

Short term NAT requirements for IPv6 transition

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

This document is an Internet-Draft. 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>.

Abstract

During the next few years, as the Ipv4 address space moves toward exhaustion, it is likely that the deployment of NAT will accelerate. By 2005, millions of NAT devices will likely be deployed on the Internet, both within enterprises and consumer households. Should those NAT devices not support either native Ipv6, or IPv6 transition mechanisms such as 6 to 4, the result would be significant delays in the deployment of IPv6.

This draft presents the requirements that NAT devices must meet in order to enable a future transition to IPv6. Rather than specifying every aspect of a NAT's operation in detail, our focus is solely on identifying those requirements that are absolutely essential to ensure compatibility with what we believe will be the most popular IPv6 transition mechanisms.

1 Introduction

During the next few years, as the Ipv4 address space moves toward exhaustion, it is likely that the deployment of NAT will accelerate. By 2005, millions of NAT devices will likely be deployed on the Internet, both within enterprises and consumer households. Should those NAT devices not support either native Ipv6, or IPv6 transition mechanisms such as 6 to 4, the result would be significant delays in the deployment of IPv6.

This draft presents the requirements that NAT devices must meet in order to enable a future transition to IPv6. Rather than specifying every aspect of a NAT's operation in detail, our focus is solely on identifying those requirements that are absolutely essential to ensure compatibility with what we believe will be the most popular IPv6 transition mechanisms.

1.01 Requirements language

In this document, the key words "MAY", "MUST", "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [[RFC2119](#)].

1.1 The case for IPv6 transition

As described in [NAT Complications] today's NAT devices are relatively successful at supporting TCP/UDP "client" applications which represented the bulk of Internet usage during the 1990s. These applications include Web browsing with HTTP and SSL, FTP, email and DNS. However, the current generation of NAT products has some unfortunate consequences on the users ability to deploy new applications, many of which follow a "peer-to-peer" model, and expect all "clients" to be also able to behave as "servers." Napster is a typical example of one such popular application: the peer-to-peer exchanges of music files cannot take place if both peers are located behind a NAT. With peer-to-peer applications such as NAPSTER now comprising more than 75 percent of Internet traffic in some locations, it has become clear that NAT devices are in danger of retarding the evolution of the Internet.

We believe that the proper solution to the NAT problem is to move towards IPv6. We realize that IPv6 cannot be turned on instantly, and thus we also believe that the interim solution will be to enable peer-to-peer applications behind NATs by deploying IPv6 on home networks, and linking these IPv6 islands together using IPv6 transition mechanisms such as 6to4. Unfortunately, not only does the continued deployment of the current generation of NAT devices make it more difficult to deploy new applications, but it will also make it difficult to handle the IPv6 transition. We therefore believe that action is required now in order to ensure that the generation of NATs that will be deployed over the next few years is IPv6-friendly.

Since we expect that the "IPv6 island" solution will eventually give way to widespread native IPv6 deployment, our approach is designed to be minimally intrusive. Rather than requiring large scale changes to NATs in the short term, we are requiring only a few modest changes to ensure IPv6 transition support. Also, rather than requiring modifications to existing host TCP/IP stacks, we only require minimal modifications to the applications.

2 Definitions

2.1 NAT

As defined in [[RFC2663](#)], Network Address Translation is a method by which IP addresses are mapped from one realm to another, in an attempt to provide transparent routing to hosts. Traditionally, NAT devices are used to connect an isolated address realm with private unregistered addresses to an external realm with globally unique registered addresses.

2.2 Global IPv4 Internet

We use the term "global IPv4 Internet" to designate the fraction of the Internet that uses globally unique IP addresses, and where connectivity to all globally unique addresses is expected.

2.3 Private network

We use the term "private network" to designate a network that uses private addresses as defined in [[RFC1918](#)], and that usually connects to the global IPv4 Internet through a NAT device.

