

6Bone Routing Practice

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

1. Introduction

The 6Bone is an environment supporting experimentation with the IPv6 protocols and products implementing it. As the network grows, the need for common operation rules emerged. In particular, operation of the 6Bone backbone is a challenge due to the frequent insertion of bogus routes by leaf or even backbone sites.

This memo identifies guidelines on how 6Bone sites might operate, so that the 6Bone can remain a quality experimentation environment and to avoid pathological situations that have been encountered in the past. It defines the 'best current practice' acceptable in the 6Bone for the configuration of both Interior Gateway Protocols (such as RIPng [RFC 2080]) and Exterior Gateway Protocols (like BGP4+ [RFC 2283]).

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119].

2. Basic principles

The 6Bone is structured as a hierarchical network with pseudo Top Level Aggregator (pTLA) sites, pseudo Next Level Aggregator (pNLA) sites and leaf sites. This topology supports the IPv6 address aggregation architecture as described in [1]. The 6Bone backbone is made of a mesh interconnecting pTLAs only. pNLAs connect to one or more pTLAs and provide transit service for leaf sites.

pTLA sites MUST use BGP4+ [RFC 2283] as the mandatory routing protocol for exchanging routing information among them.

Multi-homed sites or pNLAs SHOULD also use BGP4+. Regular sites MAY use a simple default route to their ISP.

3. Common Rules

This section details common rules governing the routing on the 6Bone. They are derived from issues encountered on the 6Bone, with respect to the routes advertised, handling of special addresses, and aggregation:

- 1) link local prefixes
- 2) site local prefixes
- 3) loopback prefix & unspecified prefix
- 4) multicast prefixes
- 5) IPv4-compatible prefixes
- 6) IPv4-mapped prefixes
- 7) default routes
- 8) Yet undefined unicast prefixes (from a different /3 prefix)
- 9) Inter site links issues
- 10) aggregation & advertisement issues

3.1 Link-local prefix

The link-local prefix (FE80::/10) MUST NOT be advertised through either an IGP or an EGP.

By definition, the link-local prefix has a scope limited to a specific link. Since the prefix is the same on all IPv6 links, advertising it in any routing protocol does not make sense and, worse, may introduce nasty error conditions.

Well known cases where link local prefixes could be advertised by mistake include:

- a router advertising all directly connected network prefixes including the link-local one.
- Subnetting of the link-local prefix.

In such cases, vendors should be urged to correct their code.

3.2 Site-local prefixes

Site local prefixes (in the FEC0::/10 range) MAY be advertized by IGPs or EGPs within a site. The precise definition of a site is ongoing work discussed in the IPng working group.

Site local prefixes MUST NOT be advertised to transit pNLAs or pTLAs.

3.3 Loopback and unspecified prefixes

The loopback prefix (::1/128) and the unspecified prefix (::0/128) MUST NOT be advertised by any routing protocol.

3.4 Multicast prefixes

Multicast prefixes MUST NOT be advertised by any unicast routing protocol. Multicast routing protocols are designed to respect the semantics of multicast and MUST therefore be used to route packets with multicast destination addresses (in the range FF00::/8).

Multicast address scopes MUST be respected on the 6Bone. Only global scope multicast addresses MAY be routed across transit pNLAs and pTLAs. There is no requirement on a pTLA to route multicast packets.

Organization-local multicasts (in the FF08::/16 or FF18::/16 ranges) MAY be routed across a pNLA to its leaf sites.

Site-local multicasts MUST NOT be routed toward transit pNLAs or pTLAs.

Obviously, link-local multicasts and node-local multicasts MUST NOT be routed at all.

3.5 IPv4-compatible prefixes

Sites may choose to use IPv4 compatible addresses (::a.b.c.d) internally. As there is no real rationale today for doing that, these addresses SHOULD

NOT be used in the 6Bone.

The `::/96` IPv4-compatible prefixes MAY be advertised by IGPs.

IPv4-compatible prefixes MUST NOT be advertised by EGPs to transit pNLAs or pTLAs.

