The Principles of Modern Attacks Analysis for Penetration Tester

Adam Ali.Zare Hudaib *adamhudaib@gmail.com*
Licensed Penetration Tester |EC-Council
Certified Ethical Hacker |EC-Council
Certified Security Analyst |EC-Council
Certified Network Analyst | WireShark University
Information & Cyber Security Expert
CEH , ECSA , LPT , WCNA
Ukraine

Abstract

Modern cyber defense requires a realistic and thorough understanding of web application security issues. Anyone can learn to sling a few web hacks, but web application penetration testing requires something deeper. Major web application flaws and their exploitation, a field-tested and repeatable process to consistently finding these flaws and convey them will be discussed in this article. Modern attacks principles will be analyzed on purpose to create the most sufficient penetration tests.

Keywords: Penetration Testing, DOS Attack, ICMP, IPv6, IPv4, NTP, Honey Pot Systems.

1. INTRODUCTION

Penetration tests (pentests) have gained recognition as a legitimate approach to identifying and then in theory, mitigating discovered weaknesses. The pentest industry even has a magazine (PenTest Magazine) and there are some tools out there that you, as an industrial control systems (ICS) cyber security professional, ought to have in your tool set like the PWN Phone or Metasploit modules from Digital Bond. With the various tools of the trade you will undoubtedly discover at least one vulnerability in your network and with that information in-hand, you may then get the resources to fix that problem. So we should analyze nodern attacks principles for efficient pentesting.

2. ANALYSIS OF MODERN ATTACKS PRINCIPLES FOR PENTESTERS

2.1 DOS Attacks and Free DOS Attacking Tools

The denial of service (DOS) attack is one of the most powerful attacks used by hackers to harm a company or organization. Don’t confuse a DOS attack with DOS, the disc operating system developed by Microsoft. This attack is one of most dangerous cyber attacks. It causes service outages and the loss of millions, depending on the duration of attack. In past few years, the use of the attack has increased due to the availability of free tools. This tool can be blocked easily by having a good firewall. But a widespread and clever DOS attack can bypass most of the restrictions. In this post, we will see more about the DOS attack, its variants, and the tools that are used to perform the attack. We will also see how to prevent this attack and how not to be the part of this attack.

What Is a Denial of Service Attack?

A DOS attack is an attempt to make a system or server unavailable for legitimate users and, finally, to take the service down. This is achieved by flooding the server’s request queue with fake requests. After this, server will not be able to handle the requests of legitimate users.

In general, there are two forms of the DOS attack. The first form is on that can crash a server. The second form of DOS attack only floods a service.

DDOS or Distributed Denial of Service Attack

This is the complicated but powerful version of DOS attack in which many attacking systems are involved. In DDOS attacks, many computers start performing DOS attacks on the same target server. As the DOS attack is distributed over large group of computers, it is known as a distributed denial of service attack.

To perform a DDOS attack, attackers use a zombie network, which is a group of infected computers on which the attacker has silently installed the DOS attacking tool. Whenever he wants to perform DDOS, he can use all the computers of ZOMBIE network to perform the attack.
In simple words, when a server system is being flooded from fake requests coming from multiple sources (potentially hundreds of thousands), it is known as a DDOS attack. In this case, blocking a single or few IP address does not work. The more members in the zombie network, more powerful the attack it. For creating the zombie network, hackers generally use a Trojan.

There are basically three types of DDOS attacks:

1. Application-layer DDOS attack
2. Protocol DOS attack
3. Volume-based DDOS attack

Application layer DDOS attack: Application-layer DDOS attacks are attacks that target Windows, Apache, OpenBSD, or other software vulnerabilities to perform the attack and crash the server.

Protocol DDOS attack: A protocol DDOS attacks is a DOS attack on the protocol level. This category includes Synflood, Ping of Death, and more.

Volume-based DDOS attack: This type of attack includes ICMP floods, UDP floods, and other kind of floods performed via spoofed packets.

There are many tools available for free that can be used to flood a server and perform an attack. A few tools also support a zombie network to perform DDOS attacks. For this post, we have compiled a few freely available DOS attacking tools.

Free DOS Attacking Tools

- *LOIC (Low Orbit Ion Canon)*

LOIC is one of the most popular DOS attacking tools freely available on the Internet. This tool was used by the popular hackers group Anonymous against many big companies' networks last year. Anonymous has not only used the tool, but also requested Internet users to join their DDOS attack via IRC.

It can be used simply by a single user to perform a DOS attack on small servers. This tool is really easy to use, even for a beginner. This tool performs a DOS attack by sending UDP, TCP, or HTTP requests to the victim server. You only need to know the URL of IP address of the server and the tool will do the rest.

You can see the snapshot of the tool above. Enter the URL or IP address and then select the attack parameters. If you are not sure, you can leave the defaults. When you are done with everything, click on the big button saying "IMMA CHARGIN MAH LAZER" and it will start attacking on the target server. In a few seconds, you will see that the website has stopped responding to your requests.

This tool also has a HIVEMIND mode. It lets attacker control remote LOIC systems to perform a DDOS attack. This feature is used to control all other computers in your zombie network. This tool can be used for both DOS attacks and DDOS attacks against any website or server.

The most important thing you should know is that LOIC does nothing to hide your IP address. If you are planning to use LOIC to perform a DOS attack, think again. Using a proxy will not help you because it will hit the proxy server not the target server. So using this tool against a server can create a trouble for you.

- *XOIC*

XOIC is another nice DOS attacking tool. It performs a DOS attack an any server with an IP address, a user-selected port, and a user-selected protocol. Developers of XOIC claim that XOIC is more powerful than LOIC in many ways. Like LOIC, it comes with an easy-to-use GUI, so a beginner can easily use this tool to perform attacks on other websites or servers.

In general, the tool comes with three attacking modes. The first one, known as test mode, is very basic. The second is normal DOS attack mode. The last one is a DOS attack mode that comes with a TCP/HTTP/UDP/ICMP Message.

It is an effective tool and can be used against small websites. Never try it against your own website. You may end up crashing your own website's server.

- *HULK (HTTP Unbearable Load King)*

HULK is another nice DOS attacking tool that generates a unique request for each and every generated request to obfuscated traffic at a web server. This tool uses many other techniques to avoid attack detection via known patterns.

It has a list of known user agents to use randomly with requests. It also uses referrer forgery and it can bypass caching engines, thus it directly hits the server's resource pool.

The developer of the tool tested it on an IIS 7 web server with 4 GB RAM. This tool brought the server down in under one minute.

- *DDOSIM—Layer 7 DDOS Simulator*

DDOSIM is another popular DOS attacking tool. As the name suggests, it is used to perform DDOS attacks by simulating several zombie hosts. All zombie hosts create full TCP connections to the target server.

This tool is written in C++ and runs on Linux systems.

These are main features of DDOSIM

1. Simulates several zombies in attack
2. Random IP addresses
3. TCP-connection-based attacks
4. Application-layer DDOS attacks
5. HTTP DDoS with valid requests
6. HTTP DDoS with invalid requests (similar to a DC++ attack)
7. SMTP DDoS
8. TCP connection flood on random port

- *R-U-Dead-Yet*

R-U-Dead-Yet is a HTTP post DOS attack tool. For short, it is also known as RUDY. It performs a DOS attack with a long form field submission via the POST method. This tool comes with an interactive console menu. It detects forms on a given URL and lets users select which forms and fields should be used for a POST-based DOS attack.

- *Tor's Hammer*

Tor's Hammer is another nice DOS testing tool. It is a slow post tool written in Python. This tool has an extra advantage: It can be run through a TOR network to be anonymous while performing the attack. It is an effective tool that can kill Apache or IIS servers in few seconds.

- *PyLoris*

PyLoris is said to be a testing tool for servers. It can be used to perform DOS attacks on a service. This tool can utilize SOCKS proxies and SSL connections to perform a DOS attack on a server. It can target various protocols, including HTTP, FTP, SMTP, IMAP, and Telnet. The latest version of the tool comes with a simple and easy-to-use GUI. Unlike other traditional DOS attacking tools, this tool directly hits the service.

- *OWASP DOS HTTP POST*

It is another nice tool to perform DOS attacks. You can use this tool to check whether your web server is able to defend DOS attack or not. Not only for defense, it can also be used to perform DOS attacks against a website.

- *DAVOSET*

DAVOSET is yet another nice tool for performing DDOS attacks. The latest version of the tool has added support for cookies along with many other features. You can download DAVOSET for free from Packetstormsecurity.

- *GoldenEye HTTP Denial Of Service Tool*

GoldenEye is also a simple but effective DOS attacking tool. It was developed in Python for testing DOS attacks, but people also use it as hacking tool.

Detection and Prevention of Denial of Service Attack

A DOS attack is very dangerous for an organization, so it is important to know and have a setup for preventing one. Defenses against DOS attacks involve detecting and then blocking fake traffic. A more complex attack is hard to block. But there are a few methods that we can use to block normal DOS attack. The easiest way is to use a firewall with allow and deny rules. In simple cases, attacks come from a small number of IP addresses, so you can detect those IP addresses and then add a block rule in the firewall.

But this method will fail in some cases. We know that a firewall comes at a very deep level inside the network hierarchy, so a large amount of traffic may affect the router before reaching the firewall.

Blackholing and sinkholing are newer approaches. Blackholing detects the fake attacking traffic and sends it to a black hole. Sinkholing routes all traffic to a valid IP address where traffic is analyzed. Here, it rejects back packets.

Clean pipes is another recent method of handling DOS attacks. In this method, all traffic is passed through a cleaning center, where, various methods are performed to filter back traffic. Tata Communications, Verisign, and AT&T are the main providers of this kind of protection.
As an Internet user, you should also take care of your system. Hackers can use your system as a part of their zombie network. So, always try to protect your system. Always keep your system up to date with the latest patches. Install a good antivirus solution. Always take care while installing software. Never download software from un-trusted or unknown sources. Many websites serve malicious software to install Trojans in the systems of innocent users.

2.2 ICMP Redirects

RFC 792 spelt out the goals and specifications of the Internet Control Message Protocol (ICMP). Basically, it is used as a means to send error messages for non-transient error conditions and to provide a way to query the network in order to determine the general characteristic of the network.

The Internet Protocol (IP) is not designed to be absolutely reliable. The purpose of the ICMP messages is to provide feedback about problems in the communication environment, not to make IP reliable. There are still no guarantees that a datagram will be delivered or a control message will be returned. Some datagrams may still be undelivered without any report of their loss. The higher level protocols that use IP must implement their own reliability procedures if reliable communication is required.

ICMP uses the basic support of IP as if it were a higher level protocol. However, ICMP is actually an integral part of IP and must be implemented by every IP module.

ICMP suppose to be a relatively simple protocol, but it can be altered to act as a conduit for evil purpose. It is therefore important to understand how this protocol can be used for malicious purposes.

This assignment examines how ICMP can be used in a non-convention way, putting itself as a potential threat. We will concentrate on the use of ICMP in a non-convention way rather than the normal use of ICMP.

Understanding ICMP

Conventionally, ICMP is provided as a means to send error messages for non-transient error conditions and to provide a way to query the network.

ICMP is used for two types of operations:

- Reporting non-transient error conditions (ICMP Error Messages).
- Query the network with request and reply (ICMP Query Messages).

Unlike TCP and UDP, ICMP has no port numbers. ICMP uses type and code to differentiate the services in the protocol.

Also in ICMP, there is no client-server concept. When an ICMP error message is delivered, the receiving host might respond internally but might not communicate back to the informer. Services and ports do not have to be activated or listening. ICMP can be broadcast to many hosts because there is no sense of an exclusion connection.

RFC 792 defined special conditions for the ICMP messages:

- No ICMP error messages are sent in response to ICMP error messages to avoid infinite repetition.
- For fragmented IP datagrams, ICMP messages are only sent for errors on fragmented zero (the first fragment).
- ICMP error messages are never sent in response to a datagram that is destined to a broadcast or a multicast address.
- ICMP error messages are never sent in response to a datagram sent as a link layer broadcast.
- ICMP error messages are never sent in response to a datagram whose source address does not represents a unique host (the source address cannot be zero, a loopback address, a broadcast address or a multicast address).
- ICMP error messages are never sent in response to an IGMP message of any kind.
- When an ICMP message of unknown type is received, it must be silently discarded.
- Routers will almost always generate ICMP messages but when it comes to a destination host, the number of ICMP messages generated is implementation dependent.

The ICMP has many messages that are identified by a “type” field. For each “type” field, there may also be a “code” field which acts as a sub-type. For example, echo reply has a type of 0 and code of 0 while echo request has a type of 0 and code of 8.

The list of ICMP types and codes is available at: target="_blank">http://www.iana.org/assignments/icmp-parameters.

Normal use of ICMP

The Internet Control Message Protocol (ICMP) is used to handle errors and exchange control messages. ICMP can be used to determine if a machine on the Internet is responding. To do this, an ICMP echo request packet is sent to a machine. If a machine receives that packet, that machine will return an ICMP echo reply packet. A common implementation of this process is the "ping" command, which is included with many operating systems and network software packages. ICMP is used to convey status and error information including notification of network congestion and of other network transport problems. ICMP can also be a valuable tool in diagnosing host or network problems.

Other RFCs have defined other functionalities for the ICMP:

- RFC 896 – Source Quench.
- RFC 950 – Address Mask Extensions.
- RFC 1191 – Path MTU Discovery.
- RFC 1256 – Router Discovery.
- RFC 1349 –Type of Service in the Internet Protocol Suite.

Use of ICMP – In a Non-Convention Way

Ping traffic is ubiquitous to almost every TCP/IP based network and sub-network. It has a standard packet format recognized by every IP-speaking router and is used universally for network management, testing, and measurement. As such, many firewalls and networks consider ping traffic to be benign and will allow it to pass through.

ICMP can be altered to act as conduit for evil purposes. Some of the ways that ICMP can be used for purposes other than the intended ones are:

- Reconnaissance
- Denial of Service
- Covert Channel

Reconnaissance

Reconnaissance is the first stage in the information gathering process to discover live hosts and some other essence information as part of most planned attack.

ICMP messages are broadly categorized into two kinds:

ICMP Messages	
ICMP Query Messages	ICMP Error Messages
□ Echo Request and Echo Reply □ Time Stamp Request and Reply □ Information Request and Reply □ Address Mask Request and Reply	□ Destination Unreachable □ Source Quench □ Redirect □ Time Exceeded □ Parameter Problem

By manipulating these ICMP messages, we are able to gather substantial information in the process of information gathering:

- Host Detection
- Network Topology
- ACL Detection
- Packet Filter Detection
- OS Fingerprinting

Host Detection and Network Topology

By using ICMP message, it allows one to identify hosts that are reachable, in particular from the Internet.

Traceroute attempts to map network devices and hosts on a route to a certain destination host. Intelligence use of it will allow one to map the topology of a network.

Access Control List (ACL) Detection

ICMP Error Messages may help to determine the kind ACL of the filtering device is being used and allow one to choose the tactics accordingly.

The idea is to manipulate the total length of the IP Header Field. A crafted packet with total length in the IP Header Filed claiming to be bigger than really what it is. When this packet reaches the host, it will try to grab the data from the area, which is not there. The host will thus issue an ICMP Parameter Problem back to the querying IP address.

If there is a packet filtering device present and we probe a targeted network with all possible combination of protocols and services, it will allow us to determine the access control list of the filtering device (which host is allowed to received what type of traffic).

The crafted packet can use ICMP, TCP or UDP as the underlying protocols.

Protocol/Port Scan

ICMP Error Messages (Protocol/Port Unreachable) are the common ways to determine what type of protocols/ports the host is running.

Nmap 2.54 beta 1 has integrated the Protocol Scan. It sends raw IP packets without any further protocol header (no payload) to each specified protocol on the target machine. If an ICMP Protocol Unreachable error message is received, the protocol is not in used.

OS Fingerprinting

Using ICMP for OS Fingerprinting requires less traffic initiation from the malicious person machine to the target host.

The idea is “Which operating system answer what kind of ICMP Query messages”.

This is possible because different OS implement differently. Some do not compliant strictly to RFC, while RFC may also optional. Fingerprinting of OS can be achieved via the following:

- Using ICMP Query Messages
- Using ICMP Error Messages

The ICMP Echo Request/Reply pair was intended to determine whether a host is alive or not. Negative response will either mean it is not alive or ICMP Echo traffic is filtered by a packet filtering device.

The ICMP Information Request/Reply pair was intended to support self-configuring systems such as diskless workstations at boot time to allow them to discover their network address.

The ICMP Time Stamp Request/Reply pair allows a host to query another for the current time. This allows a sender to determine the amount of latency that a particular network is experiencing. Most operation systems implemented the ICMP Time Stamp Request/Reply.

The ICMP Address Mask Request/Reply pair was intended for diskless systems to obtain its subnet mask in use on the local network at bootstrap time. It is also used when a host wants to know the address mask of an interface. RFC 1122 states that the Address Mask is optional.

At times, the ICMP Error Messages revealed substantial information about the host or network. For example, receiving a Protocol Unreachable will reveal that the host is alive and that particular protocol queried is not supported. By manipulating certain field in the query, we can generate several ICMP Error Messages.

There was done a comprehensive research on the use of ICMP in OS fingerprinting.

Based on the nature of the different implementation of OS, substantiate information can be gathered using different techniques in manipulating the ICMP messages and observe the response of the target host. The techniques are listed below:

a. Response on ICMP Query Messages Types on a targeted host
b. Response on ICMP Query Messages Types on a broadcast address
c. IP TTL value on the ICMP Messages (Request and Reply)
d. Response on ICMP Query Messages with Code Field ≠ 0
e. Response on the ICMP Query Messages with Precedence Bits value ≠ 0
f. Response on the ICMP Query Messages with TOS value ≠ 0
g. Response on the ICMP Query Messages with TOS unused bit = 1
h. Response on the ICMP Query Messages with Reserved Bit Flag = 1
i. Response on the ICMP Query Messages with DF set
j. ICMP Error Message echoing integrity with ICMP Port Unreachable Error Message

A detailed tabulation can be obtained in [1]. We extracted some results and conduct some fingerprint on the following operating systems:

- Solaris
- Linux
- Windows Family (Win 98/NT/2000)

Fingerprinting HPUX 10.20, Solaris and Linux

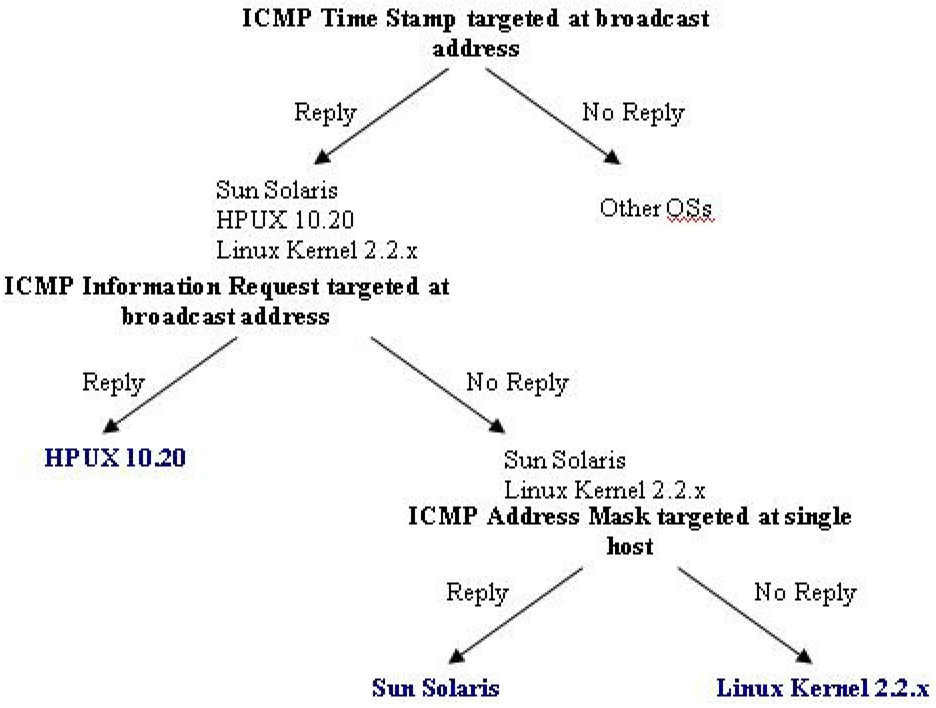

FIGURE 1: Example the technique of fingerprinting HPUX 10.20, Solaris and Linux operating systems.

Fingerprinting Windows Family (95/98/ME/NT/20000).

Using ICMP as a means to cause DoS is not new. CERT/CC has issued an advisory on Denial of Service via Ping in 1996 (CA-1996-26). Ping of Death is one of the common uses of ICMP to cause a machine to crash. Here we mentioned some other well-known DoS using ICMP as a means.

Smurf DoS

The infamous Smurf attack preys on ICMP's capability to send traffic to the broadcast address. Many hosts can listen and response to a single ICMP echo request sent to a broadcast address. This capability is used to execute a DoS attack.

The two main components to the smurf denial-of-service attack are the use of forged ICMP echo request packets and the direction of packets to IP broadcast addresses.

In the "smurf" attack, attackers are using ICMP echo request packets directed to IP broadcast addresses from remote locations to generate denial-of-service attacks. There are three parties in these attacks: the attacker, the intermediary, and the victim (note that the intermediary can also be a victim).

The intermediary receives an ICMP echo request packet directed to the IP broadcast address of their network. If the intermediary does not filter ICMP traffic directed to IP broadcast addresses, many of the machines on the network will receive this ICMP echo request packet and send an ICMP echo reply packet back. When (potentially) all the machines on a network respond to this ICMP echo request, the result can be severe network congestion or outages.

