                      6Bone Backbone Routing Guidelines

Status of this Memo

Copyright Notice

Abstract

   The 6Bone is an Ipv6 testbed to assist in the evolution and
   deployment of IPv6. Because of this, it is important that the core
   backbone of the IPv6 network maintain stability, and that all
   operators have a common set of rules and guidelines by which to
   deploy IPv6 routing equipment.

   This document provides a set of guidelines for all 6bone routing
   equipment operators to use as a reference for efficient and stable
   deployment of 6bone routing systems. As the complexity of the 6Bone
   grows,the adherence to a common set of rules becomes increasingly
   important in order for an efficient, scalable backbone to exist.

Table of Contents

1. Introduction

   The 6Bone is an IPv6 testbed to assist in the evolution and
   deployment of IPv6. Because of this, it is important that the core
   backbone of the IPv6 network maintain stability, and that all
   operators have a common set of rules and guidelines by which to
   deploy IPv6 routing equipment.

   This document provides a set of guidelines for all 6bone routing
   equipment operators to use as a reference for efficient and stable
   deployment of 6bone routing systems. As the complexity of the 6Bone
   grows,the adherence to a common set of rules becomes increasingly
   important in order for an efficient, scalable backbone to exist.

   This document uses BGP-4 with Multiprotocol Extensions for BGP-4 as
   defined [RFC 2283], commonly referred to as BGP4+, as the currently
   accepted EGP.

   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. Scope of this document

   This document is a best-practices Informational document aimed at
   IPv6 entities which operate under the 6Bone IPv6 testbed TLA
   allocation.

3. Common Rules for the 6bone

   This section details common rules governing the routing of the 6Bone.
   They are derived from the 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 and unspecified prefixes

       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) 6to4 prefixes

       11) aggregation & advertisement issues

3.1 Link-local prefixes

   This link-local prefix (FE80::/10) MUST NOT be advertised through
   either an IGP or an EGP. Under no circumstance should this prefix be
   seen in the 6Bone backbone routing table.

   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, but are not limited to:

      - 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. While
   vendors should be encouraged to fix the problem, the ultimate
   responsibility lies on the operator of that IPv6 site to correct the
   problem through whatever means necessary.

   Should a pTLA discover link-local prefixes coming from another pTLA,
   it is the responsibility of the pTLA leaking the routes to filter
   these, and correct the problem in a timely fashion. Should a pTLA
   discover that a downstream of that pTLA is leaking link-local
   prefixes, it is the pTLA's responsibility to ensure that these
   prefixes are not leaked to other pTLA's, or to other downstreams of
   that pTLA.

   Failure to filter such routes in a timely fashion may result in the
   manual shutting down of BGP4+ sessions to that pTLA, from other
   pTLA's.

   (Also, it is each pTLA, pNLA, and end-site's responsibility to not
   only filter their own BGP4+ sessions appropriately to peers, but to
   filter routes coming from peers as well, and to only allow those
   routes that fit the aggregation model, and do not cause operational
   problems).

3.2 Site-local prefixes

   Site local prefixes (in the FEC0::/10 range) MAY be advertised by
   IGP's or EGP's within a site. The precise definition of a site is
   ongoing work of the IPng working group, but should generally include
   a group of nodes that are operating under one administrator or group
   of administrators, or a group of nodes which are used for a common
   purpose.

   Site-local prefixes MUST NOT be advertised across transit pNLAs,
   pTLAs, or leaf-sites.

   Again, should site-local prefixes be leaked outside of a given site,
   it is the responsibility of the site to fix the problem in a timely
   manner, either through filters, or via other means which remove the
   operational impact that those prefixes had on the peering sites
   involved. However, every site SHOULD filter not only outbound on
   their EGP, but also inbound, in order to ensure proper routing
   announcements are not only sent, but also received.

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.

   The same responsibility lies with the party guilty of advertising the
   loopback or unspecified prefix as in Section 3.1 and 3.2.

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 of 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
   at the time of the writing of this memo.

   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.

   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 where
   a.b.c.d represents the octets of an IPv4 address) internally. As
   there is no real rationale today for doing so, these address SHOULD
   NOT be used or routed 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.

   Should ::/96 IPv4-compatible prefixes be leaked into an EGP, it is
   the responsibility of the party who is advertising the route to fix
   the problem, either through proper filters, or through other means,
   while it remains in the best interest of all particiapants of the
   6Bone to filter both outbound and inbound at their IGP borders.

3.6 IPv4-mapped prefixes

   IPv4-mapped prefixes (::FFFF:a.b.c.d where a.b.c.d represents the
   octets of 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, to aid in deployment of
   IPv6.

   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 any downstream peer (non-pTLA
   site). Transit pNLAs MAY advertise a default route to any of their
   downstreams (other transit pNLA or leaf site).

   Should a default route be redistributed into an EGP and found on any
   pTLA EGP sessions, it is the responsibility of the pTLA to fix this
   problem immediately upon realization of the route's existence, and
   the responsibility of the guilty pTLA to push the entity from which
   the default route was originated, should the default route have
   originated from downstream of a pTLA.

