

NAT Classification Test Results
draft-jennings-behave-test-results-01

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 17, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

IETF has several groups that are considering the impact of NATs on various protocols. Having a classification of the types of NATs that are being developed and deployed is useful in gauging the impact of various solutions. This draft records the results of classifying NATs.

This draft is not complete and has only a few test results but it is worth discussing all the testing we wish to do before all the test results are collected. The test results here are very old and work

is being done to update them with more current information.

This work is being discussed on the ietf-behave@list.sipfoundry.org mailing list

1. Conventions

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](#) [2].

2. Introduction

A major issue in working with NAT traversal solutions for various protocols is that NATs behave in many different ways. This draft describes the results of testing several residential style NATs.

3. Descriptions of Tests

3.1 UDP Mapping

This test sends STUN packets from the same port on three different internal IP addresses to the same destination. The source port on the outside of the NAT is observed. The test records whether the port is preserved or not and whether all the mappings get different ports.

A second set of tests checks out how the NAT maps ports above and below 1024.

Tests are run with a group of several consecutive ports to see if the NAT preserves port parity.

3.2 UDP Filtering

This test sends STUN packets from the same port on three different internal IP addresses to the same destination. It then tests whether places on the outside with 1) a different port but the same IP address and then 2) a different port and a different IP address can successfully send a packet back to the sender.

3.3 UDP Hairpin

This test sends a STUN packet from the inside to the outside to create a mapping and discover the external source address called A. It does the same thing from a different internal IP address to get a second external mapping called B. It then sends a packet from A to B and B to A and notes if these packets are successfully delivered from

one internal IP address to the other.

3.4 ICMP

A device on the inside sends a packet to an external address that causes an ICMP Destination Unreachable packet to be returned. The test records whether this packet makes it back through the NAT correctly.

3.5 Fragmentation

The MTU on the outside of the NAT is set to under 1000; on the inside it is set to 1500 or over. Then a 1200 byte packet is sent to the NAT. The test records whether the NAT correctly fragments this when sending it. Another test is done with DF=1. An additional test is done with DF=1 in which the adjacent MTU on the NAT is large enough the NAT does not need to fragment the packet but further on, a link has an MTU small enough that an ICMP packet gets generated. The test records whether the NAT correctly forwards the ICMP packet.

In the next test a fragmented packet with the packets in order is sent to the outside of the NAT, and the test records whether the packets are dropped, reassembled and forwarded, or forwarded individually. A similar test is done with the fragments out of order.

3.6 UDP Refresh

A test is done that involves sending out a STUN packet and then waiting a variable number of minutes before the server sends the response. The client sends different requests with different times on several different ports at the start of the test and then watches the responses to find out how long the NAT keeps the binding alive.

A second test is done with a request that is delayed more than the binding time but every minute an outbound packet is sent to keep the binding alive. This test checks that outbound traffic will update the timer.

A third test is done in which several requests are sent with the delay less than the binding time and one request with the delay greater. The early test responses will result in inbound traffic that may or may not update the binding timer. This test detects whether the packet with the time greater than the binding time will traverse the NAT which provide the information about whether the inbound packets have updated the binding timers.

An additional test is done to multiple different external IP

addresses from the same source, to see if outbound traffic to one destination updates the timers on each session in that mapping.

[3.7](#) Multicast and IGMP

Multicast traffic is sent to the outside of the NAT, and the test records whether the NAT forwards it to the inside. Next an IGMP Membership Report is sent from inside. The test records whether the NAT correctly forwards it to the outside and whether it allows incoming multicast traffic.

[3.8](#) Multicast Timers

The test records how long the NAT will forward multicast traffic without receiving any IGMP Membership Reports and whether receiving Reports refreshes this timer.

[3.9](#) TCP Timers

TBD: Measure time before ACK, after ACK, and after FIN and RST.

[3.10](#) TCP Port Mapping

Multiple SYN packets are sent from the same inside address to different outside IP addresses, and the source port used on the outside of the NAT is recorded.

[3.11](#) SYN Filtering

Test that a SYN packet received on the outside interfaces that does not match anything gets discarded with no reply being sent. Test whether an outbound SYN packet will create a binding that allows an incoming SYN packet.

[3.12](#) DNS

Does the DNS proxy in them successfully pass through SRV requests.

[3.13](#) DHCP

Do any DHCP options received on the WAN side get put into DHCP answers sent on lan side?