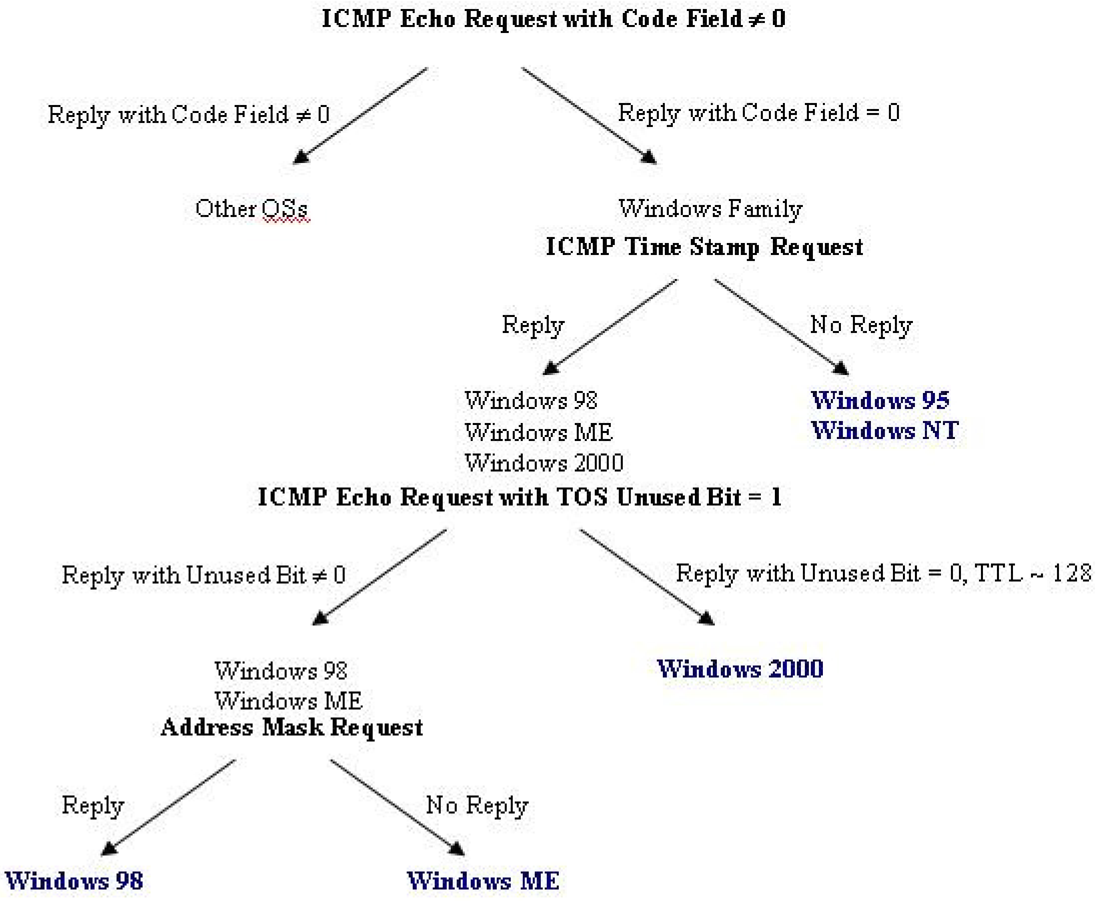

FIGURE 2: Example of fingerprinting the Windows Family.

Denial of Service (DoS)
When the attackers create these packets, they do not use the IP address of their own machine as the source address. Instead, they create forged packets that contain the spoofed source address of the attacker's intended victim. The result is that when all the machines at the intermediary's site respond to the ICMP echo requests, they send replies to the victim's machine. The victim is subjected to network congestion that could potentially make the network unusable.

More detailed description of Smurf attack can be found in [5].

Tribe Flood Network (TFN)
The Tribe Flood Network (TFN) attack is another DoS attack that uses ICMP for communication. 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.

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.

Communication from the TFN client to daemons is accomplished via ICMP Echo Reply packets. Each "command" to the daemons is sent in the form of a 16-bit binary number in the ID field of an ICMP Echo Reply packet (The sequence number is a constant 0x0000, which would make it look like the response to the initial packet sent out by the "ping" command). This is to prevent the kernel on the daemon system from replying with an ICMP Echo Reply packet. The daemon then responds (if need be) to the client(s), also using an ICMP Echo Reply packet. The payload differs with TFN, as it is used for sending command arguments and replies.

Some network monitoring tools do not show the data portion of ICMP packets, or do not parse all of the various ICMP type-specific fields, so it may be difficult to actually monitor the communication between client and daemon.

A detailed analysis of TFN can be found in [6].

WinFreeze
WinFreeze is a DoS attack against Windows.
A small exploit code that can cause a Windows 9x/NT box on the local LAN to freeze completely. The program initiates ICMP/Redirect-host messages storm that appears to come from a router (by using the router's IP). The Windows machine will receive redirect host messages causing it to change its own routing table. This will make it get stuck, or operate very slowly until a reboot is done.

Covert Channel
Many firewalls and networks consider ping traffic to be benign and will allow it to pass through. Use of ping traffic can open up covert channels through the networks in which it is allowed.

Loki
The concept of the Loki is simple: arbitrary information tunneling in the data portion of ICMP Echo Request and ICMP Echo Reply packets.

Loki exploits the covert channel that exists inside of ICMP Echo traffic. ICMP Echo packets have the option to include a data section. This data section is used when the record route option is specified, or, the more common case, (usually the default) to store timing information to determine round-trip times. Although the payload is often timing information, there is no check by any device as to the content of the data. So, as it turns out, this amount of data can also be arbitrary in content as well. Therein lies the covert channel. Most network devices do not filter the contents of ICMP Echo traffic. They simply pass them, drop them, or return them. The trojan packets themselves are masqueraded as common ICMP Echo traffic.

If a host is compromised and a Loki server is installed, it can response to traffic send to it by a Loki client.

Because the programs use ICMP Echo Reply packets for communication, it will be very difficult (if not impossible) to block it without breaking most Internet programs that rely on ICMP. With a proper implementation, the channel can go completely undetected for the duration of its existence. Detection can be difficult. If you know what to look for, you may find that the channel is being used on your system. However, knowing when to look, where to look, and the mere fact that you should be looking all have to be in place. A surplus of ICMP Echo Reply packets with a garbled payload can be ready indication the channel is in use. More information on the Loki project can be obtained in [7].

Filtering ICMP Traffic and the Challenge for the IDS
Network devices requires ICMP Messages for communications. ICMP is a protocol that is supposed to be used to alert hosts of problem conditions or exchange messages. However, using it in a malicious manner allows one to dig out host information and network topology. To use a Network Intrusion Detection System to actively monitor the network for malicious ICMP traffic is laborious. Given this, appropriate filtering of ICMP traffic should be done to minimize the potential threat.

It is therefore important to understanding how operating systems response to ICMP Messages. This will allow us to determine what type of ICMP Messages should only be allow in and out of the network. With appropriate configuration of the packet filtering device to block unnecessary ICMP Messages, potential threats resulting from ICMP Messages can be reduced. This, however, should be done wisely and selectively. For example, incoming “ICMP Error Message, Fragmentation Needed but Don’t Fragment Set”, will be necessary to inform the internal host on such errors and to adjust the datagrams accordingly.

Even with proper filtering of ICMP traffic, NIDS should still be deployed to monitor the kind of ICMP activities. The challenge of the NIDS will be have accurate signatures to detect malicious ICMP traffic.

Host-based IDS is another option. Nevertheless, it still needs “inputs” to monitor the traffic accurately.

Ultimately, human will be required to perform the final analysis of the IDS detects to determine whether detects are legitimate or hostile.

Spoofing ICMP redirect host messages with hping
An icmp redirect host message can be sent from any router on the same broadcast segment as the end host that "needs redirection". Modern network infrastructures will typically have a single router gateway address per subnet however it is possible to have more than one router in a segment making the operational case for ICMP redirect messages.

An ICMP redirect host message has ICMP type 5, code 1. The ICMP redirect network code is 0. There also exists redirect with Type of Service (ToS) for both network and host (codes 2 and 3).

With the advent of classless Internet domain routing (CIDR, RFC 1518/1519 in 1993), an end host cannot readily determine the network class and thus ICMP type 5, code 0 is basically useless. RFC 1812 additionally states that a router should not generate type 5, code 0. While working on this post, I observed that a Windows host will accept code 0 and treat it the same as code 1 adding a /32 route to the table.

Because IP source address spoofing is trivial, ICMP redirect message abuse potential exists. The only specific limitation is that the "new" destination gateway address of the redirect message must exist within the same subnet as the end host itself.

The ICMP redirect use case would most likely be employed in a network penetration testing scenario whereby extensive layer 2 security features are enabled limiting the effectiveness of layer 2 attacks such as ARP cache poisoning and rogue DHCP server use. The primary goal being to intercept traffic for a specific destination address.

The end host must be configured to accept ICMP redirect messages and update its routing table accordingly. Within Microsoft Windows, there is a registry key that enables the acceptance of ICMP redirect messages. This DWORD registry key has a default setting of 0x0001, that being the "enabled" state.

HKLM\System\CurrentControlSet\Services\Parameters\Tcpip\EnableICMPRedirect
Based on my reading, I believe some implementations of the Microsoft TCP/IP stack also read the plural form of this key "EnableICMPRedirects" rather than the singular form, so it is possible that both keys exist.

With regard to the Windows XP firewall, it will block all ICMP requests in its default configuration state. Of course, there may be site wide group policy that changes this situation for legitimate operational reasons such as multiple router gateways existing in a single segment. If you wish to experiment, and enable 'icmp redirect' from the command line, there are two useful 'netsh' commands as follows:

```
C:\> netsh firewall show icmpsetting
```

shows the current state of ICMP acceptance if any. A blank output indicates that no ICMP policies are in effect.

```
C:\> netsh firewall set icmpsetting type=5 mode=enable
```

will enable the acceptance of ICMP redirects through the firewall.

The Linux kernel has two settings that control ICMP redirect acceptance behavior. For the 'eth0' interface, these settings are as follows:
/proc/sys/net/ipv4/conf/eth0/accept_redirects
/proc/sys/net/ipv4/conf/eth0/secure_redirects

If "secure_redirects" is enabled, the Linux system will only accept ICMP redirects that are redirected to a default gateway that is already listed in the routing table. This is the default in most modern linux distributions and is an effective defense against spoofing attempts.

'accept_redirects' is enabled as the default also. If the 'secure_redirects' kernel parameter is set to 0, then the linux kernel is susceptible to an ICMP host redirect attack in the same way that a Windows system is susceptible. The one thing to note is that the linux kernel will not show the accepted route in the routing table that is listed through 'route show' or 'netstat -nr' commands, even though the route is in effect.

Our scenario below is laid out as follows:

Attacker IP Address: 172.16.235.99
Legitimate Router Gateway: 172.16.235.1
Victim IP Address: 172.16.235.100
The legitimate DNS server address is 10.1.1.1.

We can use ICMP redirect host to insert a new route table entry for the 10.1.1.1 address as follows:

```
hping -I eth-dest -C 5 -K 1 -a 172.16.235.1 --icmp-ipdst 10.1.1.1 --icmp-gw 172.16.235.99 --icmp-ipsrc 172.16.235.100 172.16.235.100
```

whereby:
-I eth-dest is the destination ethernet interface on the attacker to send the packets out of/from.
-a is the spoofed source address of the legit. router gateway
--icmp-ipdst is the new route table entry address you want to create
--icmp-gw is the new route destination address/gateway you want to create and must live within the same subnet as the victim.
--icmp-ipsrc must match the source address of the victim to pass sanity checking

If you check the route table on the victim using "netstat -nr" or "route print" after executing this command from the attacker, you should see a new route table entry. Since MS-Windows will readily accept these new route table entries, many ICMP redirects can be generated with random IPv4 prefixes to perform a denial of service against the target. A /32 host route learned via an ICMP redirect message will remain in the routing table for 10 minutes.

In this example, we assumed that the attacker was on the same subnet in order to receive / intercept the traffic. In other words, the attacker would also be a DNS server ready to serve some bogus response to the victim. The attacking host could well be a different machine on another network, but the "man in the middle" host or "router gateway" if you like needs to remain on the same subnet in order to receive the traffic.

2.3 ICMP ADDRESS MASK PING

An ICMP address mask request is an antiquated ICMP method that queries an IP gateway for an appropriate subnet mask on the current IP subnet. This ICMP type is extremely rare, and the traffic pattern is very obvious when observing network traces.

ICMP Address Mask Ping Operation
The ICMP address mask ping operates by sending an ICMP address mask request to a remote device. If this device is configured to reply to the ICMP address mask, it will send an ICMP address mask reply:

```
Source        Destination    Summary
--------------------------------------------------------------------------------------
[192.168.0.5]  [192.168.0.67] ICMP: C Get address mask
[192.168.0.67] [192.168.0.5]  ICMP: R Address mask = [255.255.255.0]
```

If the remote device isn't active or the remote device does not respond to ICMP address mask requests, no response will be seen and the ping will fail.
If the nmap is not running as a privileged user, the –PM option provides the following warning:

```
Warning: You are not root -- using TCP pingscan rather than ICMP
```

The ping process then continues with a TCP connect()-style ping.

Advantages of the ICMP Address Mask Ping A successful ICMP address mask ping can be indicative of an older or unprotected TCP/IP stack or gateway. Most modern operating systems and routers will not respond to this request, or (at the very least) they will not respond to this request from systems that are not on the same subnet. This ping could be useful as a filtering mechanism, since it would identify all systems on the network that have older or unusually open TCP/IP protocol stacks.

ICMP doesn't rely on any particular networking service or application. It's common for ICMP to respond to a request without any particular open or available ports on a system.

Disadvantages of the ICMP Address Mask Ping

The ICMP address mask request is an unusual frame, and it's rarely seen in normal network traffic. When looking at network trace files, the ICMP frames requesting address masks are very obvious.

This ICMP ping type doesn't work on most modern systems, which means that this ping will often fail. If it's important to find active systems, this method won't provide a high percentage of successful pings.

his ping type won't work at all unless the nmap user is privileged. If the nmap user isn't privileged, the ping type will change to a TCP connect()-style ping. Although there is a warning when this occurs, there's no option to stop the scan. Since this ping type doesn't accept a port number variable, this change to a TCP connect()-style ping will only run on the default port of 80. If there's an active web server on the destination station, this uncontrolled ping change will result in the initialization of an application session on the remote device.

ICMP is a difficult protocol to transmit through firewalls and packet filters. Since ICMP is often filtered, this ping has a low percentage of operation through firewalls.

When to use the ICMP Address Mask Ping The ICMP address mask ping is useful on networks that contain older operating systems or gateways.

The successful ping trace shown above was performed against a system using an older version of the VxWorks operating system.

- This address mask ping is only useful on networks that allow for the free flow of ICMP frames. If the link contains firewalls or packet filters, a better choice would be a non-ICMP-based ping type.
- If the nmap user is non-privileged, this ping type will revert to a TCP connect()-style ping. Since this ping type doesn't allow port specifications, a better ping choice for non-privileged users would be the TCP ACK ping (-PA) or the TCP SYN ping (-PS).

2.4 Cyber Project Plan

Organizations of all sizes and across all industries are potential victims of security attacks. According to the2013 Cost of a Data Breach: Global Analysis*, 35% of incidents can be attributed to the human factor.

Understanding that knowledge is power and employee education is an imperative line of defense, DDI is helping organizations build and sustain a culture of security through a comprehensive education program that encompasses engaging content/test components, corporate communications, employee incentives, and efficient delivery and tracking of employee status and testing results.

Just as organizations are forced to take precautions to protect their networks from attack, it is essential that they effectively educate their employees to help build a culture of security.

Through the DDI *Cyber Crime Prevention & Safety Program*, organizations can implement a security campaign that will help train, educate and reinforce a security aware employee base.
Generate Awareness

- Generate maximum awareness and support through a comprehensive campaign kick-off and strategic planning customized to the needs of the organization.
- Introduce the SecurED® Security Training & Awareness Video Library to all employees giving them access to content that can be viewed on any device that supports video including PC, laptop, iPad/tablets and mobile phones.
- Leverage the power of a Security Communications Kit custom designed to effortlessly engage your employees.
- Benefit from a DDI Security Ambassador to assist with planning, implementation, marketing, campaign analysis and reporting.

This fully integrated security campaign empowers organizations to invest in security awareness without burdening internal resources.

- Engaging training topics (Social Engineering, Password Development, Mobile Device Security, Safe Web-Browsing Habits, and many more)
- Security training that is informative and fun
- Keep security awareness top of mind throughout the year
- Pause and play videos from virtually any device
- Easily integrated into your existing LMS
- Gain certification and accreditation

2.5 Security Flaws in IPv6

This section covers IPv6 related security issues, which came to my attention during the research. The following issues are included:

• Neighbor Discovery Protocol issues
– Neighbor Solicitation / Neighbor Advertisement spoofing
– Redirect spoofing
– Router Solicitation / Router Advertisement spoofing
– Duplicate Address Detection attack
– Neighbor Advertisement flooding
– Router Advertisement flooding
• IPv6 smurfing
• Routing header type 0
• Implementation issues of IPv6
• Transition techniques issues
– Dual-stack network issues
– Tunneled IPv6 network issues
– Low user and administration awareness of IPv6 autoconfiguration

In this report I have proposed detection and mitigation methods for the issues, when applicable. Along with some of the covered issues I have included examples using the THC IPv6 toolkit.

Neighbor Discovery Protocol issues

RFC3756 [26] describes the IPv6 NDP as a mechanism used by nodes in an IPv6 network to learn the local topology. This includes the IP to MAC address mappings for the local nodes, the IP and MAC addresses of the routers present in the local network, and the routing prefixes served by the local routers. NDP uses five Internet Control Message Protocol version 6 (ICMPv6) packet types:

• Neighbor Solicitation (NS)
• Neighbor Advertisement (NA)
• Router Solicitation (RS)
• Router Advertisement (RA)
• Redirect

According to Scott Hogg and Eric Vyncke (“IPv6 Security” [5]) there is no authentication mechanism built into ICMPv6 and those packets can be spoofed. This is a flaw, which can allow an attacker to perform malicious activities such as traffic redirection and DoS. NDP related issues have only local impact, because routers don’t forward NDP messages. However I consider them a serious threat, because flat IPv6 networks can be much larger compared to IPv4 networks. Details about the NDP packet types and possible attacks are described below. Neighbor Solicitation / Neighbor Advertisement spoofing NS / NA packets function in a similar way as ARP in IPv4. The basic mechanisms of the attack are described by US-CERT [27]: After receiving a neighbor solicitation request from a system that is on-link and is using a spoofed IPv6 address as the source address, a router will create a neighbor cache entry.

When this entry is made, some IPv6 implementations will create a Forwarding Information Base (FIB) entry. This FIB entry may cause the router to incorrectly forward traffic to the device that sent original spoofed neighbor solicitation request. I consider the NS / NA spoofing attack similar to the well known ARP spoofing. The parasite6 tool from THC IPv6 toolkit redirects all local traffic to the attacker's system by answering falsely to NS requests. The following command is used to perform the attack on interface eth0: parasite6 eth0 IP forwarding should be enabled on the attacking machine or the redirected traffic can cause DoS. 19Router Solicitation / Router Advertisement spoofing IPv6 hosts can auto-configure when connected to a routed IPv6 network using Stateless Address Autoconfiguration (SLAAC). SLAAC is stateless compared to DHCP, because a DHCP server stores a state - the leased IP addresses. SLAAC is based on RS / RA messages exchanged between the router and hosts. RFC1256 [29] describes the method used for router discovery: Each router periodically multicasts a Router Advertisement from each of its multicast interfaces, announcing the IP address(es) of that interface. Hosts discover the addresses of their neighboring routers simply by listening for advertisements. When a host attached to a multicast link starts up, it may multicast a Router Solicitation to ask for immediate advertisements, rather than waiting for the next periodic ones to arrive. RS / RA messages enable hosts to discover the existence of local routers, but don't provide data which router is a better choice for reaching a particular destination. If a host chooses not the optimal first-hop router for that destination, it should receive an NDP Redirect from the first chosen router, suggesting a better first-hop. Some of the fields that a RA packet contains are the local network prefix, link-local address of the router and router priority. An attacker can advertise a fake router by sending spoofed RA packets. As a result a host can receive a RA for a router different from the one expected by the host or for a non existing router. The attack is illustrated on Figure 5. Figure 5: Router Solicitation / Router Advertisement spoofing 1. The host doesn't have a default router. In order to learn one, the host sends a RS packet. 2. The router replies to the host with a RA packet. 3. The attacker also replies and claims to be a router, with higher priority. 204. The host chooses the attacker as a default router. One can compare this attack to a rogue DHCP server in IPv4, because of the similar outcome. The Router Solicitation / Router Advertisement spoofing attack is covered in details in RFC3971 [30]. The THC IPv6 toolkit contains the tool fake router6, which can set any IP address as a default router, define network prefixes and DNS servers. An example of router advertisement with fake router6 on interface eth0: fake_router6 eth0 2001:610:158:1020:226:55ff:fecd:8f84/64.

Redirect spoofing Redirect is an NDP mechanism used by routers to inform a host for a better route to a particular destination. Routers should detect if a host on the local network has made an inefficient first-hop routing decision and then recommend a better first-hop. The NDP Redirect has a simple security mechanism - a copy of the packet, which caused the redirection, must be included in the NDP Redirect message. An attacker cannot blindly spoof a Redirect message, because the victim will not accept it. This can be bypassed by sending a spoofed ICMPv6 echo request to the victim, where the source IP address is set to the IP address of the router the victim is using. The victim will send an ICMPv6 echo reply to the router. The attacker can predict the content of the reply message and use it to craft an NDP Redirect message, which advertises another system as a better first-hop to the router. This attack vector is not possible in IPv4. Figures 6 and 7 illustrate the attack. Figure 6: Redirect spoofing - The attacker sends ICMP echo request message 1. The attacker sends to the victim an ICMPv6 echo request, with a spoofed source address claiming to be originating from the router. 2. The victim sends the router an ICMPv6 echo reply.

The attacker can sniff the packets and forward them to the router in order to stay undetected. The redir6 tool from the THC IPv6 toolkit is an implementation of NDP Redirect spoofing. The tool accepts the following parameters: redir6 <interface> <source-ip> <target-ip> <original-router> <new-router> [new-router-mac] Duplicate Address Detection attack In IPv6 networks it is not allowed several hosts to share an IP address. To prevent duplicate IPv6 addresses, a host must check whether an IPv6 address it intends to use is free or already used by another host. This procedure is called Duplicate Address Detection (DAD). Since no higher layer traffic is allowed until a host has obtained an IP address, DAD relies on NS and NA messages in order to check if an IP address is in use. An attacker could launch a DoS attack by responding to duplicate address detection attempts made by a newly connected host. If the attacker claims every IP address, then the host will not be able to obtain an address at all. The DAD attack doesn't have an analogue in IPv4. The attack is described in RFC3756 [26]. dos-new-ipv6 is the implementation of this attack in the THC IPv6 toolkit. The tool accepts only one parameter - the network interface: dos-new-ip6 eth0 Neighbor.

Advertisement flooding Routers can store a limited number of ND cache entries. In case a router is flooded with NA packets, the flood can result in exhaustion of the resources causing the router to crash or become slower and eventually filling up the entry table. When the table is full, a router cannot learn new ND entries or can even 22cause old (legitimate) entries to be overwritten. This attack is evaluated by Jeff S. Wheeler in his presentation "IPv6 NDP Table Exhaustion Attack" [31]. The same effect as the described NA flooding can be achieved unintentionally, in case a host is using a random IPv6 address for every outgoing TCP connection as a privacy and security mechanism. This aspect was published in the draft of RFC3041 [32].