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 the 6Bone range
   (3ffe::/16), and routing of global unicast prefixes yet undelegated
   in the range (3ffe::/16) are discussed in section 4, Routing
   policies, below.

3.9 Inter-site links

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

   Sites MAY use Other addressing schemes for Inter-site links, but
   these addresses MUST NOT be advertised into the IPv6 global routing
   table.

Prefixes for inter-site links MUST NOT be injected in the global
routing tables.

3.10 6to4 Prefixes

The 6to4 prefix, or some portion thereof, MAY be announced by any
pTLA which has a current implementation of 6to4 in their IPv6
network. However, as 6to4 implementors gain more operational
experience, it MAY be necessary to change this in some way. At the
time of the writing of this docuement, any pTLA MAY announce the 6to4
prefix into global EBGP. However, in order to announce this block,
the pTLA MUST have a 6to4 router active, sourcing this prefix
announcement.

This section subject to change, and MAY vary, depending on 6to4
progress within the NGTRANS working group.

3.11 Aggregation & advertisement issues

Route aggregation MUST be performed by any border router talking EGP
with any other IPv6 sites. More-specifics MUST NOT be leaked into or
across the IPv6 6Bone backbone.

4. Routing Policies for the 6bone

Leaf sites or pNLAs MUST only advertise to an upstream provider the
prefixes assigned by that provider. Advertising a prefix assigned by
another provider to a provider is not acceptable, and breaks the
aggregation model. A site MUST NOT advertise a prefix from another
provider to a provider as a way around the multi-homing problem.
However, in the interest of testing new solutions, one may break this
policy, so long as ALL affected parties are aware of this test, and
all agree to support this testing. These policy breaks MUST NOT
affect the 6bone routing table globally.

To clarify, if one has two upstream pNLA or pTLA providers, (A and B
for this example), one MUST only announce the prefix delegated to one
by provider A to provider A, and one MUST only announce the prefeix
delegated by one from provider B upstream to provider B. There exists
no circumstance where this should be violated, as it breaks the
aggregation model, and could globally affect routing decisions if
downstreams are able to leak other providers' more specific
delegations up to a pTLA. As the IPNG working group works through the
multi-homing problem, there may be a need to alter this rule
slightly, to test new strategies for deployment. However, in the case
of current specifications at the time of this writing, there is no
reason to advertise more specifics, and pTLA's MUST adhere to the
current aggregation model.

Site border routers for pNLA or leaf sites MUST NOT advertise
prefixes more specific (longer) than the prefix that was allocated by
their upstream provider.

All 6bone pTLAs MUST NOT advertise prefixes longer than a given pTLA
delegation (currently /24 or /28) to other 6bone pTLAs unless special
peering arrangements are implemented. When such special peering
aggreements are in place between any two or more 6bone pTLAs, care
MUST be taken not to leak the more specifics to other 6bone pTLAs not
participating in the peering aggreement. 6bone pTLAs which have such
agreements in place MUST NOT advertise other 6bone pTLA more
specifics to downstream 6bone pNLAs or leaf sites, as this will break
the best-path routing decision.

The peering agreements across the 6Bone may be by nature non-
commercial, and therefore MAY allow transit traffic, if peering
agreements of this nature are made. However, no pTLA is REQUIRED to
give or receive transit service from another pTLA.

Eventually, the Internet registries will assign prefixes under other
than the 6Bone TLA (3FFE::/16). As of the time this document was
written in 1999, the Internet registries were starting to assign /35
sub-TLA (sTLA) blocks from the 2001::/16 TLA. Others will certainly
be used in the future.

The organizations receiving prefixes under these newer TLAs would be
expected to want to establish peering and connectivity relationships
with other IPv6 networks, both in the newer TLA space and in the
6bone pTLA space. Peering between new TLA's and the current 6Bone
pTLA's MAY occur, and details such as transit, and what routes are
received by each, are outside of general peering rules as stated in
this memo, and are left up to the members of those TLA's and pTLA's
that are establishing said peerings. However, it is expected that
most of the rules discussed here are equally applicable to new TLAs.

5. The 6Bone Registry

   The 6Bone registry is a RIPE-181 database with IPv6 extensions used
   to store information about the 6Bone, and its sites. The 6bone is
   accessible at:

          <http://www.6bone.net/whois.html>)

   Each 6Bone site MUST maintain the relevant entries in the 6Bone
   registry. In particular, the following object MUST be present for all
   6Bone leaf sites, pNLAs and pTLAs:

   - IPv6-site: site description

   - Inet6num: prefix delegation (one record MUST exist for each
      delegation)

   - Mntner: contact info for site maintance/administration staff.

   Other object 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 MAY
   NOT necessarily reside in the 6Bone registry. They can be stored
   within any of the Internet registry databases (ARIN, APNIC, RIPE-NCC,
   etc.)

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 web site at <http://www.6bone.net> has various information
   and tools to help find candidate 6bone networks.

   Any site connected to the 6Bone MUST maintain a DNS server for
   forward name lookups and reverse address lookups. The joining site
   MUST maintain the 6Bone objects relative to its site, as describe in
   section 5.

