

ICMP Tunneling: Defense Against the Vulnerability

Abhishek Singh, Chenguhai Lu, Ola Nordstrom, Andre L M dos Santos
Georgia Tech. Information Security Center (GTISC)
Center for Experimental Research in Computer Systems(CERCS)
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
{abhi,nalo,lulu,andre}@cc.gatech.edu

ABSTRACT

This paper presents a systematic solution to the ICMP tunneling problem. ICMP is not multiplexed via port numbers and the data part of the ICMP packet provides considerable bandwidth for covert channels. These factors make it an integral part of various malicious software such as remote access and denial of service attack tools. These tools use ICMP to establish covert communication channels. In this paper a stateless model is proposed to prevent ICMP tunneling. A Linux kernel module was implemented to demonstrate the proposed stateless solution. The module enforces a fixed payload policy for ICMP packets and virtually eliminates ICMP tunneling which arises due to the data carrying capability of ICMP. The performance impact on end hosts and routers due to our stateless monitoring model show that it can eliminate ICMP tunneling with modest performance penalties.

Keywords

ICMP, Remote Access Tools, Netfilter

1 INTRODUCTION

The Internet and the World Wide Web (WWW) have had a phenomenal growth during the past few years, interconnecting average users and average user to expert users who are not always good neighbors. Many Internet attacks, both simulated and real, have been described by the security community and have appeared in mainstream media. Two factors that have contributed to widespread Internet attacks are the lack of security as an initial design consideration for the Internet and the average user's inadequate knowledge of threats faced every time their computer is connected to the Internet. Although the attacks vary on their form, most of them have a common goal of leaving a back door open for a future communication with a victim machine.

Internet communication is based, in addition to the Internet Protocol (IP), on three basic protocols: Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP). Firewalls, depending on the services required by their internal networks, totally block or partially filter Internet packets using one or more of these protocols. An attacker has to decide which protocol or protocols they will use to communicate with a backdoor installed on a host that has been compromised. The attacker's objective is to make the traffic generated by the backdoor appear as much as possible to be normal traffic so it is not blocked by the firewall. TCP and UDP, the most widely used protocol by application servers, are the one that can be abused with the greatest chance of allowing the traffic to pass through. TCP and UDP packets can carry information either manipulating unused parts of the packet or making the payload look legitimate. For example,[10] various header fields like ACK flags and port number can be used to establish covert communication using TCP or UDP. However, firewalls like IP Filter can be used to prevent covert channels that make use of TCP/UDP header fields for its communication. IP Filter [1] uses stateful packet filtering. The state engine not only inspects the presence of ACK flags in TCP packets but also includes sequence numbers and window sizes in its decision to block or to allow packets. IP Filter does not check the content of ICMP packets and hence fails to prevent covert channels that can arise due to misuse of the payload of ICMP packets. Therefore, although TCP and UDP continue to be a subject for studies in vulnerabilities, ICMP also provides several means for stealth traffic.

ICMP tunneling was first reported in the 1997 [2][3]. Initial version of ICMP tunneling enabled an attacker to execute remote commands and steal information from a compromised machine. Although ICMP tunneling has been used for user-user and user – machine communication, its most damaging usage has been for coordination of distributed denial of service attacks. In early February a distributed denial of service attack was launched against Yahoo, Amazon, eBay and

other popular Internet sites. The attacks on Yahoo, eBay, Amazon.com and E*Trade resulted in a loss of approximately \$1.2 billion, according to The Yankee Group, a Boston consulting firm. It is reported in [4] that almost all of the tools used on the distributed denial of service (DDOS) attacks on Yahoo, Amazon, eBay, and E*Trade internet sites, have used ICMP for covert communications between the DDOS clients and the attacker's handler program. Some of the most widely known distributed denial of service attack tools like Tribe Flood Net2K[5] and Stacheldraht rely on ICMP tunneling to establish communication channels between the compromised machines and the hacker's machine. Since ICMP tunneling is very simple to deploy and can cause a significant amount of damage it has been classified as a high risk security threat by Internet Security Services[6] and SANS [7].

The rest of the paper is organized as follows. Section 2 presents the solutions which are currently being used to prevent ICMP tunneling. Section 3 presents the experimental result which demonstrates the seriousness of ICMP tunneling. Section 4 discusses the proposed solution and its performance impact on routers and on end hosts. Finally in section 5 conclusion and future work is presented.