NA flooding is comparable to MAC flooding of a network switch, where the content addressable memory table of the switch is overfilled and the switch starts operating as a hub. The two attacks are not similar, because of their different outcome. The THC IPv6 toolkit includes the flood advertise6, which floods a target network with random NA messages. The tool requires only a network interface as a parameter: flood_advertise6 eth0 Router Advertisement flooding When receiving a RA messages from different routers and announcing different network prefixes, hosts and routers update their network knowledge according to the content of the messages. Although this activity is computation intensive, it is not likely that many routers will be sending RA messages in an average network. But if an attacker floods the local network with random RA messages, this will result in consumption of the available resources of the systems in the local network. RA flooding will make the systems unusable and unresponsive. Marc Heuse listed several operating systems with IPv6 and SLAAC enabled by default, which are known for being vulnerable to this issue [33]. The most notable one is the Microsoft Windows series, including the latest version - Windows 7.

According to the document, where Marc Heuse lists the vulnerable systems, Microsoft are aware of the RA flooding security issue, but they do not plan to release a fix for the issue. The router advertisement flooding is IPv6 specific vulnerability. This vulnerability can be exploited using the flood router6 tool from the THC IPv6 toolkit. The following command sends router advertisements on interface eth0: flood_router6 eth0 Detection of Neighbor Discovery Protocol attacks The ICMPv6 based attacks are local to a subnet. This implies that detection mechanisms cannot be centralized in a single IDS responsible for a large network. Decentralized solution with access to every subnet in a network is necessary to detect NDP attacks. NDPMon is an application, which monitors NDP traffic and can notify a network administrator if a host on the network spoofs NDP packets. The program is similar to arpwatch used for detection of ARP spoofing in IPv4. NDPMon can monitor NS and NA packets and detect if a new NA message is conflicting with a previous one, which is a sign of possible spoofed NA message. 23In order to detect a fake Router.

Advertisement, an IDS can compare the source IP address and MAC address of the router, which sent the RA, to a list of known routers. A draft for RFC [34], published in 2005, suggests changes to RFC2461 [35]. They include a method for detection of "exploitation of inherent vulnerabilities in the Neighbor Discovery processes". This method forces NDP packets to be multicast only to the host's Solicited Node Multicast group, thus allowing a security device to detect attacks. The draft proposes a solution for the NA / NS spoofing and host Redirect issue, but solution for RA / RS problems are not discussed. Proposed method for detection of Neighbor Discovery spoofing: • Neighbor Advertisements must be sent to the recipient's Solicited-node Multicast Address • Require that a node shall silently discard Neighbor Advertisements that are not addressed to the node's SNA.

Proposed method for detection of host Redirect: • Require host Redirect messages to be sent to the destination node's SNA. • Require that a node shall silently discard Host Redirection packets that are not addressed to the node's SNA. Apart from the above, I suggest the following additional measures for detection of Neighbor Advertisement flooding: • The NDP entry table of the router can be monitored. In case the table is filling up faster, than its entries are expiring, a notification can be sent to the network administrator. • If possible a list of trusted devices can be implemented on the router, which gives them a priority over the rest hosts, that send ND messages in a network. When a certain limit in the NDP entry cache table is reached, only messages from those trusted devices will be processed and the rest messages will be ignored. RA flooding can be detected in a similar way as NS flooding. I suggest monitoring the number of new RA messages. If unusual large numbers of routers advertise new prefixes, most likely the network is being flooded. The network administrator can make a list of trusted routers and when one or several not listed routers advertise themselves, a script can notify the administrator.

Mitigation of Neighbor
Discovery Protocol attacks RFC3756 [26] recommends the use of IPsec for authenticating NDP message. However the RFC doesn't provide detailed implementation instructions. Due to the requirement of manual configuration of IPsec, SEcure Neighbor Discovery (SEND) was developed (specified in RFC3971 [30]). SEND adds new options to NDP that make it more secure. The security of SEND is based on signing the 24NDP messages using RSA Public key signatures and the use of cryptographically generated addresses. NDP packets, which are not signed, are treated as unsecured. During the research I studied possible ways to reduce the likelihood of IPv6 NDP vulnerabilities from being exploited. One of the effective solutions is segmenting the network by assigning a unique prefix to every router interface or by implementing Virtual LANs. In this way an attacker won't be able to affect a large number of hosts. Another possible solution is reducing the subnet size. This can limit the number of possible hosts under the maximum capacity of the router's entry table, which will mitigate NA flooding. However this solution is not compatible with SLAAC, because SLAAC requires at least a /64 subnet. Additionally smaller subnets can allow an attacker to enumerate hosts easier. To improve the robustness against Man in the Middle attacks, the administrator can configure application and transport layer encryption (TLS, SSH tunnels, etc.), because the encryption can prevent third parties from viewing the intercepted network traffic.

For mitigation of Neighbor Advertisement flooding I suggest the size of IPv6 entry cache table to be increased to a value allowing reasonable time for reaction. This should be implemented along with other measures such as throttling the host learning speed during the attack. The throttling will allow a router to continue serving the old, known hosts and learn a limited amount of new hosts. Router Advertisement flooding can be mitigated by throttling the learning speed of hosts. 3.1.2 IPv6 smurfing The IPv4 smurf attack is a way of generating significant traffic on the victim's network. It is an amplified attack, in which an attacker sends an ICMP echo request with spoofed source address to the broadcast address. All hosts, which receive the request, will reply to the source IP, thus generating traffic and possibly cause a DoS. IPv6 does not use broadcasting as a form of communication. However, IPv6 relies on multicasting, and multicast addresses might also be used for a smurf attack. This makes the differences between IPv4 and IPv6 smurfing small. A simplified illustration of IPv6 smurfing can be seen on Figure 8. Figure 8: Smurf attack 251. The attacker sends an ICMPv6 echo request packet with spoofed source address to a multicast address. 2. The hosts, which received the request send an reply to the victim, which can overload the victim's network connection. The attacker can send packets to the link-local all nodes multicast address (FF02::1) and the link-local all routers multicast address (FF02::2) for performing the smurf attack on IPv6.

These two addresses identify the group of all nodes and routers in the scope of the local subnetwork. The THC IPv6 toolkit features the smurf6 and rsmurf6 tools. The differences between the two tools are listed below. The smurf6 tool sends ICMPv6 echo request packets with spoofed source (using the victim's IP address) to the multicast address FF02::1. The hosts on the LAN that are vulnerable to the attack send ICMPv6 echo reply packets, which flood the victim. The victim of smurf6 can be on the local subnet with the attacker or on a remote subnet. smurf6 eth0 2001:610:158:960::100 rsmurf6 uses a different approach. It sends ICMPv6 echo reply packets that are sourced from FF02::1 and destined for remote computers. If the destination system is allowed to respond to packets sourced from a multicast address, the response causes a traffic flood on the remote LAN. This attack has stronger amplification, because each packet generated by rsmurf6 can generate large amount of packets on the remote LAN. rsmurf6 eth0 2001:610:158:960::100 Most modern IPv6 implementations are protected against this vulnerability and drop multicast packets, which can cause smurfing. Scott Hogg and Eric Vyncke recommend that IPv6 hosts should not be responding to echo request packets destined to a multicast group address [5]. Possible protection against remote smurf attacks can be ingress filtering, which rejects the attacking packets on the basis of the forged source address. 3.1.3 Routing header type 0 RFC2460 [36] defines an IPv6 extension header called Routing Header. The value of this header can be set to a specific type as defined in the RFC. The header type "0" (known as RH0) forces a packet to follow a strictly predefined path between network nodes. This feature allows RH0 to be used for amplification attack.

RFC5095 [37] explains the attack: A single RH0 may contain multiple intermediate node addresses, and the same address may be included more than once in the same RH0. This allows a packet to be constructed such that it will oscillate between two RH0-processing hosts or routers many times. This allows a stream of packets from an attacker to be amplified along the path between two remote routers, which could be used to cause congestion along arbitrary remote paths and hence act as a denial-of-service mechanism.

This attack is particularly serious in that it affects the entire path between the two exploited nodes, not only the nodes themselves or their local networks. Analogous functionality may be found in the IPv4 source route option, but the opportunities for abuse are greater with RH0 due to the ability to specify many more intermediate node addresses in each packet. According to information collected after the CanSecWest/core07 talk, several major operating systems and network vendors are vulnerable to this issue [38]. The operating systems, which are not vulnerable to the RH0 amplification attack did not implement RH0 or implemented it not according to the IETF standard. Another possible malicious use of RH0 is to bypass firewalls that prohibit outside access to a host in the internal network. An attacker can send a packet with RH0 through the firewall to a router, which will redirect it to the target host. The possible security problems that RH0 can cause were considered by IETF. RH0 was deprecated with RFC5095 [37] from December 2007. The RFC recommends using ingress filtering until routers are updated. The ingress filtering should be applied as recommended in RFC2827 [39] and RFC3704 [13]. If a whole network has to be protected, the ingress filtering should be implemented on the border, where the network is connected to the outside world. 3.1.4 Implementation issues of IPv6 A large percentage of the IPv6 issues listed in the National Vulnerability Database are not related to the design of the protocol, but are a result of insecure implementations [40]. Vulnerabilities, which are exploited using flaws in the IPv6 protocol, were not considered implementation specific during the research.

Based on information collected from the 98 vulnerabilities listed by NVD between October 2002 and June 2011, one can conclude that most of the implementation issues can result in: • DoS • Security policies bypassing • Buffer overflow These vulnerabilities can occur in the network stack of the device OS / firmware or in a specific piece of software installed on the device. The vulnerabilities, which allow bypassing of the security policies, are caused by none or insufficient filtering of the IPv6 packets compared to IPv4 [40]. Two vulnerabilities listed by NVD are not IPv6 specific. They apply also to IPv4 and are caused by bugs in the products. When implementation specific vulnerabilities are discovered, usually they are fixed by the vendors in the newer versions of their products. IPv6 implementations are relatively new and are not tested in production environments as thoroughly as the IPv4 implementations. With the wider adoption of IPv6 it is likely that the bugs will be fixed and IPv6 implementation will be as good as IPv4.Operating systems with IPv6 enabled by default can be considered vulnerable, because an attacker can advertise a rogue router, which will be automatically configured on the host. According to SixXS these operating systems include the latest versions of Windows, Mac OS and most Linux distributions [41]. A penetration tester can check for implementation specific issues by identifying an affected product by its fingerprint. nmap and Nessus are tools with large fingerprint databases,

which can recognize vulnerable versions of software. Nessus is able to detect IPv6 specific issues listed in the Common Vulnerabilities and Exposures (CVE) database. Along with IPv6 specific implementation issues, there can be issues which apply only when IPv6 is used along with IPv4 during the transition period. 3.1.5 Transition techniques related issues The switch between IPv4 and IPv6 cannot happen instantaneously. A migration period is necessary, during which the two protocols will coexist allowing users to be connected to both IPv4 and IPv6 networks. During this phase, transition techniques like dual-stack, tunneling and translation will be used. However these transition mechanisms can introduce security issues, discussed in this section. During the transition phase, users and administrators have to consider both, IPv4 and IPv6 issues and combination of attacks using both protocols. For example an attacker can compromise a remote system through IPv6 vulnerability and perform ARP spoofing on the IPv4 network connection of the compromised system. Dual-stack network issues A dual-stack system can be less secure compared to a single stack (either IPv4 or IPv6), because an attacker has more possible attack vectors to exploit. Also it is more difficult for a system administrator to secure both IPv4 and IPv6 networks on adequate level. Firewalls may not be enforcing the same policy for IPv4 as for IPv6 traffic, which could be due to misconfiguration of the firewalls. It is possibility for firewalls to have more relaxed policy for IPv4 or IPv6, thus allowing unfiltered traffic to pass through. In 2007 ICANN did a survey on the of IPv6 support in commercial firewalls [42].
The results show that the support of IPv6 was low at that time and traffic could go through unnoticed. A new survey was conducted by ICANN in 2010, but the results are not published yet. These issues can be individual for every system and configuration. They can be detected by scanning the hosts and firewalls for opened port and if the same rules are enforced for both IPv4 and IPv6 networks. Tunneled IPv6 network issues A host using a tunneled IPv6 connection over a native IPv4 connection can be more vulnerable compared to a dual-stack host. Ryan Giobbi [43] shows examples how the encapsulated IPv6 traffic can pass unnoticed by firewalls creating security vulnerability. The tunneling software requires opening a port in the firewall that can be used for attacks, unless tunnel-aware firewall is in 28place. According to the RFC draft “Issues with Dual Stack IPv6 on by Default” [44] a poorly configured or implemented VPN may redirect traffic from a protected VPN network to an unprotected IPv6 interface, causing security issues. A penetration tester can detect if data, which is normally blocked by the firewall, will pass through a tunnel in the firewall. Low user and administration awareness of IPv6 autoconfiguration The operating systems, which have as a feature IPv6 enabled by default (e.g. Windows 7 and Linux distributions with kernel version higher than 2.6 [41]), can autoconfigure without the knowledge of their user or system administrator. If security mechanisms and policy are not in place to protect against IPv6 based attacks, a host might get compromised through an IPv6 network. Mitigation of transition techniques related issues I believe that administrators (and maybe users as well) should be educated about the features and required security policies of IPv6. I recommend that IPv6 is disabled if administrators or users are not planning to use it or haven’t implemented protection from threats originating from their IPv6 networks. The security policy implemented in firewalls, VPN software, or other devices, should take a stance whether it applies equally to IPv4 and IPv6 traffic. The “Issues with Dual Stack IPv6 on by Default” RFC draft [44] covers such issues and recommends the implementation of more complex techniques for mitigation: There is still a risk that IPv6 packets could be tunneled over a transport layer such as UDP, implicitly bypassing the security policy. Some more complex mechanisms could be implemented to apply the correct policy to such packets. This could be easy to do if tunnel endpoints are co-located with a firewall, but more difficult if internal nodes do their own IPv6 tunneling. A shorter transition period will minimize the time during which systems could be vulnerable to transition technique specific issues. I consider important that IPv6 is deployed fast so it can become the most used version of IP and minimize the transition period.

Traditional host discovery via network scanning won't work with IPv6, but alternative methods are available

IPv6 brings some welcome security and other features, but there are some 'gotchas' for IP professionals that may not be immediately apparent.

The next generation IPv6 protocol has been "coming soon" for the last decade and is finally nearing the point of necessity as IPv4 addresses get closer to exhaustion. Many hail it as the next great thing for security because of nifty features like native IPSec support.

But it will also bring challenges for security pros, namely in vulnerability scanning and penetration testing. With the addition of all of this new IP space afforded by IPv6, scanning each IP to determine which hosts are up, and then performing a vulnerability scan, would take years. Host discovery through traditional means of network scanning -- host by host and subnet by subnet -- will go away. Instead, new host discovery methods will need to be put in place to make vulnerability scanning more targeted.

Fortunately, there are several techniques that can be used based on existing hardware and software tools. With little to no extra cost, it's possible to determine which IPs are in use on the network without scanning. You can then feed the discovered IPs into the vulnerability scanner so the scanner can spend more time on vulnerability scanning, and not on host discovery.

When looking for other hosts on the local network segment, IPv4's ARP was the usual method, but it's going away with IPv6. ARP's functionality is being replaced by IPv6's new Neighbor Discovery Protocol (NDP) within ICMPv6. There are several different functions in NDP, but in relation to host discovery, it can be used to discover other IPv6 hosts on the local network segment.

The limitation with NDP is it only finds hosts on the local network segment, which is great if you have a flat network. But in large, geographically diverse enterprises, it doesn't work across routers. To work around the problem segmented networks pose, you can issue those commands through a host on each network segment, or from a router with access to each segment.

From a penetration testing perspective, NDP can be used once access is gained into the network through physical access or compromise of an internal host. Attackers can leverage NDP to start finding other juicy targets to compromise that could provide deeper penetration into the network.
Network flow data is a commonly overlooked source of valuable information in a network. It provides a record of all network traffic on the network without recording the content. Most business-class and practically all enterprise-class routers and layer 3 switches support exporting network flow data. Using network flow records, it is easy to identify all hosts that have been communicating on the network during any given time period.

Armed with the hosts enumerated through network flow data, vulnerability scanning can commence and target only those hosts that are communicating during the last hour.

Some commercial vulnerability scanning and vulnerability management solutions such as Tenable's Security Center already have features to scan known hosts and new hosts as they are seen for the first time. Similarly, most network behavioral analysis products use network flow data and can initiate a vulnerability scan of hosts seen for the first time or acting outside of their normal baseline of behavior.

IPv6 IP addresses also are going to cause the Domain Name System (DNS) to become increasingly important, as IP addresses become nearly impossible to remember. In Microsoft Active Directory domains, DNS records are updating dynamically, so hostnames can be provided to a scanner instead of a list of IPs.

Interrogating DNS services using tools like Fierce will also become essential for penetration testers outside the network who don't have access to NDP and network flow data.

The Dynamic Host Configuration Protocol (DHCP) will also serve as a good source of host information for those networks using DHCPv6 to allocate site-local IP addresses to their internal hosts. DHCP server logs can be queried to see which IPs have been allocated, and to limit scans to only those IPs.

Systems and network administrators who are holding off as long as possible to deploy IPv6 are sticking their heads in the sand and ignoring the reality that IPv6 is already on their network. The inclusion of a fully functional IPv6 network stack in all modern operating systems including Windows, Mac OS X, and Linux means that even though IPv6 may not be routed across the network and Internet gateway, it can still be used for attacks on local network segments.

Penetration testers have realized this for several years, and tools like the Metasploit Framework have supported IPv6 since 2008. Once a host is compromised, IPv6 can be used to connect and exploit other systems on the local network segment. The root causes: host-based firewalls may not support IPv6, or the systems administrators didn't realize services often will listen on IPv4 and IPv6 and they only bothered to lock down the IPv4 side.

Intrusion detection and prevention systems (IDS/IPS) may also be leaving security teams blind to internal attacks carried over IPv6. Commercial and open source IDS and IPS solutions are adding support for IPv6 but it doesn’t exist across the board yet. Snort, for example, is an open source IDS on which many commercial products are based and it includes support for IPv6. But support is not enabled by default, and many of the available reporting tools do not support IPv6.

Whether you think you're running IPv6 or not, it's time to prepare for the impact it will have on your security efforts now and in the near future. Vulnerability management efforts will need to adapt, and hardening of hosts will require ensuring that both IPv4 and IPv6 are both locked down.

IPv6 Security Testing and Monitoring Tools

The tools below are useful for system administrators to test and monitor the security of their IPv6 networks. They can be used for IPv6 troubleshooting, intrusion detection and security audits — or for exploiting IPv6 vulnerabilities. They have been freely available on the Internet for a long time to anyone who wants them, including crackers, spammers, black hats, white hats, and national security services.

IPv6 is already available on all modern operating systems and network devices. It can be used today by those who seek to bypass firewalls, steal data, consume resources or simply eavesdrop. Significant amounts of IPv6 traffic now circulate on networks worldwide, and the software below can be used to diagnose IPv6 security vulnerabilities.

Please be certain you have the appropriate rights and permissions to access any networks on which you use this software. IPv6Now provides these links as an IPv6 educational resource and accepts no liability for their use in any way.

Intrusion Detection and Network Monitoring

Security Onion is a Linux distribution for intrusion detection and network security monitoring. It is based on Ubuntu and contains numerous security tools. The Setup wizard builds an army of distributed sensors for an enterprise in minutes. Security Onion provides visibility into network traffic and context around alerts and anomalous events. It seamlessly weaves together three core functions: full packet capture, network-based and host-based intrusion detection systems, and powerful analysis tools.

Network Inventory and Security Auditing

Nmap (Network Mapper) is a free and open source utility for network discovery and security auditing. It uses raw IP packets to determine what hosts are available on the network, what services (application name and version) those hosts are offering, what operating systems (and OS versions) they are running, what type of packet filters/firewalls are in use, and dozens of other characteristics. It was designed to rapidly scan large networks. Nmap runs on all major computer operating systems.

Website Security Scanner

Qualsys Website Scan checks websites for vulnerabilities, hidden malware and SSL security errors (requires registration, 10 free checks).

- Supports Internet Protocol version 6 (IPv6).
- Scans web servers and apps on the Internet or inside networks.
- Detects vulnerabilities and review ways to fix them.
- Finds malware uploaded by malicious users.
- Verifies SSL is properly configured and working.

Firewall Testing

FT6, Firewall Tester for IPv6 is a tool for examining how firewalls handles IPv6.

- ICMPv6 Filtering: verifies the firewall is able to filter and forward certain ICMPv6 Messages.
- Type 0 Routing Header: Checks for Type 0 Routing Header (RH0), has been deprecated due to security concerns.
- Header Chain Inspection: Extension Header tester. ft6 sends a selection of valid and invalid packets.
- Overlapping Fragments: firewall should be able to drop overlapping fragments but still permit non-overlapping fragments.
- Tiny Fragments: no TCP or UDP header in the first fragment. Firewall must wait to decide whether to forward or drop the packet.
- Tiny Fragments Timeout: too many tiny fragments will cause the firewall's reassembly buffers to fill, can lead to DoS.
- Excessive Hop-By-Hop Options: each Hop-By-Hop Option should occur at most once in any IPv6 packet. ft6 sends a variety of duplicates.
- PadN Covert Channel: the PadN Option aligns at 8-byte boundaries. The padding bytes could be used to send messages covertly.
- Address Scopes: multicast addresses are not to be used as source addresses and link-local addresses should not be forwarded.

Troubleshooting Toolset

The IPv6 Toolkit is a set of IPv6 security/trouble-shooting tools that can send arbitrary IPv6-based packets. Supported on FreeBSD, NetBSD, OpenBSD, Linux and Mac OS.

- addr6: an IPv6 address analysis and manipulation tool.
- flow6: performs a security assessment of the IPv6 Flow Label.
- frag6: performs and assesses IPv6 fragmentation-based attacks.
- icmp6: performs attacks based on ICMPv6 error messages.
- jumbo6: assesses potential flaws in the handling of IPv6 Jumbograms.
- na6: sends and assesses Neighbor Advertisement messages.
- ni6: sends and assesses ICMPv6 Node Information messages.
- ns6: sends and assesses Neighbor Solicitation messages.
- ra6: sends and assesses Router Advertisement messages.
- rd6: sends and assesses ICMPv6 Redirect messages.
- rs6: sends and assesses Router Solicitation messages.
- tcp6: sends TCP segments and performs TCP-based attacks.
- scan6: An IPv6 address scanning tool.

Penetration Toolset

THC-IPv6 is a complete toolset to attack the inherent protocol weaknesses of IPv6 and ICMP6. Partial list of tools:

- alive6: detects all systems listening to an address.
- detect-new-ip6: detect new ip6 devices which join the network.
- exploit6: known ipv6 vulnerabilities to test against a target.
- denial6: a collection of denial-of-service tests againsts a target.

- firewall6: firewall tester, sends many different types of SYN packets.
- implementation6: performs various implementation checks on ipv6.
- parasite6: icmp neighbor solicitation/advertisement spoofer.
- redir6: redirect traffic with a clever icmp6 redirect spoofer.
- dos-new-ip6: detect new ip6 devices, cause denial-of-service.
- trace6: very fast traceroute6 which supports ICMP6 echo request and TCP-SYN.
- flood_router6: flood a target with random router advertisements.
- flood_advertise6: flood a target with random neighbor advertisements.
- fake_mipv6: steal a mobile IP to yours if IPSEC is not needed for authentication.
- smurf6: local smurfer, icmp flood attack.
- 6to4test - check an ipv4 address for dynamic 6to4 tunnel setup.
- etc. etc.