[4.](#) Observations

Several NATs attempt to use the same external port number as the internal host has used. This is referred to as port preservation. Some of the NATs that do this were found to have different

characteristics depending on whether the port was already in use or not. This was tested by running the STUN tests from a particular port on one internal IP address and then running them again from the same port on a different internal IP address. The results from the first interface, where the port was preserved, are referred to as the primary type; while the results from the second interface, which did not manage to get the same external port because it was already in use, are referred to as the secondary type. On most NATs the secondary type is the same as the primary but on some it is different; these are referred to as nondeterministic NATs, since a client with a single internal IP address cannot figure out what type of NAT it is.

There are several NATs that would be detected as address restricted by the STUN tests but are not. These NATs always use the same external port as the internal port and store the IP address of the most recent internal host to send a packet on that port. The NATs then forward any traffic arriving at the external interface of the NAT on this port to the internal host that has most recently used it. These NATs are labeled "Bad" in the result table since they do not meet the definitions of NAPT in [RFC 3022](#). Interestingly, as long as the clients behind the NAT choose random port numbers, they often do work. STUN detects these NATs as address restricted although they are really not address restricted NATs. This type of NAT is easily detected by sending a STUN packet from the same port on two different internal IP addresses and looking at the mapped port in the return. If both packets have been mapped to the same external port, the NAT is of the Bad type.

Another important aspect of a NAT for some applications is whether it can send media from one internal host back to another host behind the same NAT. This is referred to as supporting hairpin media.

It was rumored that some NATs existed that looked in arbitrary packets for either the NATs' external IP address or the internal host IP address - either in binary or dotted decimal form - and rewrote it to something else. STUN could be extended to test for exactly this type of behavior by echoing arbitrary client data and the mapped address but sending the bits inverted so these evil NATs did not mess with them. NATs that do this will break integrity detection on payloads.

To help organize the NATs by what types of applications they can support, the following groups are defined. The application of using a SIP phone with a TLS connection for signaling and using STUN for media ports is considered. It is assumed the RTP/RTCP media is on random port pairs as recommended for RTP.

Group A: NATs that are deterministic, not symmetric, and support hairpin media. These NATs would work with many phones behind them.

Group B: NATs that are not symmetric on the primary mapping. This group would work with many IP phones as long as the media ports did not conflict. This is unlikely to happen often but will occasionally. Because they may not support hairpin media, a call from one phone behind a NAT to another phone behind the same NAT may not work.

Group D: NATs of the type Bad. These have the same limitations of group B but when the ports conflict, media gets delivered to a random phone behind the NAT.

Group F: These NATs are symmetric and phones will not work.

5. Results

To help with common reporting of test results. This specification will use the following format:

Address and port mapping behavior:

"endpoint independent" |
"address dependent" |
"address and port dependent"

IP address pooling behavior:

"unsupported" |
"arbitrary" |
"paired"

Port preservation:

"yes" |
"no" |
"overloading"

Port-range preservation:

"none" |
"registered" |
"dynamic" |
"registered and dynamic"

Port-parity preservation:

"yes" |
"no"

Port-contiguity preservation:

"yes" |
"no"

Mapping refresh timer:

<seconds> |
"configurable"

Mapping outbound refresh:

"yes" |

```
"no"
Mapping inbound refresh:
  "yes" |
  "no"
Unsolicited packet filtering:
  "endpoint independent" |
  "address dependent" |
  "address and port dependent"
Filter refresh timer:
  <seconds> |
  "configurable"
Filter outbound refresh:
  "yes" |
  "no"
Filter inbound refresh:
  "yes" |
  "no"
Hairpinning behavior:
  "none" |
  "external source address and port" |
  "internal source address and port"
Fixed application level gateways:
  <DNS, FTP, etc.> |
  "none"
Configurable application level gateways:
  <DNS, FTP, etc.> |
  "none"
Mapping and filtering determinism:
  "deterministic" |
  "non-deterministic"
Supports ICMP destination unreachable:
  "yes" |
  "no"
Supports fragmentation:
  "yes" |
  "no"
Fragment receive ordering:
  "ordered" |
  "out of order" |
  "none"
Maximum transmission unit (MTU):
  <bytes> |
  "configurable"
```

OPEN ISSUE: Should this be XML? can we make these shorter?

