

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 5, 2009

D. Thaler
Microsoft
L. Zhang
UCLA
March 4, 2009

IAB Thoughts on IPv6 Network Address Translation
draft-iab-ipv6-nat-00.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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 September 5, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

There has been much recent discussion on the topic of whether the IETF should develop standards for IPv6 Network Address Translators

(NATs). This document articulates the architectural issues raised by IPv6 NATs, the pros and cons of having IPv6 NATs, and provides the IAB's thoughts on the current open issues and the solution space.

Table of Contents

1.	Introduction	3
2.	What is the Problem?	3
2.1.	Avoiding Renumbering	4
2.2.	Site Multihoming	4
2.3.	Topology Hiding	5
2.4.	Preventing Host Counting	5
2.5.	Simple Security	6
2.6.	Discussion	6
3.	Architectural Considerations of IPv6 NAT	6
4.	Solution Space	7
4.1.	Discussion	9
5.	Security Considerations	9
6.	IANA Considerations	9
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	10
	Authors' Addresses	10

1. Introduction

In the past, the IAB has published a number of documents relating to Internet transparency and the end-to-end principle, and other IETF documents have also touched on these issues as well. These documents articulate the general principles on which the Internet architecture is based, as well as the core values that the Internet community seeks to protect going forward. Most recently, [RFC 4924](#) [[RFC4924](#)] reaffirms these principles and provides a review of the various documents in this area.

Facing imminent IPv4 address space exhaustion, recently there have been increased efforts in IPv6 deployment. However, since late last year there have also been increased discussions about whether the IETF should standardize network address translation within IPv6. People who are against standardizing IPv6 NAT argue that there is no fundamental need for IPv6 NAT, and that as IPv6 continues to roll out, the Internet should converge towards reinstallation of the end-to-end reachability which has been a key factor in the Internet's success. On the other hand, people who are for IPv6 NAT believe that NAT vendors would provide IPv6 NAT implementations anyway as NAT can be a solution to a number of problems, and that the IETF should avoid repeating the same mistake with IPv4 NAT, where the lack of protocol standards led to different IPv4 implementations, making NAT traversal difficult.

An earlier effort, [[RFC4864](#)], provides a discussion of the real or perceived benefits of NAT and suggests alternatives for most of them, with the intent of showing that NAT is not required to get the desired benefits. However, it also identifies several gaps remaining to be filled.

This document provides the IAB's current thoughts on this debate. We believe that the issue in hand must be viewed from the overall architecture standpoint in order to fully assess the pros and cons of IPv6 NAT on the global Internet and its future development.

2. What is the Problem?

The discussions on the necessity for IPv6 NAT can be summarized as follows: network address translation is viewed as a solution to achieve a number of desired properties for individual networks: avoiding renumbering, facilitating multihoming, internal topology hiding, and in particular preventing host counting. We discuss below each of these perceived benefits from NAT.

2.1. Avoiding Renumbering

As discussed in [\[RFC4864\] Section 2.5](#), the ability to change providers with minimal operational difficulty is an important requirement in many networks. However, renumbering is still quite painful today, as discussed in [\[I-D.carpenter-renum-needs-work\]](#). Currently it requires reconfiguring devices that deal with IP addresses or prefixes, including DNS servers, DHCP servers, firewalls, IPsec policies, and potentially many other systems such as intrusion detection systems, inventory management systems, patch management systems, etc.

In practice today, renumbering does not seem to be a significant problem in consumer networks, such as home networks, where addresses or prefixes are typically obtained through DHCP, and are rarely manually configured in any component. However in managed networks, renumbering can be a serious problem.

We also note that many, if not most, large enterprise networks avoid the renumbering problem by using provider-independent (PI) IP address blocks. The use of PI addresses is inherent in today's Internet operations. However in smaller managed networks that cannot get provider-independent IP address blocks, renumbering remains a serious issue. Regional Internet Registries (RIRs) constantly receive requests for PI address blocks; one main reason that they hesitate in assigning PI address blocks to all users is the concern about the PI addresses' impact on the routing system scalability.