Penetration Testing
BackTrack is a Linux-based penetration testing arsenal intended for all audiences, from the most savvy security professionals to early newcomers to the information security field. BackTrack promotes a quick and easy way to find and update the largest database of security tools to-date. Our community of users range from skilled penetration testers in the information security field, government entities, information technology, security enthusiasts, and individuals new to the security community. Feedback from all industries and skill levels allows us to truly develop a solution that is tailored towards everyone and far exceeds anything ever developed both commercially and freely available.

Packet Scanning and Probing
Scapy is a powerful interactive packet manipulation program for scanning and probing. It is able to forge or decode packets of a wide number of protocols, send them on the wire, capture them, match requests and replies, and more. It can easily handle most classical tasks like scanning, tracerouting, probing, unit tests, attacks or network discovery. It also performs very well at a lot of other specific tasks that most other tools can't handle, like sending invalid frames, injecting your own 802.11 frames, combining techniques (VLAN hopping+ARP cache poisoning, VOIP decoding on WEP encrypted channel, etc).

2.6 Avoid Pay Per Click Problems

One of the most important Internet marketing tools is PPC advertising. With Yahoo and Google leading the pack, the industry as a whole has grown immensely in the past few years. PriceWaterHouseCoopers reports that in 2004 alone, Internet Advertising brought in an estimated $9 billion dollars. Anybody can use this marketing system that is really quite simple in theory.
With PPC advertising you choose "keywords/phrases," then bid how much you'd like to pay for each click. When a searcher goes to a search engine and types in one of your keyphrases, your short text ad appears, and if someone click on it your account is then charged. In a "perfect world" this is the way it would work, but thanks to unscrupulous people, there's a dirty little secret known as "click fraud."

Click fraud is simply the act of clicking on ads for the direct purpose of costing the advertiser money. It is recognized as the biggest problem today in PPC marketing. According to InternetWeek.com, 60% of people surveyed by the "Search Engine Professional Organization" have stated that fraud is a problem when it comes to PPC advertising.

PPC marketing can cost you a lot if you do not administer it right. Bad targeting plus fraud can be a costly problem. The main sources of click fraud are the following four:
1) AdSense Users:Google Adsense has a program called "Adsense" that pays website owners to run their Adwords ads and compensates them per click. Google does monitor this and it's against their terms of service to click on any of the ads on your own site. If they find a publishers doing this, they will lose their accounts, but some may still be clicking under the radar.

2) Your Own Competitors:our competitors could be clicking on your ads over a period of several days in order to deplete your ad budget. This way they neutrilize your advertizing campaigns.
3) Software:There are those who use automated clicking tools, such as robot programs, to click on PPC listings.
4) Paid Clickers:In some Asian countries, people are often paid to click on PPC ads for hours. Many don't know why they do it, and don't care. The only important issue is that they will be well rewarded for their efforts. If you do a search on any search engine you'll see plenty of sites offering to hire people for just this purpose. Type in 'earn rupees clicking ads' in Google and you get quite a few leads.

Most PPC networks have measures in place to protect you against click fraud. Yahoo's Overture tracks more than 50 data points, including IP addresses, browser info, users' session info and what they call "pattern recognition." They have a "proprietary system" in place for detecting fraud and a specialized team that monitors things and works with the advertisers to stop it.

Google offers suggestions to avoid click thru fraud, such as "using negative keywords" to keep your ads from showing up for products and services that are unrelated. They also suggest adding tracking url's to your links so you can track the traffic coming from Google. If you go through your log files, you'll be able to see your Google traffic at a glance.

If you suspect fraud, Google asks that you contact them right away, because they have a team of researchers that will investigate. They also take action to block future impressions from anyone they identify as committing click fraud. Like Overture, they also have "proprietary technology" that distinguishes between normal clicks and invalid ones. Google never bills you for any "bad clicks" that are caught by their system.

All honest website owners need to be alert to any "suspicious activity" by researching their server logs or stats. If you're experiencing a lot of clicks and no sales you'll also want to take a closer look. You need to watch for any spikes in traffic, usually on one keyword or phrase and coming from only one PPC source. You need to measure and track all of your PPC accounts closely.

A variety of new services have opened recently to help combat the click fraud problem. Some of them also offer web analytic tools that help improve your advertising productivity. You may want to look at these outside services to take care of problems for you. Here are some links:
1) Keyword Max: http://www.KeyWordMax.com
Offers up a service called "Click Auditor," which monitors the activity on your PPC accounts and alerts you to any suspicious activity. You can request a free demo at the site.
2) Click Detective: http://www.ClickDetective.com
A website monitoring service that uses sophisticated tracking mechanisms to determine whether "visitor behavior" is normal or not. Offering a 15 day free trial. Easy to use, you just copy and paste a snippet of code on your page and add a campaign ID by logging into your account.
3) Click Assurance: http://www.ClickAssurance.com
An Internet Security Firm that specializes in click fraud. They will audit your PPC accounts and go after any refunds you are due because of fraud.
4) Nami Media: http://www.NamiMedia.com
Specializes in post-click actions and landing page optimization technologies. Offers to increase sales and give marketers the ability to intelligently define landing pages to achieve business objectives. Works on ASP platform.
5) Who's Clicking Who: http://www.WhosClickingWho.comAn independent auditing service that tracks individual users for fraud. Can also detect abuse coming from proxy servers. A one month subscription is $79.00, which includes free installation and up to 50,000 transactions per month.
6) ClickLab: http://www.ClickLab.com/
his service isolates bad clicks with a scorecard based detection system. Pricing starts at $50.00 per month and is based on the number of sites you need to track and their page views.
ClickLab also offers white papers you should download while visiting, get them at the resources section.

7) Tracking ROI: http://www.TrackingROI.com/
TrackingROI’s Content Personalization System is truly a technological innovation that targets site visitors more closely. This system allows you to segment visitors to your site into groups and then provides personalized content through a Microsoft Word like tool! Your visitors can be segmented based upon each individual campaign. It offers ad tracking, site optimization, as well as click fraud control.
Click fraud isn't going away anytime soon. Most probably, it will get worse before it get's any better. It's up to you as a vigilant website owner to do what you can to keep your PPC advertising costs down. You can't stop it, but with the right tracking in place, it can be managed and controlled, and hopefully kept to a minimum.

2.7 Increasing website traffic

Google Webmaster tools is essential for any website. Google informs you of problems it identifies when indexing your site and it’s wise not to ignore this!

Here are ways to use Google Webmaster Tools to increase traffic to website.
1. Add a Sitemap to help Google index relevant pages
A sitemap is like an index in a book. It tells Google about the structure of your site. You don’t have to build this sitemap manually, there are tools available that can do it. For example, there arewordpress plugins that will build a sitemap.
If you have video on your site you also need to install a video sitemap. This allows Google to index video on your site instead of redirecting visitors off to Youtube or similar sites.
2. Optimise Existing Posts to Increase Traffic
In the traffic section under search queries Google shows you the keywords that are used to find your website and the position in the search results. You can review this, improve it and generate better results.

In the example above there are some search terms that are appearing in search results but not many people are clicking on these results. For example, Google Keyword tool appeared 2,500 times and got less than 10 clicks. I have a blog post directly relevant to this so I should be getting clicks.

Blog post – How to use Google Keyword Tool to Increase Traffic to Your Blog

Title – How to use Google Keyword Tool to Increase Traffic to your blog
So there are a few things I can do to increase traffic to this post:
a). Linking from another post to this blog post. In this post I have linked to the post using the keywords ‘Google Keyword Tool’, that will help Google understand what this post is about.
b). Change the title. The words at the start of a title are more important than the words at the end so I’ve updated the title to the following which will help:
Google Keyword Tool – How to use to increase traffic to your blog.
With the keywords at the start that will help.
c). Get links back to this post. The best way of increasing ranking will be to get links back from external (high profile) websites that have these keywords in the link!.
3. Tidy up your sitelinks
When Google starts liking your site it displays a few additional links within the search results as follows:
These site links are automatically created by Google but within webmaster tools you have some control over them. You can demote up to 100 different links so gives you a better chance of putting in the best links. By demoting links that are not useful this should help point out the better links and you’ll get more clicks.
4. Resolve any HTML Errors
Google reports on some HTML errors that will effect the indexing of your site so it’s important to resolve these issues. If you resolve these issues this will improve the chances of getting more traffic to your site.

In the above example the description tag (which is displayed when someone searches) is duplicated on some posts. This needs to be resolved. Quite often you will find some Titles that are duplicated which should be resolved also.

5. Resolve any server errors

It is not going to do you any benefit having server errors and it's likely that it will cause your harm. The following shows errors returned for content that should not be indexed. So to resolve these issues I had to block access to this directory from google.

One way to tell Google not to index content is to update the file robots.txt. This file is read by Google when it attempts to index that site. If I wanted to remove access to the plugins directory I'd enter the following: disallow: /wp-content/plugins

6. Implement Geographic Targeting

If a website has a generic top-level domain (e.g. .com, .org) then you can inform Google what is your main country that you want to target. This can improve the ranking based on this country.

You can target the main country you want to focus on in webmaster tools

Note: If you really want to target a specific country ideally you should have your website located on a server in that country also.

7. Remove Bad Links!

Google recently released an addition to webmaster which allows you to remove bad links to the page. If Google has notified you that you have some unnatural links you need to consider removing them using the disavow tool.

Generally you would only use this tool if Google reports that you have some unnatural linking on your website. If you login to Google and get a message similar to the following then using the disavow tool will help:

"We have detected that some of the links pointing to your site are using techniques outside of Google Webmaster guidelines"

So make sure to remove links to your site that Google doesn't like otherwise this will reduce the traffic you get. Deciding what links are damaging and worth deleting is another blog post!

Google provides lots of valuable information in Google webmaster tools so it's important to monitor it and make any relevant changes. Making sure all issues are resolved and it's configured correctly can help drive additional traffic to your website so it's well worth spending the time doing it.

2.8 IPv6 Risks

The IP version currently used in networks and the Internet is IP version 4 (IPv4). IPv4 was developed in the early '70s to facilitate communication and information sharing between government researchers and academics in the United States. At the time, the system was closed with a limited number of access points, and consequently the developers didn't envision requirements such as security or quality of service. To its credit, IPv4 has survived for over 30 years and has been an integral part of the Internet revolution. But even the most cleverly designed systems age and eventually become obsolete. This is certainly the case for IPv4. Today's networking requirements extend far beyond support for web pages and email. Explosive growth in network device diversity and mobile communications, along with global adoption of networking technologies, new services, and social networks, are overwhelming IPv4 and have driven the development of a next-generation Internet Protocol.

IPv6 has been developed based on the rich experience we have from developing and using IPv4. Proven and established mechanisms have been retained, known limitations have been discarded, and scalability and flexibility have been extended. IPv6 is a protocol designed to handle the growth rate of the Internet and to cope with the demanding requirements on services, mobility, and end-to-end security.

When the Internet was switched from using Network Control Protocol (NCP) to Internet Protocol (IP) in one day in 1983, IP was not the mature protocol that we know today. Many of the well-known and commonly used extensions were developed in subsequent years to meet the growing requirements of the Internet. In comparison, hardware vendors and operating system providers have been supporting IPv6 since 1995 when it became a Draft Standard. In the decade since then, those implementations have matured, and IPv6 support has spread beyond the basic network infrastructure and will continue to be extended.

It is very important for organizations to pay attention to the introduction of IPv6 as early as possible because its use is inevitable in the long term. If IPv6 is included in strategic planning; if organizations think about possible integration scenarios ahead of time; and if its introduction is considered when investing in IT capital expenditures, organizations can save considerable cost and can enable IPv6 more efficiently when it is needed.

An interesting and humorous overview of the history of the Internet can be found in RFC 2235, "Hobbes' Internet Timeline." The account starts in 1957 with the launch of Sputnik in Russia and the formation of the Advanced Research Projects Agency (ARPA) by the Department of Defense (DoD) in the United States. The RFC contains a list of yearly growth rate of hosts, networks, and domain registrations in the Internet.

Some excerpts from the RFC:
1969: Steve Crocker makes the first Request for Comment (RFC 1): "Host Software."
1970: ARPANET hosts start using Network Control Protocol (NCP).
1971: 23 hosts connect with ARPANET (UCLA, SRI, UCSB, University of Utah, BBN, MIT, RAND, SDC, Harvard, Lincoln Lab, Stanford, UIU©, CWRU, CMU, NASA/Ames).
1972: InterNetworking Working Group (INWG) is created with Vinton Cerf as Chairman to address the need for establishing agreed-upon protocols. Telnet specification (RFC 318) is published.
1973: First international connections to the ARPANET are made at the University College of London (England) and Royal Radar Establishment (Norway). Bob Metcalfe's Harvard PhD thesis outlines the idea for Ethernet. File transfer specification (RFC 454) is published.
1976: Queen Elizabeth II sends an email.
1981: Minitel (Teletel) is deployed across France by France Telecom.
1983: The cutover from NCP to TCP/IP happens on January 1.
1984: The number of hosts breaks 1,000.
1987: An email link is established between Germany and China using CSNET protocols, with the first message from China sent on September 20. The thousandth RFC is published. The number of hosts breaks 10,000.
1988: An Internet worm burrows through the Net, affecting 10 percent of the 60,000 hosts on the Internet.
1989: The number of hosts breaks 100,000. Clifford Stoll writes Cuckoo's Egg, which tells the real-life tale of a German cracker group that infiltrated numerous U.S. facilities.
1991: The World Wide Web (WWW) is developed by Tim Berners-Lee and released by CERN.
1992: The number of hosts breaks 1,000,000. The World Bank comes online.
1993: The White House comes online during President Bill Clinton's time in office. Worms of a new kind find their way around the Net—WWW Worms (W4) are joined by Spiders, Wanderers, Crawlers, and Snakes.
1994: Internet shopping is introduced; the first spam mail is sent; Pizza Hut comes online.
1995: The Vatican comes online. Registration of domain names is no longer free.
1996: 9,272 organizations find themselves unlisted after the InterNIC drops their name service as a result of their not having paid their domain name fees.
1997: The 2,000th RFC is published.

This is how far the RFC goes. But history goes on. According to http://www.internetworldstats.com/emarketing.htm, the worldwide online population reached 361 million users in 2000 (a penetration rate of 5.8%) and 587 million users in 2002. In 2003, the U.S. Department of Defense announced that they would be migrating the DoD network to IPv6 by 2008, and the Moonv6 project was started (now concluded). In 2005, Google registered a /32 IPv6 prefix, and Vint Cerf, known as "Father of the Internet," joined Google. By that time the number of Internet users had reached 1.08 billion. Today, at the time of writing in 2014, we are at approximately 2.4 billion Internet users, which corresponds to a penetration rate of 34%.

So while these numbers reflect all Internet users, independent of the IP protocol version, now we are starting to watch the growth of the IPv6 Internet. It is in its early days, but according to the growth numbers of the last two years, we expect growth to be exponential, and probably much faster than even the enthusiasts among us expect. The growth of the IPv6 Internet can be seen on the Google IPv6 Adoption statistics and the stats as of spring 2014 are shown in Figure 1-1. The stats show that in early 2011 (when the IANA IPv4 pool ran out), the percentage of native IPv6 Internet users was at approximately 0.2%. The stats also show that the percentage of users that were not native IPv6 (e.g., 6to4 or Teredo, red line) dropped to almost zero and are since then insignificant. Within one year the number of IPv6 Internet users doubled to 0.4%—a small number but still growth. In January 2013, the IPv6 Internet had crossed the 1% mark, and we entered 2014 with almost 3% IPv6 Internet users, which corresponds to approximately 72 million users. At the time of delivering this chapter, in April 2014, we were at 3.5%. The number of IPv6 Internet users currently doubles approximately every nine months.

These are just a few selected events and milestones of the Internet's history. Keep watching as more history unfolds. We are all creating it together.

Google's global IPv6 adoption statistics as of spring 2014
Figure 1-1. Google's global IPv6 adoption statistics as of spring 2014

The History of IPv6
The Internet Engineering Task Force (IETF) began the effort to develop a successor protocol to IPv4 in the early 1990s. Several parallel efforts to solve the foreseen address space limitation and to provide additional functionality began simultaneously. The IETF started the Internet Protocol Next Generation (or IPng) area in 1993 to investigate the different proposals and to make recommendations for further procedures.

The IPng area directors of the IETF recommended the creation of IPv6 at the Toronto IETF meeting in 1994. Their recommendation is specified in RFC 1752, "The Recommendation for the IP Next Generation Protocol." The Directors formed an Address Lifetime Expectation (ALE) working group to determine whether the expected lifetime for IPv4 would allow the development of a protocol with new functionality, or if the remaining time would allow only the development of an address space solution. In 1994, the ALE working group projected that the IPv4 address exhaustion would occur sometime between 2005 and 2011 based on the available statistics.

For those of you who are interested in the different proposals, here's some more information about the process (from RFC 1752). There were four main proposals: CNAT, IP Encaps, Nimrod, and Simple CLNP. Three more proposals followed: the P Internet Protocol (PIP), the Simple Internet Protocol (SIP), and TP/IX. After the March 1992 San Diego IETF meeting, Simple CLNP evolved into TCP and UDP with Bigger Addresses (TUBA), and IP Encaps became IP Address Encapsulation (IPAE). IPAE merged with PIP and SIP and called itself Simple Internet Protocol Plus (SIPP). The TP/IX working group changed its name to Common Architecture for the Internet (CATNIP). The main proposals were now CATNIP, TUBA, and SIPP. For a short discussion of the proposals, refer to RFC 1752.

The Internet Engineering Steering Group approved the IPv6 recommendation and drafted a Proposed Standard on November 17, 1994. RFC 1883, "Internet Protocol, Version 6 (IPv6) Specification," was published in 1995. The core set of IPv6 protocols became an IETF Draft Standard on August 10, 1998. This included RFC 2460, which obsoleted RFC 1883.

One of the big challenges but also one of the main opportunities of IPv6 is the fact that we can redesign our networks for the future. This is what enterprises should focus on most when planning their IPv6 integration in order to make sure they don't just copy old concepts onto a new protocol. We have to rethink our architectures. This once-in-a-lifetime opportunity can be used to get rid of a lot of legacy. An interesting RFC that helps in the process of seeing the big picture is RFC 6250, "Evolution of the IP Model." It shows how much this model has changed in the many years of operating our networks. So it helps to free our minds for hinking in new ways. One funny little quote that demonstrates what I am talking about is included below.

In this RFC there is mention of the first IP model and addressing architecture, and it quotes RFC 791, which defined IPv4 and the IPv4 address:

Addresses are fixed length of four octets (32 bits). An address begins with a one-octet network number, followed by a three-octet local address. This three-octet field is called the “rest” field.
This is how far we have come. Now project this into the future with the vast IPv6 address space in mind. Making meaningful use of the new address architecture and the enormous space will write the next chapter of the evolution of the IP model.

The vast IPv6 address space opens up serious opportunities for the re-examination of the notion of address. The IETF has only allocated 1/8th of the IPv6 address space for current use. The remaining 7/8ths of the address space is still to be allocated. In consequence we may be able to interpret new segments of the IP address space in ways that are different from topological end points. This is precisely the reason that a focus on the future of IPv6 is so important at this point in the evolution of the Internet.

What’s New in IPv6?
IPv6 is an evolution of IPv4. The protocol is installed as a software upgrade in most devices and operating systems. If you buy up-to-date hardware and operating systems, IPv6 is usually supported and needs only activation or configuration. In many cases it is activated by default. Currently available transition mechanisms allow the step-by-step introduction of IPv6 without putting the current IPv4 infrastructure at risk.

Here is an overview of the main changes:
Extended address space
The address format is extended from 32 bits to 128 bits. This provides multiple IP addresses for every grain of sand on the planet. In addition, it also allows for hierarchical structuring of the address space in favor of optimized global routing.

Autoconfiguration
Perhaps the most intriguing new feature of IPv6 is its Stateless Address Autoconfiguration (SLAAC) mechanism. When a booting device in the IPv6 world comes up and asks for its network prefix, it can get one or more network prefixes from an IPv6 router on its link. Using this prefix information, it can autoconfigure for one or more valid global IP addresses by using either its MAC identifier or a private random number to build a unique IP address. In the IPv4 world, we have to assign a unique IP address to every device, either by manual configuration or by using DHCP. SLAAC should make the lives of network managers easier and save substantial cost in maintaining IP networks. Furthermore, if we imagine the number of devices we may have in our homes that will need an IP address in the future, this feature becomes indispensable. Imagine reconfiguring your DHCP server at home when you buy a new television! Stateless Address Autoconfiguration also allows for easy connection of mobile devices, such as a smartphone, when moving to foreign networks.

Simplification of header format
The IPv6 header is much simpler than the IPv4 header and has a fixed length of 40 bytes. This allows for faster processing. It basically accommodates two times 16 bytes for the Source and Destination address and only 8 bytes for general header information.

Improved support for options and extensions
IPv4 integrates options in the base header, whereas IPv6 carries options in so-called Extension headers, which are inserted only if they are needed. Again, this allows for faster processing of packets. The base specification describes a set of six Extension headers, including headers for routing, Quality of Service, and security.

Why Do We Need IPv6?

For historic reasons, organizations and government agencies in the United States used the largest part of the allocatable IPv4 address space. The rest of the world had to share what was left over. Some organizations used to have more IPv4 address space than the whole of Asia (where more than 50% of the world's population live). This is one explanation of why the deployment of IPv6 in Asia is much more common than in Europe and the United States.

The IPv4 address space has a theoretical limit of 4.3 billion addresses. However, early distribution methods allocated addresses inefficiently. Consequently, some organizations obtained address blocks much larger than they needed, and addresses that could be used elsewhere are now unavailable. If it were possible to reallocate the IPv4 address space, it could be used much more effectively, but this process is not possible, and a global reallocation and renumbering is simply not practical. In addition to that it would not buy much, as even 4.3 billion addresses would not suffice for long at the current growth rate. We have to take into account that in the future we will need IP addresses for billions of devices. Vendors in all industries are developing monitoring, control, and management systems based on IP.

As the previous section shows, the IPv6 working group has done more than just extend the address space. For many complex networks of today and tomorrow, and for the number of IP devices of all types, the Autoconfiguration capability of IPv6 will be a necessity. The management of such services can't be accomplished with traditional addressing methods, and Stateless Address Autoconfiguration will also help to reduce administrative costs for organizations.

