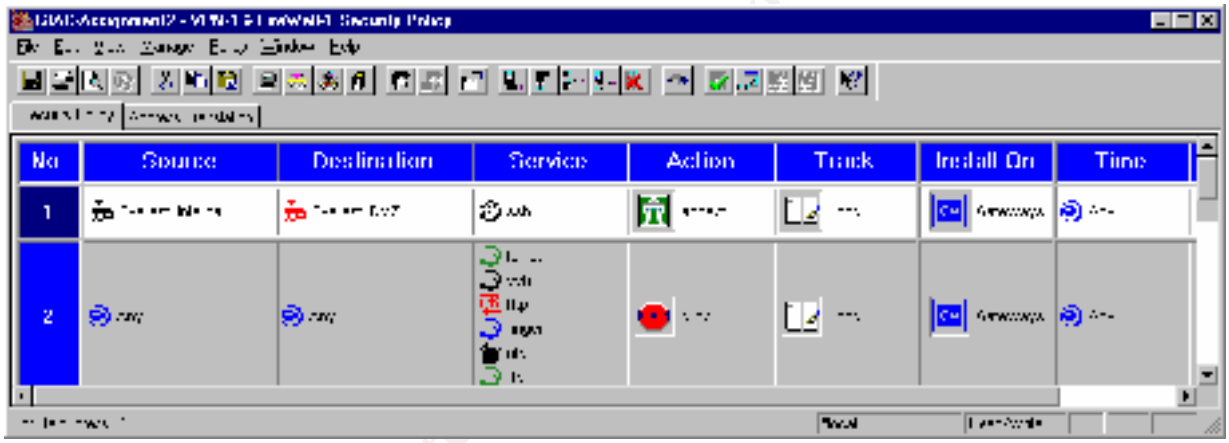

Whenever a packet comes in, the firewall scans all rules sequentially in the order in which they are defined in the ruleset. The first rule that matches source & destination addresses as well as the requested service (whatever the action) will be applied to the packet. So, in this case, any packet from the internal network to the DMZ on port 22 (ssh) will be allowed to go through by rule number 1. Any other packet on port 22 either coming from somewhere else than the internal network, or going to somewhere else than the DMZ, will be dropped by rule 2 (as well as any packets for the other listed services).

Now, as far as ssh and the other listed services are concerned, our networks are secure:
1) Nobody can connect to our internal network from any location
2) ssh to the DMZ is only allowed from our internal network, that means from a 10.x.x.x address. But since we implemented the ACLs on our Internet router to block anything with such an address to come from the Internet, we are sure that packets showing 10.x.x.x as their source address are really coming from our internal network.

An important thing to note is that, even though packets must be transmitted back and forth between the different segments, it is not necessary to add a rule allowing ssh traffic from the DMZ to the internal network, because CheckPoint FW1 uses stateful inspection. This means that once a TCP packet is allowed to go through, the inspection module will remember the source, destination and port numbers used by this connection. When the reply comes back, all parameters will match, and the inspection module will then know that this packet is part of an already established connection (or being established): it will be allowed even if no

specific rule matches the exact source, destination and service of the packet. If ever another rule was added to allow ssh from the DMZ to the internal network in this case, it would allow ssh connections to be initiated from the DMZ towards the internal network, which is something that must be avoided. Of course this is not relevant to UDP where no connecti ons are established.

We can follow the same path and gradually extend our ruleset in order to cover more services:

2.4. Netbios

Windows services will not be included in the "permit" list because we have chosen to make all our servers on the DMZ Unix -based. Additionally, such services are very insecure and provide easy ways to compromise systems. They should always be confined to a firewalled network.
Everything can be blocked altogether. Here is what the rule would look like:

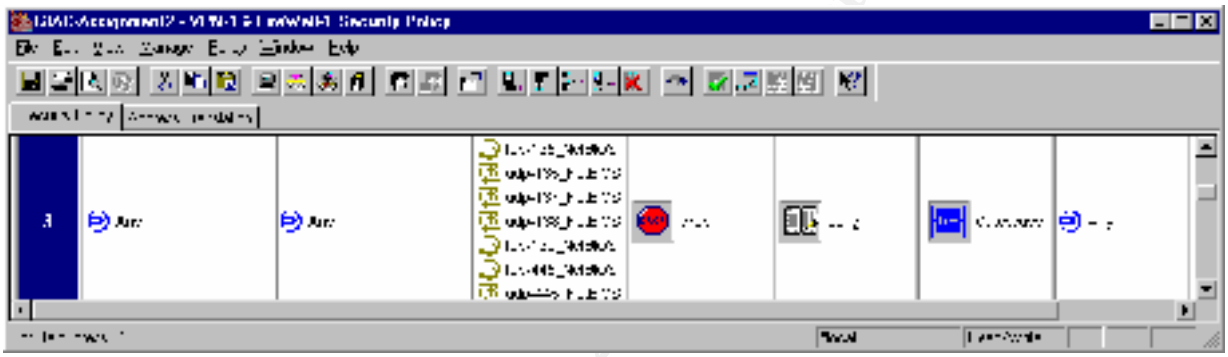

2.5. X Windows and other Unix serv ices (RPC, NFS, rlogin)

This is the same thing as in the previous case, we want to block all these services altogether, they are very unlikely to be needed from the outside of our network.
For X-Windows, we have to possibility to exclude an entire range i nstead of having to specify each individual port:

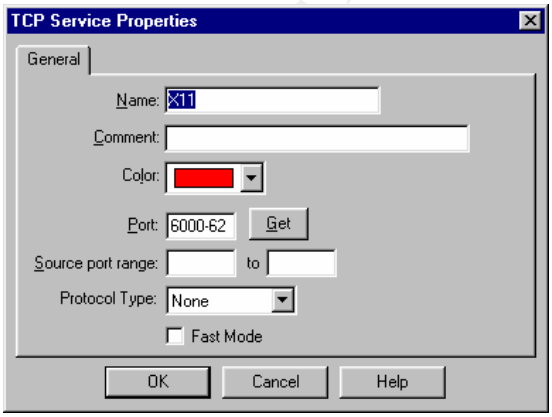

The rule is:

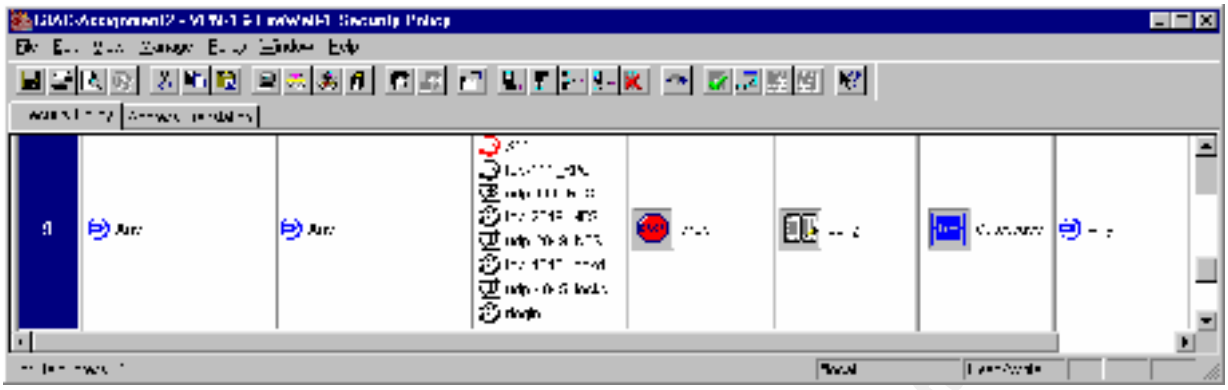

2.6. Mail, DNS, Web, ftp

These services are to be considered much more carefully because some of them need to be available from outside our network, or need to access the outside world. We w ill then have to carefully design our ruleset for these services.

The exact requirements are:

- mail: all internal hosts connect to the internal mail relay. All external (Internet) hosts need to be able to access the mail relay on the DMZ. Then the 2 mail r elays (internal and DMZ) can send SMTP traffic to each other.
- - web: all external hosts need to be able to access the company's server on the DMZ, and the internal proxy needs to be able to access the external web server as well as all other web servers on the Internet.
- - ftp: the company's customers and suppliers need to be able to access the company's ftp server to download patches, updates etc...
- dns: the internal server will resolve all internal names and forward other queries to the external DMZ server. The external server will answer all queries from the internet about servers placed on the DMZ (and only about these !).

2.6.1. HTTP

For http access, the allowed ports will be the most common ports on which web sites are listening:

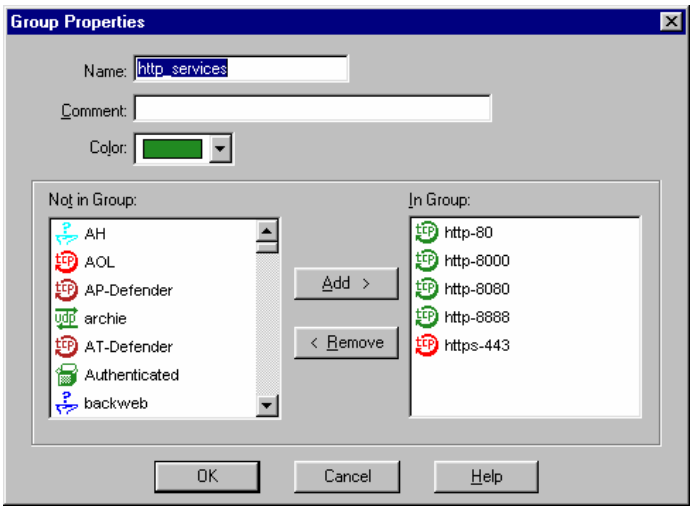

This is the correspondi ng rule:

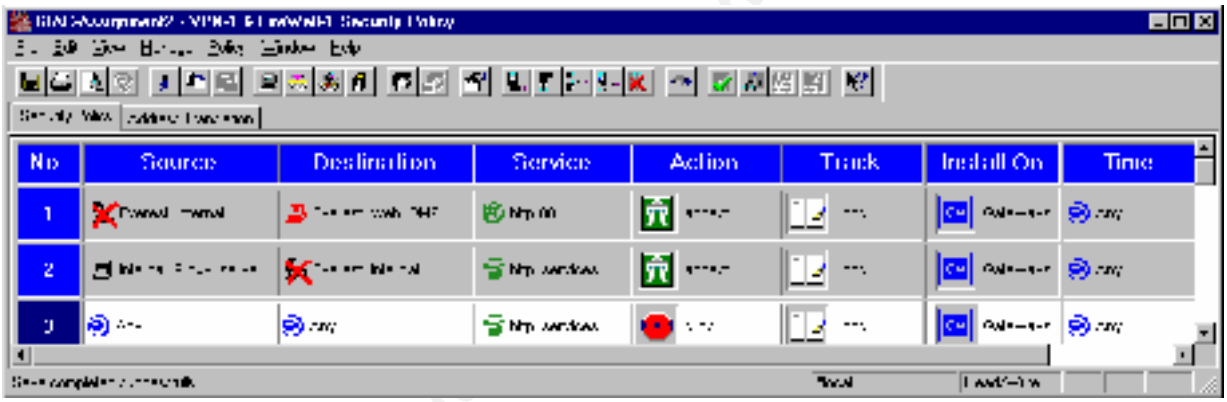

Rule number 1 provides http 80 service from all external (Internet) hosts to our company's public web server on the DMZ (which listens on port 80 only). That little red cross on the icon "Everest_Internal" in the first column "Source" is a nega tion, it describes all objects which are not part of this group. In this case, it means everything but Everest's internal network, that is the Internet (to be precise, this would represent the Internet plus the DMZ, but of course the firewall cannot preven t any server on the DMZ from connecting to the web server since they are on the same segment).
Rule number 2 provides http services from the Internal proxy server to all external web sites. Why is the origin of all HTTP connections limited to the proxy ? Because it is better to have all connections to the Internet to originate from 1 secure machine only, and the proxy will provide an easy content -filtering implementation (for example deny pornographic & racist sites), as well as protection from malicious J ava or active-X content, which would not be more difficult to implement if all internal hosts could connect directly . Finally, rule number 3 denies all other kind of http traffic.

These 3 new rules have been inserted at the very top of the ruleset for per formance reasons. As HTTP is likely to be the most used protocol in terms of numbers of packet, it is better to have the firewall find the match for these packets as early as possible, without having to go through the entire ruleset. However rule 3 can lik ely be pushed down.