2.2. Site Multihoming

Another important requirement in many networks is site multihoming. A multihomed site essentially requires that its prefixes be present in the global routing table to achieve the desired reliability in its Internet connectivity as well as load balancing. In today's practice, multihomed sites with PI addresses announce their PI prefixes to the global routing system; multihomed sites with provider-allocated (PA) addresses also announce the PA prefix they obtained from one provider to the global routing system through another provider, effectively disabling provider-based prefix aggregation. This practice makes the global routing table scale linearly with the number of multihomed user networks.

This issue was identified in [\[RFC4864\] Section 6.4](#). Unfortunately, no solution except NAT has been deployed today that can insulate the global routing system from the growing number of multihomed sites, where a multihomed site simply assigns multiple IPv4 addresses, one from each of its providers, to its exit router which is a IPv4 NAT box. Using address translation to facilitate multihoming support has

one unique advantage: there is no impact on the routing system scalability, as the NAT box simply takes one address from each provider, and the multihomed site does not inject its own routes into the system. Intuitively it also seems straightforward to roll the same solution into multihoming support in the IPv6 deployment.

However it is important to point out that a multihomed site announcing its own PI prefix(es) achieves important benefits that NAT-based multihoming support does not provide. Using PI addresses, end-to-end communications can be preserved in face of connectivity failures of individual providers, as long as the site remains connected through at least one operational provider. Announcing PI prefixes also gives a multihomed site the ability to perform traffic engineering and load balancing. While the users gain these benefits from PI-based multihoming, we also note that, in today's routing system, these gains are at the cost of the increased routing table size for all providers.

2.3. Topology Hiding

It is perceived that a network operator may want to hide the details of the network topology, the size of the network, the identities of the internal routers, and the interconnection among the routers. This desire has been discussed in [[RFC4864](#)] Sections [4.4](#) and [6.2](#).

However it remains an open question of whether one can successfully hide the network topology through address translation. Even if one can hide the actual addresses of internal hosts through address translation, this does not necessarily prove sufficient to hide internal topology. It may be possible to infer some aspects of topological information from passively observing packets, for example based on packet timing, the Hop Limit field, or other fields in the packet header. An even bigger open question is what quantifiable benefits topology hiding can bring, which will determine how much cost one is willing to pay to achieve it.

2.4. Preventing Host Counting

As a specific measure of topology hiding, preventing an external entity from counting the number of hosts on one's network is another perceived benefit of NAT. Like topology hiding in general, it remains to be seen how important this issue may really be in practice.

Even where host counting is a concern, it is worth pointing out some of the challenges in preventing it. In [[Bellovin](#)], Bellovin showed how one can successfully count the number of hosts behind a NAT box. Although that paper uses the IPid field to count IPv4 hosts, the

point applies more generally, in that if any fields, such as fragment ID, TCP initial sequence number, or ephemeral port number, are chosen in a predictive fashion (e.g., sequentially), then an attacker may correlate packets or connections coming from the same host.

2.5. Simple Security

It is commonly perceived that a NAT box provides one level of protection because external hosts cannot directly initiate communication with hosts behind a NAT. As discussed in [\[RFC4864\] Section 2.2](#), the act of translation does not provide security in itself, but rather the lack of pre-established or permanent filtering state. The stateful filtering function can provide the same level of protection without requiring a translation function. For further discussion, see [\[RFC4864\] Section 4.2](#).

2.6. Discussion

At present, the primary benefits one may receive from deploying NAT appear to be avoiding renumbering and supporting multihoming without impacting routing scalability. Topology hiding and host counting may be concerns to some, however solving them calls for further research as address translation may or may not be an effective solution. Furthermore, their quantifiable benefit is yet to be understood, and one must compare that benefit against the potential costs.

