TCP Maintenance and Minor
Extensions (tcpm)
Internet-Draft
Expires: June 8, 2005

F. Gont
UTN/FRH

December 8, 2004

ICMP attacks against TCP
draft-gont-tcpm-icmp-attacks-02.txt

Abstract

This document discusses the use of the Internet Control Message Protocol (ICMP) to perform a variety of attacks against the Transmission Control Protocol (TCP) and other similar protocols. It proposes several counter-measures to eliminate or minimize the impact of these attacks.

Table of Contents

1. Introduction

Recently, awareness has been raised about several threats against the TCP [1] protocol, which include blind connection-reset attacks [12]. These attacks are based on sending forged TCP segments to any of the TCP endpoints, requiring the attacker to be able to guess the four-tuple that identifies the connection to be attacked.

While these attacks were known by the research community, they were considered to be unfeasible. However, increases in bandwidth availability, and the use of larger TCP windows [13] have made these attacks feasible. Several general solutions have been proposed to either eliminate or minimize the impact of these attacks [14][15][16]. For protecting BGP sessions, specifically, a counter-measure had already been documented in [17], which defines a new TCP option that allows a sending TCP to include a MD5 [18] signature in each transmitted segment.

All these counter-measures address attacks that require an attacker to send spoofed TCP segments to the attacked host. However, there is still a possibility for performing a number of attacks against the TCP protocol, by means of ICMP [2]. These attacks include, among others, blind connection-reset attacks.

This document aims to raise awareness of the use of ICMP to perform a number of attacks against TCP, and proposes several counter-measures that can eliminate or minimize the impact of these attacks.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [3].

2. Background

2.1 The Internet Control Message Protocol (ICMP)

The Internet Control Message Protocol (ICMP) is used in the Internet Architecture to perform the fault-isolation function, that is, the group of actions that hosts and routers take to determine that there is some network failure [19].

When an intermediate router detects a network problem while trying to forward an IP packet, it will usually send an ICMP error message to the source host, to raise awareness of the network problem. In the same way, there are a number of cases in which an end-system may generate an ICMP error message when it finds a problem while processing a datagram. These error messages are notified to the corresponding transport-protocol instance.

When the transport protocol is notified of the error condition, it will perform a fault recovery function. That is, it will try to survive the network failure.

In the case of TCP, the typical fault recovery policy is as follows:

- o If the network problem being reported is a hard error, abort the corresponding connection.

- o If the network problem being reported is a soft error, just record this information, and repeatedly retransmit the segment until either it gets acknowledged, or the connection times out.

Some stacks honor hard errors only for connections in any of the synchronized states (ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK or TIME-WAIT).

2.1.1 ICMP for IP version 4 (ICMP)

[2] specifies the Internet Control Message Protocol (ICMP) to be used with the Internet Protocol version 4 (IPv4). It defines, among other things, a number of error messages that can be used by end-systems and intermediate systems to report network errors to the sending host.

The Host Requirements RFC [4] states that ICMP error messages of type 3 (Destination Unreachable) codes 2 (protocol unreachable), 3 (port unreachable), and 4 (fragmentation needed and DF bit set) should be considered hard errors. Thus, any of these ICMP messages could elicit a connection abort.

The ICMP specification also defines the ICMP Source Quench message (type 4, code 0), which is meant to provide a mechanism for flow control and congestion control. The Requirements for IP Version 4 Routers RFC [5], however, states that experience has shown this ICMP message is ineffective for handling these issues.

[6] defines a mechanism called "Path MTU Discovery" (PMTUD), which makes use of ICMP error messages of type 3 (Destination Unreachable), code 4 (fragmentation needed and DF bit set) to allow hosts to determine the MTU of an arbitrary internet path. For obvious reasons, those systems implementing the PMTUD do not treat ICMP error messages of type 3 code 4 as hard errors.

Appendix D of [7] provides information about which ICMP error messages are produced by hosts, intermediate routers, or both.

2.1.2 ICMP for IP version 6 (ICMPv6)

[8] specifies the Internet Control Message Protocol (ICMPv6) to be used with the Internet Protocol version 6 (IPv6) [9].

Even though ICMPv6 didn't exist when [4] was written, one could extrapolate the concept of "hard errors" to ICMPv6 Type 1 (Destination Unreachable) codes 1 (communication with destination administratively prohibited) and 4 (port unreachable). Thus, any of these messages could elicit a connection abort.

