International Journal of Communication Systems and Network Technologies
Vol.8, No.1, Dec 2019

Case Study on IPv6 Routing Protocols Over the IPv4: Designing Needs

Sanjay Sharma
Shri Guru Ram Rai ITS, Dehradun

Abstract- A modified next-generation Internet Protocol, known first as IPng and then as IPv6, has been under development by the Internet Engineering Task Force (IETF) for several years to replace the current Internet Protocol known as IPv4. This paper describes the reasons behind the need for IPv6. Of major importance during the development of IPv6 has been how to do the transition away from IPv4, towards IPv6. The work on powerful transition strategies, secure tools and flexible mechanisms has been a part of the basic IPv6 design effort from the beginning. The current transition efforts, taking place at the IETF IPng Transition Working Group (ngtrans) will continue until it is clear that the transition will be successful.

Keywords- IETF, IPv4, IPv6, Routing, Protocols

I. Introduction

A big question mark arises how to migrate IPv6 to IPv4 and IPv4 to IPv6. A combination of techniques can be used purposely, such as tunneling, path MTU, discovery DSN looksup, IP spoofing, hop limit manipulation and IPv6 header modifying. These transition design efforts resulted in a basic Transition Mechanism specification for IPv6 hosts and routers [4] that specifies the use of a Dual IP layer providing complete support for both IPv4 and IPv6 in hosts and routers, and IPv6-over-IPv4 tunneling, encapsulating IPv6 packets within IPv4 headers to carry them over IPv4 routing infrastructures.

For smooth transition from IPv4 to IPv6, more and more web servers are implemented and configured with both IPv4 and IPv6 protocol stacks. By choosing dual-stack web sties as our data sources we can gain better understanding of IPv6 performance as well as its distinctive problems by comparison with its IPv4 counterpart.

Of great concern to transition strategy planners is how to provide connectivity between IPv6-enabled end-user sites (also known as routing domains) when they do not yet have a reasonable (or any) choice of Internet Service Provider (ISP) that provides native IPv6 transport services. One way to provide IPv6 connectivity between end-user sites (when native IPv6 service does not exist) is to use IPv6-over-IPv4 encapsulation (tunneling) between them, similar to the technique currently used in the 6bone [5] IPv6 testbed network. This requires complexity for both end-user sites, and the networks providing the tunneling service (for instance, the 6bone backbone ISPs) in creating, managing, and operating manually configured tunnels.

The "6to4" transition mechanism, "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels" [6], provides a solution to the complexity problem of using manually configured tunnels by specifying a unique routing prefix for each end-user site that carries an IPv4 tunnel endpoint address.

It should also be noted that each end-user site with as little as a single IPv4 address has a unique, routable, IPv6 site routing prefix thanks to the 6to4 transition mechanism.

II. Connecting IPv6 Routing Domains

To determine successful connectivity between both the end users, tunneling can be used similar to technique currently used in 6bone IPv6 testbed networks. An ordered set of

links through interior routers exchange routing information through an interior gateway protocol (IGP) whereas exterior routers use an exterior Gateway Protocol (EGP) that identifies and connects both the protocols.

Compilation of connectivity can be done by arranging tunnels directly with each IPv6 site to which

DOI- 10.18486/ijcsnt.2019.8.1.01
ISSN: 2053-6283

connectivity is needed, but more typically, it is done by arranging a tunnel into a larger IPv6 routing infrastructure that could guarantee connectivity to all IPv6 end-user site networks (See Figure 1). The 6bone IPv6 testbed was the first IPv6 routing infrastructure to provide worldwide IPv6 connectivity (starting in 1996), while more recently (late 1999) networks providing production IPv6 Internet service are also interconnected to provide this connectivity. In fact, the 6bone and production IPv6 routing infrastructures are well interconnected to guarantee worldwide IPv6 connectivity.