2 ICMP Tunneling Vulnerability

Most types of ICMP packets like *echo_request* / *echo_reply* (commonly used for ping) have the capability of carrying data in its payload. This data carrying capability of ICMP can be used to establish covert channels. Various malicious applications like "Loki" use the data carrying capability of ICMP to establish covert channels. The use of ICMP for covert communication presents one big advantage over the use of TCP or UDP. ICMP packets use fewer parameters than TCP or UDP. For example, ICMP does not use port numbers. Port number gives an additional parameter for firewalls to filter suspicious traffic. The first 64 bits of the original IP datagram's data are used by the host to match ICMP error messages to the appropriate application level process. The simplicity and lack of parameters used in ICMP have made it popular for hackers designing tools that require covert communication.

Many solutions have been proposed to prevent ICMP tunneling. Some of the proposed solutions currently being used to prevent ICMP tunneling are:

- Disable all ICMP traffic. Disabling all ICMP traffic prevents covert communication using ICMP packets. However, ICMP messages are required to check the status of a network and communicate IP packet errors. Therefore,

disabling all ICMP messages prevents covert communication, but also prevents users on one side of the network protected by the firewall to check the status of machines on the other side or receive valid ICMP packets related to transient network problems. This limitation is not acceptable in some environments and is not a general solution.

- Disable part of the ICMP traffic allowed by a firewall. For example, disable incoming *echo_request*, while allowing outgoing *echo_request*. If naively implemented, policies like this will still allow covert communication, limiting only which host needs to start a communication. In addition, outgoing ICMP packets could be used to establish a unidirectional channel to send compromised information. A modified server can periodically send ICMP packets containing sensitive information. For example, an e-commerce web server could switch to ICMP mode of communication for 10 second after every 30 minutes. During these 10 seconds it could send to a receiving machine, credit card information that has been captured during the previous 30 minutes.
- Limit the size of ICMP packets. Large ICMP packet can be seen as suspicious by an IDS system that could inspect the ICMP packet and raise an alarm. However, since there are legitimate uses for large ICMP packets it is difficult to determine if a large ICMP packet is malicious. For example, large *echo_request* packets are used to check if a network is able to carry large packets. Differentiating legal from illegal large packets is even more difficult if covert communication is encrypted. An IDS needs to be able to determine if a packet is encrypted or not. Distinguishing encrypted from non-encrypted packet still remains an open interesting research problem.
- Preserve the state of ICMP packet to check for covert channel. Some firewalls like Raptor use state preservation technique to prevent ICMP tunneling. A daemon called pingd runs as an application process on the firewall. If the firewall is the target of ping then pingd responds to the client normally. If the firewall is not the target of ping then pingd will construct a new echo request with a new sequence number, new time to live, and a new payload (with new checksum). When the reply is received it is ensured that the data is the same as what had been sent, the sequence number and responders IP address are checked. After a successful check the firewall transmits

the *echo_reply* to the original client. Although state preserving technique can easily prevent ICMP tunneling, it is an expensive process and hence application level personal firewalls do not use state preserving techniques. Application level personal firewalls[8] running on personal computers only monitor inbound and outbound Internet traffic and alerts the user when an application is attempting access their pc or their machine is trying to access something on the Internet. As per the configuration rule, personal firewalls can either block ICMP or allow it irrespective of the fact that the ICMP is being used for tunneling. Protecting personal machines from ICMP tunneling is very important since there are a large number of personal machines connected to high speed links which in turn can be used to launch DDOS attacks.

Although some of the solutions presented above can be acceptable, ideally the prevention of ICMP misuse due to their data carrying capabilities should be able to provide the following.

- It should enable users to use ICMP messages for administrative purposes freely.
- It should allow large size of ICMP so users can find out if the network can carry large size of data packets.
- It should be able to handle encrypted packets.
- It should be able to prevent personal machines that are not behind powerful state preserving firewalls from being used as DDOS slaves.

