Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 30, 2007

M. Blanchet
Viagenie

February 26, 2007

IPv6 Routing Policies Guidelines
draft-ietf-v6ops-routing-guidelines-01.txt

Status of this Memo

Copyright Notice

Abstract

Guidelines on how to manage and filter IPv6 routes are needed for operators of networks, either providers or enterprises. It describes IPv6 routes from the protocol point of view. It does not discuss operational or policy issues such as the maximum length of prefixes to filter. This document is a followup on RFC2772 work but for the production IPv6 Internet. RFC2772 is obsoleted.

Table of Contents

1. Introduction

To maintain stability, efficiency and scalability of the IPv6 Internet, guidelines for routing policies are needed for operators deploying IPv6 networks. Prior experience on IPv6 routing guidelines on the 6bone[RFC2772], practical deployment of the IPv6 internet and IPv6 specifications were used as input to this document.

This document first describes the different types of addresses and then summarizes the suggested policies in RPSL.

"Advertisement" in this document refers to the prefix advertisement, not the next-hop.

2. Address Types

2.1. Node-scoped Unicast

The node-scoped unicast addresses[RFC4291] such as the loopback (::1/128), the unspecified (::/128) must not be advertised in an IGP or EGP and should be filtered out when received.

2.2. IPv4-Mapped Addresses

IPv4-mapped addresses (::FFFF:0:0/96) [RFC4291] must not be advertised and should be filtered out.

2.3. Link-scoped Unicast

The link-scoped unicast[RFC4291] routes (fe80::/10) must not be advertised in an IGP or EGP and should be filtered out when received.

2.4. Site-scoped Unicast

The site-scoped unicast routes, known as Unique-local[RFC4193], (fc00::/7) may be advertised in an IGP. It must not be advertised in an EGP connected to the global Internet and should be filtered out when received. However, it may be advertised in an EGP between two networks sharing a private interconnect, but must not be advertised outside the scope of these networks. When advertised in an EGP, these routes should be of length /48 or smaller.

2.5. Global Unicast

The global unicast routes (2000::/3) [RFC4291] may be advertised in an IGP or EGP.

A minimal EGP routing policy should filter out routes that exceed a maximum length. Determining the maximum length of a global Internet route is outside the scope of this document.

A finer EGP routing policy may use only the allocated address space from IANA to registries as specified in http://www.iana.org/assignments/ipv6-unicast-address-assignments. This would result in better filtering since the non-allocated prefixes will be filtered out.

An even finer EGP routing policy may use only the assigned address space from registries to providers as available in the registries databases. This would result in the best filtering since the non-assigned prefixes will be filtered out. However, this requires the synchronization of the filters with the registries databases.

2.5.1. Documentation Prefix

The 2001:0db8::/32 prefix[RFC3849] is used for documentation purposes and must not be advertised in an IGP or EGP and should be filtered out when received.

2.5.2. 6to4

The 6to4[RFC4291][RFC3056] prefix (2002::/16) may be advertised in an IGP or EGP, when the site is running a 6to4 relay or offering a 6to4 transit service. However, the provider of this service should be aware of the implications of running such service[RFC3964], which includes some specific filtering rules for 6to4.

2.5.3. Teredo

The Teredo[RFC4380] prefix (2001::/32) may be advertised in an IGP or EGP, when the site is running a Teredo relay or offering a Teredo transit service.

2.5.4. 6bone

The 6bone experimental network used some experimental allocations, such as 5f00::/8[RFC1897] and 3ffe::/16[RFC2471] that were later returned to IANA[RFC3701]. These prefixes should not be advertised in an EGP unless IANA reallocates them subsequently.

2.6. Default Route

The default unicast route (::) may be advertised in an IGP. It must not be advertised in an EGP unless it has been requested by the recipient.

2.7. Multicast

Multicast addresses (ff00::/8) [RFC4291] have a 4 bits scope in the address field. Only addresses having the 'E' value in the scope field are of global scope, all other values are local or reserved. Therefore, only ffXe:: routes may be advertised outside an organisation network, where X may be any value.

Multicast routes must not appear in unicast routing tables.

2.8. Unknown addresses

Any non listed address above must not be advertised and should be filtered out. Future work might reserve additional address space for protocol use which might require specific routing guidelines. The reader should refer to newer versions of the normative references in this document to verify the existence of newer protocol address space.

3. Implementing routing policies

This document focuses on protocol addresses and their use in the networks. It does not discuss any allocation policies and their impact on the routing policies, such as /48 Micro-allocations for infrastructure providers or maximum length of a unicast prefix. As such, to implement a complete routing policy, one should augment these guidelines with the current registry allocation policies and by appropriate ingress filtering techniques[RFC3704].

4. RPSL Implementation

The Route Policy Specification Language(RPSL) [RFC4012] used in route registries supports the policies described in this document and should be considered to manage route policies.

The following RPSL code implements the policies described in this document. This code should be considered as an example and should be adapted to the target usage.

```
route-set: rs-exclude
mp-members: ::1/128, ::/128, ::ffff:0:0/96^+, fe80::/10^+,
 2001:0db8::/32^+

route-set: rs-ula
mp-members: fc00::/7^+

route-set: rs-global-unicast
mp-members: 2000::/3^+

route-set: rs-6to4
mp-members: 2002::/16^+

route-set: rs-teredo
mp-members: 2001::/32^+

filter-set: fltr-v6egp
mp-filter: NOT (rs-exclude AND rs-ula) AND rs-global-unicast

filter-set: fltr-v6igp
mp-filter: NOT rs-exclude AND rs-global-unicast
```

5. Security Considerations

This document list guidelines that should improve the security of networks by the filtering of invalid routing prefixes.

6. Acknowledgements

Florent Parent, Pekka Savola, Tim Chown, Alain Baudot, Stig Venaas, Vincent Jardin, Olaf Bonness, David Green, Gunter Van de Velde, Michael Barnes, Fred Baker, Edward Lewis, Marla Azinger, Brian Carpenter, Mark Smith and Kevin Loch have provided input and suggestions to this document.

7. References

7.1. Normative References

[RFC1897] Hinden, R. and J. Postel, "IPv6 Testing Address Allocation", RFC 1897, January 1996.

[RFC2471] Hinden, R., Fink, R., and J. Postel, "IPv6 Testing Address Allocation", RFC 2471, December 1998.

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

[RFC3849] Huston, G., Lord, A., and P. Smith, "IPv6 Address Prefix Reserved for Documentation", RFC 3849, July 2004.

[RFC4012] Blunk, L., Damas, J., Parent, F., and A. Robachevsky, "Routing Policy Specification Language next generation (RPSLng)", RFC 4012, March 2005.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

7.2. Informative References

[RFC2772] Rockell, R. and B. Fink, "6Bone Backbone Routing Guidelines", RFC 2772, February 2000.

[RFC3701] Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout", RFC 3701, March 2004.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, December 2004.

Author's Address

Marc Blanchet
Viagenie

Email: Marc.Blanchet@viagenie.ca

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).