ICMPv6 defines the "Packet Too Big" (type 2, code 0) error message, that is analogous to the ICMP "fragmentation needed and DF bit set" (type 3, code 4) error message. For IPv6, intermediate systems do not fragment IP packets. Thus, there's an implicit "don't fragment" bit set in every IPv6 datagram sent on a network. Therefore, hosts do not treat ICMPv6 "Packet Too Big" messages as a hard errors, but use them to discover the MTU of the corresponding internet path, as part of the Path MTU Discovery mechanism for IP Version 6 [10].

Appendix D of [7] provides information about which ICMPv6 error messages are produced by hosts, intermediate routers, or both.

2.2 Handling of ICMP errors

The Host Requirements RFC [4] states that a TCP instance should be notified of ICMP error messages received for its corresponding connection.

In order to allow ICMP messages to be demultiplexed by the receiving host, part of the orignal packet that elicited the message is included in the payload of the ICMP error message. Thus, the receiving host can use that information to match the ICMP error to the instance of the transport protocol that elicited it.

Neither the Host Requirements RFC nor the original TCP specification [1] recommend any security checks on the received ICMP messages. Thus, as long as the ICMP payload contains the correct four-tuple that identifies the communication instance, it will be processed by the corresponding transport-protocol instance, and the corresponding action will be performed.

Therefore, an attacker could send a spoofed ICMP message to the attacked host, and, as long as he is able to guess the four-tuple that identifies the communication instance to be attacked, he can use ICMP to perform a variety of attacks.

As discussed in [12], there are a number of scenarios in which an

attacker may be able to know or guess this four-tuple. Furthermore, it must be noted that most Internet services use the so-called "well-known" ports, so that only the client port would need to be guessed. In the event that an attacker had no knowledge about the range of port numbers used by clients, this would mean that an attacker would need to send, at most, 65536 packets to perform any of the attacks described in this document.

It is clear that security checks should be performed on the received ICMP error messages, to mitigate the impact of the attacks described in this document.

3. ICMP attacks against TCP

ICMP messages can be used to perform a variety of attacks. These attacks have been discussed by the research community to a large extent.

Some TCP/IP implementations have added security checks on the received ICMP error messages to minimize the impact of these attacks. However, as there has not been any official proposal about what would be the best way to deal with these attacks, these security checks have not been widely implemented.

Section 4 of this document discusses the constraints in the general counter-measures that can be implemented against the attacks described in this document. Section 5 proposes several general conter-measures that apply to all the ICMP attacks described in this document. Finally, Section 6 and Section 7 discuss a variety of ICMP attacks that can be performed against TCP, and propose attack-specific counter-measures that eliminate or mitigate them. These attack-specific counter-measures are meant to be additional counter-measures to the ones proposed in Section 5. In particular, all TCP implementations SHOULD perform the TCP sequence number checking described in Section 5.1.

4. Constraints in the possible solutions

For ICMPv4, [2] states that the internet header plus the first 64 bits of the packet that elicited the ICMP message are to be included in the payload of the ICMP error message. Thus, it is assumed that all data needed to identify a transport protocol instance and process the ICMP error message is contained in the first 64 bits of the transport protocol header. [4] states that "the Internet header and at least the first 8 data octets of the datagram that triggered the error" are to be included in the payload of ICMP error messages, and that "more than 8 octets MAY be sent", thus requiring implementations to include more data from the original packet than that required by

the original ICMP specification. The "Requirements for IP Version 4 Routers RFC" [5] states that ICMP error messages "SHOULD contain as much of the original datagram as possible without the length of the ICMP datagram exceeding 576 bytes".

Thus, for ICMP messages generated by hosts, we can only expect to get the entire IP header of the original packet, plus the first 64 bits of its payload. For TCP, that means that the only fields that will be included are: the source port number, the destination port number, and the 32-bit TCP sequence number. This clearly imposes a constraint on the possible checks we can perform, as there is not much information avalable on which to perform these security checks. While there exists a proposal to recommend hosts to include more data from the original datagram in the payload of ICMP error messages [20], and some TCP/IP implementations already do this, we cannot yet propose any work-around based on checks performed on any data past the first 64 bits of the payload of the original IP datagram that elicited the ICMP error message. Thus, the only check that can be performed on the ICMP error message is that of the TCP sequence number contained in the payload.

As discussed above, for those ICMP error messages generated by routers, we can expect to receive much more octets from the original packet than just the entire IP header and the first 64 bits of the transport protocol header. Therefore, not only can hosts check the TCP sequence number contained in the payload of the ICMP error message, but they could perform further checks such as checking the TCP acknowledgement number, as discussed in Section 5.2.