3 ICMP Tunneling : Case Study

Before presenting the solution that can eliminate ICMP tunneling, it is important to investigate the seriousness of the ICMP tunneling. A remote access tool was used as the test bed for our experimentation. Remote Access Tools allows a user to access data and gain control over a remote computer. BO2K[14] which is one of the remote access tools was used to test ICMP tunneling. The communication infrastructure of BO2K was moved to ICMP. An ICMP *echo_request* contained the remote command issued by a client of BO2K and an *echo_reply* contained the information from the machine which was running BO2K server. Strong authentication and encryption was also implemented to demonstrate the seriousness of the exploit. The Bellovin and Merritt[9] key exchange known for its strong authentication and establishment of session keys was implemented in the BO2K communication protocol. The session key generated by the Bellovin and Merritt protocol was then

used to encrypt the data using 3DES. The server and the client of the modified BO2K was then installed a pair of personal computers. The machines were connected via the Internet and were using the personal firewall as their application level firewall.

3.1 Results

Since ICMP uses raw sockets for its communication, root (administrator) privileges are required. For the experimentation purposes it was assumed that the root privileges can be obtained. More details about getting the root permission can be found in [13]. However in most of the commonly used operating systems [11] like Windows ME/98 getting root permission is not required. While performing this attack it was also assumed that the firewall is customized to allow ICMP packets.

Personal firewall used in the machines was set up such that it raise an alarm in the form of a pop up window each and every time it sees an incoming or an outgoing data packet. Various options provided by the personal firewall can be seen from figure 1.

- Incoming Connection Alert / Outgoing Connection Alert – indicates whether the requested connection is outgoing (from the local network) or incoming (eg. from the Internet)
- Time – exact time and date, when the connection was requested
- Remote – information about the remote end-node (IP address, port and communication protocol)
- Details – detailed information about the connection
- Details about application – information about the local application taking part in the communication (as a client or server). For example in figure 4 internet explorer is trying to access the internet.
- Permit – let the packet pass through
- Deny – stop (filter) the caught packet

For an incoming connection, the pop up window displays the IP address of the machine initiating the connection and the protocol, which is being used for the connection. For an outgoing connection, the pop up window shows the IP address of the destination machine along with the protocol used for the communication. The server of BO2K was installed on a compromised machine and the client of BO2K was installed on a user's machine. Both the server and the client machines were running the Windows ME operating system.

Figure 2.1 and 3.1 are the snapshots of the connection initiation steps. Figure 2.2 shows the firewall of the client when it tries to initiate the connection. "Protocol number:0" in figure 2.2 and 3.2 identifies the IP packet as an ICMP packet. The IP address of the compromised machine was passed as a parameter to the client program. Figure 3.1 shows the incoming connection as seen by the compromised machine. It can be inferred from figure 2.2 and figure 3.2 that the firewalls see the IP packet sent by the kernel. The ICMP mode of communication provides false information to the firewall. Even though ICMP packets are being sent by the application, the firewall infers them to be sent by the operating system. After exchanging the session keys, the user at the remote machine issues a command



Figure 1. Options provided by the Personal Firewall.

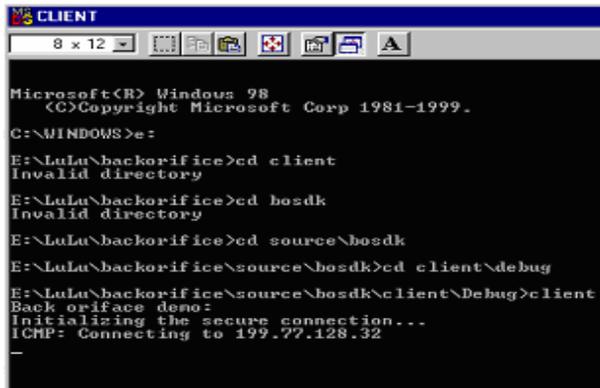


Figure 2.1 Client Initiating Connection



Figure 2.2 Firewall of the client while initiating Connection.