The extended address space and the restoration of the original end-to-end model of the Internet allows for the elimination of Network Address Translation (NAT), in which a single or a few public IPv4 address(es) are used to connect a high number of users with private addresses to the Internet by mapping the internal addresses to the public address(es). NATs were introduced as a short-term fix for solving the address space limitations with IPv4, since IPv6 was not ready yet (refer to RFC 1631; the original NAT specification was obsoleted by RFC 3022 in 2001). NATs have become pretty common in IPv4 networks, but they create serious disadvantages in management and operation: in order to do the address mapping, NATs modify end node addresses in the IP header. Very often, Application Level Gateways (ALG) are used in conjunction with NAT to provide application-level transparency. There is a long list of protocols and applications that create problems when used in a NAT environment. IPsec and peer-to-peer applications are two well-known examples. Another known issue with NAT is the overlapping of private address space when merging networks, which requires either the renumbering of one of the networks or the creation of a complex address-mapping scheme. The amplification of limited address space, the primary benefit of NAT, is not needed with IPv6 and therefore is not supported by design.

By introducing a more flexible header structure (Extension headers), the protocol has been designed to be open and extensible. In the future, new extensions can easily be defined and integrated in the protocol set. Based on the fact that IPv4 has been in use for almost 30 years, the development of IPv6 was based on the experience with IPv4 and focused on creating an extensible foundation; you can expect it to last a long time.

Broadband penetration rates in many countries continue to accelerate and, in some cases, have reached 65% or more. This level of always-on connectivity with substantial bandwidth capacity means that there is greater opportunity for devices to be connected. And many consumer electronic manufacturers have taken advantage of this. Online gaming is no longer the sole purview of games on PCs. Gaming stations, such as Sony's PlayStation 4, Xbox One, or Nintendo Wii U, have added capabilities to take them online. Many telecommunication carriers are providing television-type services (movies, audio content, etc.) over their IP networks. Even appliances, such as refrigerators, stoves, water heaters, and bathtubs, are getting connected. While it may seem rather silly to network-enable a bathtub, many of these devices are being connected to facilitate things such as power management, remote control, and troubleshooting, and for telemetry/monitoring purposes. We are entering the age of smart buildings and smart cities. The end result of this network-enablement process is a greater number of devices that need addressing, many of which will not have standard user interfaces. In these cases, the IPv6 address space, coupled with features such as Neighbor Discovery, Stateless Autoconfiguration, and Mobile IPv6, will help to usher in a new era of computerization in the home, but hopefully without the enormous deployment headache that it would cause if it were attempted with the current protocol.

The growth of the wireless industry (both cellular and wireless networks) has been nothing short of phenomenal. In more and more countries the number of cell phones actually exceeds the number of people. In this world of continuous reachability and reliance on the ability to access information at any time, the mobility requirements for end users have become exceptionally important. From the carriers' perspective, especially those supporting multiple media access types (e.g., 3G, WiMax, LTE), leveraging IP as the method of transporting and routing packets makes sense. Smartphones access the Internet, play games with other users, make phone calls, and even stream video content. Instead of supporting all of these functions using different transport protocols and creating intermediary applications to facilitate communications, it is far more efficient to leverage the existing network infrastructure of the Internet and a company's network. We will see later that from a technical perspective, Mobile IPv6 is very elegant in its design, supporting mobile users in a highly efficient manner and providing the overlay mechanisms for users to maintain their connections when moving between networks, even if those networks do not use the same type of media access.

There still remain some questions about the value of IPv6 to the enterprise, and it is worth conceding that each organization needs to evaluate the benefits and best timing of IPv6 for their own internal use. In many instances, organizations can find clever ways to use IPv6 to solve "pain" issues without migrating their entire network. Adoption can occur in an incremental fashion with a plan that minimizes integration pain but also ensures that everything is ready when the time comes to "flip the switch." As many case studies show, well-planned introduction costs substantially less than you would expect; the main cost-saving aspect is the fact that the advance planning lets you use all your refresh cycles, which minimizes cost. The step-by-step introduction allows you to learn as you go, thereby saving a lot of money and headaches, and you can do it without putting the current IPv4 infrastructure at risk.

But with all these thoughts and considerations, let's not forget the most essential advantage of IPv6. With its new structure and extensions, IPv6 provides the foundation for a new generation of services. There will be devices and services on the market in the near future that cannot be developed with IPv4. This opens up new markets and business opportunities for vendors and service providers alike. The first-mover opportunities are substantial, as are the opportunities to extend current product life cycles by refreshing their technology with IPv6. On the other hand, it means that organizations and users will require such services in the mid-term. It is therefore advisable to integrate the new protocol carefully and in a nondisruptive manner, by taking one step at a time to prepare the infrastructure for these new services. This protects you from having to introduce a business-critical application based on IPv6 at unreasonably high cost with no time for thorough planning.

Common Misconceptions

When considering all these advantages, maybe the question should be: "Why not IPv6?" When talking to customers, we often find that they share a similar set of misconceptions preventing them from considering IPv6. Here are the most common ones:

"The introduction of IPv6 puts our current IP infrastructure—our networks and services—at risk."
This concern is unsubstantiated. A major focus in IPv6's development was to create integration mechanisms that allow both protocols to coexist peacefully. You can use IPv6 both in tandem with and independently of IPv4. It is possible to introduce IPv6 and use it for access to new services while retaining IPv4 to access legacy services. This not only ensures undisrupted access

to IPv4 services, but it also allows a step-by-step introduction of IPv6. I discuss these mechanisms in Chapter 7. Your biggest risk is to not take advantage of all the opportunities IPv6 offers. You can only use these opportunities if you plan while there is time.

“The IPv6 protocol is immature and hasn’t proven that it stands the test of time or whether it is capable of handling the requirements.”
This was a concern of many people back in 2006 when we published the second edition of this book. Now in 2014 this is not true anymore. Many ISPs and organizations are deploying IPv6, vendors are getting up to speed, and the working groups have developed and optimized mechanisms that help with the integration. There is no technical reason not to do IPv6.

“The costs of introducing IPv6 are too high.”
There will certainly be costs associated with adopting IPv6. In many cases, newer networks will find that the level of IPv6 support in their current infrastructure is actually high. Regardless, the transition will necessitate some hardware and software costs. Organizations will need to create new designs, review current concepts, train their IT staff, and may need to seek outside expertise in order to take full advantage of IPv6.

However, the cost savings associated with IPv6 are becoming easier to define. Networks based on IPv4 are becoming increasingly more complex. New IT services such as VoIP, instant messaging, video teleconferencing, IPTV, and unified communications are adding layers of middleware and complexity. Merging organizations or those conducting B2B transactions are implementing NAT overlap solutions that have high management costs and are difficult to troubleshoot. And a growing market of mobile devices and network appliances requires robust access models that are expensive and difficult to implement in an IPv4 world. In all of these cases, IPv6 presents a cleaner and more cost-effective model in the long run than IPv4 can provide. And the fact is that an investment in IPv4 is an investment in an end-of-life technology, while an investment in IPv6 is an investment in the future technology.

“With Stateless Address Autoconfiguration, we will not be able to control or monitor network access.”
While this statement may generally be true for networks that widely utilize Stateless Address Autoconfiguration, administrators will have a choice about their level of control. DHCPv6 as defined in RFC 3315 has been extended to support two general modes of operation, Stateful and Stateless. Stateful mode is what those who currently utilize DHCP (for IPv4) are familiar with, in which a node (DHCP client) requests an IP address and configuration options dynamically from a DHCP server. DHCPv6 also offers a Stateless mode in which DHCPv6 clients simply request configuration options from a DHCPv6 server and use other means, such as Stateless Address Autoconfiguration, to obtain an IPv6 address.

“Our Internet Service Provider (ISP) does not offer IPv6 services, so we can’t use it.”
You do not have to wait for your ISP to use IPv6 in your corporate or private network. If you want to connect to the global IPv6 Internet, you can use one of the transition mechanisms and tunnel your IPv6 packets over the IPv4 infrastructure of your ISP. This may be doable for smaller organizations. On the other hand, at the time of writing in 2014, you could expect a large ISP targeting enterprise customers to support IPv6. And this should be your standard requirement in any renewal of contract and SLAs (Service Level Agreements). If your ISP does not provide IPv6 services, consider finding a new provider.

“It would be too expensive and complex to upgrade our backbone.”
The transition mechanisms make it possible to use IPv6 where appropriate without dictating an order of upgrade. Usually for the backbone it is advisable to wait for the regular life cycle, when hardware needs to be exchanged anyway. Make sure to choose hardware that supports performance IPv6 routing. In the meantime, you can tunnel your IPv6 packets over the IPv4 backbone. Networks that use MPLS have an easy way to tunnel IPv6 packets over their IPv4 MPLS backbone. Read more about it in Chapter 7. More and more organizations are considering

migrating their backbone and data centers to IPv6 only with the next refresh or redesign cycle, because it substantially reduces operational cost. In this scenario we will start to tunnel IPv4 packets over IPv6 backbones. IPv4 as a service is the new keyword.

“It would be too complex and expensive to port all of our applications to IPv6.”
The effort necessary to port applications to run over IPv6 is often much lower than expected. If an application is well written, it may simply run over IPv6 without modification. Instead of assuming that it won’t work, test it to find out. For applications that need modifications that are not yet available, or for applications in which porting does not make sense, there are mechanisms available that support IPv4 applications in IPv6 networks and IPv6 applications in IPv4 networks. Alternatively, you can run a dual-stack network, in which you use IPv4 to access IPv4 applications and IPv6 to access IPv6 applications. In any case it is recommendable for enterprise customers to start the planning process early and provide good labs for the application teams to test their applications before there is time pressure.

“We have enough IPv4 addresses; we don’t need IPv6.”
True—if you have enough IPv4 addresses, there may be no immediate need to integrate IPv6 today. But ignoring IPv6 for this reason is a perspective that assumes that your network stands completely isolated from the rest of the world, including your vendors, partners, and customers. IPv6 adoption is further along in Asia and Europe than in the United States, so even though you may have adequate address space for your operations in Denver, interconnecting with a partner organization in Tokyo may eventually become complicated if you do not support IPv6. Plus, the assumption that IPv6 is about address space only doesn’t account for the advanced features that IPv6 brings to the table.

When Is It Time for IPv6?
The answer in 2014 is now! If the rest of the world moves to IPv6 while you insist on continuing to use IPv4, you will exclude yourself from global communication and reachability. The risks if you wait too long include losing potential customers and access to new markets and the inability to use new IPv6-based business applications.

There is a golden rule in IT: “Never touch a running system.” As long as your IPv4 infrastructure runs well and fulfills your needs, there is no reason to change anything. But from now on, whenever you invest in your infrastructure, you should consider IPv6. An investment in the new technology gives it a much longer lifetime and keeps your network state-of-the-art.

These are the main indicators that it may be time for you to consider switching to or integrating IPv6:
You need to extend or fix your IPv4 network or NAT implementation.
You are running out of address space.
You want to prepare your network for applications that are based on advanced features of IPv6.
You need end-to-end security for a large number of users and you do not have the address space, or you struggle with a NAT implementation.
You need to replace your hardware or applications that are at the end of their life cycles. Make sure you buy products that support IPv6 adequately, even if you don’t enable it right away.
You want to introduce IPv6 while there is no time pressure.
The following provisions can be taken in order to prepare for IPv6 adequately:
Build internal knowledge, educate IT staff, and create a test network.
Include IPv6 in your IT strategy.
Design future-proof network, security, and service concepts while you have time.
Create integration scenarios based on your network and requirements.
Put IPv6 support on all of your hardware and software purchasing guidelines. Be specific about which features (RFCs) must be supported. Don’t forget to add IPv6 requirements to outsourcing and service contracts, as well as SLAs.
Compel your vendors to add IPv6 support to their products.

If you do this, you can determine the right moment for the introduction of IPv6 in your network. You can also assess whether a further investment in your IPv4 infrastructure makes sense or whether introducing IPv6 would be a better way to go.

There will be no “flag day” for IPv6 like there was for the 1983 move from NCP to IPv4. Probably there will be no killer application either, so don’t wait for one. Or as some people like to say, the killer application for IPv6 is the Internet. IPv6 will slowly and gradually grow into our networks and the Internet. Taking a step-by-step approach to IPv6 may be the most cost-efficient way to integrate it, depending on your requirements. This method does not put your current infrastructure at risk or force you to exchange hardware or software before you are ready, and it allows you to become familiar with the protocol, to experiment, to learn, and to integrate what you’ve learned into your strategy.

You may want to enable IPv6 in your public services first. Due to the lack of IPv4 addresses, ISPs that want to grow their customer base (and who does not want to do that?) make use of NAT-type mechanisms to extend their IPv4 address space. This includes CGN (Carrier Grade NAT), which means multiple customers share one single public IPv4 address and sit behind multiple layers of NAT.

These users may have a bad user experience accessing your IPv4 website, and for e-commerce or other more complex services it may even fail. The users will not know that it is the provider’s CGN causing the issue and will blame your website for their problems. If you provide your website dual-stack, these users can access it over IPv6 and bypass the IPv4 NATs.

IPv6 Status and Vendor Support

As previously mentioned, IPv6 is implemented in most up-to-date versions of routing and operating systems. For standard applications, assume that IPv6 support has already been added or will be added with their next major release at the latest. For creating an IPv6 integration plan for your corporate network, you will need to assess the status and degree of IPv6 support with each vendor individually. Many vendors have an information site that can often be found at http://www.<vendor>.com/ipv6.

It can be said that IPv6 support up to the network layer is mature, tested, and optimized. This includes routing, transition mechanisms, DNS, and DHCPv6.

Development is most active in the security, transition mechanism, IPv4/IPv6 MIB integration, and Mobile IPv6 areas. More work needs to be done in the areas of network management and firewalls. Vendors such as Cisco, Checkpoint, Juniper, and many others are working on these areas. The application area is continuously developing, and new applications will appear on the market that will make use of the advanced features of IPv6. Thanks to the transition mechanisms, you can still use IPv4 applications in IPv6 networks.

Predictions about when the world will end are about as consistent as the predictions about when IPv4 internet addresses will finally run out, but some IT security professionals say that is really the least of our worries.

A much bigger concern, they say, should be the security holes that will open up in many business organisations as the world moves over to internet protocol version six (IPv6).

This is an important aspect of the changeover that has been lost in all the hype around how IPv4 is about to run out of IP addresses assigned to each internet-connected device because of the explosion of internet users, devices and web services.

IPv6 will solve this problem because it provides over four billion times more addresses than IPv4, but in solving that problem, it could expose businesses to cyber attacks as hackers use IPv6 to bypass security controls and filters designed and configured for IPv4 traffic.

Although the move to IPv6 could be completed as soon as 2011 in China, this will take at least two more years in the US and elsewhere, so the security threat is a much more immediate and pressing problem than ensuring networks are ready for IPv6 traffic.

IPv6 attacks likely to increase with adoption

The number of IPv6 attacks is relatively small, but as we see a wider adoption to IPv6, we are much more likely to see an increase in attacks as well as a greater focus from attackers, says Raj Samani, chief technology officer, EMEA, McAfee.

Danger lurks where companies are adopting IPv6 because of its greater speed and efficiency without ensuring that their network defences are updated accordingly, says James Lyne, director of technology strategy at security firm Sophos.

But perhaps an even bigger danger is where companies are using IPv6 without being aware of it, because the latest versions of most network hardware devices and operating systems are IPv6-enabled by default.

"Any business that is using Windows Server 2008, Windows 7 or even Mac OS X and a growing number of applications, including Skype, could be using IPv6 without even knowing it," says James Lyne.

Security researchers have already seen widespread malware with IPv6-based command-and-control capabilities. Given the relative lack of attention paid to IPv6, this technique can bypass existing protection such as non-IPv6 enabled firewalls completely.

IPv6 uses a completely different scheme of IP addresses, which effectively means that the concept of a network border no longer exists as it is possible to have a single IP address that will work anywhere in the world.

The hierarchy has been redesigned, says Lyne, so the danger is that businesses will implement IPv6 in much the same way they did IPv4.

"All they will succeed in doing is solving the problem of too few IP addresses, while opening up a host of security vulnerabilities and without getting any of the benefits such a massive performance gains from IPv6's ability to handle much bigger data packets," he says.

Without careful planning, Lyne warns that businesses could end up accidentally running IPv4 and IPv6 in parallel, effectively nullifying security measures they have put around either protocol.

Security advantages of IPv6

Lyne is critical of supporters of IPv6 for selling IPv6 only in terms of additional IP addresses and performance gains, instead of the inherent security benefits, such as internet protocol security (IPsec) which was originally developed for IPv6 and back-engineered for use with IPv4.

"IPsec, which is optional in IPv4, is an integral and mandatory part of IPv6, making man-in-the-middle attacks much more difficult for hackers," he says.

Encryption is also mandatory, which automatically ensures a higher level of data protection than IPv4. Unlike its predecessor, IPv6 was built from the ground up to be capable of end-to-end encryption.

The encryption and integrity checking used in current VPNs is a standard component in IPv6, available for all connections and supported by all compatible devices and systems.

IPv6 is also much stronger from a security point of view for mobile devices, says Lyne, because each device gets a consistent IP address which enables businesses to define a security policy for each device that will apply wherever that device is used.

The abundance of IP addresses makes it possible to allocate businesses their own blocks of IP addresses, which in turn delivers another security benefit. With such blocks of IP under their control, says Lyne, businesses can apply security policies to all corporate IP addresses, making the process much more manageable.

The availability and abundance of global IPv6 addresses enable a business to create specific services for targeted users, ranging from customers, partners or employees from remote sites.

"Each service can be guarded by fine-grained security and access policy containers, thus simplifying the implementation and maintenance of external facing services," says Qing Li, chief scientist and senior technologist at Blue Coat Systems.

Security challenges of IPv6
While having a large number of IP addresses will benefit companies from a management point of view, it will also benefit cyber criminals. Not only will criminals be able to switch IP addresses frequently - making it difficult to track and trace them - but many existing security controls that rely on blacklisting malicious IP addresses will cease to be effective.

This is a problem, says Lyne, as he estimates around 90% of web filtering tools used by business today rely on blacklists. Once the world has moved to IPv6, criminals will be able to rotate IP addresses very quickly, which will severely challenge the effectiveness of blacklisting, and even grey- and whitelisting, he says.

Not only is older technology a potential security threat, so too is an older skill set.

"It is important to remember however, that the majority of security professionals and networking engineers are most familiar with protecting IPv4 networks and aware of the signs so as we move across to IPv6 a real risk is the relative skills shortage," says McAfee's Raj Samani.

Blue Coat's Qing Li points out that many IT managers have not had the opportunity to develop working knowledge of the technology nor have they gone through a transition like this in the past.
"As a result, there is a potential to create security holes during the transition process. The most likely place for this to occur is in the creation of usage and security policies for IPv6. Not all of the existing corporate policies and rules that are implemented in IPv4 environments can simply be translated syntactically for IPv6 environments. Instead, they need to be rewritten. The lack of operational expertise makes it more likely that an IT manager will inadvertently create a security hole while writing those new policies," he says.

Also, in the traditional IPv4 infrastructure it is common to find network address translation (NAT) devices, which obscure an internal network's structure, but a NAT that performs the same type of duty is rarely found in IPv6 networks.

"Consequently, IT managers have been mostly managing private addresses that will eventually be translated into a single public address. Now the IT staff is faced with public address management at a grand scale and must figure out how to prevent internal users from creating secure tunnels to the outside, which may create corporate liability," says Li.

Avoid the security pitfalls of IPv6
Lyne says the switch over to IPv6 is an important opportunity to avoid the mistakes that were made with the implementation of IPv5 and SSL. Stricter IP address allocation processes that require proof that applicants represent a legitimate business, for example, could help address the problem of rapid IP address switching. But, he says, in the absence of any single recognised

internet authority, there is the risk that IPv6 implementation will lack co-ordination and, like IPv4 and SSL, will be determined organically and therefore lack the joined-up thinking required to ensure it is done in a way that makes the protocol as secure as it can be, with as few vulnerabilities as possible that can be exploited by criminals.

One of the challenges IPv6 poses to security suppliers is that they will have to re-write firewalls, but again, without any single organisation setting the agenda for how IPv6 will be deployed, says Lyne, the exact approach and requirements of doing this will be constantly changing as the situation evolves. This inability of security suppliers to anticipate how IPv6 will work in practice, is likely to create further opportunities for cyber criminals.

The lack of ownership by any single organisation is also one of the biggest reasons, says Lyne, that IPv6 adoption has been relatively slow, despite its speed, efficiency and security advantages over IPv4. But while the business world has been standing still, the cyber criminal world has been moving forward to apply the speed and efficiency benefits to their botnets or networks of hijacked computers. "Cyber criminals have long being capitalising on the fact that few people are filtering IPv6 traffic or even know how to," says Lyne.

While the mandatory encryption of IPv6 traffic is a good thing that will reduce the seriousness of data breaches that occur, it is a double-edged sword, as it also presents a challenge to government organisations who, once the transition to IPv6 is complete, will find their network traffic monitoring capabilities severely diminished.

The way ahead for IPv6 users

In the transition period, Lyne advises businesses turn off IPv6 until they are thoroughly prepared for the security implications of the new protocol and have updated all security filters and controls in their networks. Only switch IPv6 on, he says, once the controls are in place.

In terms of the technical concerns linked to IPv6 attacks facing companies, CIO'S should look out for rogue IPv6 devices, built-in ICMP and multicast, rogue IPv6 traffic and tunnels, says McAfee's Raj Samani.

There is no instant switch to the new protocol, says Lyne, so partial adoption means using tunnelling technologies to transport IPv6 over IPv4, and this kind of workaround is another potential source of confusion, misconfiguration and security gaps.

It is important businesses understand if their web security solution can rate and analyse IPv6 content because, without that ability, users will be vulnerable to attacks.

"The larger malware attack surface created by IPv6 also demands a real-time defence. In IPv4, we are already seeing very dynamic malware attack with the malware deliverable changing URLs more than 1,500 times in a single day, and we expect this trend to accelerate with the adoption of IPv6 and increase in number of available addresses," says Blue Coat's Qing Li.

He believes waiting a week or even 24 hours to analyse requests and update databases will leave users exposed to malware.

To truly protect their users, businesses need a web security defence that can analyse requests as they are made and deliver immediate protection when a new threat is discovered. Since individual users in a business may now be assigned a global IPv6 address and can create encrypted tunnels, it is important for the IT manager to have visibility into this encrypted traffic to eliminate security threats.

Business has been largely ignoring the inevitable transition to IPv6 since the early 1980s, and although IPv6 will not be the dominant standard tomorrow, businesses need to start planning today for the skills and hardware they will need to make the transition securely, says Lyne.

Businesses need to be more proactive on this issue, he says and challenge network hardware suppliers now about their strategies for IPv6. "Businesses should also ensure that any hardware they buy from now is IPv6 compatible, so when the time comes, they are ready.

"We are going to have to do this, so we may as well make use of the opportunity to apply the lessons learned from IPv4 and make serious advances in terms of security," says Lyne.

The transition to IPv6 is not something to be taken lightly and will require considerable effort, preparation and consideration because if done incorrectly or incompletely, the transition to IPv6 could leave gaping security holes in corporate networks.

