UNIRAS (UK Gov CERT)

Advisory Type: Alert

Id: 20050412-00308, Ref: 12/05 **Date: 12 April 2005 Time: 12:55**

Title: NISCC Vulnerability Advisory ICMP - 532967

Abstract: ICMP is the control protocol for IP (Internet Protocol), a core network protocol used in the majority of networked computer systems today. Most vendors include support for this protocol in their products and may be impacted to varying degrees. Furthermore any network service or application that relies on a long-lived TCP connection will also be impacted if the host processes ICMP messages in accordance with RFC 1122. For the Source Quench attack the severity will depend on the throughput of the TCP connection; the application may well become unusable.

Vendors affected: multiple

Operating Systems affected: multiple

Applications/Services affected: multiple

```
Title
=====
NISCC Vulnerability Advisory ICMP - 532967

Detail
======

NISCC Vulnerability Advisory 532967/NISCC/ICMP

Vulnerability Issues in ICMP packets with TCP payloads


Version Information
- -------------------

Advisory Reference 532967/NISCC/ICMP
Release Date         12 April 2005
Last Revision         12 April 2005
Version Number         1.0


What is Affected?
- -----------------

The vulnerabilities described in this advisory affect the TCP (Transmission Control
Protocol) by using Internet Control Message Protocol (ICMP) messages that comply with the
Internet Engineering Task Force's (IETF's) Requests For Comments (RFCs) for ICMP,
including
RFC 792 "Internet Control Message Protocol: DARPA Internet Program Protocol
Specification"
(for IP Version 4), RFC 1122, "Requirements for Internet Hosts -- Communication Layers"
and
potentially RFC 2463 "Internet Control Message Protocol (ICMPv6) for the Internet
Protocol
Version 6 (IPv6) Specification" (for IP Version 6). The original TCP specification is
provided in RFC 793.

ICMP is the control protocol for IP (Internet Protocol), a core network protocol used in
```

the
majority of networked computer systems today. Most vendors include support for this protocol
in their products and may be impacted to varying degrees. Furthermore any network service or
application that relies on a long-lived TCP connection will also be impacted if the host
processes ICMP messages in accordance with RFC 1122. For the Source Quench attack the
severity will depend on the throughput of the TCP connection; the application may well
become unusable.

Severity
- --------

The impact of the ICMP TCP reset vulnerability (called "the TCP blind connection-reset
vulnerability" in this advisory) varies by vendor and application, but in some deployment
scenarios it is likely to be rated medium to high. Please see the 'Vendor Information'
section below for further information. Alternatively contact your vendor for product
specific information.

If exploited, the TCP blind connection-reset vulnerability could allow an attacker to create
a denial-of-service condition against existing TCP connections, resulting in premature
session termination. The resulting session termination will affect the application layer,
the nature and severity of the effects being dependent on the application layer protocol.
The primary dependency is on the tolerance of the network service or application to the loss
of a TCP connection.

The Border Gateway Protocol (BGP) is judged to be potentially most affected by this
vulnerability. BGP relies on a persistent TCP connection between BGP peers; resetting the
connection can result in medium term unavailability due to the need to rebuild routing
tables and route flapping. Route flapping may result in route dampening (suppression) if the
route flaps occur frequently within a short time interval. The overall impact on BGP is
likely to be low to moderate based on the likelihood of successful attack, but could be high
if an ICMP implementation does not perform any checks on the ICMP payload.

If an access control list is applied at routers to block packets of ICMP Type 3 codes 2, 3
and 4 then the impact will be low as this measure will successfully mitigate the
vulnerability. Anti-spoofing measures can also be of benefit if the attacker spoofs the IP
address from where the ICMP packet is sent. Anti-spoofing measures include access control
lists to block non-routable IP addresses (see RFCs 1918 and 3330) and Unicast Reverse Path
Forwarding (URPF) that checks the consistency of the source IP address with the interface on
which the packets are received. See NISCC Technical Note 06/02 "Response to Distributed
Denial-of-Service (DDoS) Attacks" for further details.

