INTERNET DRAFT C. Huitema, Microsoft
<draft-huitema-natreq4udp-00.txt>
Expires August 13, 2001 February 14, 2001

Short term NAT requirements for UDP based peer-to-peer applications

Status of this memo

Abstract

During the next few years, as the IPv4 address space moves toward exhaustion, it is likely that the deployment of NAT will accelerate. This draft presents the requirements that NAT devices must meet in order to enable use of UDP by peer-to-peer applications. The requirement can be summed up by the need to avoid gratuitous filtering and too short timers.

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.

This draft presents the requirements that NAT devices must meet in order to enable UDP based peer-to-peer applications during the transition to IPv6. Rather than specifying every aspect of a NATÆs 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.

1.1 Requirements language

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

1.2 Enabling peer-to-peer applications

As described in [NAT Complications] todayÆs NAT devices are relatively successful at supporting TCP/UDP ôclientö applications which represented the bulk of Internet usage during the 1990s. These applications include Web browsing with HTTP and SSL, FTP, email and DNS. The current generation of NAT products has some unfortunate consequences on the users ability to deploy new applications, many of which follow a ôpeer-to-peerö model, and expect all ôclientsö to be also able to behave as ôservers.ö 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 that during this deployment even the IPv6 enabled hosts will have to continue using IPv4 to reach those hosts that are not yet IPv6 enabled. We believe that very simple restrictions on the way NAT perform mappings for UDP applications can greatly ease the deployment of those peer-to-peer applications that rely on UDP.

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 ôglobal IPv4 Internetö 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 ôprivate networkö 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.

3 Model, requirements

Experience shows that the implementers of NAT products can adopt widely different treatments of UDP packets:

* Some implement the simplest solution, which is to map an internal UDP port, defined by an internal address and a port number on the corresponding host, to an external port, defined by a global address managed by the NAT and a port number valid for that address. In this simple case, the mapping is retained as long as the port is active, and is removed after an inactivity timer. As long as the mapping is retained, any packet received by the NAT for the external port is relayed to the internal address and port.

* Some implement a more complex solution, in which the NAT device not only establishes a mapping for the UDP port, but also maintain a list of external hosts to which traffic has been sent from that port. The packets originating from third party hosts to which the local host has not yet sent traffic are rejected.

* Instead of keeping just a list of authorized hosts, some NAT devices keep a list of authorized host and port pairs. UDP packets coming from remote addresses are rejected if the internal host has not yet sent traffic to the outside host and port pair.

* Finally, some NAT device map the same internal address and port pair to different external address and port pair, depending on the address of the remote host.

The most complex implementations break many application scenarios that use ôtriangular routingö: a Host A sends a packet to a server B; the server B forwards the query to a secondary server C; C responds directly to A. Such scenarios are quite common in real-time communication applications.

The goal of this document is to enable easy deployment of UDP-based peer-to-peer applications in private networks that are connected to the ôglobal IPv4 Internetö through a NAT device.

4 Description of the solution

The following requirements apply to NAT devices establishing mappings of private hosts addresses and UDP ports to global host addresses and ports:

1) The NAT MUST maintain UDP mappings for at least MIN_INTERVAL minutes. By default, MIN_INTERVAL = 3.

2) The NAT SHOULD use the same external address and external port for all current UDP session involving a given internal address and a given port, regardless of the third party address and port

involved in the session,

3) The NAT MUST NOT filter incoming traffic on the basis of the third party address and port.

The intent of these requirements is for a NAT to only implement the simple mapping of an internal address and port pair to a corresponding external address and port pair. The NAT SHOULD retain that mapping as long as the port is active; an inactivity timer of at least INACTIVITY_MIN minutes is required, where by default INACTIVITY_MIN = 2. NATs SHOULD NOT filter UDP traffic based on its external origin, and SHOULD use the same mapping for all external destinations. These rules have the following results:

1) If an internal host A sends UDP packets from its internal address and port AI:PA to an external host B through a NAT N, the packets will appear to B to originate from the address and port NA:PNA, i.e. the address and port chosen by N for this session.

2) If A sends UDP packets from the same address and port AI:PA to another host C, the packets will appear to C to come from the same address and port, NA:PNA.

3) If an external D sends a packet to the address and port pair NA:PNA before the mappings have expired, the packets will be forwarded by N to the internal address and port AI:PA.

A discussion of the rationale for this choice is provided in the next section.

5 Discussion of the solution

The rationale for the complex filtering is that by being as restrictive as possible into what the NAT will accept, they offer maximum protection to the internal host against external attacks. However, the protection value is somewhat limited, since port mappings are chosen at random and are hard to guess by third parties. If these parties can observe the port number by getting a copy of a packet, they can also get the address of the authorized peer, and spoof that address in their attacks. Indeed, the port could also be found through a port scan, but there should be other ways to protect against port scans than disabling peer-to-peer applications. In any case, the attacks that can be exercised through access to a single application port should be minimal, if the application is well programmed.

While the complex options may provide some limited additional security, they also disable many application scenarios, which is not a good trade-off. On balance it is better to be less restrictive.

6 Future Work

In order to enable hosts in private domains to receive TCP connections, we must provide a well adopted standard. This work is undertaken by the MIDCOM working group.

7 Security Considerations

Making UDP ports available is arguably more risky than restricted mappings. The trade-off between security and connectivity is discussed in section 5.

8 IANA Considerations

None.

9 Copyright

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

10 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.

11 References

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

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

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

12 Authors' Addresses

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

Email: huitema@microsoft.com