Figure 3.1 Server receiving request



Figure 3.2 Firewall of the Server Receiving Request

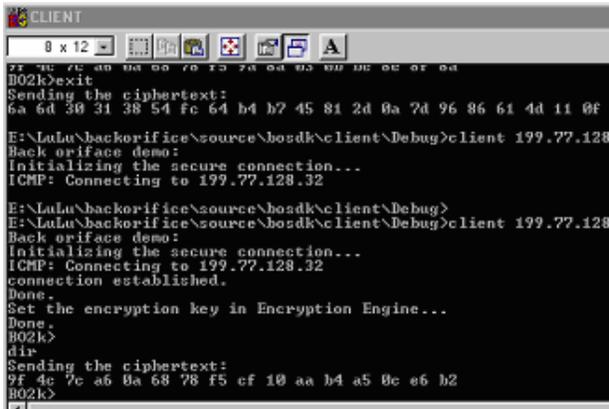


Figure 4.1 Client requesting for directory Listing



Figure 4.2 Firewall of the client when it requests for directory listing

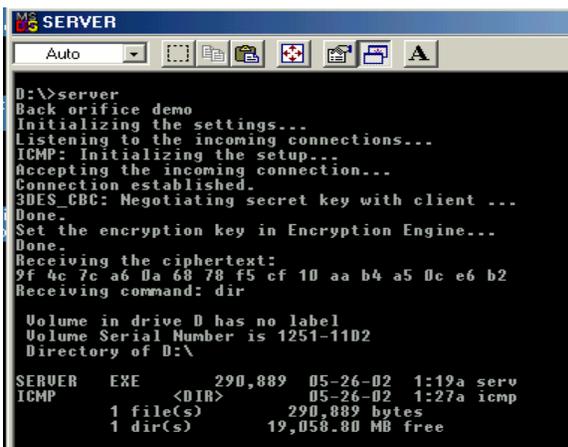


Figure 5.0 Server Executing dir command

to the compromised machine. Figure 4.1 shows the command “dir” to view the file listing at the server issued at the client. This command is encrypted as “9f 4c 7c a6 0a 68 78 f5 cf 10 aa b4 a5 0c e6 b2” before being sent to the compromised machine. Again it can be seen in figure 4.2 that the outgoing command is sent by using ICMP. In figure 5 it can be seen that the compromised machine receives the encrypted command over ICMP channel, decrypts it and executes it. Thus from the screen shots of the case study, it can be inferred that the ICMP mode of communication provides the same data carrying capability as TCP/UDP mode of communication. Even though ICMP packets are being used by the application, the personal firewall infers these ICMP packets as the control messages issued by the operating system.

4.0 Proposed Solution

It becomes difficult to stop ICMP tunneling since the personal firewall infers them to be control messages issued by the operating system. Strong encryption and authentication in addition to ICMP tunneling only makes the matter worse. Thus the solution to prevent ICMP tunneling should

- Enable the administrators to use ICMP messages freely.
- Enable large size of ICMP to pass through.
- Should work even in the case of strong authentication and encrypted payload.
- Should work even when the machine is behind application level personal firewall.

As compared to the expensive state preserving model used by industrial firewalls, a simple state less model is proposed to prevent tunneling. The stateless model caters to the above mentioned requirements. The state less model requires a common agreement upon the allowable payload of ICMP messages. The policy should be implemented in the ICMP protocol implementation of the kernel. It should be enforced either when an ICMP packet is going up the network stack or when it is going down the stack.

Because of the freely available source code for Linux, it was chosen for the implementation of the proposed module. The algorithm used by the module is shown in figure 9. This module first scans the content of the payload of the ICMP packet against a predefined set. If any malicious content is found in the payload an alarm is raised. The module then zeros out the entire data carrying field irrespective of the ICMP type. The unused portion of ICMP packets varies with the ICMP type. The Netfilter framework was used to extend existing functionality within the kernel. Netfilter[15] is a set of hooks inside the Linux 2.4.x kernel's network stack

which allows kernel modules to register callback functions called when a network packet traverses one of the predefined hooks. The handlers will execute as if they were part of the packet processing pipeline directly. Netfilter allows a module to register the *IP_POST_ROUTING* hook which is called after a packet has been through the routing table and right before it is delivered to the outgoing interface (typically an Ethernet device). The ICMP monitor (*icmp_mon*) is registered with the *IP_POST_ROUTING* handler, and thus allowing *icmp_mon* to process all locally outbound packets as well as forward packets since all packets go through the post routing hooks. The *icmp_mon* module at runtime calls *icmp_mon_erase* and *icmp_mon_scan* to perform a combination of actions.