2.6.2. SMTP

For SMTP, we would have the following rules:

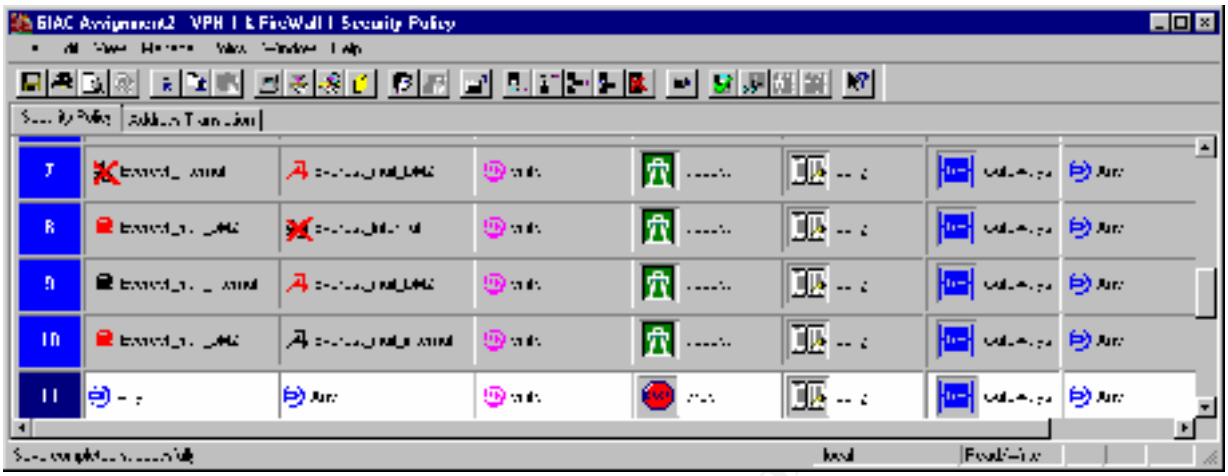

Rules 7 and 8 allow bi-directional traffic between the Internet and the mail relay on the DMZ (bi-directional because either side needs to be able to initiate a connection). The same principle is used in rules 9 and 10. As usual, rule 11 drops all traffic which has not been allowed by preceding rules.

2.6.3. DNS

For DNS, it is possible to use the integrated function of Checkpoint FW-1 as shown below:

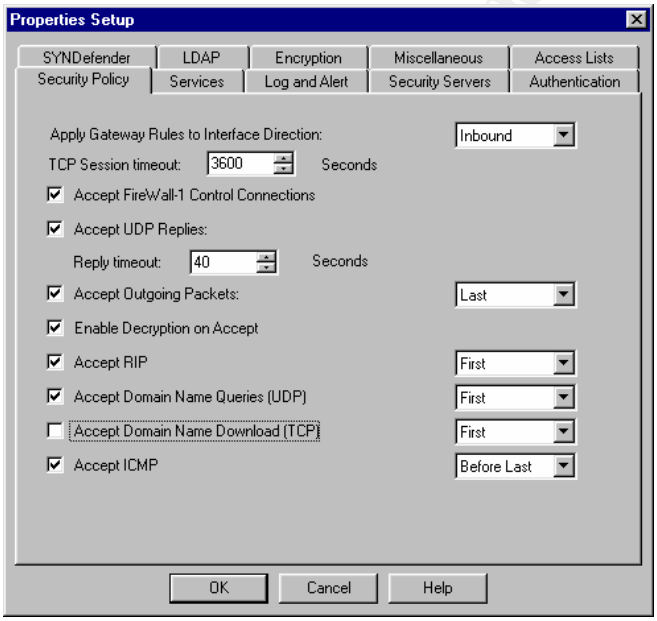

DNS requests are accepted over udp but not over tcp (no zone transfers, no long DNS requests). Since udp is a connectionless protocol, no state table similar to that used for tcp can be maintained by the firewall. However, the "Accept UDP replies" option provides a kind of equivalent. The firewall does not log port numbers, but

will remember queries between 2 hosts and will allow a UDP reply if it comes within 40 seconds of the request being sent.

2.6.4. ftp

A rule preventing ftp from any to any was already added to the ruleset. In order to allow customers to access the public ftp server, the following rule can be added before it:

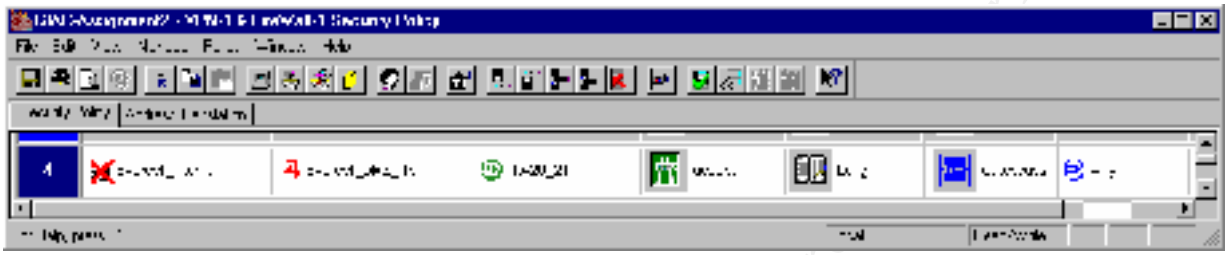

2.7. ICMP

All ICMP messages will be blocked. The reason for that is that ICMP messages can be used, for example, to modify the routing of packets sent by hosts accepting ICMP_REDIRECT. ICMP messages can also be used to organize DOS (Denial of Service) attacks.
However, they are also very useful when it comes to basic network troubleshooting, so at least the management station should be able to use them, both to the DMZ and the Internet.

Some common ICMP types of packets are defined on the Firewall GUI this way:

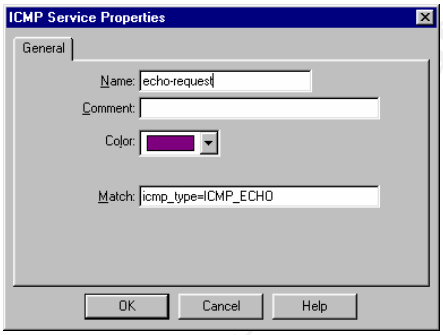

The other types of messages will be defined the same way, using the strings ICMP_ECHOREPLY, ICMP_REDIRECT, ICMP_TIMXCEED and ICMP_UNREACH for reply, redirect, t ime exceeded and destination unreachable, respectively.