There is a potential impact on other application protocols such as DNS (Domain Name System)
and SSL (Secure Sockets Layer) in the case of zone transfers and ecommerce transactions
respectively, but the sessions can be restarted without medium term unavailability problems.
In the case of DNS the TCP connections are short lived, so the chance of the vulnerability
being exploited is lower than for long lived TCP connections. In the case of SSL it may be
difficult to guess the source IP address because it could be dynamically allocated home user
address (in the case of Internet banking).

The severity of the related issue of slowing down routers that use Path MTU discovery is
also likely to be moderate to high in some vendors' products because RFC 792 and RFC 1191
("Path MTU discovery") do not specify checking of sequence numbers.

The severity of the spoofing of ICMP Source Quench packets is likely to be moderate to low
because the support of routers for Source Quench as a means of congestion control has been
deprecated for ten years. RFC 1812 section 5.3.6 states: "As described in Section [4.3.3.3],
this document recommends that a router SHOULD NOT send a Source Quench to the sender of the
packet that it is discarding. ICMP Source Quench is a very weak mechanism, so it is not
necessary for a router to send it, and host software should not use it exclusively as an

indicator of congestion." On the other hand, RFC 1122 section 4.2.3.9 states that "TCP MUST
react to a Source Quench by slowing transmission on the connection", a statement honoured in
 a number of TCP implementations. To mitigate Source Quench attacks using spoofed IP addresses in the payload, ICMP Source Quench (ICMP Type 4) messages should not be allowed through routers or through firewalls at the organisational perimeter. It is reasonable to allow routers to block Source Quench packets if their use is deprecated.

Summary
- -------

The first issue described in this advisory is the practicability of resetting an established
TCP connection by sending suitable ICMP packets that simulate a hard error condition in an
existing TCP connection. Hard error conditions are defined in RFC 1122 section 4.2.3.9 and
include a number of common ICMP types. An ICMP error packet records the IP header of the
packet causing the error as well as the first 64 bits of the TCP header, which consists of
the source and destination ports and the sequence number. Many ICMP implementations only
check the IP addresses and TCP ports at either end of the connection; they do not check
whether the sequence number of the packet is within an acceptable range (see 'Details'
section below for characterisation of this range).

It is thus possible in some implementations for an attacker to reset an existing TCP
connection by sending a suitably crafted ICMP packet with the correct IP addresses and TCP
ports. The target of these denial-of-service attacks is any TCP connection, especially one
for which the source port can be identified or guessed. Moreover any application protocol
which relies on long term TCP connections and for which the source and destination IP
addresses and TCP ports are known or can be easily guessed will be vulnerable to
denial-of-service attacks.

A related, subsidiary issue is the potential ability to slow down traffic through hosts that
 use Path MTU discovery (defined in RFC 1191) by sending forged ICMP Type 3 Code 4
("Fragmentation Needed and Don't Fragment was Set") packets that report a (false) low
"next-hop MTU" to a host using the Path MTU discovery mechanism.

The third issue described in this advisory is the practicability of slowing the traffic
between two hosts by sending ICMP Source Quench packets to an endpoint of the session. The
technique used is to send a Source Quench packet including the details of the TCP connection
to be targeted. According to RFC 1122, the Source Quench packet will limit the rate of the
TCP connection.

It is possible to apply the TCP blind connection-reset vulnerability to ICMP Version 6
packets (the control protocol of IP Version 6) by equating hard errors to ICMPv6 Type 1
(Destination Unreachable) codes 1 (communication with destination administratively
prohibited) and 4 (port unreachable). The Path MTU discovery attack could be affected by
sending an ICMPv6 Type 2 code 0 ("Packet Too Big") packet, which does not describe a hard
error but is used to determine end to end path MTU. Source quench is not defined in RFC 2463
for ICMPv6. Since IP Version 6 was not defined when RFC 1122 was written, the discussion in
this advisory will concentrate on IP Version 4. Further details of how the attacks apply to
IP Version 6 is available in Fernando Gont's Internet Draft "ICMP attacks against TCP".