3. Architectural Considerations of IPv6 NAT

The discussions on IPv6 NAT often refer to the wide deployment of IPv4 NAT, where people have both identified tangible benefits and gained operational experience. However the discussions so far seem mostly focused on the potential benefits that IPv6 NAT may, or may not, bring. Little attention has been paid to the bigger picture, as we elaborate below.

When considering the benefits that IPv6 NAT may bring to a site that deploys it, we must not overlook a bigger question: if one site deploys IPv6 NAT, what is the potential impact it brings to the rest of the Internet that does not do IPv6 NAT? This important question does not seem to have been addressed, or addressed adequately.

We believe that the discussions on IPv6 NAT should be put in the context of the overall Internet architecture. The foremost question is not how many benefits one may derive from using IPv6 NAT, but more fundamentally, whether a significant portion of parties on the Internet are willing to deploy IPv6 NAT, and hence whether we want to make IP address translation a permanent block in the Internet

architecture.

One may argue that the answers to the above questions depend on whether we can find adequate solutions to the renumbering and multihoming problems. It is worthwhile pointing out that IPv6 NAT is not the only solution to these two problems. Renumbering can be avoided by allocating to users provider-independent addresses. Multihoming is already a pervasive practice today, not some new feature to be supported in the future, and NAT-based multihoming has serious limitations as discussed earlier. The real issue is not multihoming per se, but the need for a scalable routing system.

If the answer to the above two questions is no, then non-IPv6-NAT parts of the world should **not** be affected by those sites that want to deploy IPv6 NAT. More specifically, IPv6 users should be able to reach each other directly, without having to worry about address translation boxes between the two ends. IPv6 application developers in general should be able to program based on the assumption of end-to-end reachability, without having to address the issue of traversing NAT boxes. Similarly, network operators should be able to run their networks without the added complexity of NATs, which can bring not only the cost of additional boxes, but also increased difficulties in network monitoring and problem debugging.

Given the diversity of the Internet user populations and the diversity in today's operational practice, it is conceivable that some parties may have a strong desire to deploy IPv6 NAT, and the Internet should accommodate different views that lead to different practices (i.e., some using IPv6 NAT, others not).

If we accept the view that some, but not all, parties want IPv6 NAT, then the real debate should not be on what benefits IPv6 NAT may bring. As every coin has two sides, it is undeniable that network address translation can bring certain benefits to its users. However the real challenge we should address is how to design IPv6 NAT in such a way that it can hide its impact within some localized scope. If IPv6 NAT design can achieve this goal, then the Internet as a whole can strive for (re-installing) the end-to-end reachability model.

4. Solution Space

From an end-to-end perspective, the solution space for renumbering and multihoming can be broadly divided into three classes:

1. Endpoints get a stable, globally reachable address: In this class of solution, end sites use provider-independent addressing and hence endpoints are unaffected by changing ISPs. For this to be a complete solution, provider-independent addressing must be available to all managed networks (i.e., all networks that use manual configuration of addresses or prefixes in any type of system). However in today's practice, assigning provider-independent addresses to all networks, including small ones, raises concerns with the scalability of the global routing system. This is an area of ongoing research and experimentation. In practice, operators have also been developing short-term approaches to resolve today's gap between the continued routing table growth and limitations in existing router capacity [[NANOG](#)].
2. Endpoints get a stable but non-globally-routable address on physical interfaces but a dynamic, globally routable address inside a tunnel: In this class of solution, hosts use locally-scoped (and hence provider-independent) addresses for communication within the site. As a result, managed systems such as routers, DHCP servers, etc. all see stable addresses. Tunneling from the host to some infrastructure device is then used to provide the host with globally routable addresses which may change, but address changes are constrained to systems that operate over or beyond the tunnel, including DNS servers and applications. These systems, however, are the ones that often can already deal with changes today using mechanisms such as DNS dynamic update. However, if endpoints and the tunnel infrastructure devices are owned by different organizations, then solutions are harder to incrementally deploy due to the incentive and coordination issues involved.
3. Endpoints get a stable address which gets translated in the network: In this class of solution, end sites use non-globally-routable addresses within the site, and translate them to globally routable addresses somewhere in the network. In general, this causes the loss of end-to-end transparency which is the subject of [[RFC4924](#)] and the documents it surveys. If the translation is reversible, and the translation is indeed reversed by the time it reaches the other end of communication, then end-to-end transparency can be provided. However if the two translators involved are owned by different organizations, then solutions are harder to incrementally deploy due to the incentive and coordination issues involved.