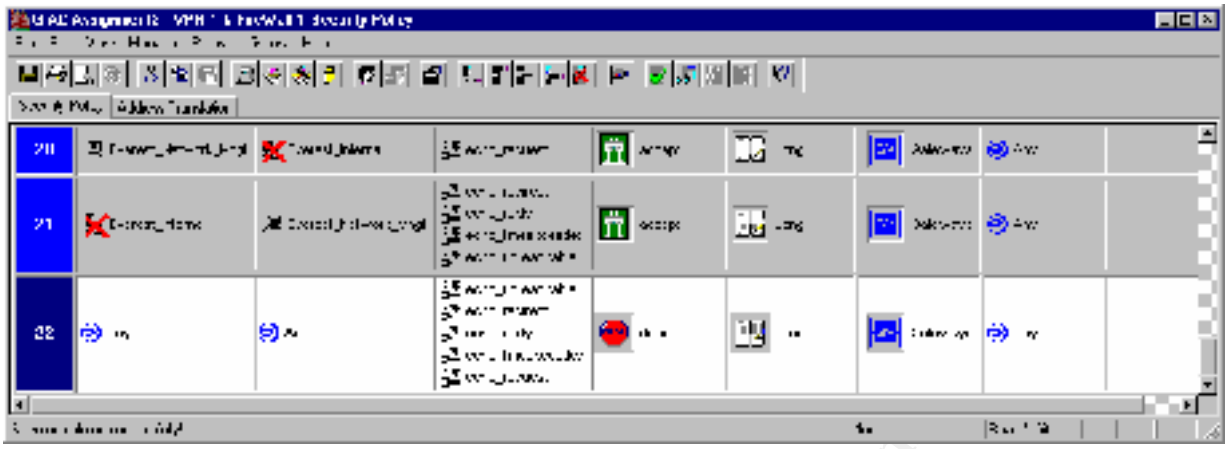

2.8. SNMP

SNMP is a protocol which is used for network management. It allows network devices to be remotely managed (read and write of configuration parameters), and also for the transmission of alarm s (traps). The most widely version used at present, version 2, is essentially insecure because it transmits all passwords (called "strings" in the SNMP world) in clear text. So SNMP traffic from and to the Internet must be filtered. However it can provide useful services, so it will be allowed between the internal network and the Internet router.
SNMP uses ports 161 (TCP, UDP) for its GET and SET commands (read, write) and 162 (TCP, UDP) for alarm (trap) transmission.

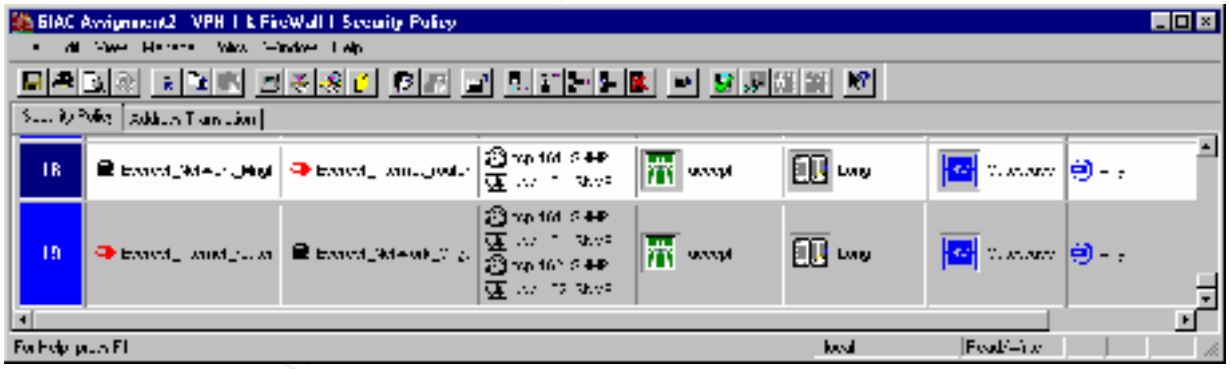

2.9. Management of the Firewall

The rules shown below allow necessary services in order to manage the firewall remotely.

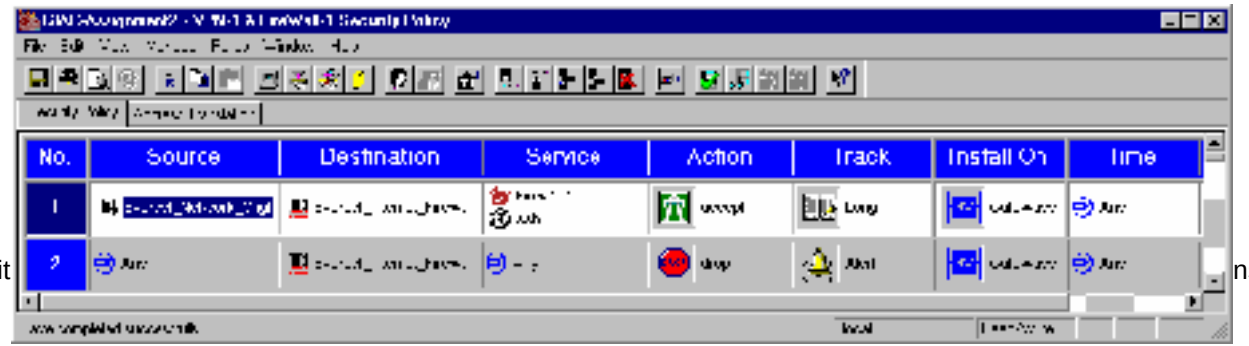

2.10. *Other services (small services, LPD, NNTP, syslog, BGP, SOCKS)*

The rule to block these services is shown below:

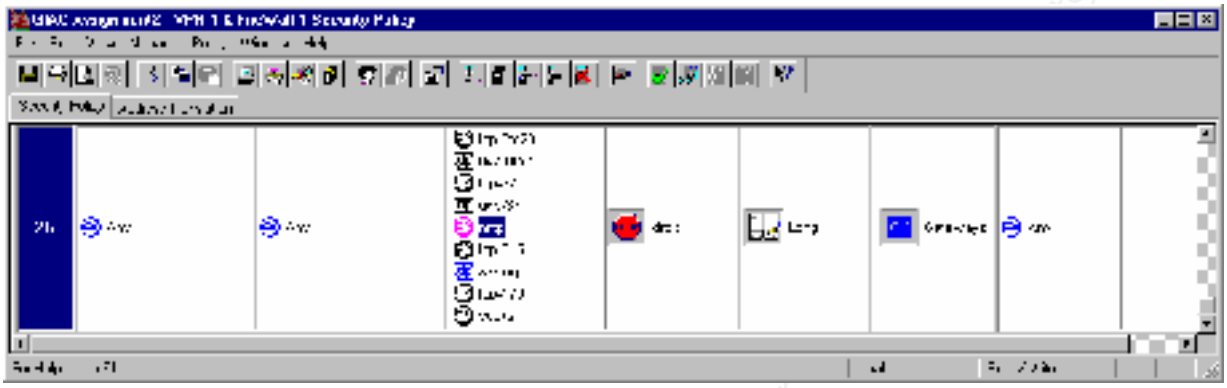

2.11. *Network Address Translation*

This rule defines the NAT operations which have to be performed for packets leaving the internal network, as presented in the assignment 1.

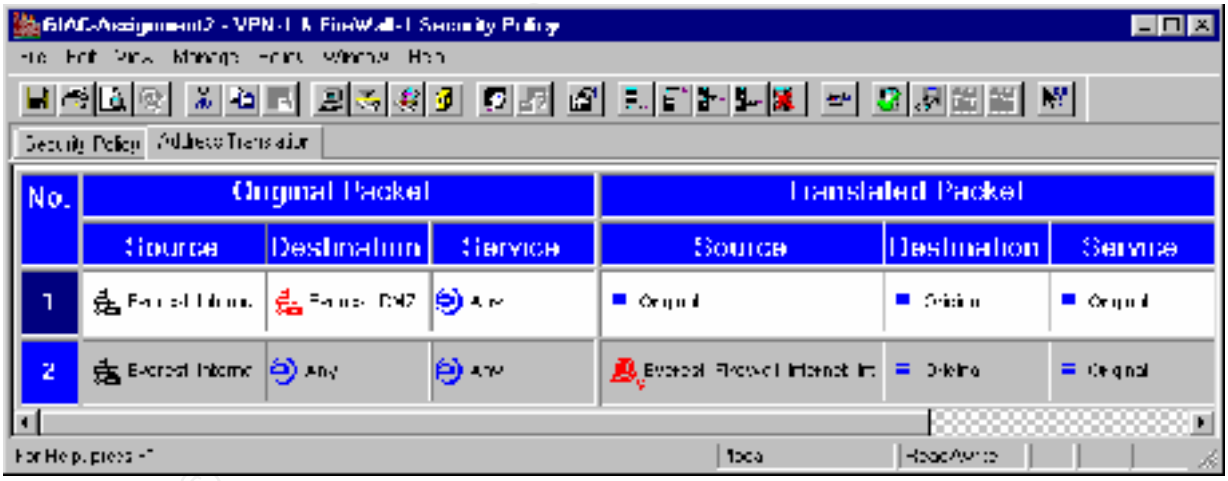

All packets going to the Internet will have their source address replaced with the address of the public Internet interface of the fire wall. However such an address translation is not required for the DMZ (because no routing is done on the DMZ, all packets must transit via the firewall which knows where to route to 10. addresses).

So NAT rule number 2 implements the translation for all p ackets sent to the Internet, and rule number 1, which takes precedence, prevents any translation taking place for the traffic to the DMZ.

2.12. Remote Connectivity.

The aim is to restrict users on the road to access the company's intranet server as well as th e internal mail relay. The remote access server (dial -in) as well as the VPN server can be configured in order to allow only HTTP and SMTP traffic so that users are limited to the intranet server and the internal mail server.

2.13. Wrap-up and Installation of the ruleset

Since the Firewall will drop by default all services which are not specifically allowed, it is not necessary to keep in our ruleset all rules which deny services. It is possible to safely delete these rules.

In order to install the ruleset, t he command to use under Solaris is "fw load *file firewall_name*", *file* being the file in which the ruleset was saved and *firewall_name* the name of the firewall.

3. Assignment 3

The scope of the assessment will include critical host systems, the local ne tworks, and all external connectivity to the site.
The assessment will consist of a penetration test of the network from outside and inside the network perimeters, assessment of the physical security of the devices, and assessment of systems security of v arious hosts (for example UNIX, Novell, NT…) using automated tools and manual methods.

The penetration test will involve a security scan of external connect points from the internet to identify accessible hosts and their respective services. Attempts will then be made to enter accessible hosts (i.e., those with access services turned -on) by exploiting discovered vulnerabilities and/or by using password cracking programs. A similar test will be conducted from inside the network. The host based assessment wi ll involve a manual review and an automated scan of various host security related settings and parameters to ensure that secondary defenses are in place in case a perpetrator was able to circumvent host login controls.

3.1. Assessment planning

The planning phase is used to gather preliminary information about the target network before the actual audit takes place. The types of activities required in this phase are things like coordinating the points of contact for the audit, establishing any ground rules suc h as whether this audit will be conducted in person or remotely over the WAN, and making arrangements for technical interviews with persons knowledgeable about the target network.

The planning phase consists of the following tasks:
- Obtain approval to ga in access to the different network segments
- Interview the local contacts to learn what the requirements are/should be for the target network
- Schedule the assessment once the necessary information regarding the target network is acquired and it is known what data is needed to ensure proper protection.

Requirements

Hardware: the hardware requirements for conducting this assessment consists in a simple mid -range laptop. Ideally, the laptop would have two OSes installed on, a version of Windows (NT is pre ferable, but 95/2K would do as well) and a version of Linux.

Software: the software to be used will be either commercial tools or open source software, as described further.

Training: at a minimum, auditors will need good theoretical knowledge of network security principles. A good command of the tools used for the assessment is required. At best, tools should have been extensively used on tests systems in order to get familiar with them.

Infrastructure: the infrastructure required for the assessment sho uld be relatively painless to find. The auditors should make sure they can have access to all parts of the networks via switch and/or hub connections, as well as a connection to the internet.