2.4 Global IPv6 Internet

We use the term "global IPv6 Internet" to designate a network of nodes that are use globally unique IPv6 addresses, and where connectivity to all globally unique addresses is expected.

3 Model, requirements

The goal of this document is to enable easy deployment of IPv6 in private networks that are connected to the "global IPv4 Internet" through a NAT box. The connection can be performed in one of two ways:

* The gateway device can support native IPV6 so that it performs IPV6 routing functions in parallel to the IPv4 Network Address Translation function.

* Hosts inside the private network can set up automatic tunnels to reach the IPV6 Internet, using the 6to4 transition mechanism [[RFC3056](#)].

The preferred solution is to "upgrade" the connection device, so that it performs IPV6 routing functions in parallel to the IPv4 address translation function. The second solution is to allow hosts inside the private network to set up automatic tunnels to reach the global IPv6 internet using the 6to4 technology.

The goal of supporting these mechanisms is to enable the interim

deployment of "peer-to-peer" applications behind NAT. These applications could possibly be built using either TCP or UDP. The ideal solution would be to enable hosts in private networks to publish a set of global IP addresses and port at which they can receive TCP connection requests and UDP datagrams.

4 Description of the solution

The proposed solution provides IPv6 support either through a local implementation or through a transparent relay.

Our assumption is that IPv6 will be initially deployed by means of "tunnels" over the existing IPv4 infrastructure. The "6to4" strategy [[RFC3056](#)] allows users to transform a single IPv4 address into an IPv6 prefix; an almost unlimited number of stations can then obtain globally routable addresses using this prefix.

In a NAT environment, this can be instantiated handled in two ways:

- 1) An IPv6 capable NAT implements the "6to4 router" functionality, and appears as an IPv6 router to the local PCs,
- 2) A NAT that has no knowledge of IPv6 transparently passes the IPv4 packets that encapsulate IPv6 traffic to a local PC, which will in turn act as an IPv6 router for the local network.

We discuss the requirements for each of these approaches in turn.

4.1 Requirements for NATs implementing native IPv6

A NAT device can support IPv6 by providing the 6to4 relay functionality. In this case, the NAT will construct a 6to4 prefix from one of the global IPv4 addresses that it manages, and will advertise this prefix to the local network. A NAT that has obtained connectivity to the global Internet by other means than "6to4", will advertise the correspondent IPv6 prefix to the local network.

If the NAT product chooses to implement IPv6, it should do so according to the relevant IETF standards, including the IPv6 Specification [[RFC2460](#)], Neighbor Discovery [[RFC2461](#)], IPv6 Stateless Address Autoconfiguration [[RFC2462](#)], and ICMPv6[RFC 2463].

4.2 Requirements for NATs supporting transparent IPv6 tunneling

However, in the short term it is likely that NAT devices will only fully implement IPv4, and thus will not be capable of routing IPv6 natively. In this case, it is required that the NAT device enables the hosts behind the NAT to utilize IPv6. This can be done if they can support transparent tunneling of IPv6 packets.

The tunneling of IPv6 packets into IPv4 is defined in [[RFC2893](#)]. The

tunneled packets are identified by the protocol type "41" in the IPv4 header. NAT devices that do not implement the "6to4" relaying functions will **MUST** provide the transparency relay function as follows:

- 1) Define a local variable, LOCAL-IPV6-ROUTER holding the IPv4 address of the "local IPv6 router." This variable is initially set to the null address, 0.0.0.0.
- 2) When a packet is sent from a local host for a remote destination, specifying protocol type 41, copy the address of the local host into the LOCAL-IPV6-ROUTER variable. Replace the source address by the external address of the NAT.
- 3) When a packet is sent from a remote source to the global address managed by the NAT for protocol type 41, check the value of the LOCAL-IPV6-ROUTER variable. If the value is null, the NAT **MUST** reject the packet. Otherwise, the NAT **MUST** replace the destination address by the value of the LOCAL-IPV6-ROUTER variable, and relay the packet to the corresponding local host.

