Francois Le Faucheur,
Cisco Systems, Inc.

IETF Internet Draft
Expires: December, 2002
Document: draft-lefaucheur-bgp-tunnel-transition-00.txt June, 2002

Operational Environments and Transition Scenarios for "Connecting IPv6 Islands across IPv4 Clouds with BGP"

Status of this Memo

Abstract

This document describes the common operational environments of IPv4 Service Providers wanting to add IPv6 services to their service portfolio but not wanting (yet) to upgrade their IPv4 backbone to IPv6 routing.

Two main transition scenarios are identified.

We recommend that the "MP-BGP over IPv6" and "MP-BGP over IPv4" approaches defined in [BGP-TUNNEL] be respectively used for each of the two transition scenarios.

1. Introduction

Many IPv4 Service Providers are considering/willing to complement their service portfolio by some IPv6 services. We first describe such operational environments in more details and then discuss the two main transition scenarios identified for such operational environments.

2. Operational Environments

A Service Provider (SP) runs an IPv4 backbone and offers IPv4 services. This Service Providers makes extensive use of BGP for exchange of IPv4 reachability information:

- IPv4 connection with upstream/peer Internet Service Providers (eBGP)
- IPv4 connection with some end-users IPv4 sites (eBGP)
- Distribution of IPv4 reachability information inside the SP's backbone (iBGP over TCP/IPv4).

The Service Provider is obviously very familiar with BGP. The Service Provider may also be very familiar with MP-BGP (e.g. support of inter-domain IPv4 Multicast service or support of MPLS VPN service [MPLS-VPN]).

This environment can be illustrated in the following way:

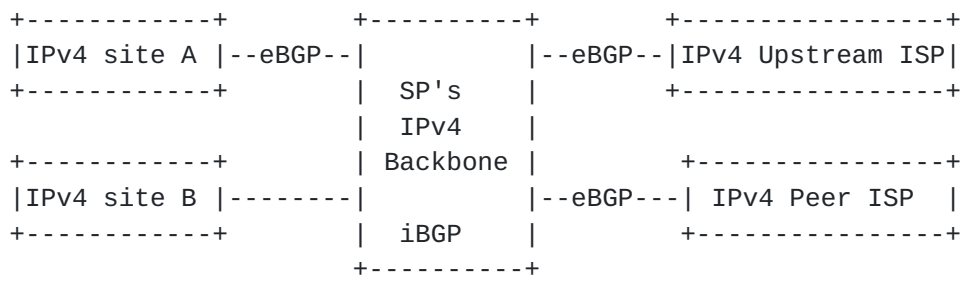

Now, the SP wants to broaden his/her service offering by adding some IPv6 services such as IPv6 connectivity across multiple IPv6 sites of an end user, and/or global IPv6 connectivity (i.e. access to the IPv6

Internet).

All the exchange of IPv6 routing information at the edge of the SP backbone uses standardized native IPv6 routing:

- the SP provides global IPv6 connectivity through his/her IPv6 customer relationship with an upstream ISP, or by peering relationships with other IPv6 ISPs in the default free routing zone (DFZ). Such peering uses MP-eBGP ([MP-BGP], [IPv6-MP-BGP)] for IPv6. It is used by the Service Provider to advertise IPv6 reachability of its IPv6 allocated prefix and to receive reachability for the IPv6 internet.
- An IPv6 routing protocol (IGPv6 or MP-eBGP for IPv6) may run between IPv6 Site and the SP's network so that the IPv6 Site advertises its IPv6 reachability and receives IPv6

reachability information from the SP. Alternatively, static IPv6 routes and/or a default IPv6 route may be used to control such reachability.

But the SP does not yet want to upgrade his/her backbone to IPv6. So native IPv6 routing is NOT used inside the SP's backbone.