- *icmp_mon_erase*; Zeros out unused portions of ICMP messages.
- *icmp_mon_scan*; Scans packets for predefined strings

icmp_mon_scan raises an alarm when it detects some suspicious content in the data portion of the packet. The packet is then forwarded to the *icmp_mon_erase*. Irrespective of the fact that it is carrying some malicious strings the data field of ICMP is filled with zeros. Thus the proposed solution can stop the encrypted traffic. Since the packet is analyzed when it is outbound our solution will work in the case of packets which are locally generated or are forwarded on behalf of another machine. Hence this patch can work in machines which are acting as a gateway, router, sensor or as a host.

```

IP_MON (ICMP packet)
Begin
    input : ICMP Packet
    for every ICMP packet
        Begin
            icmp_mon_scan (data portion of ICMP
            packet)
        Begin
            If the data field matches with the signature
            Raise an alarm.
        End
        icmp_mon_erase{data portion of the ICMP
        packet}
        Begin
            Fill in the data portion with zeros.
        End
    End
End
End

```

Figure 9.0 Algorithm for the Stateless Model

The module can be compiled directly into a monolithic kernel. The filter hooks are registered when the modules init routines are called at the boot time if it was compiled into the kernel. The proposed module can not be bypassed by user space applications like since it does not interact with any. Iptables is the userspace command which is used to modify Linux firewalling rules. Iptables is typically used to add and remove filtering rules pertaining to regular TCP/IP packets. However *icmp_mon* is not part of iptables, the only way to disable the scanning is to remove the module (using *rmmod*) which is not possible in this scenario since the module is statically compiled into a monolithic kernel. When using a monolithic kernel the machine would have to be rebooted into a different kernel to disable the *icmp_mon* from operating. This is not trivial to do on a compromised host machine. However if the proposed patch is running on a Gateway or at a sensor which is monitoring network traffic then it becomes difficult for an adversary to bypass. The proposed solution will make it difficult for an adversary to setup a covert communication channel on a compromised machine since packets are scanned and erased by the kernel. The stateless model implemented in the kernel benefits from the fact that there is no way to turn off this functionality by simply terminating or modifying an application, which could easily be done on a compromised machine if packets were queued to a user-space intrusion detection application.

The proposed stateless model of scanning and erasing the data fields should be implemented in the kernel used by hosts, sensors, gateways and routers so as to completely eliminate ICMP tunneling.

4.1 Results of the Proposed Solution.

The *icmp_mon* scanning and erasing times were tested on *ICMP ECHO* request packets. The routing machine was a Dual Pentium III 450 MHz machine with 512 MB of RAM. Internally all network packets are time stamped upon arrival in the Linux kernel. The time from arrival up to when the *icmp_mon* routines are referred as the “Kernel Time”, or specifically the time the packet has been traversing inside the kernels packet processing. The total time spent inside the *icmp_mon* routine, ignoring the overhead caused by the netfilter hooks will be referred to as the “erase/scan time” depending on the operation performed.

