Network Working Group	R. Denis-Courmont
Internet-Draft	Nokia
Intended status: Informational	February 18, 2009
Expires: August 22, 2009	

Problems with IPv6 source address selection and IPv4 NATs
draft-denis-v6ops-nat-addrsel-00

Status of This Memo

Copyright Notice

Abstract

This memo details a problem and potential solution, when using the IPv6 source address selection algorithm with private IPv4 address space.

1. Introduction

When a host initiates an IP communication flow with a remote host, a pair of local and remote IP addresses to use must be chosen. If either or both hosts is assigned multiple IP addresses, an address selection mechanism is required. That can happen, for instance, if either or both hosts are dual-stacked. The default address selection scheme[RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.) was specified to address this problem. One fundamental design assumption of this scheme is the ability to determine a scope for any address (IPv4 or IPv6). To that end, static scoping rules were defined. This memo explains why and how the current rules are inadequate when Network Address Translation (NAT) is involved, which is a common occurence in modern-day IPv4 deployments.

2. IPv4 address scopes

As defined in [RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.), a unicast IPv4 address has one of three scopes:

> **Link-local scope:** Loopback addresses (127.0.0.0/8) and autoconfigured addresses (169.254.0.0/16).
>
> **Site-local scope:** Private addresses as defined in [RFC1918] (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, “Address Allocation for Private Internets,” February 1996.).
>
> **Global scope:** All other unicast addresses.

The address scopes are supposed to be universal; and hence they are statically defined. Furthermore, per [RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.), scope matching rules (Rule 2) are normally applied before any other rule, except for the identical address rule (Rule 1). In other words, apart from the corner case whereby the local and remote hosts are one and the same, the scope matching rule always "wins".

3. NAT and address scope

When it crosses a NAT, either the source or destination address of a packet will change. As a consequence, the scope of that address might change as well. In any case, the result of the source address selection scheme could be different when the original address is substitued with the translated address.

In fact, many real-world NAT deployments use private addresses on one side of the NAT, and public addresses on the other side. This is probably the most common scenario with IPv4 network in SOHO environment: a single public IPv4 address is provisioned to a customer, and all hosts on the customer network "share" that address using a NAT function within the Customer Premises Equipment (CPE).

[RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.) assumes that a source address with a small scope cannot reach a destination address with a larger scope. However, if private IPv4 addresses and a NAT are used to reach public IPv4 addresses, then this assumption does not hold. In other words, the private IPv4 addresses behind NATs effectively have a global scope, provided that the protocols above the IP network layer can cope with network address translation.

4. Applicability to IPv6 transition mechanisms

[RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.) states that "the use of transitional addresses when native addresses are available [should be avoided]". Indeed, transitional addresses and transition mechanisms in general tend to be less reliable than native connectivity, including native IPv4 connectivity.

However, in a typical IPv4 NAT'ed private address deployments, if IPv6 transition mechanisms are available, a dual-stack host will typically have the following addresses. They are the candidate source addresses:

*a link-local IPv6 address (autoconfigured),

*a site-local scope private IPv4 address (e.g. assigned by DHCPv4),

*a global scope transitional IPv6 address, such as Teredo[RFC4380] (Huitema, C., “Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs),” February 2006.), or 6to4[RFC3056] (Carpenter, B. and K. Moore, “Connection of IPv6 Domains via IPv4 Clouds,” February 2001.) (e.g. if the CPE is a 6to4 gateway).

If the destination host is also dual-stacked, then it will typically have two public addresses (though the number is not relevant). They are the candidate destination addresses:

*a global native IPv6 address (e.g. from DNS AAAA record),

*a global IPv4 address (e.g. from DNS A record).

Because the candidate source IPv4 address have a smaller scope (site-local) than the candidate destination IPv4 address (global), it will be eliminated. The address selection algorithm will always select the IPv6 address pair:

- the transitional IPv6 address as source,

- the global native IPv6 address as destination.

Thus, the transitional (IPv6) address will be used instead of the native (IPv4) address, even though that should have been avoided.
There is no way to override this result with a compliant implementation of source address selection. In particular, the policy table does not affect this result, because the scope rules preempt the policy table rules.

5. Solutions

5.1. Changing the private IPv4 address scope

Several operating system vendors appear to work around this issue by assigning a global scope to IPv4 address. Thus, rule 2 is no longer discriminating against the IPv4 address pair.
In that case, provided the policy table has separate labels for transitional addresses, the IPv4 addresses pair will be selected. IPv4 addresses normally all have the same label.
Note that the default policy table has a separate label for 6to4 addresses. However, as it predates Teredo, it lacks a distinct label for the Teredo prefix, 2001:0:/32. An adequate extra label would be as follow:
Prefix: 2001:0:/32, Precedence: 5, Label: 5

5.2. Address selection parameter for NAT

With the previous solution, IPv4 is always selected. This is a potential drawback if the upper-layer protocol combination is not NAT-friendly.
As an alternative, a "translation-friendly" source address selection parameter could be specified, as in [RFC5014] (Nordmark, E., Chakrabarti, S., and J. Laganier, “IPv6 Socket API for Source Address Selection,” September 2007.). However, a default value will be needed

for the many existing applications that would fail to set this parameter.

6. IPv6 Address Translation

The implications of IPv6 Address Translation and protocol translation are left beyond the scope of this document. However, it can only be recommended that RFC3484 be taken into account when designing such translation systems.

7. Security Considerations

TBD.

8. IANA Considerations

This document raises no IANA considerations.

9. References

9.1. Normative References

[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, “Address Allocation for Private Internets,” BCP 5, RFC 1918, February 1996 (TXT).
[RFC3484]	Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” RFC 3484, February 2003 (TXT).
[RFC5014]	Nordmark, E., Chakrabarti, S., and J. Laganier, “IPv6 Socket API for Source Address Selection,” RFC 5014, September 2007 (TXT).

9.2. Informative References

[RFC3056]	Carpenter, B. and K. Moore, “Connection of IPv6 Domains via IPv4 Clouds,” RFC 3056, February 2001 (TXT).
[RFC4380]	Huitema, C., “Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs),” RFC 4380, February 2006 (TXT).

Author's Address

	Rémi Denis-Courmont
	Nokia Corporation
	P.O. Box 407
	NOKIA GROUP 00045
	FI
Phone:	+358 50 487 6315
EMail:	remi.denis-courmont@nokia.com