3.2. Assessment implementation

The assessment will be divided in 4 main parts:
- the network review
- the host/application review
- the remote access review
- the physical security and process review

3.2.1. Network Review

The network review, or Scan, consists of using a set of software programs that will attempt to connect to every host on the target networks in a variety of ways to determine amongst other things what the OS type is, if the host is running Internet services, and any potential vulnerabilities.
These scanners can usually compare the results from a scanned host against the contents of an internal database to determine the results of the scan. The results from a network scan consisting of many hosts can then be compiled into a report depicting the hosts consisting of the weakest links in terms of overall network security. These weak links may need to be patched or upgraded, etc…
Due to the amount of time that can be required to scan large numbers of hosts on a network, the information gathered during the interviews conducted at the planning phase can help reduce the number of security holes that are scanned for, such that if it is known that there is no software on a network capable of running the Sun NIS service, then it can be eliminated from the list of services to scan for. The other way round, if a network doesn't contain any Windows -based systems, all Microsoft services can be eliminated as well.

3.2.1.1. NMAP

An NMAP port scan is used to identify which network services are available on a target machine. This information could then be used to exploit vulnerabilities specific to the service running on the port.

The SCAN should be done both from the Internet (from a segment which is located beyond the firewall and the Internet router) in order to test that no unexpected ports are left open on the Firewall, and also directly locally on the DMZ in order ensure that only the expected services are running on these high risk machines.

Here are two examples of a NMAP scan directed at the mail relay on the DMZ. The first one is made from within the DMZ, the other one from behind the firewall. These examples were produced using NMAP on a Solaris workstation.

NMAP from the DMZ:

```
> nmap –P0 157.159.12.12
Starting nmap V. 2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Interesting ports on  (157.159.12.12):
(The 1496 ports scanned  but not shown below are in state: closed)
Port       State       Service
21/tcp     open        ftp
23/tcp     open        telnet
25/tcp     open        smtp
79/tcp     open        finger

Nmap run completed  -- 1 IP address (1 host up) scanned in 1 second
```

This output shows that some unneeded services are still running on the server. Apart from SMTP, which is the only required service, ftp, telnet and finger are still listening on their respective ports. These services should be turned off.

NMAP from behi nd the firewall:

> nmap –P0 157.159.12.12
In this case, NMAP should fail. It is possible to have a closer look at what happens with a log at the firewall level (with the command "fw log –ftn"):

```
17:21:31 drop Everest_FW >qfe0 proto tcp src 157.159.11.10 d     st 157.159.12.12 service ftp  -data s_port 49548 len 44 rule 5
17:21:31 drop Everest_FW >qfe0 proto tcp src 157.159.11.10 dst 157.159.12.12 service ftp s_port 49549 len 44 rule 8
17:21:34 drop Everest_FW >qfe0 proto tcp src 157.159.11.10 dst 157.159.12.12 se       rvice smtp s_port 49550 len 44 rule 17
17:21:34 drop Everest_FW >qfe0 proto tcp src 157.159.11.10 dst 157.159.12.12 service telnet s_port 49551 len 44 rule 11
```

The format of the log is:
Time action host interface protocol source -address destination -address service source -port packet -length firewall -rule

For example, the second line of the log shows:

At 17:21:31, the Everest firewall dropped a packet.
The tcp packet was coming from 157.159.11.10/port 49549, going to 157.159.12.12/port 21 (ftp), length was 44 bytes
It was dropped according to the policy defined in rule number 8.

More information about NMAP can be obtained at http://www.insecure.org/nmap

3.2.1.2. ISS Internet Scanner

Internet Scanner is a commercial product for network vulnerability assessment and audit analysis. It is designed to identify and address technical vulnerabilities. It provides detailed technical policies and performs scheduled and selective vulnerability probes and assessments of communication services, operating systems, applications, and routers.
It searches for those vulnerabilities used by unscrupulous attackers to probe, investigate, and attack networks.

Being a commercial product, it is quite easy to install and use. More information can be obtained from http://www.iss.net.

Here is an example of a report that was created by Internet Scanner on the SMTP server on the DMZ:

IP Address	157.159.12.12
DNS Name	mail.everest.com
Vulnerability Name	STP verify (VRFY) command can be used to validate u sers
Severity	Low

Description:
The SMTP VRFY command is enabled. The VRFY (Verify) command allows an attacker to determine if an account exists on a system, providing significant assistance to a brute force attack on user accounts. VRFY provides additio nal information about users on the system, such as if they exist and their full names.

Fix

If you are running Sendmail, add the line Opnovrfy to your Sendmail configuration file, usually located in /etc/sendmail.cf. For other mail servers, contact your v endor for information on how to disable the verify command. Newer versions of Sendmail are available at http://www.sendmail.org or from ftp://ftp.cs.berkeley.edu/ucb/sendmail.

3.2.2. Host/Application review

3.2.2.1. ISS System Sc anner

System Scanner is a tool that provides a host -based security assessment that targets security weaknesses undetectable through network scanning. While the Internet Scanner determines vulnerabilities by scanning devices at the network level, System Sc anner detects internal vulnerabilities at the system level through a variety of cross platform agents that reside on each system. These agents allow a security policy to be implemented and controlled across an entire enterprise from a centralized point. It can also be used to lock down the configuration with a digital fingerprint - making it easier to detect unauthorized tampering. System Scanner agents are available for both Windows NT and many UNIX platforms.

Here is an example of a report that was creat ed by System Scanner on a NT server.

Check Risk level Brief description: reg -sys-log-03 Medium Recent events overwritten if the system log is full
Description
Is the Overwrite Events Older Than period less than the number of days specified for the syste m log in the site security policy?
Consequences
If the system is allowed to overwrite log events when the log becomes full, vital audit information may be lost.
Remedy
If "Do Not Overwrite Events" has not been selected for the system log then the "Over write Events Older Than" should be set to minimum number of days specified by the site's security policy.
The minimum number of days for which events are to be retained in the system log can be set with the Event Viewer administrative tool via the Log => Log Settings menu option, selecting the System Log, clicking the "Overwrite Events Older than" option and entering the required number of days.
Vulnerability detail
Events older than 7 days are overwritten if the System Log becomes full rather than e vents older than 365 days.

Check Risk level Brief description: reg -tcpip-02 Medium TCP filtering not enabled
Description
Has the 'Enable Security' option been selected in the TCP/IP Protocol Properties window?
Consequences
If TCP/IP Security is not e nabled, then the server does not filter incoming packets. If filtering is enabled, then it is possible to specify which TCP ports, UDP ports, and IP Protocols can be used.
Remedy
It is recommended that filtering is enabled.
This can be achieved by check ing the 'Enable Security' box on the 'Advanced IP Addressing' window accessed from the 'Network' option within 'Control Panel'.
Vulnerability detail

Check Risk level Brief description: iis -data-patch1 Low IIS files accessible (infocomm.dll out of date)
Description
Does the installed version of infocomm.dll allow unauthorized users access to IIS files?
Consequences
The installed version of infocomm.dll allows unauthorized access. Without a patch, unauthorized users can access files on the Microsoft Intern et Information Server.