Concerning routing scalability, although there is no immediate danger, routing scalability has been a long time concern in operational communities, and an effective and deployable solution must be found. We observe that the question in hand is not about whether some parties can run NAT, but rather, whether the Internet as a whole would be willing to rely on NAT to curtail the routing

scalability problem, and whether we have investigated all the potential impacts of doing so to understand its cost on the overall architecture. If effective solutions can be deployed in time to allow assigning provider-independent IPv6 addresses to all user communities, the Internet can avoid the complexity and fragility and other unforeseen problems introduced by NAT.

4.1. Discussion

As [[RFC4924](#)] states:

A network that does not filter or transform the data that it carries may be said to be "transparent" or "oblivious" to the content of packets. Networks that provide oblivious transport enable the deployment of new services without requiring changes to the core. It is this flexibility that is perhaps both the Internet's most essential characteristic as well as one of the most important contributors to its success.

We believe that providing end-to-end transparency, as defined above, is key to the success of the Internet. While some fields of traffic (e.g., Hop Limit) are defined to be mutable, transparency requires that fields not defined as such arrive un-transformed. Currently, the source and destination addresses are defined as immutable fields, and are used as such by many protocols and applications.

Each of the three classes of solution can be defined in a way that preserves end-to-end transparency. We strongly encourage the community to consider end-to-end transparency as a requirement when proposing any solution. Solutions can then be compared based on other aspects such as scalability and ease of deployment.

5. Security Considerations

[Section 2](#) discusses potential privacy concerns as part of the Host Counting and Topology Hiding problems.

6. IANA Considerations

[RFC Editor: please remove this section prior to publication.]

This document has no IANA Actions.

7. References

[7.1.](#) Normative References

[7.2.](#) Informative References

[Bellovin]

Bellovin, S., "A Technique for Counting NATted Hosts",
Proc. Second Internet Measurement Workshop ,
November 2002,
<<http://www.cs.columbia.edu/~smb/papers/fnat.pdf>>.

[I-D.carpenter-renum-needs-work]

Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering
still needs work", [draft-carpenter-renum-needs-work-02](#)
(work in progress), February 2009.

[NANOG]

"Extending the Life of Layer 3 Switches in a 256k+ Route
World", NANOG 44 , October 2008, <[http://www.nanog.org/
meetings/nanog44/presentations/Monday/
Roisman_lightning.pdf](http://www.nanog.org/meetings/nanog44/presentations/Monday/Roisman_lightning.pdf)>.

[RFC3041]

Narten, T. and R. Draves, "Privacy Extensions for
Stateless Address Autoconfiguration in IPv6", [RFC 3041](#),
January 2001.

[RFC4864]

Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and
E. Klein, "Local Network Protection for IPv6", [RFC 4864](#),
May 2007.

[RFC4924]

Aboba, B. and E. Davies, "Reflections on Internet
Transparency", [RFC 4924](#), July 2007.

Authors' Addresses

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone: +1 425 703 8835
Email: dthaler@microsoft.com

Lixia Zhang
UCLA Computer Science Department
3713 Boelter Hall
Los Angeles, CA 90095
USA

Phone: +1 310 825 2695
Email: lixia@cs.ucla.edu