The new environment can be illustrated in the following way:

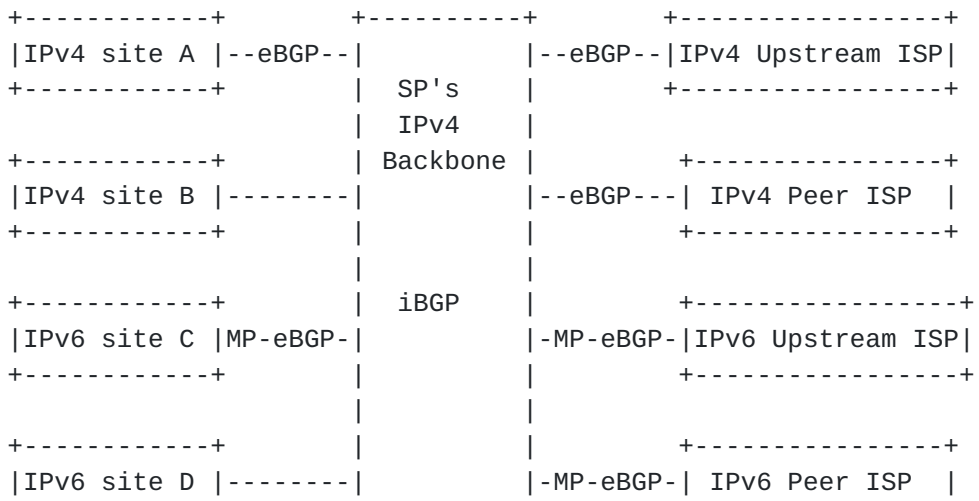

In such an environment, the SP needs to find a transition solution until he/she is ready to upgrade the backbone to native IPv6 routing. The transition solution needs to:

- distribute IPv6 reachability information over his/her non-IPv6 backbone
- tunnel IPv6 traffic inside his/her IPv4 backbone.

We feel such operational environments are very common and require clearly documented transition approaches.

3. Transition Scenarios

We identify two main transition scenarios for such environments.

3.1. Use of existing IPv6 Tunneling Techniques

In this scenario, the SP's operational constraints are such that:

- it is acceptable/desirable to establish a set of MP-iBGP peerings (MP-iBGP mesh or MP-iBGP Route Reflector structure) over TCP/IPv6, in addition to the existing set of iBGP peerings (iBGP mesh or iBGP Route Reflector structure) over TCP/IPv4.

AND,

- it is acceptable/desirable to use one of the existing NGTRANS tunneling techniques ([6to4], [ISATAP],...) to achieve IPv6 reachability of the MP-BGP Next Hops over the IPv4 backbone. Note that this tunneling technique is ONLY needed for IPv6 reachability of the MP-BGP Next Hops at the edge of the SP network and is NOT needed for IPv6 reachability of the IPv6 Internet or IPv6 end user sites. The latter can be achieved via native IPv6 MP-iBGP routing among the MP-BGP Next Hops once those are granted IPv6 reachability. Note that this means that the inherent characteristics of such existing

methods (e.g. use of special IPv6 addresses as with ISATAP, need for some configuration,...) are acceptable/desirable.

AND,

- The potential optimization of tunneling overhead achievable in some situations is not acceptable/desirable for tunneling of all the IPv6 traffic. For example, if the IPv4 backbone is also running MPLS, use of existing NGTRANS tunneling results in every IPv6 packets being effectively prepanded with both an IPv4 and an MPLS header.

In this scenario, transition can be supported using purely existing IPNG and NGTRANS techniques combined in the following manner:

- one existing NGTRANS tunneling technique ([6to4], [ISATAP],...) is used to provide IPv6 reachability among MP-iBGP Next Hops at the edge of the SPs' network
- this reachability is used to build a set TCP/IPv6 connections for MP-iBGP peering among these MP-iBGP Next Hops (and possibly Route Reflectors), in addition to the existing set of TCP/IPv4 connections for the iBGP peerings.
- existing MP-iBGP ([MP-BGP], [IPv6-MP-BGP])is used among those MP-iBGP Next Hops at the edge of the SP's network to establish MP-IBGP peerings and to exchange all the IPv6 reachability information. This is in addition to the existing iBGP peerings.
- IPv6 packets are tunneled between the MP-iBGP routers at the edge of the IPv4 backbone using the selected existing NGTRANS tunneling technique and by performing this tunneling as if the packet was addressed to the MP-iBGP Next Hop Router advertised for the relevant IPv6 prefix. In other words, the tunnel IPv4 destination and the IPv4 tunnel header are not selected directly based on the destination IPv6 address but selected based on the IPv6 address of the MP-iBGP Next Hop router advertised in MP-iBGP for the relevant IPv6 prefix.