For ICMPv6, the payload of ICMPv6 error messages includes as many octets of the IPv6 packet that elicited the ICMPv6 error message as will fit without making the resulting ICMPv6 packet exceed the minimum IPv6 MTU (1280 octets) [8]. Thus, further checks (as those described above) can be performed on the received ICMP error messages.

5. General counter-measures against ICMP attacks

There are a number of counter-measures that can be implemented to eliminate or mitigate the attacks discussed in this document. Rather than being alternative counter-measures, they can be implemented together to increase the protection against these attacks.

5.1 TCP sequence number checking

TCP SHOULD check that the TCP sequence number contained in the payload of the ICMP error message is within the range SND.UNA =< SEG.SEQ < SND.NXT. This means that the sequence number should be

within the range of the data already sent but not yet acknowledged. If an ICMP error message doesn't pass this check, it SHOULD be discarded.

Even if an attacker were able to guess the four-tuple that identifies the TCP connection, this additional check would reduce the possibility of considering a spoofed ICMP packet as valid to Flight_Size/2^^32 (where Flight_Size is the number of data bytes already sent to the remote peer, but not yet acknowledged [21]). For connections in the SYN-SENT or SYN-RECEIVED states, this would reduce the possibility of considering a spoofed ICMP packet as valid to 1/2^^32. For a TCP endpoint with no data "in flight", this would completely eliminate the possibility of success of these attacks.

5.2 TCP Ackowledgement number checking

As discussed in Section 4, for those ICMP error messages that are generated by intermediate routers, additional checks can be performed. TCP SHOULD check that the TCP Acknowledgement number contained in the payload of the ICMP error message is withing the range SEG.ACK <= RCV.NXT. This means that the TCP Acknowledgement number should correspond to data that have already been acknowledged.

This would reduce the possibility of considering a spoofed ICMP packet as valid by a factor of two.

5.3 Port randomization

As discussed in the previous sections, in order to perform any of the attacks described in this document, an attacker needs to guess (or know) the four-tuple that identifies the connection to be attacked. Randomizing the ephemeral ports used by the clients would make it harder for an attacker to perform any of the attacks discussed in this document.

[22] discusses a number of algorithms to randomize the ephemeral ports used by clients.

Also, a proposal exists to enable TCP to reassign a well-known port number to a random value [23].

5.4 Authentication

Hosts could require ICMP error messages to be authenticated [7], in order to act upon them. However, while this requirement could make sense for those ICMP error messages sent by hosts, it would not be feasible for those ICMP error messages generated by intermediate routers.

[7] contains a discussion on the authentication of ICMP messages.

5.5 Filtering ICMP errors based on the ICMP payload

As discussed in Section 4, the source address of ICMP error messages does not need to be spoofed to perform the attacks described in this draft. Thus, simple filtering based on the source address of ICMP error messages does not serve as a counter-measure against these attacks. However, a more advanced packet filtering could be used as a counter-measure. Systems performing such advanced filtering would look at the payload of the ICMP error messages, and would perform ingress and egress packet filtering based on the source IP address of the IP header contained in the payload of the ICMP error message. As the source IP address contained in the payload of the ICMP error message does need to be spoofed to perform the attacks described in this document, this kind of advanced filtering would serve as a counter-measure against these attacks.

6. Blind connection-reset attacks

6.1 Description

The Host Requirements RFC [4] states that a host SHOULD abort the corresponding connection when receiving an ICMP error message that indicates a hard error.

Thus, an attacker could use ICMP to perform a blind connection-reset attack. That is, even being off-path, an attacker could reset any TCP connection taking place. In order to perform such an attack, an attacker would send any ICMP error message that indicates a "hard error", to either of the two TCP endpoints of the connection. Because of TCP's fault recovery policy, the connection would be immediately aborted.

As discussed in Section 2.2, all an attacker needs to know to perform such an attack is the socket pair that identifies the TCP connection to be attacked. In some scenarios, the IP addresses and port numbers in use may be easily guessed or known to the attacker [12].

Some stacks are known to extrapolate ICMP errors across TCP connections, increasing the impact of this attack, as a single ICMP packet could bring down all the TCP connections between the corresponding peers.

There are some points to be considered about this type of attack:

- The source address of the ICMP error message need not be forged. Thus, simple filtering based on the source address of ICMP packets

would not serve as a counter-measure against this type of attack.

- o Even if TCP itself were protected against the blind connection-reset attack described in [12] and [14], the type of attack described in this document could still succeed.