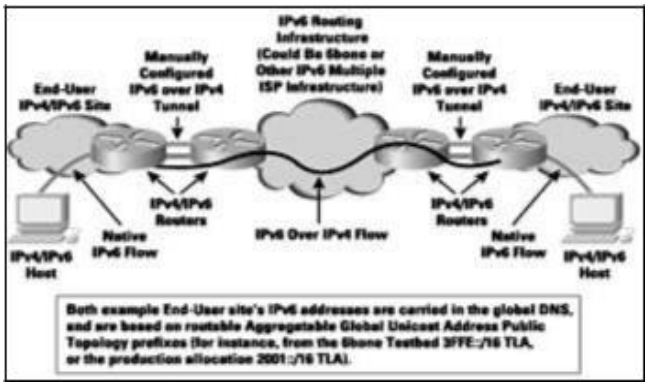

Figure 1: Configured Tunnel Overview

The 6to4 mechanism addresses many of the practical difficulties with manually configured tunneling. These are discussed below.

- The end-user site network staff must choose an IPv6 Internet service to tunnel to. This entails a process of at least three parts:
- Finding candidate networks when the site's choice of IPv4 service does not provide IPv6 service (either tunneling or native),
- Determining which ones are the best IPv4 paths to use, so that an IPv6-over-IPv4 tunnel doesn't inadvertently follow a very unreliable or low-performance path.
- Making arrangements with the desired IPv6 service provider for tunneling service, a scenario that may at times be difficult if the selected provider is not willing to provide the service, or if for other administrative/ cost reasons it is difficult to establish a business relationship.
- Clearly it is easiest to use the site's own service provider, but in the early days of IPv6 transition this will often not be an option.
- An IPv6-over-IPv4 tunnel must be built by the selected provider, and a peering relationship must be established with the selected provider. This requires

  establishing a technical relationship with the provider and working through the various low-level details of how to configure tunnels between two routers, including answering the following questions:

  - Are the site and provider routers compatible early on in this process?
- - What peering protocol will be used (presumably an IPv6-capable version of the Border Gateway Protocol Version 4 [BGP4]), and are the versions compatible and well debugged?

Have all the technical tunnel configuration issues between the site and provider been addressed?

Clearly it is easiest to perform all these steps if they are taken with the site's own IPv4 service provider.

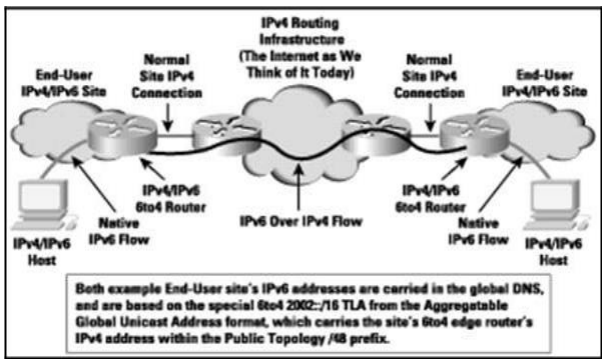

Figure 2: Configured Tunnel Overview

The 6to4 mechanism addresses many of the practical difficulties with manually configured tunneling. These are discussed below.

- The end-user site network staff must choose an IPv6 Internet service to tunnel to. This entails a process of at least three parts:
- Finding candidate networks when the site's choice of IPv4 service does not provide IPv6 service (either tunneling or native),
- Determining which ones are the best IPv4 paths to use, so that an IPv6-over-IPv4 tunnel doesn't inadvertently follow a very unreliable or low-performance path.
- Making arrangements with the desired IPv6 service provider for tunneling service, a scenario that may at times be difficult if the selected provider is not willing to provide the service, or if for other administrative/ cost reasons it is difficult to establish a business relationship.
- Clearly it is easiest to use the site's own service provider, but in the early days of IPv6 transition this will often not be an option.

! An IPv6-over-IPv4 tunnel must be built by the selected provider, and a peering relationship must be established with the selected provider. This requires
establishing a technical relationship with the provider and working through the various low-level details of how to configure tunnels between two routers, including answering the following questions:
- Are the site and provider routers compatible early on in this process?

! - What peering protocol will be used (presumably an IPv6-capable version of the Border Gateway Protocol Version 4 [BGP4]), and are the versions compatible and well debugged?