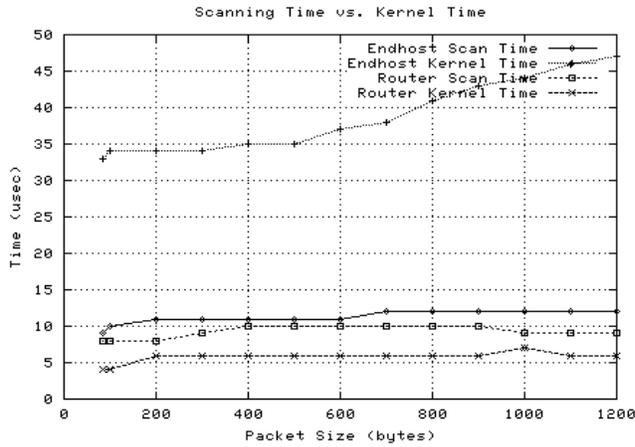


Figure 7.0 Scanning Vs. Kernel Time for Different Size Packets

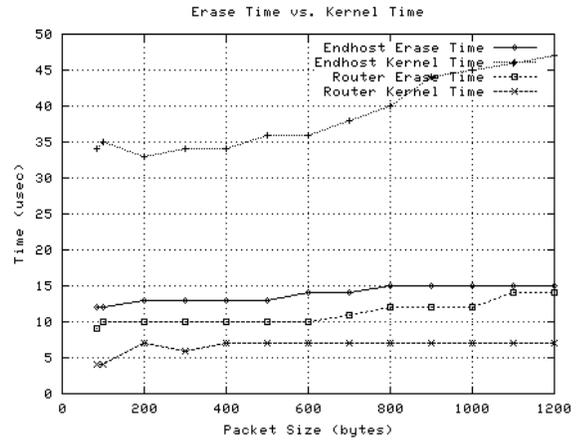


Figure 9.0 Erase time Vs. Kernel Time for Variable Size packets

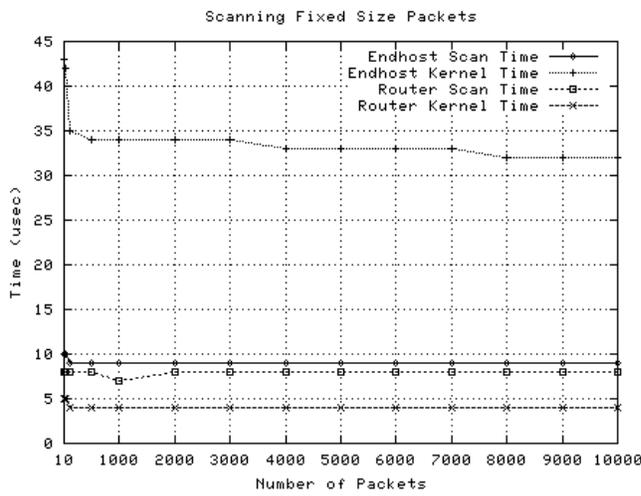


Figure 8.0 Scanning Vs. Kernel Time for Fixed Size packets

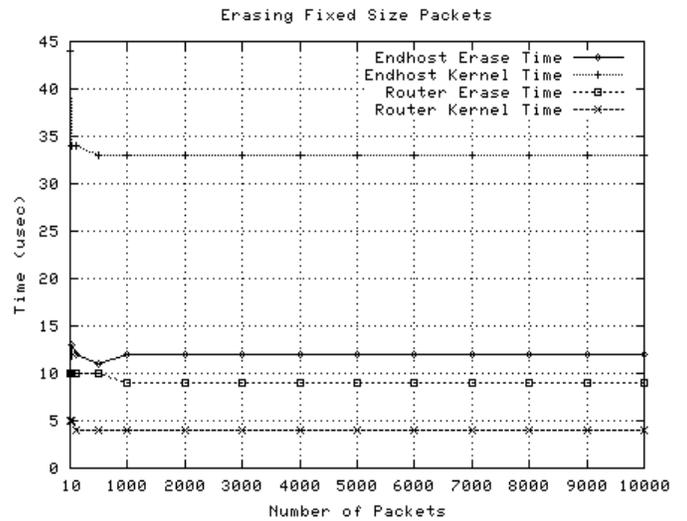


Figure 10 Erasing Time Vs. Kernel Time for Fixed Size packets

The first test was implemented in the kernel of a host machine. In the second test the proposed module was implemented in the kernel of a machine acting as a router. For each of these tests the time to process packets of fixed size in the kernel was compared with the scan and erase time. The tests on host machine and router were again repeated with the packets of different lengths. All the tests were performed using the standard Unix “ping” command; it was typically run as follows “ping -c Count -s Size -f host”, where the Count and Size were either fixed or varied depending on the test. The “-f” option was used to flood ping, the *icmp_mon* module would sum the processing times for a given number of packets (corresponding to the Count given to the ping command) then the mean time was calculated after the given number of packets had passed through. The `gettimeofday()` call was used to get the time, on the Intel platform the granularity available is microseconds.

Graph shown in figure 7.0 presents the performance cost associated with the string scanning routines of the *icmp_mon*. Size of the packet was increased from 100 bytes to 1200 bytes and time to process 1000 packets of each size was measured and plotted in the graph shown in figure 7.0. This test was performed both for the router and for the end host. The simple words searched for were “passwd”, “root”, “tmp”, “etc”, “ls”, and “dir”. They were never present in the ping packets thus the scans never raised any alarms. As per the graph for the router the scan time is more than the time to process packet. The machine is simply receiving the packets in one interface and sending them out another, which makes the kernel time incredibly small. In the case of the end host the time to process a packet increases with the packet size, however the processing time is small as compared to the kernel time. The scan time shows a marginal increase, however it remains constant when the size of packet is increases from 700 bytes to 1200 bytes.