For example, the product datasheet for a given NAT device might include the following complete description of its NAT behavior:

Address and port mapping behavior:	endpoint independent
IP address pooling behavior:	paired
Port preservation:	yes
Port-range preservation:	registered and dynamic
Port-parity preservation:	yes
Port-contiguity preservation:	yes
Mapping refresh timer:	configurable
Mapping outbound refresh:	yes
Mapping inbound refresh:	yes
Unsolicited packet filtering:	endpoint independent
Filter refresh timer:	configurable
Filter outbound refresh:	yes
Filter inbound refresh:	yes
Hairpinning behavior:	external
Supports ICMP destination unreachable:	yes
Supports fragmentation:	yes
Fragment receive ordering:	out of order

The following table shows the results from several NATs. The NATs tested include some random ones the author had lying around as well as every NAT that could be purchased in February 2004 in the San Jose Fry's, Best Buy, CompUSA, and Circuit City. Clearly this is not a very good approximation to a random sample. It is clear that the NATs widely purchased in the US are different from what are available in Japan and Europe.

In the following table the Prim column indicates the primary type of the NAT. A value of Port indicates port restricted, Cone is a full cone, Bad is described in the next section, Symm is Symmetric, and Addr is Address restricted. The Hair column value of Y or N indicates whether the NAT will hairpin media. The Pres column indicates whether the NAT attempts to preserve port numbers. The Sec column indicates the secondary type of the NAT, and a value of Same indicates it is the same as the primary type. The Grp indicates the group that this NAT falls into.

Vendor	Model	Firmware	Prim	Sec	Hair	Pres	Grp
Airlink	AS0H04P	V1.01.0095	Port	Symm	N	Y	B
Apple	Air Base	V5.2	Cone	Same	Y	N	A
Belkin	F5D5321	V1.13	Port	Same	N	N	B
Cisco	IOS		Port	Symm			-
Cisco	PIX		Port	Same			-
Corega	BAR Pro2	R1.00 Feb 21 2003	Cone				-
DLink	DI-604	2.0 Jun 2002	Cone	Same	N	N	B
DLink	DI-704P	2.61 build 2	Cone	Same	Y	N	A
Dlink	DI-804	.30, Tue, Jun 24 20	Cone	Same	Y	N	A
Hawkings	FR24	6.26.02h Build 004	Bad	Same	Y	Y	D
Linksys	BEFSR11		Port				B
Linksys	BEFSR11 V2	1.42.7, Apr 02 200	Port				B
Linksys	BEFSR41	v1.44.2	Port				B
Linksys	BEFSR81	2.42.7.1 June 2002	Addr	Same	N	Y	B
Linksys	BEFSRU31		Port				B
Linksys	BEFSX41	1.44.3, Dec 24 200	Port				B
Linksys	BEFVP41	1.41.1, Sep 04 200	Port				B
Linksys	BEFW11S4	1.45.3, Jul 1 2003	Port				B
Linksys	WRT54G	1.42.2	Port	Symm	N	Y	B
Linksys	WRT55AG	1.04, Jun.30, 2003	Port				B
Linksys	WRV54G	2.03	Port	Same	N	Y	B
Microsoft	MN-700	02.00.07.0331	Cone	Same	N	N	B
Netgear	FVS318	V1.4 Jul. 15 2003	Port	Same	N	N	B
Netgear	RP114	3.26(CD.0) 8/17/20	Cone				-
Netgear	RP614	4.00 April 2002	Cone	Same	Y	N	A
NetworkEver	NR041	Version 1.0 Rel 10	Symm	Same	N	N	F
NetworkEver	NR041	Version 1.2 Rel 03	Bad	Same	Y	Y	D
SMC	2804WBRP-G	v1.00 Oct 14 2003	Port	Symm	Y	Y	B
SMC	7004ABR	V1.42.003	Port	Same	N	N	B
SMC	7004VBR	v1.03 Jun 12, 2002	Cone				-
Toshiba	WRC-1000	1.07.03a-C024a	Port	Cone	N	Y	B
umax	ugate-3000	2.06h	Port				-
US Robotics	USR8003	1.04 08	Cone	Same	N	N	B
ZOT	BR1014	Unknown	Bad	Same	N	Y	D

Since this testing was done, some additional testing and shopping sprees in France and Taiwan have provided the following results.