This assumes that the NAT manages a single external address. Where it manages several addresses, it the NAT **SHOULD** pick one of these addresses as the preferred address for IPv6, and behave as if only this address is available.

5 Discussion of the solution

Implementing an IPv6 relay in the NAT device obviously enables local hosts to use IPv6. Transparent IPV6 tunneling also enables IPv6, if one of the hosts is designated as the IPv6 route implementing [[RFC3056](#)], or if several hosts cooperate using the service.

5.1 Single 6to4 router

The simplest way to deploy IPv6 behind a NAT providing transparent tunneling is to select one of the local hosts to act as a 6to4 relay. This host will have to discover the global address used by the local NAT, construct a 6to4 prefix based on that address, and act as an IPv6 router for the local network. The global address can be discovered either through an interaction with the local NAT, or with the help of a server that has access to the global IPv4 Internet; specifying these mechanisms is outside the scope of this memo.

The selected router will have to ensure that it sends at least one IPv6 packet to an external target before it can receive tunneled packets. Sending one packet will ensure that the address of the selected router will be copied by the NAT in the LOCAL-IPV6-ROUTER variable, and that the selected router will receive the IPv6 packets sent by external hosts.

5.2 Cooperation between multiple hosts

It is possible to avoid the selection of a specific router by letting several hosts on the private network act as cooperating 6to4 relays. Each of these hosts will discover the global address used by the local NAT, construct a 6to4 prefix based on that address, and act as an IPv6 router for the local network. Any of these routers may receive tunneled packets; they must all be ready to relay over the private network the packets that are bound to other hosts.

According to the specified algorithm, tunneled IPv6 packets will be forwarded to the last router that sent an encapsulated IPv6 packet to an external node. If all active routers send packets at regular intervals, this ensures that the packets will be sent by the NAT to an active router, rather than possibly being sent in a black hole created by a failing router.

6 Security Considerations

6.1 The generic security risks of 6to4 tunneling and the appropriate protections are discussed in [[RFC3056](#)]. The transparent IPv6 tunneling option introduces an additional vulnerability, since a rogue host on the private network could send tunneled packets at regular intervals, be perceived by the NAT as the selected router, and use this in a denial of service attack. To protect against this vulnerability, the administrators of private networks must ensure that the local hosts adopt proper behavior.

7 IANA Considerations

None.

8 Copyright

The following copyright notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the applicable copyright for this document.

Copyright (C) The Internet Society XXX 0, 0000. All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be

followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

9 Intellectual Property

The following notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the position of the IETF concerning intellectual property claims made against this document.

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use other technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

10 References

[RFC2460] S. Deering, R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. [RFC 2460](#), December 1998.

[RFC2461] T. Narten, E. Nordmark, W. Simpson. Neighbor Discovery for IP Version 6 (IPv6). [RFC 2461](#), December 1998.

[RFC2462] S. Thomson, T. Narten. IPv6 Stateless Address Autoconfiguration. [RFC 2462](#), December 1998.

[RFC2463] A. Conta, S. Deering. Internet Control Message Protocol (ICMPv6) for the Internet. [RFC 2463](#), December 1998.

[RFC2119] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. [RFC 2119](#), March 1997.

[RFC2663] P. Srisuresh, M. Holdrege. IP Network Address Translator (NAT) Terminology and Considerations. [RFC 2663](#), August 1999.

[RFC2893] R. Gilligan, E. Nordmark. Transition Mechanisms for IPv6 Hosts and Routers. [RFC 2893](#), August 2000.

[RFC1918] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot & **E. Lear**. **Address Allocation for Private Internets**. [RFC 1918](#), February 1996.

[RFC3056] B. Carpenter, K. Moore. Connection of IPv6 Domains via IPv4 Clouds. [RFC 3056](#), February 2001.

[NAT Complications] M. Holdrege, P. Srisuresh. Protocol Complications with the IP Network Address Translator. Work in Progress.

11 Authors' Addresses

Christian Huitema
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Email: huitema@microsoft.com