See Microsoft Security Bulletin MS98 -003 for details.

Remedy
Apply the patch indicated in Microsoft Security Bulletin MS98 -003.
Apply the patch as follows:
1. Open your web browser.
2. Enter the URL: http://www.microsoft.com/securit y/.
3. On the left column, under Security Bulletins, click security bulletins by date.
4. Click MS98 -003.
5. Under What Microsoft is Doing, select the patch appropriate for your system.
6. Follow the instructions.
Vulnerability detail
C:\WINNT\system32\inetsrv\infocomm.dll (4.0.1381.110) found

3.2.2.2. Password cracking

Time schedule permitting, it can be very useful to use Crack (for UNIX) and/or L0phtCrack (for NT) to identify password vulnerabilities.

L0phtCrack 2.5 (for NT) (see http://www.l0pht.com/l0phtcrack/)

L0phtCrack is designed to recover passwords for Windows NT. NT does not store the actual passwords but a cryptographic hash of the passwords. L0phtCrack can take the hashes of passwords and generate the cleartext passwords. It can recover passwords directly from the registry, from the file system and backup tapes, from repair disks, or by recovering the passwords as they traverse the network. Once it has extracted the password hashes, it goes to work computing the passwords, which is called cracking. It uses three different methods:
- dictionary attack: l0phtCrack tests all the words in a dictionary against the password hashes.
- hybrid crack method: it is an extension of the dictionary method which w orks by adding numeric and symbol characters to dictionary words. These are the types of passwords that will pass through many password filters and policies but yet still are easily crackable.
- brute force method; this method will always recover the passw ord no matter how complex. It is just a matter of time. Most passwords can be cracked in a matter of days, however complex they are. This is usually much shorter than the time most administrators set their password policy expiration time to.

The password hashes can be obtained:
a) From the Registry: If you have administrator rights, you can use the Tools Dump Passwords from Registry command on the L0phtCrack menu to retrieve the hashes
b) From the SAM File: since the operating system holds a lock on the SA M file where the password hashes are stored, it is not possible to just read them from this file while the operation system is running. However, sometimes a backup of this file is made on tape or on an Emergency Repair Disk or in the repair directory of the system hard drive. Also, another operating system such as DOS can be booted from a floppy and the password hashes can be read directly from the file system. This is especially useful if you have physical access to the machine and it has a floppy drive.
c) From SMB Packet Capture: this is a capture of the encrypted hashes over the network. Your machine must have 1 or more Ethernet devices to access the network.
d) Using PWDUMP2: Todd Sabin has released a free utility that can dump the password hashes on a local machine if the SAM has been encrypted with the SYSKEY utility that was introduced in Service Pack 3 (refer to http://www.webspan.net/~tas/pwdump2/) .

The default options for cracking are to run a dictionary attack, then a hybrid attack, and then the brute force attack. L0phtCrack runs these attacks on the password hashes in succession by default. You can select more details about the cracking attack in the Tools Options dialog box.

During any crack attack the L0phtCrack window displays status information to show the progress of the attack. During dictionary attacks the number of dictionary words tried is displayed along with the percentage complete.

Crack (For Unix)

The "Crack" utility is a tool to break Unix passwords, based on the same principles as L0phtcrack. All it requires is a copy of the /etc/passwd and /etc/shadow files, which are merged using the *shadmrg.sv* included in the distribution. Only root access to the server is required for this operation.

3.2.3. Remote Access

RAS (Remote Access Service) access to the internal network can potentially be a very vulnerable service. This service could potentially allow unauthorized users to connect to resources on your network without requiring a physical presence to do so. These types of services include, but are not limited to, remote control products like Symantec's pcANYWHERE, Virtual Private Network (VPN) access, or regular dialup access through hardware like Cisco's AS5300. Remote access can also be obtained by either a modem directly connected to a computer and a phone line at the host site, or via the internet.

3.2.3.1. Internet

All access from the Internet is going through the firewall, which has been assessed already.

3.2.3.2. Remote Access Server

Normally, this type of remote connectivity solution is very secure. In order to break into the system, an attacker would need to know which phone number(s) to use. Once this information is obtained and (s)he manages to call the server and get a login prompt, it is necessary to guess a login and password with no or little possibility to use backdoor and/or bugs. It is not easily possible to eavesdrop on other connections to steal a valid user's login and password as phone connections are circuit based. The only "easy" attack would be a brute-force attack.
If time and costs are limited, the assessment of the remote access server can safely be ignored as such solutions have been on the market for some time now and are relatively secure. Not being connected to the Internet, they are also relatively more secure than more exposed devices.

3.2.3.3. VPN

VPN servers are relatively new on the market, and are more exposed than the remote access servers. However this type of solution also provides a very secure solution for connectivity. The type of attacks that can be used against these servers is also based on brute force.

3.2.3.4. Other types of remote access

One of the most common types of attacks targeted at corporate networks are conducted via hosts which have dual connections: a LAN connection to the corporate network, and also a modem connection directly to an ISP, in order to get un-firewalled Internet access for whichever purpose. Such hosts presents a risk in that they can provide the ability to circumvent existing perimeter security devices. In order to detect the existence of such hosts, it is possible to use ToneLoc.

ToneLoc is a tool that allows to check usage of analog lines in any given environment. Of course, a list of analog lines should be obtained in advance to make the test more efficient. The test should be conducted at off peak times to minimize the distraction to valid analog users.

To scan a single number, from the command prompt, type *ToneLoc ### - ### - ####* (for example, type *ToneLoc 123 – 456 - 7899* to scan a the number (123) 456-7899 for a carrier signal).
To scan multiple consecutive numbers, *type ToneLoc ### - XXXX /R:XXXX-XXXX*. (for example executing *ToneLoc 474-XXXX /R:9000-9999* dials 1000 numbers, from 474-9000 to 474-9999 (randomly).
Results will be stored in a file called "Tone.log". Potentially rogue modems will be identified as "*CARRIER *" in this file

Once all the carriers have been identified by ToneLoc, the vulnerability of each of them can be identified by attempting to connect to it using VT100 or similar tools (Dial-up Networking, Hyper Terminal). This is done to verify the existence of a modem. Comparing the list of carriers found by ToneLoc to a list of known fax machines is an easy way to eliminate carriers identified from the list of potential modems.