Graph shown in figure 10 shows the time taken for scanning the ICMP payloads of fixed size of 84 bytes. For this test the number of packets was increased from 10 to 10000 and the scan time was plotted. A 2 microsecond increase was observed when the number of packets increased from 1000 to 10000. As the number of back to back packets increased (due to the ping flood) the kernel time per packet decrease. The kernel time goes down because the flood of packets allows the kernel routines to grab multiple packets back to back without returning to a different routine (thus losing time due to context switching). When the machine is acting as a router scan time is again more than the time to process a packet by the kernel. The scan time shows a marginal increase as the number of packet increases.

The result show in figure 9 gives the average time spent inside the kernel and the *icmp_mon* erase routines when the patch is running on a router and on an end host. This test was again performed by keeping the number of ICMP packet to 1000 packets for each size and increasing the size of packets from 84 bytes to 1200 bytes. As the packet size increase time to process packet inside the kernel for the end host increases. This increase is almost linear. Size has no effect on the processing time for the end routers. The time is contrasted against and increasing packet size, starting

5.0 Conclusion and Future Work.

ICMP tunneling is becoming a vital part of the malicious software. To prevent ICMP tunneling kernel modifications are necessary to enforce a fixed predefined payload policy for ICMP packets. If the proposed solution becomes an integral part of the kernels which runs as host, gateways, and DMZ routers then it will be impossible to establish ICMP tunnels. Another way to remove ICMP tunneling is to simply truncate the data field of ICMP. However truncation of the data field will require amendments in the RFC [11][12] which supports data field for ICMP. Scanning and erasing of the ICMP data field is compliant with RFC and prevents ICMP tunneling irrespective of the type of firewall used.

The results show that simply marking out unused and potentially dangerous portions of ICMP packets is a straightforward task and requires little overhead on a modest system. Simple string scans are also not costly and can be done to test for unencrypted covert communication. This is highly recommended for the end hosts where it offers minimal overhead on the system. For the routers it is expensive, however weighted against the potential security risks the marginal overhead is worth the security benefits. Moreover if the proposed solution becomes an integral part of the kernels in operating systems like Solaris, Windows, Linux etc. which runs at host machines then the routers will not have to examine all the ICMP packets which it comes across. It can adopt some sort of probabilistic scheme to check some ICMP packets and allow the other packets to pass through. This will reduce the load on the routers. As a part of the future work various schemes will be explored such that even for the routers, scanning and erasing becomes a cheap operation. Implementing state preserving technique used by raptor inside the kernel and studying its performance will also be the focus of future research.

REFERENCE

- [1] Guido Van Rooij, “ Real Stateful TCP Packet Filtering in IP Filtering”.
- [2] ICMP attacks Illustrated, http://www.sans.org/rr/threats/ICMP_attacks.php
- [3] CERT Advisory Denial of Service attack tools , <http://www.cert.org/advisories/CA-1999-17.html>
- [4] Loki ICMP tunneling back door , http://www.iss.net/security_center/static/1452.php
- [5] Intrusion Detection FAQs , http://www.sans.org/resources/idfaq/icmp_misuse.php
- [6] Personal Firewalls : What are they, how do they work? http://www.sans.org/rr/homeoffice/personal_fw.php
- [7] Steven M. Bellovin and Michael Merrit. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In proceeding of the IEEE Symposium on Research in Security and Privacy, May 1992.
- [8] Craig H. Rowland ,Covert Channels in the TCP/IP Protocol Suite, http://www.firstmonday.dk/issues/issue2_5/rowland
- [9] Postel, J.(ed) , “ Internet Protocol – DARPA Internet Program Protocol Specification,” RFC 791, USC/ Information Sciences Institute, September 1981
- [10] Postel J. (ed), “Internet Control Message Protocol – DARPA Internet Program Protocol Specification,”, RFC 792, September 1981.
- [11]Root Exploit and Dos in the Linux Kernel, <http://linux.oreillynet.com/pub/a/linux/2001/10/22/insecurities.html>
- [12] Backorifice SDK Documents, <http://bo2k.sourceforge.net/indexnews.html>NIST, US
- [13]Netfilter firewalling, NAT and packet mangling for Linux 2.4, <http://www.netfilter.org/>