3.6 IPv4-mapped prefixes

IPv4-mapped prefixes (`::FFFF:a.b.c.d` where `a.b.c.d` is an IPv4 address) MAY be advertised by IGPs within a site. It may be useful for some IPv6 only nodes within a site to have such a route pointing to a translation device.

IPv4-mapped prefixes MUST NOT be advertised by EGPs.

3.7 Default routes

6Bone core pTLA routers MUST be default-free.

pTLAs MAY advertise a default route to their pNLAs. Transit pNLAs MAY do the same for their leaf sites.

3.8 Yet undefined unicast prefixes

Yet undefined unicast prefixes from a format prefix other than `2000::/3` MUST NOT be advertised by any routing protocol in the 6Bone. In particular, RFC 2471 test addresses MUST NOT be advertised on the 6Bone.

Routing of global unicast prefixes outside of the 6Bone range (`3FFE::/16`) is discussed in section 4, Routing policies, below.

3.9 Inter-site links

Global IPv6 addresses MUST be used for the end points of the inter-site links. In particular, IPv4 compatible addresses MUST NOT be used for tunnels.

Prefixes for those links MUST NOT be injected in the global routing tables.

3.10 Aggregation & advertisement issues

Route aggregation MUST be performed by any border router.

Sites or pNLAs MUST only advertise to their upstream provider the prefixes assigned by that ISP unless otherwise agreed.

Site border router MUST NOT advertise prefixes more specific than the /48 ones allocated by their ISP.

pTLA MUST NOT advertise prefixes longer than 24 to other pTLAs unless special peering agreements are implemented. When such special peering agreements are in place between any two or more pTLAs, care MUST be taken not to leak the more specific prefixes to other pTLAs not participating in the peering agreement.

4. Routing policies

6Bone backbone sites maintain the mesh into the backbone and provide an as reliable as possible service, granted the 6Bone is an experimentation tool. To achieve their mission, 6Bone backbone sites MUST maintain peerings with at least 3 (three) other back bone sites.

The peering agreements across the 6Bone are by nature non-commercial, and therefore SHOULD allow transit traffic through.

Eventually, the Internet registries will assign other TLAs than the 6Bone one (currently 3FFE::/16). The organizations bearing those TLAs will establish a new IPv6 network, parallel to the 6Bone. The 6Bone MIGHT interconnect with this new IPv6 Internet, but transit across the 6Bone will not be guaranteed. It will be left to each 6Bone backbone site to decide whether it will carry traffic to or from the IPv6 Internet.

5. The 6Bone registry

The 6Bone registry is a RIPE-181 database with IPv6 extensions used to store information about the 6Bone. Each 6Bone site MUST maintain the relevant entries in the 6Bone registry (whois.6bone.net). In particular, the following objects MUST be present:

- IPv6-site: site description
- Inet6num: prefix delegation
- Mntner: coordinate of site maintenance staff

Other objects MAY be maintained at the discretion of the sites, such as routing policy descriptors, person or role objects. The Mntner object MUST make reference to a role or person object, but those must not necessarily reside in the 6Bone registry, they can be stored within any of the Internet registry databases (RIPE, InterNIC, APNIC,

6. Guidelines for new sites joining the 6Bone

New sites joining the 6Bone should seek to connect to a transit pNLA or a pTLA within their region, and preferably as close as possible to their existing IPv4 physical and routing path for Internet service. The 6Bone registry is available to find out candidate ISPs.

Any site connected to the 6Bone MUST maintain a DNS server for forward name looking and reverse address translation. The joining site MUST maintain the 6Bone registry objects relative to its site, and in particular the IPv6- site and the MNTNER objects.

The upstream ISP MUST delegate the reverse address translation zone in DNS to the joining site. The ISP MUST also create 6Bone registry objects reflecting the delegated address space (inet6num:).

Up to date information about how to join the 6Bone is available on the 6Bone Web site at <http://www.6bone.net>.