More information about ToneLoc can be found at this site:
http://www.ntsecurity.net/security/tools/def.htm

3.2.4. Physical Security and Process review

This is a review of the physical security of the site and supporting documents to identify environment related vulnerabilities, as well as validate adherence to established security practices and policies.

3.2.4.1. Physical Walkthroughs

The assessment team will conduct a tour of the computer rooms and the whole facility to identify physical security controls in place. Items of interest in this overview are things like building security (controlled access to computer rooms, monitoring, etc…) and equipment protection with passwords and such.

3.2.4.2. Policies and Processes - Supporting Documents

This is a review of all supporting documents in order to validate adherence to established security practices, standards and procedures. It can involve user interviews to corroborate compliance with established security standards and procedures. Users can be questioned on such topics as:

- do they know if a security policy is in place ?
- if yes, do they know where to find the relevant policy documentation ? Have they read it and do they understand it ?
- what would they do in cases where their anti-virus software alerted them of a suspicious file, or if a security agent was told that "an engineer" from an external company requires immediate access to the server room ?

3.3. Perimeter Analysis

Depending on the type of vulnerabilities found, some of the recommendations for improvement listed below may apply.

3.3.1. SMTP connectivity

One of the weaknesses of the SMTP implementation of the network is that mail delivery from the Int ernet to the local mail server requires an incoming from the DMZ, that is from a network which is not entirely trusted because partly exposed to the Internet.
The architecture could be modified a little in order to replace SMTP with POP3. The advantage wit h this solution is that it would eliminate the need for connections originating from the mail relay on the DMZ towards the internal mail server.
In order to implement this solution, the required actions are:

1) install a POP -3 server on the DMZ mail relay
2) replace the former rule allowing SMTP from the DMZ mail relay to the internal one with one allowing POP-3 connections from the inside as shown below (rule 12 is disabled):

However this solution is not ideal yet because POP -3 transmits password without encr yption. In order to implement an even more secure solution, it is possible to configure POP -3 over SSL though.

3.3.2. Accounts and Passwords

Depending on the number of passwords which were guessed by the password cracking programs, and also on the time which w as required to crack the simplest ones, it may be required to strengthen the password policies and also educate the users, that is teach them what a good password is.

Checking the passwords of each account is also a good opportunity to review the set of e xisting accounts. There should be as few generic accounts as possible, preferably none. A generic account available to several users prevents accountability for those users, and also offers a security risk since most generic accounts will also end up with an easily guessed password.
There should be the ability to set account expirations, as well as password expirations. There should also be a formal process in place to delete normal accounts for users that are no longer needed. An account for a former employee, for example, should be deleted or disabled to prevent access of a potentially disgruntled employee back into the network.
Another point of interest is to make local login names, remote login names and email addresses all different, and not necessarily based on people's real names. For example, Joe H. Average could have joe.h.average@everest.com as an email address, jhaverage as a local login and jhav4321 for remote access.

3.3.3. Remote access (VPN and Dial -in server + Authentication server)

In order to improve the security of the remote access infrastructure, it is possible to isolate the 3 devices (VPN server, dial-in server and authentication server) on a separate segment which would be connected to the mai n segment either through the internal firewall or through the internal router.
That would add an additional layer of defense against any internal threats such as internal users trying to break into the LDAP server to steal other people's credentials (user names and passwords).

3.3.4. Intrusion Detection Systems (IDS)

If some servers contain critical business information that would jeopardize the company's existence if ever that information was compromised, an IDS can play a critical role in the maintenance of t he overall security of a network. IDS's should not considered as perfect tools and should always be supplemented with a combination of other approaches for attaining a level of security. Some such commercial IDS products are available from companies such a s ISS, Axent, BlackIce, Network Associates, etc., but there are also freely available IDS tools as well, such as Shadow.

3.3.5. Network address translation for the DMZ

If one is concerned with some of the hosts on the DMZ being compromised, it is also possible to implement address translation on the DMZ:

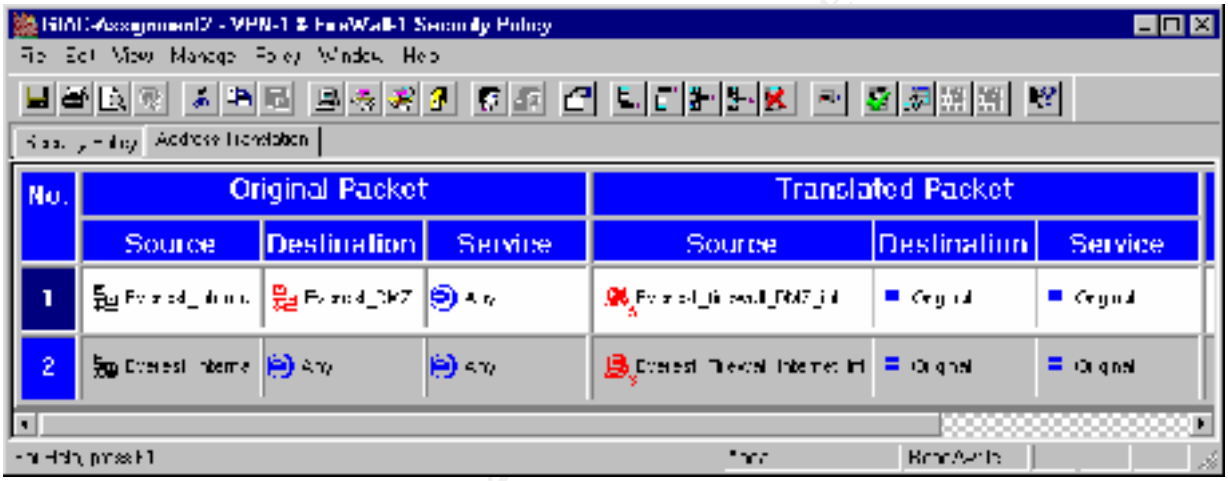

The traffic would use the Firewall's DMZ interface's address as a source. This would hide the internal address range from an eavesdropper who could have compromised and taken control of a host on the DMZ.