6.2 Attack-specific counter-measures

6.2.1 Changing the reaction to hard errors

As discussed in Section 6.1, hosts MUST NOT extrapolate ICMP errors across TCP connections.

An analysis of the circumstances in which ICMP messages that indicate hard errors may be received can shed some light to minimize (or even eliminate) the impact of blind connection-reset attacks.

ICMP type 3 (Destination Unreachable), code 2 (protocol unreachable)

For ICMP messages of type 3 (Destination Unreachable) code 2 (protocol unreachable), specifically, the Host Requirements RFC states that even those transport protocols that have their own mechanisms to indicate that a port is unreachable MUST accept these ICMP error messages for the same purpose. That is, they MUST abort the corresponding connection when an ICMP port unreachable message is received.

This ICMP error message indicates that the host sending the ICMP error message received a packet meant for a transport protocol it does not support. For connection-oriented protocols such as TCP, one could expect to receive such an error as the result of a connection establishment attempt. However, it would be strange to get such an error during the life of a connection, as this would indicate that support for that transport protocol has been removed from the host sending the error message during the life of the corresponding connection. Thus, it would be fair to treat ICMP protocol unreachable error messages as soft errors (or completely ignore them) if they are meant for connections that are in synchronized states. For TCP, this means one would treat ICMP port unreachable error messages as soft errors (or completely ignore them) if they are meant for connections that are in the ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK or TIME-WAIT states.

ICMP type 3 (Destination Unreachable), code 3 (port unreachable)

This error message indicates that the host sending the ICMP error

message received a packet meant for a socket (IP address, port number) on which there is no process listening. Those transport protocols which have their own mechanisms for notifying this condition should not be receiving these error messages. However, the Host Requirements RFC [4] states that even those transport protocols that have their own mechanism for notifying the sender that a port is unreachable MUST nevertheless accept an ICMP Port Unreachable for the same purpose. For security reasons, it would be fair to treat ICMP port unreachable messages as soft errors (or completely ignore them) when they are meant for protocols that have their own mechanism for reporting this condition.

ICMP type 3 (Destination Unreachable), code 4 (fragmentation needed and DF bit set)

This error message indicates that an intermediate node needed to fragment a datagram, but the DF (Don't Fragment) bit in the IP header was set. Those systems that do not implement the PMTUD mechanism should not be sending their IP packets with the DF bit set, and thus should not be receiving these ICMP error messages. Thus, it would be fair for them to completely ignore this ICMP error message. On the other hand, and for obvious reasons, those systems implementing the Path-MTU Discovery (PMTUD) mechanism [6] should not abort the corresponding connection when such an ICMP error message is received.

ICMPv6 type 1 (Destination Unreachable), code 1 (communication with destination administratively prohibited)

This error message indicates that the destination is unreachable because of an administrative policy. For connection-oriented protocols such as TCP, one could expect to receive such an error as the result of a connection-establishment attempt. Receiving such an error for a connection in any of the synchronized states would mean that the administrative policy changed during the life of the connection. Therefore, while it would be possible for a firewall to be reconfigured during the life of a connection, it would be fair, for security reasons, to ignore these messages for connections that are in the ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK or TIME-WAIT states.

ICMPv6 type 1 (Destination Unreachable), code 4 (port unreachable)

This error message is analogous to the ICMP type 3 (Destination unreachable), code 3 (Port unreachable) error message discussed above. Therefore, the same considerations apply.

6.2.2 Delaying the connection-reset

An alternative counter-measure could be to delay the conenction reset. Rather than immediately aborting a connection, a TCP could abort a connection only after an ICMP error message indicating a hard error has been received a specified number of times, and the corresponding data have already been retransmitted more than some specified number of times.

For example, hosts could abort connections only after a fourth ICMP error message indicating a hard error is received, and the corresponding data have already been retransmitted more than six times.

The rationale behind this proposed fix is that if a host can make forward progress on a connection, it can completely disregard the "hard errors" being indicated by the received ICMP error messages.

While this counter-measure could be useful, we think that the counter-measure discussed in Section 6.2.1 is more simple to implement and provides increased protection against this type of attack.

7. Blind throughput-reduction attacks

The following subsections discuss a number of attacks that can be performed against TCP to reduce the throughput of a TCP connection. While these attacks do not reset the corresponding TCP connection, they may reduce their throughput to such an extent that they may become practically unusable.