Details
- -------

532967/NISCC/ICMP/1
CVE number: CAN-2004-0790

RFC 1122 section 4.2.3.9 refers to some ICMP message types as representing hard errors.
These message types are ICMP Type 3 (Destination Unreachable) codes 2 (Protocol
Unreachable), 3 (Port Unreachable) and 4 (Fragmentation Needed and Don't Fragment was
Set).
For hard errors, according to RFC 1122 the TCP implementation should abort the
connection.
Thus by sending an ICMP Type 3, code 2, 3 or 4 with the IP header and the first 64 bits

of
the header of a TCP as its payload (the source and destination TCP ports and the sequence number), the receiving TCP implementation would reset the existing connection if it did not
check that the sequence number was as expected.

It should be noted that the current RFCs do not recommend that TCP implementations check the
sequence number. In section 5.1 of his Internet Draft "ICMP attacks against TCP" Fernando Gont states: "TCP SHOULD check that the sequence number in the TCP header 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."
Here SND.UNA is the oldest unacknowledged sequence number, SEG.SEQ is the sequence number contained in the payload of the ICMP error message, and SND.NXT is the next sequence number
to be sent. There is a dependency on the TCP window size as without scaling (see RFC 1323)
the range of unacknowledged sequence numbers can be in the range SND.UNA to SND.NXT-1.

532967/NISCC/ICMP/2
CVE number: CAN-2004-1060

In the case where a host complies with RFC 1191 ("Path MTU Discovery"), it is possible to use the blind connection-reset attack with a ICMP Type 3 Code 4 packet and the addition of
the "next-hop MTU" field in the ICMP header set to a value of 68 (octets) to slow down the
transmission rate for traffic from the host.

NISCC/532967/ICMP/3
CVE number: CAN-2004-0791

RFC 1122 section 4.2.3.9 states "TCP MUST react to a Source Quench by slowing transmission
on the connection. The RECOMMENDED procedure is for a Source Quench to trigger a "slow start," as if a retransmission timeout had occurred." Thus by sending an ICMP Type 4 (Source
Quench) packet to a host with the IP header and the first 64 bits of the header of a TCP as
its payload, the receiving TCP implementation would rate limit the existing connection if it
did not check that the sequence number was expected. As noted above, the current RFCs do not
recommend that TCP implementations check the sequence number.

Mitigation
- ----------

The main impact of the TCP blind connection-reset vulnerability is on applications that are
intolerant of loss of a TCP session, such as BGP. In this case applying an access control list to block ICMP Type 3 code 2, 3 and 4 packets to BGP routers will be an effective mitigation, as will the use of anti-spoofing measures in the case that the ICMP packet inducing the reset is sent from a spoofed IP address.

In the case of hosts which use Path MTU discovery, the only mitigation is to disable the use
of the Path MTU discovery mechanism until the vendor provides a security patch.

The Source Quench vulnerability can be mitigated by applying an access control list blocking
ICMP Type 4 packets on routers and by blocking ICMP Type 4 packets to corporate networks at
the organisational boundary. Again, anti-spoofing measures such as URPF and access control
lists blocking private, non-routable IP addresses at routers will also provide some protection if the source IP address of the ICMP packet is spoofed.