The corporate perspective on IPv6

The Corporate IT Forum says members are taking a "wait and see" approach to IPv6 to ensure that the software and hardware are mature before implementation.

An IT head in the food distribution sector commented that he expects issues to be resolved by the time the company considers IPv6 in detail. A security and architecture manager in the hospitality industry says he does not percieve IPv6 to pose a significant threat due to any inherent risks in the standard.

One area starting to raise concerns, however, is risk due to a lack of comprehensive support for IPv6 security. Members are questioning whether the support in firewalls, ISPs, application proxies and the like is mature enough to be able to trust their own countermeasures to be doing a good enough role protecting the enterprise.

Fundamentally all bolt-ons are vulnerable and IPv6 is just that. One IT manager from the professional services sector is more concerned that organisations will become lazy due to IPv6 being perceived as more secure than IPv4.

With this level of complacency, it is unsurprising IPv6 has seen little press or uptake to date, says the Corporate IT Forum.

When quizzed further about the issues, members were not forthcoming leading the Forum to conclude that this has yet to reach the forefront of a team still challenged with user awareness, malware and phishing scams.

When people hear about IPv6-specific security issues, they frequently tend to rate this as an argument in favour of delaying or avoiding IPv6 deployment on their enterprise or campus network. Even without IPv6 being consciously deployed, however, some of the IPv6-related security issues were already introduced to most networks many years ago. The reason for this is simple: IPv6 is implemented in all common operating systems and enabled by default. We introduced hosts with these operating systems on our networks several years ago – be they clients on the office network or servers in a data centre or DMZ.

Since most, if not all, of today’s company networks are IPv6-enabled to a certain degree, they are attackable over IPv6. To make things worse, in contrast to IPv4, IPv6 brings along different kinds of autoconfiguration functionality, which can be misused. Network operators and security people who have neither basic IPv6 experience nor measures in place to detect IPv6-related attacks run a risk, and this risk is permanently increasing as the bad guys have already started to use IPv6. Bad guys are usually early adopters.

Let us look at some possible scenarios:

Rogue IPv6 router attracts traffic

In this scenario, the attacker has access to a local network segment. He is either an insider or has otherwise gained physical access to the network. Maybe he managed to compromise and take over one system on the network segment. He can now send periodic IPv6 Router Advertisements (ICMPv6 Type 134 messages) to the local network.

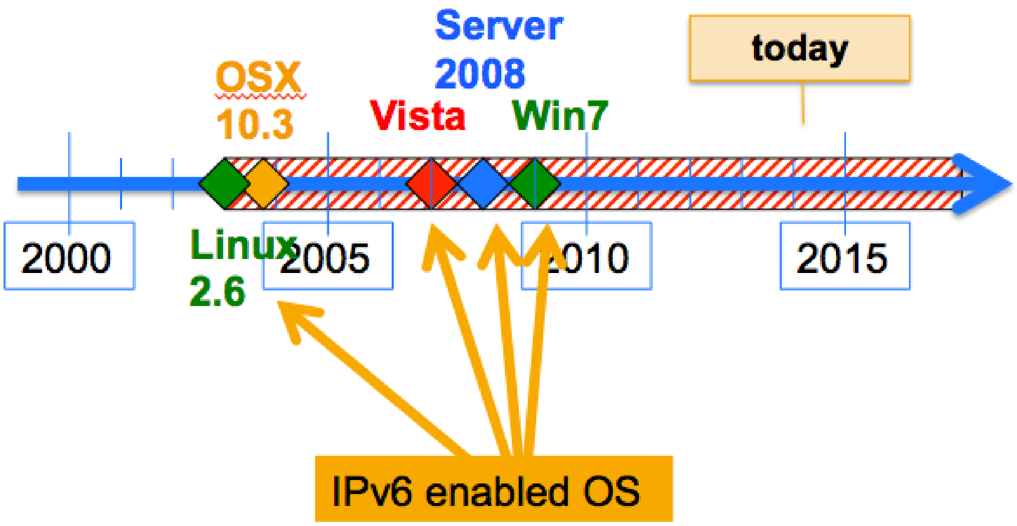

FIGURE 3: IPv6 in the OS history.

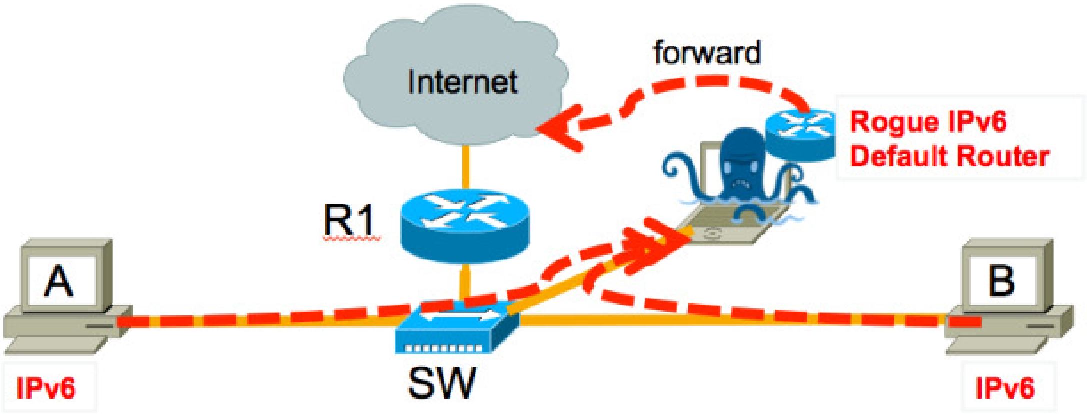

FIGURE 4: Scenario 1.

Attacker bypasses IP based access control

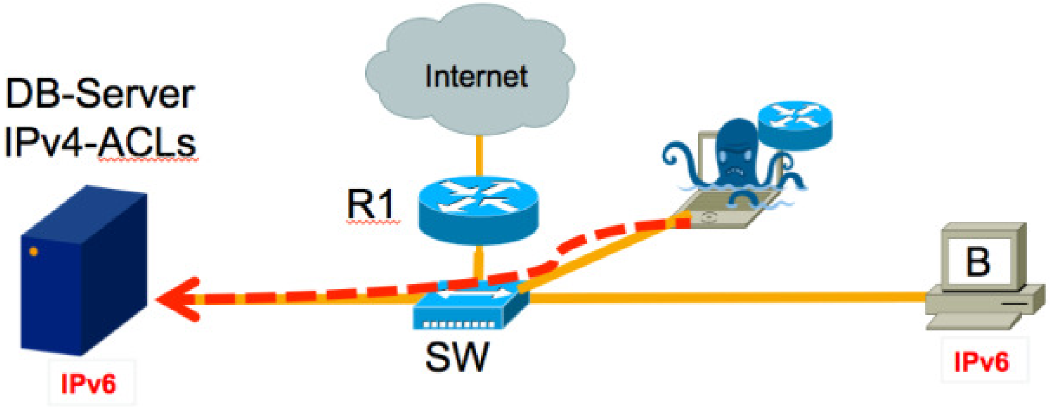

FIGURE 5: Scenario 2.

All nodes on the network that have IPv6 Stateless Address Autoconfiguration (SLAAC) enabled – which is the default – will consequently configure global routable IPv6 addresses and a default IPv6 route to the attacker's node.

As a result, all former IPv4-only nodes on the network are now dual-stack nodes and have native IPv6 connectivity, with the attacker serving as the next hop.

Here the attacker wants to access a resource (in this case a database server) protected by IP-based access control (IPv4 ACLs). He again sends Router Advertisements in order to configure IPv6 addresses on the desired machine. He might then be able to access the system via IPv6 because this "second door" is not protected by ACLs. This example applies to all devices protected by IP-based ACLs with only weak alternative access protection or none at all "because nobody can access it anyway".

Client bypasses firewall with IPv6 tunnel

In the third scenario, we have a user on the local network whose network knowledge is up to date. He wants to access a resource that is blocked according to the local firewall policy. Since this resource is available over IPv6 as well, he tells his web browser to use IPv6. The client node does not have native IPv6 connectivity, so it tries one of its IPv6 tunnel mechanisms that are enabled and can configure automatically. Given that the local firewall does not filter IPv6 traffic that has been tunnelled over IPv4 to an exterior tunnel gateway, this user can access the otherwise blocked resource.

Food for thought

These are just three examples that show how IPv6 can affect your network security, even though you have never consciously deployed IPv6. There are a number of others. The following questions may help you to assess whether your network security is at risk:

- Do you see IPv6 traffic on your network? (Monitoring)
- Are you sure your firewalls filter (tunnelled) IPv6 traffic?
- Do you have enough knowledge about IPv6 and its specific attacks to detect them?
- Do you rely on IP-based ACLs – which are ineffective for IPv6?

If you did not answer "yes" to the first three questions and "no" to the last one, then it seems you have a new item to add to your "to do" list.

2.9 NTP Attacks

In this article I am going to illustrate how NTP is vulnerable to attacks like replay-delay attacks, MITM, and a very recent attack termed as NTP DdoS (which is a kind of amplification attack used to flood the intended target with a response from the NTP server that can be 350 times bigger than the original request), and how the NTP security model addresses some of these concerns and future design considerations.

Network Time Protocol (NTP): NTP is used to synchronize the time of the computer within a few milliseconds of Coordinated Universal Time (UTC). It can be implemented in various models like client-server and peer-to-peer. The current version of NTP is ntpv4 and uses the User Datagram Protocol (UDP) and port number 123. In a client-server model, the client sends a packet mode 4 to the server and the server responds back with a packet mode 3 and provides synchronization to them. NTP uses a hierarchical model of time sources. Each layer is termed as stratum, with stratum 0 being the parent of whole layer. This layer is comprised of timekeeping devices such as atomic clocks, GPS clocks and radio clocks. Further down, the next layer is named as stratum 1, which is comprised of computer systems whose devices are synchronized to a few microseconds with their attacked stratum 0 devices. Then a stratum 2 layer, which is connected over a network to stratum 1 devices. Only stratum 0 to 15 are valid.

NTP Security Model

Intruders can play with this protocol to clog the network with big response packets (recent DDoS amplification attack), disrupt some time-dependent critical service, etc. There are various types of attacks that are possible on NTP. Some of them are discussed below:

- A replay attack in which an intruder replays one or more packets.
- Man in the middle attack (MITM) in which an intruder can intercept the packets between authentic client and server.
- A delay attack in which packets between client and server are delayed for a constant or variable time but left unchanged.
- A DDoS attack in which an attacker finds a vulnerable machine, makes it a bot master, and infects other vulnerable systems with malware. NTP DDoS is a type of reflective DDoS attack in which an attacker sends spoofed SYN packets so that when the server replies to the spoofed packet, replies goes to the spoofed IP in the SYN packet. In DDoS, amplification factor is used by attackers to increase the traffic volume in an attack. Results have shown that in an NTP DDoS attack, an attacker who has 1 GB of bandwidth can generate the attack with amplification factor close to 250 GB. More on this will be covered later in this article.

NTP Hierarchical Security Model

So let's first understand the security model of NTP. The below section covers the security layers in NTP, the various attacks possible, and how each layer protects against the respective attack projected at.

Let's discuss these layers further.

On-Wire Protocol Layer: The underlying protocol which is used to transfer packets between client and server. To detect duplicate packets and bogus packets, the wire protocol uses a 64 bit timestamp in the NTP packet, which is very unlikely for an intruder to guess. Detection of a duplicate packet is called a loopback test, and whenever a duplicate packet is found, it is discarded. Thus this layer protects attacks such as replay attacks.

Message Digest Layer: Packets between client and server can be intercepted and changed, except the transmit timestamp. To protect against these attacks, NTP has a message digest layer that uses symmetric key cryptography to compute a message digest. The message digest is computed using algorithms such as MD5 using a secret key and appending a Message Authentication Code (MAC) along with NTP and extension protocol. The MAC consists of a 32 bit key ID followed by a message digest. An algorithm such as MD5 then computes the message hash of the message digest and concatenates that with the NTP packet header. When packets are transmitted, the digest is computed and inserted in the MAC, and when packets are received, the digest is compared with the digest in the MAC, and the packet is only accepted when two digests are equal. However when a large number of clients are required, this is not suitable.

Autokey Sequence Key: In order to provide authenticity of NTP packets, an auto sequence layer is used to provide authenticity using public key cryptography and also digital signatures which are used only in responses from server to client. In client-server mode, the server distributes a unique cookie per client. The server computes the client cookie as the MD5 hash of the autokey with client and IP addresses of server, a key ID of zero and the server cookie. On receiving a request, the server returns the encrypted client cookie and the responses are signed using the server private sign key. The client on the other side decrypts the client cookie and verifies it using the server public key, which is contained in the certificate. For subsequent requests, both client and server calculate and verify the message digest. So an intruder cannot forge a packet until the possession of client cookie, which is further bound to a server signature.

However, a vulnerability known as cookie snatcher exists, where an intruder is able to wiretap the client cookie request to learn the client and server IP addresses. The intruder then launches a client cookie request using its own public encryption key. Using the client cookie, the middleman can masquerade as the legitimate server and inject bogus packets acceptable to the client.

The Autokey Protocol Layer: This protocol layer is used to retrieve certificates and identity keys. There are different key pairs, such as:

- Host Keys: These keys are used to encrypt client cookie.
- Sign Keys: These keys are used on certificates to verify signatures on extension fields. Changing the sign keys will set up the need to regenerate all the certificates.
- Identity keys: these keys are used to authenticate sign keys to avoid masquerade attacks.

Below is the process defined for client-server mode.

Client-Server Mode

Client-server mode is the most deployed for a time server and thousands of clients. All the exchanges must be in proper order. The exchange begins when the client sends an association request including its X.509 distinguished name along with its available cryptographic options. The response from the stateless server sends an association response including its X.509 distinguished name and available cryptographic options. The client, after selecting from the available options, sends a certificate request specifying the server distinguished name. The server sends the certificate signed by the server private sign key for positive identification. This design is vulnerable to the cookie snatcher attack.

So, we have seen that signing and message digest are even susceptible to an attack known as cookie snatcher attack. Let’s understand what exactly a cookie snatcher attack is and a design consideration to stop it.

The cookie snatcher attack happens at the on-wire protocol layer, thus exposing client cookie and message digest. This attack can be defended by adopting an agreement scheme the same way as the TLS handshake protocol, in which client sets message digest key as a nonce and encrypts the key using server public sign key on its certificate and sends that encrypted key during the exchange. Server in turn then decrypts it with its private key and saves the digest key for future use.

NTP DDoS Amplification Attack

Recently, an attack termed as NTP DDoS amplification attack has been realized worldwide. So what is an amplification factor?

Amplification Factor: Amplification is the factor by which the volume of an attack increases. An attacker can trigger an NTP attack by a initiating a request message which is of few bytes in size and because a command like monlist (defined later) can make the server to respond with a packet of very very big size. Recent attacks have been seen to have an amplificatory factor of 350+.

How NTP DDoS Attack Works

NTP protocol by design uses UDP to operate, which does not require any handshake like TCP, thus no record of the request. So, NTP DDoS amplification attack begins when an attacker crafts packets with a spoofed source IP to make the packets appear to be coming from the intended target and sends them to NTP server. Attacker initially crafts the packet of few bytes, but NTP responds with a large amount of data thus adding to amplification of this attack.

MONLIST command: It is a NTP protocol command which has very little use, but it is this command which is the main culprit for this attack. However, the use of MONLIST command is to give details of the last 600 clients that have connected to the NTP time service. Below is the command syntax:

Resolution or Mitigation Steps for DDoS Amplification Attack

Below are some of the ways in which NTP DDoS attack can be mitigated:

Update the NTP server: All versions of ntpd server prior to 4.2.7 are vulnerable to this attack by default. So upgrade the ntpd servers to the latest version.

For ntpd servers which cannot be upgraded to the latest version, disable the monlist functionality on NTP servers by adding the noquery directive to the restrict default line in ntp.conf like below:

- restrict default kod nomodify notrap nopeer noquery
- restrict -6 default kod nomodify notrap nopeer noquery

Like other DDoS attacks, attacker usually spoofs the source IP address in this attack as well. So organizations should restrict forged traffic. This mitigation technique is listed in BCP 38. Analysis has shown that attackers place their system inside network range 204.69.207.0/24. So by restricting transit traffic which originates from downstream network to known, this problem of address spoofing can be virtually eliminated. Take a look at the picture below:

For example, in the above figure, the ISP provides connectivity to Attacker (204.69.207.0/24) and all traffic originating from Attacker's machine is filtered on the ingress link to router 2, thus providing connectivity to Attacker's machine and filtering traffic that have originated with source address other than mentioned the range 204.69.207.0/24 thus prohibiting the attack from using an invalid IP address.

ISPs should regularly check their exposed services from a security perspective to close any vulnerable services before they are abused by attackers.

Organizations should implement tighter Access Control Lists (ACLs) in their public facing transit edges across layer 3 devices like switches and hardware based routers. This will help to mitigate the malicious traffic (NTP DDoS) from reaching targeted resources; however since the attack amplification factor is so large, mostly transit edge devices get clogged.

Disable the remote execution of monlist command. For example, most Unix /Linux distributions support monlist from a localhost and not from a remote host. Check this monlist execution with following command:
As part of Rapid7 Labs' Project Sonar, among other things, we scan the entire public IPv4 space (minus those who have opted out) looking for listening NTP servers. During this research we discovered some unknown NTP servers responding to our probes with messages that were entirely unexpected. This lead to the writing of an NTP fuzzer in Metasploit in the hopes of understanding what NTP implementations would respond in this or other anomalous manner in various configurations. This, in turn, resulted in finding six previously unpublished vulnerabilities in NTP Project's NTP implementation. One of these is similar in terms of severity to the NTP MON_GETLIST amplification vulnerability described in CVE-2013-5211 that was the source of record-sized DRDoS attacks in late 2013 and early 2014. All NTP instances vulnerable to CVE-2013-5211 are likely also vulnerable to these six new vulnerabilities, putting the number of public, vulnerable systems at approximately 65,000 based on a recent analysis.

Background
To fully grasp these vulnerabilities it is important to have a brief understanding of the technology in question (NTP), the vulnerability type (traffic amplification) and the attacks that frequently result from the abuse of these vulnerabilities (DRDoS).

NTP is the Network Time Protocol and serves to keep the clock of a computer system in sync. Properly synchronized clocks play a critical role in logging, authentication, cryptography and general system sanity, and as such NTP can be found in some manner in nearly all environments. NTP has been evolving for over 30 years and has seen four revisions its protocol. While there are numerous NTP implementations for both clients and servers, the NTP software provided by the Network Time Foundation's Network Time Protocol project powers the vast majority.

A traffic amplification vulnerability occurs when the number or size of any resulting responses is greater than that of the initiating request. These types of vulnerabilities are nearly exclusively limited to just UDP protocols and see frequent enough abuse to justify a notice from US-CERT. When discussing traffic amplification vulnerabilities, an amplification factor is used to describe the relationship between the total size or total number of responses as compared to that of the original request. For example, a vulnerability where a single 1-byte UDP message results in 3 responses of arbitrary size can be said to have a 3x packet amplification factor. Similarly, a vulnerability where a single 8-byte UDP message results in an 800-byte response can be said to have a 100x bandwidth amplification factor.

Distributed Reflective Denial of Service (DRDoS) attacks abuse traffic amplification flaws to overwhelm third-party targets. In a typical attack, an attacker will forge UDP packets with a source address of their intended target and a destination address of the system vulnerable to traffic amplification. With enough traffic amplifiers, because the number or size of each resulting response is larger than that of the forged request, a target can very quickly become overwhelmed with the responses coming from the affected UDP service. While DRDoS attacks continue to be effective, they generally only exploit vulnerabilities where the traffic amplification factor is large enough to overwhelm the target.

This particular disclosure describes traffic amplification vulnerabilities in the Network Time Protocol project's NTP implementation that could be used in DRDoS attacks.

2.10 Network Logging and Monitoring Tools

- Api Monitor v2 Alpha 13 – Portable [DL] – Latest Version (rohitab.com)
- Handy and customizable Api Monitor with advanced filtering capabilities. Standalone version for 32/64bit systems.
- Filemon 7.04 for Nt/Xp/… – Filemon for Nt/Xp/… on Amd64 – Filemon source code
- The famous file monitor
- Ice Sword v1.22
- An effective tool against rootkits, with a lot of additional functions like process dumper/killer/explorer, raw disk access monitor and much more.
- Process Explorer 16.04
- Process Explorer shows you information about which handles and DLLs processes have opened or loaded.
- ProcMon 3.1
- Monitor file system, Registry, process, thread and DLL activity in real-time aka Regmon + Filemon
- Regmon 7.04 for Nt/Xp/… – Regmon for Nt/Xp/… on Amd64 – Regmon source code
- The famous registry monitor
- Spy++ v11.00.50727
- Spying tool with point-and-click Handle/ID grabbing

Extromatica Network Monitor
Extromatica Network Monitor is a network monitoring application created and maintained by Extromatica company. It is designed to monitor network hardware, servers and network services for faults and performance degradation. It alerts users when things go wrong and again when they get better. This software supports a variety of real-time notification mechanisms, including SMS.

The development of this software started in 1999 as internal project by Maxim Perenesenko and Yuri Zaitsev. After 2 years of development if was released as Network Eagle Monitor.[1] It took one more year until first stable release in 2002.

Now, this software is maintained by Extromatica company and has name Extromatica Network Monitor.

- Testing IP channel between monitoring system and another computer or network device with ICMP protocol.
- Checking accessibility of TCP based services like (SMTP, POP3, HTTP, NNTP, FTP, SSH), and so on.

- Checking availability and responsiveness of FTP servers.
- Checking availability and responsiveness of HTTP/HTTPS servers.
- Checking content of Web page by searching for specified substring.
- Monitoring free/used space on a disk or a network share.
- Testing local or remote (Windows share) directory for changes.
- Periodically running external commands or batch scripts and check the return code.
- Accessibility testing of ODBC data source or native MS SQL data source. Running SQL query and check the return result as an option.
- Executing script tests. They can be written on Visual Basic Script, JavaScript or other languages supported by the operating system (ActiveScript technology). For example: Active Python, ActivePerl.
- Monitoring a process either on the local or remote machine by its Process identifier or process name.
- Monitoring local or remote Windows Event Log for specified messages.
- Communicating with a Windows machine to determine if a specified Windows Service is running and responding.
- Monitoring the content of specified file for changes by calculating MD5 hash or searching for substring. This check can test files inside archives.
- Monitoring the status of local or network printer. You can track more than twenty events, such as out of paper, jammed paper, out of toner and much more
- Testing Remote Access Service connection.
- Monitoring various parameters of SNMP enabled computer or device.
- LDAP server accessibility testing and LDAP directory content checking.
- System performance - monitoring CPU loading, memory loading of local or remote Windows computer.