7.1 ICMP Source Quench attack

7.1.1 Description

The Host requirements RFC states hosts MUST react to ICMP Source Quench messages by slowing transmission on the connection. Thus, an attacker could send ICMP Source Quench (type 4, code 0) messages to a TCP endpoint to make it reduce the rate at which it sends data to the other party. While this would not reset the connection, it would certainly degrade the performance of the data transfer taking place over it.

7.1.2 Attack-specific counter-measures

The Host Requirements RFC [4] states that hosts MUST react to ICMP Source Quench messages by slowing transmission on the connection. However, as discussed in the Requirements for IP Version 4 Routers

RFC [5], research seems to suggest ICMP Source Quench is an ineffective (and unfair) antidote for congestion. Thus, we recommend hosts to completely ignore ICMP Source Quench messages.

7.2 ICMP attack against the PMTU Discovery mechanism

7.2.1 Description

When one IP host has a large amount of data to send to another host, the data will be transmitted as a series of IP datagrams. It is usually preferable that these datagrams be of the largest size that does not require fragmentation anywhere along the path from the source to the destination. This datagram size is referred to as the Path MTU (PMTU), and is equal to the minimum of the MTUs of each hop in the path [6].

A technique called "Path MTU Discovery mechanism" (PMTUD) lets IP hosts determine the Path MTU of an arbitrary internet path. [6] and [10] specify the PMTUD mechanism for IPv4 and IPv6, respectively.

The PMTUD mechanism for IPv4 uses the Don't Fragment (DF) bit in the IP header to dynamically discover the Path MTU. The basic idea behind the PMTUD mechanism is that a source host assumes that the MTU of the path is the MTU of its first hop, and sends all its datagrams with the DF bit set. If any of the datagrams is too large to be forwarded without fragmentation by some intermediate router, the router will discard the corresponding datagram, and will return an ICMP "Destination Unreacheable" (type 3) "fragmentation neeed and DF set" (code 4) error message to sending host. This message will report the MTU of the constricting hop, so that the sending host reduces the assumed Path-MTU.

For IPv6, intermediate systems do not fragment packets. Thus, there's an "implicit" DF bit set in every packet sent on a network. If any of the datagrams is too large to be forwarded without fragmentation by some intermediate router, the router will discard the corresponding datagram, and will return an ICMPv6 "Packet Too Big" (type 2, code 0) error message to sending host. This message will report the MTU of the constricting hop, so that the sending host can reduce the assumed Path-MTU accordingly.

As discussed in both [6] and [10], the PMTUD can be used to attack TCP. An attacker could reduce the throughput of a TCP connection by forging ICMP "Destination Unreachable, fragmentation needed and DF set" packets (or their IPv6 counterpart), and making these packets report a low MTU.

For IPv4, this reported Next-Hop MTU could be as low as 68 octets, as

[11] requires every internet module to be able to forward a datagram of 68 octets without further fragmentation. For IPv6, the reported Next-Hop MTU could be as low as 1280 octets (the minimum IPv6 MTU) [9].

Thus, this attack could considerably readuce the throughput that can be achieved with the attacked TCP connection.

7.2.2 Attack-specific counter-measures

An analogous counter-measure to that described in Section 6.2.2 could be implemented to greatly minimize the impact of this attack.

For IPv4, this would mean that upon receipt of an ICMP "fragmentation needed and DF bit set" error message, TCP would just record this information, and would honor it only when it had received a specified number of ICMP "fragmentation needed and DF bit set" messages, and provided the corresponding data had already been retransmitted a specified number of times.

For IPv6, the same mechanism would be implemented ICMPv6 "Packet Too Big" error messages.

Henceforth, we will refer to both ICMP "fragmentation needed and DF bit set" and ICMPv6 "Packet Too Big" messages as "ICMP Packet Too Big" messages.

To implement the proposed fix, two new parameters would be introduced to TCP: MAXPKTTOOBIG, and MAXSEGREXMIT. MAXPKTTOOBIG would specify the number of times an ICMP "Packet Too Big" must be received before it can be honored to change the Path-MTU. MAXSEGREXMIT would specify the number of times a given segment must be retransmitted before an ICMP "Packet Too Big" error message can be honored.

Two variables would be needed to implement the proposed fix: npkttoobig, and nsegrexmit. npkttoobig would be initialized to zero, and would be incremented by one everytime a valid ICMP "Packet Too Big" error message is received. It would be reset to zero everytime an ICMP "Packet Too Big" error message is honored to change the assumed Path-MTU for given internet path. nsegrexmit would be initialized to zero, and would be incremented by one everytime the corresponding segment is retransmitted.