   The upstream provider MUST delegate the reverse address translation
   zone in DNS to the joining site, or have an agreement in place to
   perform primary DNS for that downstream. The provider MUST also
   create the 6Bone registry inet6num object reflecting the delegated
   address space.

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

7. Guidelines for 6Bone pTLA sites

   The following rules apply to qualify for a 6Bone pTLA allocation. It
   should be recognized that holders of 6Bone pTLA allocations are
   expected to provide production quality backbone network services for
   the 6Bone.

   1. The pTLA Applicant must have a minimum of three (3) months
      qualifying experience as a 6Bone end-site or pNLA transit. During
      the entire qualifying period the Applicant must be operationally
      providing the following:

   a. Fully maintained, up to date, 6Bone Registry entries for their
      ipv6-site inet6num, mntner, and person objects, including each
      tunnel that the Applicant has.

   b. Fully maintained, and reliable, BGP4+ peering and connectivity
      between the Applicant's boundary router and the appropriate
      connection point into the 6Bone. This router must be IPv6
      pingable. This criteria is judged by members of the 6Bone
      Operations Group at the time of the Applicant's pTLA request.

   c. Fully maintained DNS forward (AAAA) and reverse (ip6.int)
      entries for the Applicant's router(s) and at least one host
      system.

   d. A fully maintained, and reliable, IPv6-accessible system
      providing, at a mimimum, one or more web pages, describing the
      Applicant's IPv6 services. This server must be IPv6 pingable.

2. The pTLA Applicant MUST have the ability and intent to provide
   "production-quality" 6Bone backbone service. Applicants must
   provide a statement and information in support of this claim.
   This MUST include the following:

   a. A support staff of two persons minimum, three preferable, with
      person attributes registered for each in the ipv6-site object
      for the pTLA applicant.

   b. A common mailbox for support contact purposes that all support
      staff have acess to, pointed to with a notify attribute in the
      ipv6-site object for the pTLA Applicant.

3. The pTLA Applicant MUST have a potential "user community" that
   would be served by its becoming a pTLA, e.g., the Applicant is a
   major provider of Internet service in a region, country, or focus
   of interest. Applicant must provide a statement and information in
   support this claim.

4. The pTLA Applicant MUST commit to abide by the current 6Bone
   operational rules and policies as they exist at time of its
   application, and agree to abide by future 6Bone backbone
   operational rules and policies as they evolve by consensus of the
   6Bone backbone and user community.

   When an Applicant seeks to receive a pTLA allocation, it will apply
   to the 6Bone Operations Group (see section 8 below) by providing to
   the Group information in support of its claims that it meets the
   criteria above.

8. 6Bone Operations Group

   The 6Bone Operations Group is the group in charge of monitoring and
   policing adherence to the current rules. Membership in the 6Bone
   Operations Group is mandatory for, and restricted to, sites connected
   to the 6Bone.

   The 6Bone Operations Group is currently defined by those members of
   the existing 6Bone mailing list who represent sites participating in
   the 6Bone. Therefore it is incumbent on relevant site contacts to
   join the 6Bone 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 for the 6bone

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

   The following is in support of policing adherence to 6Bone rules and
   policies:

   1. Each pTLA site has committed to implement the 6Bone's rules and
      policies, and SHOULD try to ensure they are adhered to by sites
      within their administrative control, i.e. those to who prefixes
      under their respective pTLA prefix have been delegated.

   2. When a site detects an issue, it SHOULD first use the 6Bone
      registry to contact the site maintainer and work the issue.

   3. If nothing happens, or there is disagreement on what the right
      solution is, the issue SHOULD be brought to the 6Bone Operations
      Group.

   4. When the problem is related to a product issue, the site(s)
      involved SHOULD be responsible for contacting the product vendor
      and work toward its resolution.

   5. When an issue causes major operational problems, backbone sites
      SHOULD decide to temporarily set filters in order to restore
      service.

10. Security Considerations

   The result of incorrect 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
   rules and policies, routing on the 6Bone will be less sensitive to
   denial of service attacks due to misleading routes.

   The 6Bone is an IPv6 testbed to assist in the evolution and
   deployment of IPv6. Therefore, denial of service or packet disclosure
   are to be expected. However, it is the pTLA from where the attack
   originated who has ultimate responsibility for isolating and fixing
   problems of this nature. It is also every 6Bone site's responsibility
   to safely introduce new test systems into the 6Bone, by placing them
   at a strategically safe places which will have minimal impact on
   other 6Bone sites, should bugs or misconfigurations occur.

11. References

   [RFC 2373] Hinden, R. and S. Deering, "IP Version 6 Addressing
               Architecture", RFC 2373, July 1998.

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

   [RFC 2546] Durand, A. and B. Buclin, "6Bone Routing Practice", RFC
               2546, March 1999

   [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.

12. Authors' Addresses

    Rob Rockell
    EMail: rrockell@sprint.net

    Bob Fink
    EMail: fink@es.net

13. Full Copyright Statement

Acknowledgement