This transition approach is referred to as the "MP-BGP over IPv6" approach and is further documented in [BGP-TUNNEL].

3.2. "Auto tunneling"

In this scenario, the SP's operational constraints are such that:

- it is desirable to use the existing set of iBGP peerings (iBGP mesh or iBGP Route Reflector structure) over TCP/IPv4

to distribute reachability of all services (IPv4, IPv6 , and VPN-IPV4 if MPLS VPN services is also supported).

OR,

- a form of transparent automatic tunneling into IPv4 is desirable so that no configuration is required on the MP-BGP Next Hop routers at the edge of the SP network to activate/control tunneling, nor any constraints imposed on the IPv6 addresses used for these MP-BGP Next Hops.

OR,

- The potential optimization of tunneling overhead achievable in some situations is desirable for tunneling of all the IPv6 traffic. For example, if the IPv4 backbone is also running MPLS, IPv6 packets should only be prepanded by an MPLS header and not by an IPv4 header plus an MPLS header.

In this scenario, transition can be supported using existing IPNG and NGTRANS techniques combined in the following manner:

- The existing set of TCP/IPv4 connections and iBGP peering is used to also advertise IPv6 reachability via MP-iBGP ([MP-BGP], [IPv6-MP-BGP]) among the MP-iBGP Next Hops at the edge of the network (and possibly Route Reflectors)
- An MP-iBGP Next Hop advertising IPv6 reachability information uses the IPv4-mapped IPv6 address format ([V6ADDR]} to convey its IPv4 address as the Next Hop Address.
- The MP-iBGP Next Hops receiving IPv6 reachability information, use the IPv4 address contained in this IPv4-mapped address, as the destination address of the IPv4 tunnel to tunnel the corresponding IPv6 packets into, thus achieving automatic tunneling.

This transition approach is referred to as the "MP-BGP over IPv4" approach and is further documented in [BGP-TUNNEL].

4. Recommendation on Transition

We recommend that the "MP-BGP over IPv6" and "MP-BGP over IPv4" approaches defined in [BGP-TUNNEL] be respectively used for each of the two main transition scenarios described above.

We also observe that these approaches could also naturally accommodate a possible additional transition step which would be to complement the SP's service offering by an IPv6 MPLS VPN service [IPv6-MPLS-VPN] by simply supporting the IPv6-VPN address family in addition to the IPv6 address family in the same MP-iBGP peerings.

5. Security Considerations

No new security considerations are raised by this document. Those are the same as the ones of BGP and can be addressed through the corresponding mechanisms.

6. Acknowledgments

We thank Jeremy De Clercq and Benoit Lourdelet for their review and suggestions.

References

[MP-BGP] T. Bates et al, "Multiprotocol Extensions for BGP-4", RFC2858.

[IPv6-MP-BGP] Marques, P., and et.al , Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing, RFC 2545, March 1999.

[BGP-TUNNEL]. Ooms et al, Connecting IPv6 Islands across IPv4 Clouds with BGP, draft-ietf-ngtrans-bgp-tunnel-04.txt, January 2002.

[6TO4] B. Carpenter, K. Moore, "Connection of IPv6 domains via IPv4 Clouds", RFC3056, February 2001.

[ISATAP] F. Templin, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), draft-ietf-ngtrans-isatap-02.txt (work in progress).

[V6ADDR] Deering, S., and R. Hinden, "IP Version 6 Addressing Architecture", draft-ietf-ipngwg-addr-arch-v3-07.txt (work in progress).

[MPLS-VPN] Rosen E., Rekhter Y., Brannon S., Chase C., De Clercq J., Hitchin P., Marshall , Srinivasan V., "BGP/MPLS VPNs", draft-ietf-ppvpn-rfc2547bis-00.txt (work in progress).

[IPv6-MPLS-VPN] Nguyen T., Gastaud G., De Clercq J., Ooms D.,"BGP-MPLS VPN extension for IPv6 VPN over an IPv4 infrastructure",

draft-ietf-ppvpn-bgp-ipv6-vpn-01.txt> (work in progress).

Author's Address:

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3
400, Avenue de Roumanille
06410 Biot-Sophia Antipolis
France
Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com