Alerts and actions

- Displaying a popup-window with information about events.
- Executing external program.
- Playing a sound file.
- Sending an e-mail message with the information about event.
- Writing event to Syslog.
- Sending user-defined message to the Windows Event Log.
- Executing script alerts. Alerts can be written on Visual Basic Script, JavaScript or other languages supported by the operating system (ActiveScript technology).
- Rebooting local or temote computer.
- Changing running state of Windows service, local or remote.
- Running user-defined SQL query.
- Sending SMS messages.

2.11 Microsoft Product Activation

Microsoft Product Activation is a DRM technology used by Microsoft Corporation in several of its computer software programs, most notably its Windows operating system and its Office productivity suite. The procedure enforces compliance with the program's end-user license agreement by transmitting information about both the product key used to install the program and the user's computer hardware to Microsoft, inhibiting or completely preventing the use of the program until the validity of its license is confirmed.

The procedure has been met with significant criticism by many consumers, technical analysts and computer experts, who argue that it is poorly designed, highly inconvenient and ultimately does nothing to prevent software piracy. The process has been successfully circumvented on multiple occasions.

Process[edit]

Before activation

The Activation Wizard in Office 2010

When installing a retail copy of Windows or Office, the user is asked to input a unique product key supplied on a certificate of authenticity included with the program, which is later verified during activation.[1] Immediate activation is not required following installation, but the program must be activated within a specific period of time in order to continue to function properly. Throughout this grace period, the user will be periodically reminded to activate the program, with warnings becoming more frequent over time.

Certain versions of Windows and Office are available under a volume license, where a single product key is used for multiple installations. Programs purchased under this license must still be activated, with the exception of Windows XP and all versions of Office released prior to Office 2010. Businesses using this licensing system have the option of using Microsoft's activation servers or creating and managing their own.

If Windows is pre-installed on a computer by an original equipment manufacturer (OEM), the operating system is automatically activated without the need for interaction from the user.[10] In this case, the copy of Windows installed does not use the product key listed on the certificate of authenticity, but rather a master product key issued to OEMs called a System Locked Pre-installation (SLP) key. On each boot, Windows confirms the presence of specific information stored in the BIOS by the manufacturer, ensuring the activation only remains valid on that computer, even if the product key is used on another machine.

After grace period

If activation is not performed within the grace period or fails because of an illegal or invalid product key, the following restrictions will be imposed on the user:

In Windows XP, Windows Server 2003 and Windows Server 2003 R2, after a grace period of 30 days, the operating system cannot be used at all until the activation process is completed successfully.

In Windows Vista, after a grace period of 30 days, the operating system will boot only into a reduced functionality mode. The reduced functionality varies based on whether the operating system is simply out of grace or has undergone a failed activation. In the former case, built-in games and premium features like Windows Aero are disabled, and the system is rebooted every hour; in the latter case, certain premium features are disabled and some content is not available from Windows Update.

In Windows Vista SP1, Windows 7, Windows Server 2008 and Windows Server 2008 R2, after a grace period of 30 days (60 days for Windows Server 2008), the operating system will add a text message in the bottom-right hand corner of the screen stating that the copy of Windows in question is not valid, set the desktop background to black, allow only critical and security updates to be downloaded from Windows Update and give periodic reminders to activate the operating system. However, the operating system otherwise functions normally.

In Windows 8, Windows 8.1, Windows Server 2012 and Windows Server 2012 R2, the 30-day grace period has been removed. If the operating system is not activated, there is a watermark showing the edition of Windows (although it does not show to activate) on desktop, personalization features in PC Settings like changing the lock screen is disabled. Entire Screen notification appears periodically. However, the operating system otherwise functions normally.

In Office XP, Office 2003, Office 2007, Office 2010, and Office 2013, after a grace period of 30–60 days for Office 2010 and 14–60 days in Office 2013 or opening the program 25 times for

Office 2007 and 50 times for Office 2003 and XP, the programs will enter a reduced functionality mode, where files can be viewed but not edited.

When activation takes place, the program saves a record of the verification data in the user's computer. If the system is booted up with significant hardware changes, the application will likely require reactivation to prevent the same copy of the program being installed on two different systems.

During activation
Activation is performed with a utility supplied with Windows and Office called the Activation Wizard. It can be performed either over the Internet or by telephone.[1] When activating over the Internet, the Activation Wizard automatically transmits and receives verification data to and from Microsoft servers, completing the process without any interaction by the user.[22] Activation by telephone requires that a user and a Microsoft agent verbally exchange activation information. In this case, an installation ID is generated, which is then read to the agent. The agent verifies the information and replies with a confirmation ID, which is then typed into the Activation Wizard.
The Activation Wizard generates verification data primarily based on information about hardware in the computer. In Windows XP, information about the following eight categories of hardware is included:
Display adapter
SCSI adapter
IDE adapter
Network adapter MAC address
RAM amount range (e.g. 0-512 MB)
Processor type and serial number
Hard drive device and volume serial number
Optical drive (e.g. DVD-ROM)

The verification data is also based on the product key entered during activation. In some cases, the product key is checked against a list of known illegally distributed keys.[23].

Certain retail copies of Windows and Office sold in certain countries classified as emerging markets have geographical activation restrictions, which only allow the user to activate the product within the indicated region.

After activation
If activation completes successfully, the user can continue to use the application without any further issues or impediments.

2.12 Half-open Ports

In a port scan based on SYN packets, the scanner machine sends out SYN packets to the different ports of a remote machine.

When the scanner machine receives a SYN+ACK packet in return for a given port, the scanner can be sure that the port on the remote machine is open. It is the “duty” of a good port-scanner to immediately send back to the target machine an RST packet in response to a received SYN+ACK packet so that the half-open TCP circuit at the target is closed immediately.

TCP SYN Scan: “half-open” scan, look for SYN-ACK, then send RESET, in this case the target system will not record the attempted connection. It is faster than the TCP connect scan.

2.13 Honey Pot Systems

Honey Pot Systems are decoy servers or systems setup to gather information regarding an attacker or intruder into your system. It is important to remember that Honey Pots do not replace other traditional Internet security systems; they are an additional level or system.

Honey Pots can be setup inside, outside or in the DMZ of a firewall design or even in all of the locations although they are most often deployed inside of a firewall for control purposes. In a sense, they are variants of standard Intruder Detection Systems (IDS) but with more of a focus on information gathering and deception.

An example of a Honey Pot systems installed in a traditional Internet security design:
A Honey Pot system is setup to be easier prey for intruders than true production systems but with minor system modifications so that their activity can be logged of traced. The general thought is that once an intruder breaks into a system, they will come back for subsequent visits. During these subsequent visits, additional information can be gathered and additional attempts at file, security and system access on the Honey can be monitored and saved.

Generally, there are two popular reasons or goals behind setting up a Honey Pot:
Learn how intruders probe and attempt to gain access to your systems. The general idea is that since a record of the intruder’s activities is kept, you can gain insight into attack methodologies to better protect your real production systems.

Gather forensic information required to aid in the apprehension or prosecution of intruders. This is the sort of information often needed to provide law enforcement officials with the details needed to prosecute.

The common line of thought in setting up Honey Pot systems is that it is acceptable to use lies or deception when dealing with intruders. What this means to you when setting up a Honey Pot is that certain goals have to be considered.

Those goals are:
The Honey Pot system should appear as generic as possible. If you are deploying a Microsoft NT based system, it should appear to the potential intruder that the system has not been modified or they may disconnect before much information is collected.

You need to be careful in what traffic you allow the intruder to send back out to the Internet for you don’t want to become a launch point for attacks against other entities on the Internet. (One of the reasons for installing a Honey Pot inside of the firewall!).

You will want to make your Honey Pot an interesting site by placing "Dummy" information or make it appear as though the intruder has found an "Intranet" server, etc. Expect to spend some time making your Honey Pot appear legitimate so that intruders will spend enough time investigating and perusing the system so that you are able to gather as much forensic information as possible.

Some caveats exist that should be considered when implementing a Honey pot system. Some of the more important are:
The first caveat is the consideration that if the information gathered from a Honey Pot system is used for prosecution purposes, it may or may not be deemed admissible in court. While information regarding this issue is difficult to come by, having been hired as an expert witness for forensic data recovery purposes, I have serious reservations regarding whether or not all courts will accept this as evidence or if non-technical juries are able to understand the legitimacy of it as evidence.

The second main caveat for consideration is whether hacking organizations will rally against an organization that has set "traps" and make them a public target for other hackers. Examples of this sort of activity can be found easily on any of the popular hacker’s sites or their publications.

Levels or Layers of Tracking

The information provided on an intruder depends on the levels of tracking that you've enabled on your Honey Pot. Common tracking levels include the firewall, system logs on the Honey Pot and sniffer-based tools.

Firewall Logs

Firewalls are useful as part of the overall Honey Pot design for many reasons. Most firewalls provide activity-logging capabilities which can be used to identify how an intruder is attempting to get into a Honey Pot. I liken firewall logs to router logs; they can both be set to trap and save packets of a pre-determined type. Remember that when setting up the firewall, you would normally want to log ALL packets going to the Honey Pot system, as there should be no legitimate reason for traffic going to or from the Honey Pot.

Reviewing the order, sequence, time stamps and type of packets used by an intruder to gain access to you Honey Pot will help you identify the tools, methodology being used by the intruder and their intentions (vandalism, data theft, remote launch point search, etc.). Depending on the detail capabilities of logging on your firewall you may or not be able to gain considerable information from these logs.

Another useful function of many firewalls is their notification capabilities. Most firewalls can be configured to send alerts by email or pager to notify you of traffic going to or from your Honey Pot. This can be extremely useful in letting you review intruder activity WHILE it's happening.

System Logs

Unix and Microsoft NT seem to have the lion share of the Internet server markets. Luckily, both operating systems have logging capabilities built into their operating systems, which help identify what changes or attempts have been made. It should be noted that out-of-the box, Unix offers superior logging capabilities as compared to Microsoft NT.

Some of their out-of-the box logging capabilities include:
Microsoft NT
Security – Available from Event Viewer
User Management – Needs to be enabled through User Manager
Running Services – Netsvc.exe needs to be manually run and compared to baseline.

Unix
User activity logs – utmp, wtmp, btmp, lastlog, messages
Syslogd – An important option is that it can log to a remote server! The range of facilities and priorities available through syslogd is very good.
There are also several tools available that greatly increase the information that can be gathered. Many of the Unix tools are public domain, while many of the Microsoft NT tools are not.

Sniffer Tools

Sniffer tools provide the capability of seeing all of the information or packets going between the firewall and the Honey Pot system. Most of the sniffers available are capable of decoding common tcp packets such as Telnet, HTTP and SMTP. Using a sniffer tool allows you to interrogate packets in more detail to determine which methods the intruder is trying to use in much more detail than firewall or system logging alone.

An additional benefit to sniffer tools is that they can also create and store log files. The log files can then be stored and used for forensic purposes.

Honey Pot Solutions

Implementation of a Honey Pot solution as part of a security system first involves the decision of whether to purchase a commercial solution or decide to develop your own.

Building a Honey Pot
There is a variety of public domain tools and software available that can be useful to help you setup a Honey Pot as well as many sites dedicated to helping guide you through the process. Most tools seem to have originated on the Unix platform, while many have been ported to Microsoft NT.

What you will need to create or develop your own Honey Pot system are a minimum of the following components and considerable configuration time:
A Workstation or PC. It appears as though an Intel-based workstation is fine.
An operating system. I prefer BSD Unix or RedHat as there are more tools available for the Unix platform than NT.

Sniffer package.
Commercial Honey Pot Systems
There are a variety of commercial Honey Pot systems available. The operating systems most widely supported are Microsoft NT and Unix. As many of the commercial product have been released in the past 12 – 18 months, some of them are still in relatively early versions. I tried to find information regarding market share but wasn't able to find any published statistics.

2.14 DNS Hacking
DNS is a naming system for computers that converts human readable domain names e.g. (infosecinstitute.com) into computer readable IP-addresses. However some security vulnerabilities exist due to misconfigured DNS nameservers that can lead to information disclosure about the domain. This forms an important step of the Information Gathering stage during a Penetration test or Vulnerability assessment. In this article we will look at the following areas.
DNS Basics
Resource records and the Zone file
DNS Lookup and Reverse DNS Lookup
Understanding Wildcard Entries
DNS Zone transfer
DNS Bruteforcing

1) DNS Basics-DNS converts human readable domain names into IP-addresses. This is because domain names are much easier to remember than IP-addresses. This process may take place through a local cache or through a zone file that is present on the server. A zone file is a file on the server that contains entries for different Resource Records (RR). These records can provide us a bunch of information about the domain. We will look more into Resource Records and the zone file in the next section.

So Let's understand how DNS resolution works. Let's say the user opens up the browser and types in infosecinstitute.com. It is now the responsibility of the DNS resolver in the user's operating system to fetch the IP address. It first checks it's local cache to see if it can find a record for the queried domain name. A cache usually contains a mapping of IP-addresses to hostnames which are saved during recent lookups so that the resolver does not have to fetch the IP address again and again. If it can't find the IP address in it's cache it queries the DNS server to see if it has a record for it. A DNS server is usually given to you by the ISP or you can manually set up a DNS server for yourself.If it still can't find the IP Address then it goes through a process or recursive DNS query in which it queries different nameservers to get the IP-address of the domain. As soon as it finds the IP-address it returns the IP-address back to the user and also caches it for it's future use.

Let's do a quick demo. We are going to use the "nslookup" utility for this demo. Just type in the commands as shown in the figure below.
Nslookup Intro

a) In the second line we set the type = a . This means that we are querying for the A records which will return us an IP-address in return for the domain we query. We will look more into records in the next section.

b) As soon as we type in google.com we get an output showing the server and an IP-address#port. This server is basically the current DNS server that will be serving our request. In this case it is 10.0.1.1 and the port no is 53. This is because DNS uses UDP port 53 to serve its requests. We can also set the current DNS server by using the command “server Ip-address”.

c) The third line in the output shows “Non-authoritative answer”. This basically means that our DNS server queried an external DNS server to fetch the IP-address. Below we can see all the IP-addresses associated with google.com. This is usually the case with large organizations. They use multiple servers to serve the request as one server is generally not capable of handling all the requests.

2) Resource Records and the Zone file -A Zone file is basically a text file present on the server hosting the domain that contains entries for different resource records. Each line is represented by a different record .In some cases these records may exceed one line and hence must be enclosed within a parantheses. Each zone file must start with a Start of Authority (SOA) record containing an authoritative nameserver for the domain (for e.g. ns1.google.com for google.com) and an email address of someone responsible for the management of the nameserver.

Different types of Resource Records exist within a Zone file. However we are going to discuss some of the important ones

A Records- Maps an IP Address to a hostname.For e.g. 74.125.236.80 for google.com.

NS Records-Delegates a given zone to use the given authoritative nameserver. For e.g. ns1.google.com is an authoritative nameserver for google.com.

MX Records-This basically tells us which server is responsible for receiving mails sent to that domain name.

TXT Records-This consists of arbitrarily human readable text in a record.

CNAME Records- Gives an alias of one name to another.

Let’s do a demo to make this clear. I have purposely added some records in my website searching-eye.com for this article,so they may not be available when you perform this, however you can try the same exercise on other domains, type in the commands as shown in the figure below.

Nslookup Detail
a) In the first command in nslookup I set the type to A which means I want IP-address for a particular domain. I type in the domain name as the second command and get the corresponding IP-address for it.

b) In the third command i set the type to NS as i am interested in finding the nameservers for searching-eye.com. Type in the domain name as the fourth command and we get the corresponding nameservers for the domain searching-eye.com. Note that finding the nameservers can give us some information about the hosting provider of the domain. Some large organizations use their own nameservers e.g. ns2.google.com.

c) I now set the current server to one of the nameservers, this is because I am interested in finding the latest information about the domain. Note that querying from your own dns server may not give you the accurate information every time. I set the type to MX and again type in the domain name. What we get is a list of mail servers responsible for handling emails sent to that domain. The number before them denotes the priority with which to fetch mails. Lower the number, higher the priority.

d) Next i set the type to CNAME and type in a subdomain, i get a canonical name as infosecinstitute.com. This means any request to the queried domain (in this case prateek.searching-eye.com) will be redirected to infosecinstitute.com.
I will take this moment to introduce DIG which is a handy little tool, we can
also do the same queries using DIG. Let's search for MX records in the same domain. I would suggest you try querying for the other domains yourself.
Dig Intro

3) DNS Lookup and Reverse DNS Lookup
DNS Lookup-Let's perform a DNS Lookup ourselves for infosecinstitute.com. We will do this by traversing the entire DNS hierarchy from the root servers to the top level domain.

4) Understanding Wildcard Entries-
WildCard – A wildcard entry is used to provide responses for subdomains that do not exist. For e.g. let's say we have a domain example.com. If we set a wildcard record for *.example.com and give it the value example.com then the requests for all the non-existent subdomains of example.com (for e.g. abcd.example , blah.example.com) will point to example.com. In the information gathering stage of a penetration test of a website, it is important to identify the subdomains and the IP-addresses corresponding to them. Introducing a Wildcard feature reduces this to a small extent.

Bypassing Wildcard entries – In case wildcard entries are set on a particular domain, they could be bypassed to reveal information about it's subdomains. This is done by brute forcing the subdomains. We have a wordlist in which we contain the subdomain names we want to test the domain against. Then we do a ping of all these subdomains, if these domains resolve to an IP-address different than the host IP-address, then we can very surely say that this subdomain actually exists. However before performing a brute force it would be better to actually check if Wildcard entries are enabled or not. For that we can ping some random subdomains for e.g. 434234.example.com and see if it's IP-address is the same as the host IP-address(in this case example.com). If this is the case for some random subdomains, then we can clearly say that Wildcard entries are enabled for this domain. We will perform a demo of this in the coming section.

5) DNS Zone Transfer-We saw in the previous exercises that every domain has some authoritative name servers associated with it. For eg in the case of google.com, the nameservers were ns1.google.com to ns4.google.com .These Nameservers are used for handling requests related to the domain google.com. Let's say we have a domain example.com and it has it's two nameservers as ns1.example.com and ns2.example.com. Usually a big organization will have more than one nameservers so that if one goes down for some time, the other one is ready to back it up and handle the requests. Usually one of these servers will be the Master server and the other one will be the slave server. Hence to stay in sync with each other, the slave server must query the Master server and fetch the latest records after a specific period of time. The Master server will provide the slave server with all the information it has. This is basically what is called a "Zone Transfer". It's like asking the nameserver "Give me everything you have". A properly configured nameserver should only be allowed to serve requests of Zone transfer from other Nameservers of the same domain. However if the server is not configured properly it will serve all requests of Zone transfer made to it without checking the querying client. This leads to leakage of valuable information. DNS Zone transfer is sometimes referred through it's opcode mnemonic AXFR.

A zone transfer reveals a lot of information about the domain. This forms a very important part of the "Information Gathering" stage during a penetration test, vulnerability assessment etc. We can figure out a lot of things by looking at the dump.For e.g. we can find different subdomains. Some of them might be running on different servers.Those server may not be fully patched and hence be vulnerable.From this point, we can start thinking about Metasploit ,Nessus,Nmap etc and do a full vulnerability assessment of the domain. Hence this kind of information increases our attack vector by a fair amount, an amount which cannot be ignored.

To protect your nameservers from leaking valuable information, one must allow zone transfer to other nameservers of the same domain only. For e.g. ns1.example.com should allow zone transfer to ns2.example.com only and discard all the other requests.

6)DNS Bruteforcing-DNS Zone transfers may not work all the time. In fact, it will not work most of the time. Most of the DNS servers are properly configured and do not allow zone transfers to every client. Well what do we do then ? Simple answer, the same thing we do when nothing works, BRUTE FORCE it ! Basically we have a wordlist containing a huge list of hosts. We first check for wildcard entries by checking if a random subdomain for e.g. 132qdssac.example.com resolves to the same IP-address as example.com .If this is the case, we know Wildcard entries are set. We then query the domain by using each of the word in our wordlist. For e.g. if one of the entries in the wordlist file is “ads” , then we make a query for ads.example.com. If it resolves to a different IP-address then we are sure that this subdomain actually exists. Hence we now have information about the name of subdomain and it’s IP-address. If wildcard entries are not set , we do the same thing and see if we get response from any subdomain we query. If we get a response back, we could be sure that the subdomain actually exists. In the end what we get is a bunch of information about the domain.

Let’s see this through a demo. We will again use the tool “Fierce”. Fierce is a very handy tool for DNS Analysis and it is something everyone should have in their armory. Fierce will first check if Zone transfers are allowed or not, if zone transfers are allowed, it will dump all the information and exit happily, otherwise it will brute force it. We need to supply Fierce with a wordlist containing a list of all the possible subdomain names (for e.g. hosts,ads,contracts). Fierce comes with an inbuilt wordlist file “hosts.txt” and we will be using the same for our demo.
Fierce Bruteforcing
Fierce Bruteforcing2
As we can see, Fierce dumps out information about the subdomains of google.com

2.15 Routing Protocol Attacks

SRP is a secure on-demand source routing protocol for ad-hoc networks proposed in [9]. The design of the protocol is inspired by the DSR protocol [7], however, DSR has no security mechanisms at all. Thus, SRP can be viewed as a secure variant of DSR. SRP tries to cope with attacks by using a cryptographic checksum in the routing control messages (route requests and route replies).

This checksum is computed with the help of a key shared by the initiator and the target of the route discovery process; hence, SRP assumes only shared keys between communicating pairs.
In SRP, the initiator of the route discovery generates a route request message and broadcasts it to its neighbors. The integrity of this route request is protected by a Message Authentication Code (MAC) that is computed with a key shared by the initiator and the target of the discovery. Each intermediate node that receives the route request for the first time appends its identifier to the request and re-broadcasts it. The MAC in the request is not updated by the intermediate nodes, as by assumption, they do not necessarily share a key with the target. When the route request reaches the target of the route discovery, it contains the list of identifiers of the intermediate nodes that passed the request on. This list is considered as a route found between the initiator and the target. The target verifies the MAC of the initiator in the request. If the verification is successful, then it generates a route reply and sends it back to the initiator via the reverse of the route obtained from the route request. The route reply contains the route obtained from the route request, and its integrity is protected by another MAC generated by the target with a key shared by the target and the initiator. Each intermediate node passes the route reply to the next node on

the route (towards the initiator) without modifying it. When the initiator receives the reply it verifies the MAC of the target, and if this verification is successful, then it accepts the route returned in the reply.

The basic problem in SRP is that the intermediate nodes cannot check the MAC in the routing control messages. Hence, compromised intermediate nodes can manipulate control messages, such that the other intermediate nodes do not detect such manipulations. Furthermore, the accumulated node list in the route request is not protected by the MAC in the request, hence it can be manipulated without the target detecting such manipulations.