Thus, the Path-MTU for a given internet path would be changed only when a ICMP "Packet Too Big" is received, provided npkttoobig >= MAXPKTTOOBIG and nsegrexmit >= MAXSEGREXMIT.

The rationale behind this proposed fix is that if there is progress

on the connection, ICMP "Packet Too Big" messages must be a false claim.

MAXPKTTOOBIG and MAXSEGREXMIT might be a function of the Next-Hop MTU claimed in the received ICMP "Packet Too Big" message. That is, higher values for MAXPKTTOOBIG and MAXSEGREXMIT could be required when the received ICMP "Packet Too Big" message claims a Next-Hop MTU that is bellow some specified value.

As discussed in Section 7.2.1, hosts should impose lower limits in the reported Next-Hop MTU values they honor. For example, for IPv4 this lower limit could be safely raised to 296 octets, the MTU for Point-To-Point (low delay) links [6]. This lower limit could be probably raised to a higher value, such as 500 octets.

A mechanism that allows hosts to determine the Path-MTU without the use of ICMP has been is described in [24].

8. Future work

The same considerations discussed in this document should be applied to other similar protocols.

9. Security Considerations

This document describes the use of ICMP error messages to perform a number of attacks against the TCP protocol, and proposes a number of counter-measures that either eliminate or reduce the impact of these attacks.

10. Acknowledgements

This document was inspired by Mika Liljeberg, while discussing some issues related to [25] by private e-mail. The author would like to thank James Carlson, Alan Cox, Juan Fraschini, Markus Friedl, Guillermo Gont, Vivek Kakkar, Michael Kerrisk, Mika Liljeberg, David Miller, Eloy Paris, Kacheong Poon, Andrew Powell, and Pekka Savola for contributing many valuable comments.

The author wishes to express deep and heartfelt gratitude to Jorge Oscar Gont and Nelida Garcia, for their precious motivation and guidance.

11. References

11.1 Normative References

[1] Postel, J., "Transmission Control Protocol", STD 7, RFC 793,

September 1981.

[2] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[4] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.

[5] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[6] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.

[7] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[8] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.

[9] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[10] McCann, J., Deering, S. and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.

[11] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

11.2 Informative References

[12] Watson, P., "Slipping in the Window: TCP Reset Attacks", 2004 CanSecWest Conference , 2004.

[13] Jacobson, V., Braden, B. and D. Borman, "TCP Extensions for High Performance", RFC 1323, May 1992.

[14] Stewart, R., "Transmission Control Protocol security considerations", draft-ietf-tcpm-tcpsecure-02 (work in progress), November 2004.

[15] Touch, J., "ANONsec: Anonymous IPsec to Defend Against Spoofing Attacks", draft-touch-anonsec-00 (work in progress), May 2004.

[16] Poon, K., "Use of TCP timestamp option to defend against blind spoofing attack", draft-poon-tcp-tstamp-mod-01 (work in progress), October 2004.

[17] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

[18] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.

[19] Clark, D., "Fault isolation and recovery", RFC 816, July 1982.

[20] Gont, F., "Increasing the payload of ICMP error messages", (work in progress) draft-gont-icmp-payload-00.txt, 2004.

[21] Allman, M., Paxson, V. and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.

[22] Larsen, M., "Port Randomisation", draft-larsen-tsvwg-port-randomisation-00 (work in progress), October 2004.

[23] Shepard, T., "Reassign Port Number option for TCP", draft-shepard-tcp-reassign-port-number-00 (work in progress), July 2004.

[24] Mathis, M., "Path MTU Discovery", draft-ietf-pmtud-method-03 (work in progress), October 2004.

[25] Gont, F., "TCP's Reaction to Soft Errors", draft-gont-tcpm-tcp-soft-errors-01 (work in progress), October 2004.

Author's Address

Fernando Gont
Universidad Tecnologica Nacional
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
EMail: fernando@gont.com.ar

Appendix A. Changes from draft-gont-tcpm-icmp-attacks-01

- The document was restructured for easier reading.
- A discussion of ICMPv6 was added in several sections of the document
- Added Section 5.2
- Added Section 5.5
- Added Section 7.2
- Fixed typo in the ICMP types, in several places
- Fixed typo in the TCP sequence number check formula
- Miscellaneous editorial changes

Appendix B. Changes from draft-gont-tcpm-icmp-attacks-00

- Added Section ChangingHandling
- Added a summary of the relevant RFCs in several sections
- Miscellaneous editorial changes

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.