Vendor	Model	Firmware	Prim	Sec	Hair	Pres	Grp
Netgear	MR814v2	Version 5.01	Bad	Same	Y	Y	D
Cisco	PIX 515	6.3(3)	Port	Same	N	N	B
Dynex	DX-E401	1.03	Cone	Same	Y	N	A
Asante	FR1004	R1.13 V2	Cone	Same	N	N	B
Linksys	BEFSR81	2.42.7.1	Addr	Note 1	N	Y	B
Lanner	BRL-04FPU		Cone	Same	N	N	
AboCom	CAS3047		Port	Same	N	Y	
Lemmel	LM-IS6400B		Port	Same	N	Y	

The NAT with a secondary type of "Note 1" is particularly weird. The primary connection is address restricted. If a second host uses this same port, it also gets an Address Restricted, but when a third host uses this same port, it gets Symmetric.

Another good source of information for behavior of various NATs is the NATCECK [\[9\]](#) and STUNT [\[8\]](#) web pages.

Open Issue: How should we arrange all the results? There are going to be too many to put it as one row per device.

[6.](#) Discussion

It is clear from discussions with various vendors and watching how tests have changed over the years that symmetric is becoming less common. This change is being driven primarily by the desire to make online gaming work; many games use methods similar to STUN for NAT traversal. The only symmetric NAT found was an old device. More recent versions of the software on the same device were not symmetric. It is clear that other symmetric NATs are deployed, but it is hard to find them.

[7.](#) Security Concerns

It is often assumed that symmetric NATs are more secure than port restricted NATs. This is not true - they are identical from a security point of view. They both only allow a packet to come inside the NAT if the host inside has previously sent to the exact same external IP and port. One can argue that cone is less secure than port restricted, but this is not true if the attacker can spoof the IP address, which is fairly easy to do in many cases. What level of security can be expected from NATs at all is a strange and curious topic. With all the NATs, if you allow packets out, packets can come in, so don't be surprised if NATs provide less security than anticipated.

8. Open Issues

The hairpin media tests were done by having a single host use STUN to find a public address on the NAT and then send media to itself and see if it was received. It is possible that NATs might not hairpin media to the same host but would hairpin media to another host behind the same NAT. It is possible that because of this, the hairpin results reported here might be wrong.

This sample set of NATs is very US-centric: D-Link, Linksys, and Netgear dominate the US consumer market. It would be good to get more results from other places.

These test results should be verified by another group. This has not been done yet.

This draft should be moved to be consistent with the classification in [\[11\]](#).

9. Acknowledgments

Many people and several mailing lists have contributed to the material on understanding NATs in this document. Many thanks to Larry Metzger, Dan Wing, and Rohan Mahy. The STUN server and client is open source and available at <http://sourceforge.net/projects/stun>, and thank you to Jason Fischl who runs the public STUN server at larry.gloo.net. Thanks to Yutaka Takeda who tested and found bugs and Christian Stredicke for getting people thinking. Thanks to Francois Audet for catching mistakes, verifying several results, and finding the very strange non-deterministic nature in the BEFSR81.

The work of the various people on STUN Client and Server [\[6\]](#), NATCECK [\[10\]](#), and STUNT [\[7\]](#) has greatly helped this work.

10. References

10.1 Normative References

- [1] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", [RFC 3489](#), March 2003.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

10.2 Informative References

- [3] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", [RFC 3424](#), November 2002.
- [4] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [5] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [6] Jennings, C., "STUN Client and Server: <http://www.vovida.org/applications/downloads/stun>", February 2005.
- [7] Guha, S. and P. Francis, "STUNT <http://nutss.gforge.cis.cornell.edu/stunt.php>", February 2005.
- [8] Guha, S. and P. Francis, "STUNT Results <http://www.guha.cc/saikat/stunt-results.php>", February 2005.
- [9] Ford, B. and D. Andersen, "Nat Check Results <http://bgp.lcs.mit.edu/~dga/view.cgi>", February 2005.
- [10] Ford, B. and D. Andersen, "Nat Check <http://midcom-p2p.sourceforge.net>", February 2005.
- [11] Audet, F. and C. Jennings, "NAT Behavioral Requirements for Unicast UDP", [draft-ietf-behave-nat-udp-00](#) (work in progress), January 2005.
- [12] Wing, D., "IGMP Proxy Behavior", [draft-wing-behave-multicast-00](#) (work in progress), October 2004.
- [13] Sivakumar, S., "NAT Behavioral Requirements for TCP", [draft-sivakumar-behave-nat-tcp-req-00](#) (work in progress), January 2005.

Author's Address

Cullen Jennings
Cisco Systems
170 West Tasman Drive
Mailstop SJC-21/2
San Jose, CA 95134
USA

Phone: +1 408 421 9990
Email: fluffy@cisco.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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