Solution
- --------

General solutions to the vulnerabilities are described in section 5 of Fernando Gont's Internet Draft "ICMP attacks against TCP" (see

http://www.ietf.org/internet-drafts/draft-gont-tcpm-icmp-attacks-03.txt). These solutions include:

- - Checking that the TCP sequence number is within the range of the data already sent but not
yet acknowledged

- - (On routers) checking that the TCP acknowledgement number is in the range of the last sequence number acknowledged to the next sequence number expected

- - Randomising source (ephemeral client) port numbers (making the port number more difficult
to guess)

- - Providing authentication mechanisms for ICMP messages to ensure that an ICMP message is
processed only if it is correctly authenticated

- - Ingress and egress filtering on the IP addresses and TCP ports in the payload of ICMP packets

- - Changing the behaviour of a TCP implementation when the ICMP message is not expected to
cause a soft rather than a hard error

- - Delaying the TCP connection reset until an ICMP error message indicating a hard error has
been received a specified number of times

- - To protect against resets against Path MTU Discovery, delay the handling of a ICMP Type 3
Code 4 ("packet too big" error) packets in the case where the Path MTU has already been negotiated, i.e. where larger packets sizes have been already sent and acknowledged

- - Changing the handling of ICMP hard error messages for connections in synchronized states
to ignore the messages during the life of a connection or to treat them as soft errors

- - Ignoring Source Quench ICMP packets

Please refer to the 'Vendor Information' section of this advisory for implementation specific remediation.

Credits
- -------

This issue was discovered by Fernando Gont (UTN/FRH) and is the subject of an Internet draft
by the same author; "ICMP attacks against TCP" (see
http://www.ietf.org/internet-drafts/draft-gont-tcpm-icmp-attacks-03.txt). The latest version
of the Internet Draft can be obtained from
http://www.gont.com.ar/drafts/icmp-attacks-against-tcp.html.

NISCC wishes to thank Fernando Gont for bringing this vulnerability to our attention and for
agreeing to allow NISCC to co-ordinate the disclosure of this issue. NISCC also wishes to thanks Fernando for his comments on this advisory.

Vendor Information
- ------------------

A list of vendors affected by this vulnerability is not currently available. Please visit the web site (http://www.niscc.gov.uk/niscc/docs/re-20050412-00303.pdf?lang=en) in order to check for updates.

Contact Information
- -------------------

The NISCC Vulnerability Management Team can be contacted as follows:

Email   vulteam@niscc.gov.uk
           Please quote the advisory reference in the subject line

Telephone +44 (0)870 487 0748 Ext 4511
           Monday - Friday 08:30 - 17:00

```
Fax    +44 (0)870 487 0749

Post    Vulnerability Management Team
           NISCC
           PO Box 832
           London
           SW1P 1BG

We encourage those who wish to communicate via email to make use of our PGP key. This is
available from http://www.niscc.gov.uk/niscc/publicKey2-en.pop.

Please note that UK government protectively marked material should not be sent to the
email
address above.

If you wish to be added to our email distribution list, please email your request to
uniras@niscc.gov.uk.


What is NISCC?
- --------------

For further information regarding the UK National Infrastructure Security Co-Ordination
Centre, please visit the NISCC web site at: http://www.niscc.gov.uk.

Reference to any specific commercial product, process or service by trade name, trademark
manufacturer or otherwise, does not constitute or imply its endorsement, recommendation,
or
favouring by NISCC. The views and opinions of authors expressed within this notice shall
not
be used for advertising or product endorsement purposes.

Neither shall NISCC accept responsibility for any errors or omissions contained within
this
advisory. In particular, they shall not be liable for any loss or damage whatsoever,
arising
from or in connection with the usage of information contained within this notice.

C 2005 Crown Copyright

<End of NISCC Vulnerability Advisory>
```

Acknowledgements

UNIRAS wishes to acknowledge the contributions of NISCC Vulnerability Team for the information contained in this Briefing.

Updates

This advisory contains the information released by the original author. Some of the information may have changed since it was released. If the vulnerability affects you, it may be prudent to retrieve the advisory from the canonical site to ensure that you receive the most current information concerning that problem.

Legal Disclaimer

FIRST

UNIRAS is a member of the Forum of Incident Response and Security Teams (FIRST) and has contacts with other international Incident Response Teams (IRTs) in order to foster cooperation and coordination in incident prevention, to prompt rapid reaction to incidents, and to promote information sharing amongst its members and the community at large.