- Have all the technical tunnel configuration issues between the site and provider been addressed?

3. 6to4 Eliminates Complex Tunnel Management

The 6to4 transition mechanism provides a solution to the complexity problem of building manually configured tunnels to an ISP by advertising a site's IPv4 tunnel endpoint (to be used for a dynamic tunnel) in a special external routing prefix for that site. Thus one site trying to reach another will discover the 6to4 tunnel endpoint from a Domain Name System (DNS) name to address lookup and use a dynamically built tunnel from site to site for communication. (See Figure 2.) The tunnels are transient, in that there is no state maintained for them, lasting only as long as a specific transaction uses the path. A 6to4 tunnel also bypasses the need to establish a tunnel to a wide-area IPv6 routing infrastructure, such as the 6bone.

The specification of a 48-bit external routing prefix in the IPv6 Aggregatable Global Unicast Address Format (AGGR) [7] (see Figure 3) that provides just enough space to hold the 32 bits required for the 32-bit IPv4 tunnel endpoint address (called V4ADDR in Figure 3) makes this setup possible.

Thus, this prefix has exactly the same format as normal prefixes assigned according to the AGGR. Within the subscriber site, it can be used exactly like any other valid IPv6 prefix, for instance, for automated address assignment and discovery according to the normal IPv6 mechanisms for this.

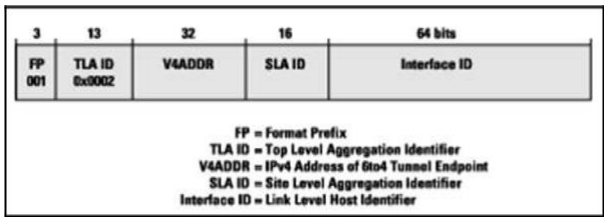

Figure 3: 6to4 Prefix Format

4. The Simplest Use of 6to4

The simplest scenario for 6to4 is when several sites start to use IPv6 alongside IPv4, and have no native IPv6 ISP service available. Thus, each site identifies a router to run dual stack (that is, IPv4 and IPv6 together) and 6to4 tunneling, ensuring that this router has a globally routable IPv4 address (that is, not in private IPv4 address space).

It is assumed that this new 6to4 router is reachable by IPv6-capable hosts within the site. Although the various ways in which these hosts may be reached are not discussed in detail here, they include using IPv6-enabled site IPv4 routers, operating special IPv6-only routers in parallel with site IPv4 routers, using the "6over4" mechanism [8] , and employing other tunneling methods.

A new 6to4 site advertises the 6to4 prefix to its site via the Neighbor Discovery (ND) protocol [9], which will cause IPv6 hosts at this site to have their DNS name/address entries to include the 6to4 prefix for the site in them.

In operation, when one IPv6-enabled host at a 6to4 site tries to access an IPv6-enabled host by domain name at another 6to4 site, the DNS will return both an IPv4 and an IPv6 IP address for that host, indicating that it is reachable by both IPv4 and IPv6. The requesting host selects the IPv6 address, which will have a 6to4 prefix, and sends a packet off to its nearest router, eventually reaching its site boundary router, which we assume has 6to4 service as well.

5. Sending and Receiving Rules for 6to4 Routers

When the requesting site's 6to4 router sees that it must send a packet to another site (that is, there is a nonlocal destination), and that the next hop destination prefix contains the special 6to4 Top Level Aggregation (TLA) value of 2002::/16, the IPv6 packet is encapsulated in an IPv4 packet using an IPv4 protocol type of 41, as defined in the Transition Mechanisms RFC [4]. The source IPv4 address will be the one in the requesting site's 6to4 prefix (which is the IPv4 address of the outgoing interface to the Internet on the 6to4 router, and contained in the source 6to4 prefix of the IPv6 packet), and the destination IPv4 address will be the one in the next hop destination 6to4 prefix of the IPv6 packet.

When the destination site's 6to4 router receives the IPv4 packet, and recognizes that it has an IPv4 protocol type of 41, IPv4 security checks are made and the IPv4 header is removed, leaving the original IPv6 packet for local forwarding.