7. Guidelines for 6Bone pTLA sites

6Bone pTLA sites are altogether forming the backbone of the 6Bone. In order to ensure the highest level possible of availability and stability for the 6Bone environment, a few constraints are placed onto sites wishing to become or stay a 6Bone pTLA:

1. The site MUST have experience with IPv6 on the 6Bone, at least as a leaf site and preferably as a transit pNLA under an existing pTLA.
2. The site MUST have the ability and intent to provide "production-like" 6Bone backbone service to provide a robust and operationally reliable 6Bone backbone.
3. The site MUST have a potential "user community" that would be served by becoming a pTLA, e.g., the requester is a major player in a region, country or focus of interest.
4. Must commit to abide by the 6Bone backbone operational rules and policies as defined in the present document.

When a candidate site seeks to become a pTLA site, it will apply for it to the 6Bone Operations group (see below) by bringing evidences it meets the above criteria.

8. 6Bone Operations group

The 6Bone Operations group is the body in charge of monitoring the adherence to the present rules, and will take the appropriate actions to correct deviations. Membership in the 6Bone Operations group is mandatory for, and restricted to, any site connected to the 6Bone.

The 6Bone Operations group is currently defined by those members of the existing 6Bone mailing list, i.e., 6bone@isi.edu, who represent sites participating on the 6Bone. Therefore it is incumbent on relevant site contacts to join the mailing list. Instructions on how to join the list are maintained on the 6Bone web site at <http://www.6bone.net>.

9. Common rules enforcement

Participation in the 6Bone is a voluntary and benevolent undertaking. However, participating sites are expected to adhere to the rules described in this document, in order to maintain the 6Bone as quality tool for experimenting with the IPv6 protocols and products implementing them.

The following processes are proposed to help enforcing the 6Bone rules:

- Each pTLA site has committed when requesting their pTLA to implement the rules, and to ensure they are respected by sites within their administrative control (i.e. those to who prefixes have been delegated).
- When a site detects an issue, it will first use the 6Bone registry to contact the site maintainer and work the issue.
- If nothing happens, or there is disagreement on what the right solution is, the issue can be brought to the 6Bone Operations group.
- When the problem is related to a product issue, the site(s) involved is responsible for contact the product vendor and work toward its resolution.
- When an issue causes major operational problems, backbone sites may decide to temporarily set filters in order to restore service.

10. Security Considerations

The result of bogus entries in routing tables is usually unreachable sites. Having guidelines to aggregate or reject routes will clean up the routing tables. It is expected that using these guidelines, routing on the 6Bone will be less sensitive to denial of service attacks due to misleading routes.

The 6Bone is a test network. Therefore, denial of service, packet disclosure, are to be expected.

11. Acknowledgements

This document is the result of shared experience on the 6Bone. Special thanks go to Bob Fink for the hard work make to date to direct the 6Bone effort, to David Kessens for the 6Bone registry, and to Guy Davies for his insightful contributions.

12. References

- [1] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [RFC 2471] Hinden, R., Fink, R. and J. Postel (deceased), "IPv6 Testing Address Allocation", RFC 2471, December 1998.
- [RFC 2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC 2283] Bates, T., Chandra, R., Katz, D. and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 2283, March 1998.
- [RIPE-181] Bates, T., Gerich, E., Joncheray, L., Jouanigot, J., Karrenberg, D., Terpstra, M. and J. Yu, Representation of IP Routing Policies in a Routing Registry. Technical Report ripe-181, RIPE, RIPE NCC, Amsterdam, Netherlands, October 1994.

13. Authors' Addresses

Alain Durand
Institut d'Informatique et de Mathematiques Appliquees de Grenoble
IMAG BP 53
38041 Grenoble CEDEX 9 France

Phone : +33 4 76 63 57 03
Fax : +33 4 76 51 49 64
EMail: Alain.Durand@imag.fr

Bertrand Buclin
AT&T International S.A.
Route de l'aeroport 31, CP 72
CH-1215 Geneve 15 (Switzerland)

Phone : +41 22 929 37 40
Fax : +41 22 929 39 84
EMail: Bertrand.Buclin@ch.att.com

14. Full Copyright Statement

Copyright (C) The Internet Society (1999). 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 assigns.

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.