2.16 Cross-Site Scripting

Cross-Site Scripting (also known as XSS) is one of the most common application-layer web attacks. XSS vulnerabilities target scripts embedded in a page which are executed on the client-side (in the user's web browser) rather than on the server-side. XSS in itself is a threat which is brought about by the internet security weaknesses of client-side scripting languages such as HTML and JavaScript. The concept of XSS is to manipulate client-side scripts of a web application to execute in the manner desired by the malicious user. Such a manipulation can embed a script in a page which can be executed every time the page is loaded, or whenever an associated event is performed.

XSS is the most common security vulnerability in software today. This should not be the case as XSS is easy to find and easy to fix. XSS vulnerabilities can have consequences such as tampering and sensitive data theft.

Key Concepts of XSS
XSS is a Web-based attack performed on vulnerable Web applications.
In XSS attacks, the victim is the user and not the application.
In XSS attacks, malicious content is delivered to users using JavaScript.

Explaining Cross-Site Scripting
An XSS vulnerability arises when Web applications take data from users and dynamically include it in Web pages without first properly validating the data. XSS vulnerabilities allow an attacker to execute arbitrary commands and display arbitrary content in a victim user's browser. A successful XSS attack leads to an attacker controlling the victim's browser or account on the vulnerable Web application. Although XSS is enabled by vulnerable pages in a Web application, the victims of an XSS attack are the application's users, not the application itself. The potency of an XSS vulnerability lies in the fact that the malicious code executes in the context of the victim's session, allowing the attacker to bypass normal security restrictions.

Cross-Site Scripting Video
XSS Attack Examples
Reflective XSS
There are many ways in which an attacker can entice a victim into initiating a reflective XSS request. For example, the attacker could send the victim a misleading email with a link containing malicious JavaScript. If the victim clicks on the link, the HTTP request is initiated from the victim's browser and sent to the vulnerable Web application. The malicious JavaScript is then reflected back to the victim's browser, where it is executed in the context of the victim user's session.

Persistent XSS
Consider a Web application that allows users to enter a user name which is displayed on each user's profile page. The application stores each user name in a local database. A malicious user notices that the Web application fails to sanitize the user name field and inputs malicious JavaScript code as part of their user name. When other users view the attacker's profile page, the malicious code automatically executes in the context of their session.

Impact of Cross-Site Scripting

When attackers succeed in exploiting XSS vulnerabilities, they can gain access to account credentials. They can also spread Web worms or access the user's computer and view the user's browser history or control the browser remotely. After gaining control to the victim's system, attackers can also analyze and use other intranet applications.

By exploiting XSS vulnerabilities, an attacker can perform malicious actions, such as:

- Hijack an account.
- Spread Web worms.
- Access browser history and clipboard contents.
- Control the browser remotely.
- Scan and exploit intranet appliances and applications.
- Identifying Cross-Site Scripting Vulnerabilities
- XSS vulnerabilities may occur if:
- Input coming into Web applications is not validated
- Output to the browser is not HTML encoded

2.17 Session Fixation Bugs

Improper handling of session variables in asp.NET websites is considered a serious threat and opens various doors to malicious hackers. For instance, a session variable could be manipulated in a way to subvert login authentication mechanisms. However, this article illustrates a session fixation bug in a .NET website by demonstrating various live scenarios which usually leads to a website becoming vulnerable, in terms of session hijacking. Moreover, the article circulates detailed information about exploiting vulnerable websites, as well as recommendations of practices for protecting them against session fixation attacks.

Internal Session Fixation

A session fixation attack allows spoofing another valid user and working on behalf of their credentials. It typically fixates another person's session identifier to breach currently happening communication. The asp.NET base website usually keeps session variables to track the user by creating a cookie called asp.NET_SessionId in the browser. A session variable is typically used to record a currently logged-in user, and such a cookie value is validated on each round-trip to ensure that the data being served is specific to that user. Here the following image is describing the process of cookies-based authentication, where the user performs the login operation to a vulnerable website, and in return, the server issues this particular user a cookie token value for session management.

Websites usually engage session management to construct a user-friendly environment. But this mechanism is vulnerable to some extent, because session IDs present an attractive target for attackers, as they are stored on a server and associated with respective users by unique session identifier values. There are a couple of approaches applied by the attacker to perform a session fixation attack, depending on the session ID transport mechanism (cookies, hidden fields, and URL arguments) and the loopholes identified on the target system.

The mechanics of session management is that the server generates a unique session identifier value during user authentication, and sends this session ID back to the client browser and makes sure that this same ID will be sent back by the browser along with each forthcoming request. Hence, such a unique session ID value thereby becomes an identification token for users, and servers can use them to maintain session data.

An asp.NET_SessionID cookie is only configured by the server whenever working on behalf of any page request of the website. So when the login page is first accessed, the asp.NET_SessionID cookie value is set by the client browser and server uses this cookie value for all subsequent requests. Even after authentication is successful and logged out, the asp.NET_SessionID value does not change. This results in the possibility of a session fixation attack, where a hacker can potentially sniff the traffic across the wire or physically access the victim machine in order to get the stored cookie values in the browser and fix a victim's session by accessing the login page, even if they don't have the actual user name or password.

The following image shows the real time session fixation attack scenario where a potential hacker sits somewhere in the network and intercepts the traffic happening between a server and client. Here, the hacker employs a packet sniffer to capture a valid token session and then utilizes the valid token session to gain unauthorized access to the web server. Finally, the hacker successfully accesses the asp.NET_SessionID value and logs in successfully to the website's sensitive zone.

Vulnerable Code Scenario

Session fixation bugs usually occur on websites which manipulate sensitive data while transacting or incorporating with the login page to authenticate valid users with correct user name and password. This paper illustrates this crucial bug in detail by presenting this vulnerable login authentication code as follows:

```
if (txtUsr.Text.Equals("frank") && txtPwd.Text.Equals("password"))
{
   Session["LIn"] = txtU.Text.Trim();
   Server.Transfer("<atitle="Home"href="http://resources.infosecinstitute.com/">Home</a>.aspx");
}
else
{
  lblMessage.Text = "Wrong username or password";
}
```

When a user browses this website and enters the valid credentials for authentication, the internal mechanism flashes the server message that either the user name and password are correct or incorrect as follows:

The user typically assumes that this transaction is safe and there are fewer possibilities of other website-related attacks, but still, a couple of serious attacks such as spoofing, replay and session hijacking attacks could be possible, even if managing the user name and password correctly. We shall see this in a forthcoming segment of this article.

Stealing Cookies

Valid session IDs are not only recognized to be identification tokens, but also employed as an authenticators. Users are authenticated based on their login credentials (e.g. user names and passwords) and are issued session IDs that will effectively serve as temporary static passwords for accessing their sessions, which makes session IDs a very appealing target for attackers. The moment a user enters his credentials on login to authenticate, these data are stored in the session and cookies are generated in the client browser. The user is typically over-confident that when he is logged out, all the data would be scrubbed automatically and the session is terminated, but unfortunately, the cookie values are not deleted from the client browser, even if the session is ended, and such cookie values could be exploited by a hacker to breach into the website's sensitive zone, without being aware of user name and password.

As the following figure shows, when a user is logged in, the browser shows cookie values which are generated during authentication as:

Now log out and refresh the page. It is generally assumed that cookie values should be wiped-out automatically at the time of ending the current session, but even after proper sign-out from the current session, which is performing Session.Abandon(), Session.Clear() implicitly, the browser is still showing the previous session's generated cookies values as follows:

Hence, revealing cookie values even without being logged in could be considered a serious threat and opens the doors to a session hijacking attack. A malicious hacker could directly access the sensitive zone of a website without being logged in by adding such retrieved cookie details manually to the browser. Here, the attacker typically uses this technique to inject the stolen cookies in the browser to hijack the someone else's current session as follows:

Defense (Securing Cookies)

Countermeasures combine several approaches to overcome such session hijacking attacks. For instance, making cookie values bullet-proof by HttpOnly, explicitly removing session cookie values, employing HTTPS/ TLS (via Secure Attribute) and proper configuration. This section fixes the session hijacking vulnerability in the aforesaid code, where cookie values are not discarded even after logout, by generating another cookie having a unique value which is compared to the session value at each round-trip. Resemblance of both of these values could allow the user to enter into the website's sensitive zone; otherwise the user is redirected to the login page. This generates a unique value never to be duplicated and there is a very low chance that the value of the new GUID is all zeroes or equal to any other GUID. Hence, such an applied random token ensures protection against a CSRF attack in a website.

```
protected void Page_Load(object sender, EventArgs e)
 {
       if (Session["LIn"] != null && Session["AuthToken"] != null
                   && Request.Cookies["AuthToken"] != null)
       {
          if (!Session["AuthToken"].ToString().Equals(
                   Request.Cookies["AuthToken"].Value))
          {
             lblMessage.Text = "You are not logged in.";
          }
          else
          {
             ..
          }
       }
       ..
 }
```

This time in the sign-in button, another unique value GUID is generated and stored with the session variable AuthToken which is added to cookies later as follows:

```
protected void btnLogin_Click(object sender, EventArgs e)
{

      if (txtUsr.Text.Equals("frank") && txtPwd.Text.Equals("password"))
      {
         Session["LIn"] = txtU.Text.Trim();
         string guid = Guid.NewGuid().ToString();
         Session["AuthToken"] = guid;
         // now create a new cookie with this guid value
         Response.Cookies.Add(new HttpCookie("AuthToken", guid));
      }
    ..
}
```

Finally, the logout button has the code to expire the session cookie values explicitly, which removes them from the client browser permanently. Here, we shall have to remove both session asp.NET_SessionId and AuthToken variables as follows:

```
protected void btnLogout_Click(object sender, EventArgs e)
 {
       Session.Clear();
       Session.Abandon();
       Session.RemoveAll();

       if (Request.Cookies["asp.NET_SessionId"] != null)
```

```
        {
            Response.Cookies["asp.NET_SessionId"].Value = string.Empty;
            Response.Cookies["asp.NET_SessionId"].Expires = DateTime.Now.AddMonths(-20);
        }

        if (Request.Cookies["AuthToken"] != null)
        {
            Response.Cookies["AuthToken"].Value = string.Empty;
            Response.Cookies["AuthToken"].Expires = DateTime.Now.AddMonths(-20);
        }
}
```

Okay, now browse the website again and login with the correct credentials and compare the output in the firebug with the previous output shown in figure 1.5. Here another session value AuthToken with new cookies is generated as follows:

Thereafter, sign out from the current session as earlier and refresh the page and notice the cookies section in the firebug again. Bingo! This time the browser doesn't retain any previously stored cookie values. Hence, making cookie values bullet-proof ensures to protect against session fixation attack.

2.18 Session Fixation Bugs

The pwdencrypt() vulnerability is a classic buffer overflow exploit. The attacker can make use of an undocumented, little understood function that is accessible by default to every authenticated user to pass any code of their choosing to the operating system in the context of a trusted service. This vulnerability is particularly dangerous in that some system administrators would initially dismiss its severity. They may reason that to exploit it would be difficult since the attacker must use a legitimate connection to SQL server in order to initiate the exploit. However, there are many ways to acquire such a connection through insiders, weak credentials, development systems, linked servers, default passwords, through SQL injection, or through other vulnerabilities such as CVE-2001-0344 which allows local users to gain connections by re-using cached sa connections.

Vulnerabilities like this combined with the lack of documentation, support, toolkits or training specific to SQL Server security create a dangerous environment ripe for targeting in a variety of attacks.

At the time of this writing (Jan 27th, 2003), the W32/SQL Slammer worm has just infected over 35,000 servers in under 24 hours using a vulnerability that has been patched since July 24th, 2002. The patch for the pwdencrypt() vulnerability has existed since July 10th, 2002.

2.19 Meeting Point for Hacker - Penetration Tester

Attackers are becoming more clever and their attacks more complex. To keep up with the latest attack methods, you need a strong desire to learn, the support of others, and the opportunity to practice and build experience.

Pentester should know advanced penetration concepts and providing an overview to prepare students for what lies ahead. Pentester should be able access, manipulate, and exploit the network, perform attacks against NAC, VLANs, OSPF, 802.1X, CDP, IPv6, VOIP, SSL, ARP, SNMP, and othersShould know how perform penetration testing against various cryptographic implementations, network booting attacks, escaping Linux restricted environments such as chroot, and escaping Windows restricted desktop environments. Learn to identify privileged programs, redirect the execution of code, reverse-engineer programs to locate vulnerable code, obtain code execution for administrative shell access, and defeat modern operating system controls such as ASLR, canaries, and DEP using ROP and other techniques. Know local and remote exploits as well as client-side exploitation techniques.

3. Conclusions

We presented our analysis for modern attacks from penetration testing side. We found serious logic flaws in advanced attacks mechanisms. We discussed the weaknesses and ways of its protection.

SSL/TSL has been around for many years without any major modifications. This protocol was considered to be secure. The CRIME, BREACH,BEAST, Heartbleed attacks proved that in one very specific use case it can be compromised. While this use case can be avoided and SSL/TSL re-secured, will this have an effect on the thoughts of SSL/TSL security as a whole. People tend to lose faith in security protocols as soon as the simplest attack is successful. Will this be the end to SSL/TSL, or will users still have faith in the non-compressed version, that has yet to be broken, or will they run to a new protocol to be positive that they are secure? This will only be answered in time.

Our results provide insights into modern DDoS attacks and help us to understand how such attacks are carried out nowadays.

We believe that our study takes some steps in the security problems. In this article, we learned about the denial of service attack and tools used to perform the attack. DOS attacks are used to crash servers and disrupt service. Sony has faced this attack for a long time and lost millions of dollars. It was a big lesson for other companies who rely on server-based income. Every server should set up a way to detect and block DDOS attacks. The availability of free tools makes it easier to perform DOS attack against a website or server. Although most of these tools are only for DOS attacks, a few tools support a zombie network for DDOS attacks. LOIC is the most used and most popular DOS attacking tool. In the past few years, it has been used many times by hackers against big company’s network, so we can never deny the possibility of attack.

So, every company should take care of it and set up good level of protection against DOS attack.
Also DNS protocol is a very critical component of the Internet as it resolves IP-address into hostnames and makes life a lot easier for us. However, if the nameservers are not properly configured they might leak out the whole DNS server database to any malicious hacker. Even if the servers are properly configured, they can be brute forced to leak information about their mail servers, IP addresses, etc. It is therefore important to properly configure your DNS servers and be aware of the security issues with DNS. This article has explained the session fixation attack on asp.NET website in detail by giving the real time code scenario, and also pinpoints the common glitches committed by programmer at time of coding of sensitive parts like login pages. We have seen how a potential hacker can access the cookies values stored in the client browser in order to execute a session hijacking attack and breach into the sensitive zone of a website, even without being aware or having a real user name and password. Finally, we have come to an understanding to secure or make bullet-proof the cookie session values to protect our website from session fixation attack.

4. REFERENCES

[1] Fraser, B “Site Security Handbook”. Internet: http://www.ietf.org/rfc/rfc2196.txt?number=2196.

[2] Herzog, Pete “The open source security testing methodology manual”. Internet: http://www.ideahamster.org/osstmm.htm.

[3] Kaye, Krysta “Vulnerability Assessment of a University Computing Environment”. Internet: http://rr.sans.org/casestudies/univ_comp.php.

[4] “Risk Assessment Tools and Practices for Information System Security”. Internet: http://www.fdic.gov/news/news/financial/1999/FIL9968a.html.

[5] Klikushina, Natalya "Firewall Penetration". Internet: http://shrike.depaul.edu/~mchen/420/natalya.html.

[6] "Nmap Free Stealth Security Scanner". Internet: http://nmap.org.

[7] Corcoran, Tim "An Introduction to NMAP". Internet: http://rr.sans.org/audit/nmap2.php.

[8] "Quality Security Tools". Internet: http://nmap.org/tools.html.

[9] "Internet Security Systems". Internet: http://www.iss.net.

[10] Kurtz, George and Prosise, Chris "Security Strategies" Information Security Magazine September 00(also available at Internet: http://www.infosecuritymag.com/articles/september00/features3.shtml).

[11] Antionline.com. Internet: http://www.antionline.com/index.php?action=forums.

[12] Moyer, Philip "Penetration Testing: Issues for Management". Computer Security Institute's Alert Magazine March 1998 (also available at Internet: http://www.gocsi.com/penet.htm).

[13] McClure, Stuart; Scambray, Joel; Kurtz, George Hacking Exposed Berkley, Osborne 1999.

[14] DOS Attacks and Free DOS Attacking Tools. Internet: http://resources.infosecinstitute.com/dos-attacks-free-dos-attacking-tools/.

[15] KoonYaw Tan Intrusion Detection FAQ: How can attacker use ICMP for reconnaissance? Internet: http://www.sans.org/security-resources/idfaq/icmp_misuse.php

[16] Ofir Arkin, ICMP Usage in Scanning – The Complete Know How. Internet: http://www.sys-security.com/html/papers.html

[17] Stephen Northcutt and Judy Novak, Network Intrusion Detection .

[18] ICMP Parameters Internet: http://www.iana.org/assignments/icmp-parameters .

[19] RFC 792 Internet Control Message Protocol . Internet: http://www.ietf.org/rfc/rfc0792.txt .

[20] Craig Huegen, The Latest in Denial of Service Attacks: 'Smurfing': Description and Information to Minimize Effects. Internet: http://www.pentics.net/denial-of-service/white-papers/smurf.cgi .

[21] David Dittrich, The "Tribe Flood Network" Distributed Denial of Service Attack Tool. Internet: http://staff.washington.edu/dittrich/misc/tfn.analysis .

[22] Loki Project. Internet: http://www.phrack.org/show.php?p=49&a=6 .

[23] RFC 1122 Requirements for Internet Hosts – Communication Layers. Internet: http://www.ietf.org/rfc/rfc1122.txt.

[24] SING utility. Internet: http://sourceforge.net/projects/sing/ .

[25] HPING2 utility. Internet: http://sourceforge.net/projects/hping2/ .

[26] NMAP. Internet: http://www.insecure.org/nmap/.

[27] Spoofing ICMP redirect host messages with hping. Internet: http://blog.packetheader.net/2010/05/spoofing-icmp-redirect-host-messages.html.

[28] Icmp address mask ping. Internet: http://www.networkuptime.com/nmap/page4-8.shtml.

[29] Internet: http://www.ddifrontline.com/security-awareness-education/cyber-crime-security-prevention-

[30] IPv6 Security Testing and Monitoring Tools . Internet: http://ipv6now.com.au/testing.php.

[31] Vesselin Hadjitodorov Security of IPv6 and DNSSEC for penetration testers. Internet: http://www.delaat.net/rp/2010-2011/p40/report.pdf

[32] Tech Insight: Retooling Vulnerability Scanning, Penetration Testing For IPv6. Internet: www.darkreading.com/vulnerabilities---threats/tech-insight-retooling-vulnerability-scanning-penetration-testing-for-ipv6/d/d-id/1134284?

[33] Avoid Pay Per Click Problems. Internet: http://www.internetworldstats.com/articles/art090.htm.

[34] 7 Ways to Use Google Webmaster Tools to Increase Traffic To Your Website. Internet: http://www.razorsocial.com/google-webmaster-tools-7-ways-to-increase-traffic-to-your-website/.

[35] Warwick Ashford, IPv6: The security risks to business. Internet: http://www.computerweekly.com/feature/IPv6-The-security-risks-to-business.

[36] IPv6 Essentials, 3rd Edition by Silvia Hagen Published by O'Reilly Media, Inc., 2014. Intenret: https://www.safaribooksonline.com/library/view/ipv6-essentials-3rd/9781449335229/ch01.html.

[37] Frank Herberg IPv6 insecurities on “IPv4-only” networks. Internet: http://securityblog.switch.ch/2014/08/26/ipv6-insecurities-on-ipv4-only-networks/.

[38] Network Time Protocol (NTP): Threats and Countermeasures. Internet: https://isc.sans.edu/forums/diary/NTP+reflection+attack/17300/.

[39] Internet: http://resources.infosecinstitute.com/network-time-protocol-ntp-threats-countermeasures/.

[40] Internet: http://tools.ietf.org/html/bcp38.

[41] Internet: http://en.wikipedia.org/wiki/Network_Time_Protocol#cite_note-29.

[42] Internet: http://www.eecis.udel.edu/~mills/security.html.

[43] R7-2014-12: More Amplification Vulnerabilities in NTP Allow Even More DRDoS Attacks. Internet: https://community.rapid7.com/community/metasploit/blog/2014/08/25/r7-2014-12-more-amplification-vulnerabilities-in-ntp-allow-even-more-drdos-attacks.

[44] Extromatica Network Monitor. Internet: http://en.wikipedia.org/wiki/Extromatica_Network_Monitor.

[45] Logging and Monitoring Tools. Internet: https://quequero.org/downloads/logging-and-monitoring-tools/.

[46] Internet: http://en.wikipedia.org/wiki/Microsoft_Product_Activation.

[47] Internet: http://www.sans.org/security-resources/idfaq/honeypot3.php.

[48] DNS Hacking (Beginner to Advanced). Internet: http://resources.infosecinstitute.com/dns-hacking/.

[49] Internet: http://searchsecurity.techtarget.com/tip/Routing-protocol-security.

[50] Ta Vinh Thong, Attacks against secure routing protocols. Internet: http://crysys.hu/members/tvthong/links/adhocAttacks.pdf.

[51] Hunting Session Fixation Bugs. Internet: http://resources.infosecinstitute.com/hunting-session-fixation-bug/

[52] Use offense to inform defense. Find flaws before the bad guys do. Internet: http://pen-testing.sans.org/resources/papers/gcih/port-1433-vulnerability-unchecked-buffer-password-encryption-procedure-104360.

[53] Microsoft. “Security Tools and Checklists.” Internet: http://www.microsoft.com/technet/security/tools/tools.asp.

[54] Microsoft. “SQL2000 C2 Admin and User Guide”, November 2, 2002. Internet: http://www.microsoft.com/Downloads/details.aspx?displaylang=en&FamilyID=71C146F3-9907-40CDBABF-3506ECD33254.

[55] Rakhmanoff, Martin. jimmers@yandex.ru. June 14, 2002. Internet: http://online.securityfocus.com/archive/1/276953.

[56] CERT: VU#225555. July 29, 2002. Internet: http://online.securityfocus.com/advisories/4308.

[57] Rakhmanoff, Martin. jimmers@yandex.ru. SecuriTeam. 10/22/2002. Internet: http://www.securiteam.com/windowsntfocus/6O00L0K5PC.html.

[58] Microsoft, “SQL Server Documentation Chapter 11”. Internet: http://www.microsoft.com/technet/prodtechnol/sql/proddocs/diag/part3/75528c11.asp?.

[59] Anley, Chris. “Advanced SQL Server Injection in SQL Server Applications” Internet: http://www.nextgenss.com/papers/advanced_sql_injection.pdf.

[60] Litchfield, David. “Threat Profiling SQL Server”, July 20, 2002. Internet: http://www.nextgenss.com/papers/tp-SQL2000.pdf

[61] Nolan, Patrick. Incidents.org “Slapper Worm Update.” Jan 25, 2003. Internet: http://isc.incidents.org/analysis.html?id=180.