DOI- 10.18486/ijcsnt.2019.8.1.01
ISSN: 2053-6283

The sending rule above is the only modification to IPv6 forwarding, because the receiving rule was already specified for the basic IPv6 Transition Mechanism mentioned earlier [4]. Along with advertisement of the 6to4 prefix by appropriate entries in the DNS, any number of sites can interoperate without manual tunnel configuration.

It is not necessary to operate an exterior routing protocol (for instance, BGP4+) for 6to4 simple scenarios because the IPv4 exterior routing protocol is handling this function. Also, no new entries in IPv4 routing tables result from the use of 6to4.

6. The Return Path and Source Address Selection

Packets must flow in both directions to be useful; thus it is essential that IPv6 packets sent use a packet with a 6to4 prefix as a source address when talking to a site with a 6to4 prefix; in other words, the destination must have a 6to4 prefix. In the simple example given above, this is not an issue because both sites have only IPv4 connectivity, so they have 6to4 prefixes for their site to communicate with. DNS lookups for host systems at these sites will return only one IPv6 address, which will be the one with a 6to4 prefix. Source address selection is thus not an issue.

As we will soon see, source address selection is an issue for more complex 6to4 usage scenarios; therefore, some source address selection algorithm is necessary in IPv6 hosts. The exact form and method of the algorithm to use is under active study at the IETF IPv6 (ipng) working group [10], and an algorithm is likely to be chosen in early 2000. Mean-while, for the purposes of understanding 6to4, it is sufficient to realize that when a 6to4 connected sending site is sending to a destination site using that site's 6to4 prefix, the sending host must guarantee that the source IPv6 address uses the sending site's 6to4 prefix.

7. More Complex 6to4 Usage Scenarios

Several more interesting 6to4 usage scenarios exist when a site has both 6to4 connectivity and native IPv6 connectivity. The simplest of these is when such a site is trying to reach another site that has only 6to4 con-nectivity, in which case the source address selection algorithm mentioned above is essential to ensure that the site's 6to4 IPv6 address is chosen. No destination selection is required because there is only one choice, that is, 6to4.

Similarly, when a site that has only 6to4 connectivity tries to reach a site with both 6to4 and native IPv6 connectivity, some host rule for choosing among multiple destination addresses must result in the 6to4 address being chosen, because only a local 6to4 IPv6 source address is available. Of course source selection is not an issue in this case because there is only the 6to4 IPv6 address to use.

Another variation of these scenarios is when a site with 6to4 and native IPv6 connectivity is trying to reach another site that has only native IPv6 connectivity, making a source address selection algorithm essential to make sure the site's native IPv6 address is chosen. No destination selection is required, because there is only one choice, that is, the native IPv6 address.

Similarly, when a site that has only native IPv6 connectivity tries to reach a site with 6to4 and native IPv6 connectivity, a host rule is essential for choosing among multiple addresses to ensure that a native IPv6 address is chosen, because only a local native IPv6 source address is available. Again, source selection is not an issue in this case because only the native IPv6 address can be used.

An interesting choice develops in the situation when both sites have 6to4 and native IPv6 connectivity as both 6to4-to-6to4 and native IPv6-to-native- IPv6 connections are a possibility. Current thinking as of the writing of this article is to prefer the native IPv6 connection.

8. The 6to4 Relay

Thus the 6to4-only site will try to send a packet to the native IPv6-only site by forwarding an encapsulated (tunneled) IPv6 packet to the 6to4 relay, which removes the IPv4 header (decapsulates) and forwards the packet on to the IPv6-only site.

Potentially, multiple 6to4 relays are needed, one for each separate IPv6 routing realm (collection of IPv6 routing ISPs). In practice, it is expected that all native IPv6 ISP services will be interconnected even if the use of inter-IPv6-ISP manually configured tunnels are required to do so. This is currently the case as of early 2000, because all 6bone 3FFE::/16 TLA networks and all production 2001::/16 subTLA networks are interconnected with each other.

It is expected that native IPv6 service providers will choose to operate 6to4 relays as a simple extension of their service. There are no special rules or exceptions to 6to4 as described here for this to happen because the 6to4 relay is simply operated as part of an end-user site that belongs to the IPv6 ISP.

The most interesting, and most complex, 6to4 scenario is that of sites with only 6to4 connectivity communicating with sites with only native IPv6 connectivity. This is

DOI- 10.18486/ijcsnt.2019.8.1.01
ISSN: 2053-6283

accomplished by the use of a 6to4 relay that supports both 6to4 and native IPv6 connectivity (Figure 4). The 6to4 relay is nothing more than an IPv4/IPv6 dual-stack router.

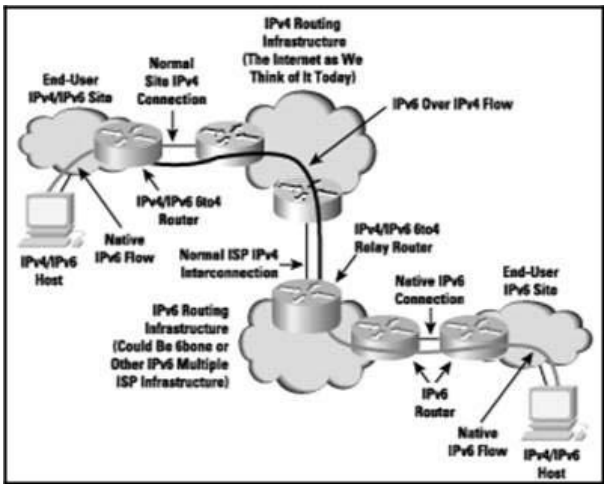

Figure 4: The 6to4 Relay

The 6to4 relay advertises a route to 2002::/16 for itself into the native IPv6 infrastructure it is attached to. The native IPv6 network operators must filter out and discard any 6to4 (2002:...) prefix advertisements longer than /16. In addition, the 6to4 relay may advertise into its 6to4 connection whatever native IPv6 routes its policies allow, which the 6to4 router at the 6to4-only site picks up with either a BGP4+ peering session, or with a default route, to the 6to4 relay.

9. Other Issues

Several other 6to4 issues are presented below for completeness.

! The IPv6 Maximum Transmission Unit (MTU) size could prove too large for some intermediate IPv4 link when a 6to4 tunnel is in use, thus IPv4 fragmentation will occur. Though undesirable, fragmentation is not disastrous, so the IPv4 "Do Not Fragment" bit should not be set in the IPv4 packet carrying the 6to4 tunnel.

! How sites move IPv6 packets internal to a site is not important to the 6to4 process. For illustrative purposes in this article, it is generally assumed that native IPv6 transmission exists within a site. This may not be strictly true because "6over4," manual tunnels, and other methods of moving IPv6 packets could be in use. Nonetheless, it is not important to the 6to4 processes described here.

! Security issues with the 6to4 mechanism are not discussed here. The reader is referred to the current 6to4 draft for an explanation of these issues [6] .

! 6to4 sites with IPv6 connectivity must not inject their 6to4 prefix into the IPv6 routing infrastructure via the native IPv6 connection.

References

[1]. Fink, R., "IPv6-What and Where It Is," The Internet Protocol Journal , Volume 2, No. 1, March 1999.

[2]. IPng and IPv6 information, including formal specifications, can be found at: http://playground. sun.com/pub/ipng/html

[3]. "The Case for IPv6," an Internet Draft of the IAB, can be found at: http://www.6bone.net/misc/case-for-ipv6.html

[4]. IETF IPv6 Transition Working Group (ngtrans) information, including status of all its current projects, can be found at: http://www.6bone.net/ngtrans/

[5]. "Transition Mechanisms for IPv6 Hosts and Routers," RFC 1933, can be found at: http://www.ietf.org/ rfc/rfc1933.txt

[6]. The 6bone IPv6 Testbed Network is explained at: http://www.6bone.net

[7]. "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels" ("6to4"), an Internet Draft of the IETF

DOI- 10.18486/ijcsnt.2019.8.1.01
ISSN: 2053-6283