BT Diamond IP

Service Provider IPv6 Deployment Strategies

By Tim Rooney

Director, Product Management

BT Diamond IP

www.btdiamondip.com

Contents

Service Provider IPv6 Deployment Strategies

By Tim Rooney, Director, Product Management

Introduction

With available public IPv4 address space diminishing rapidly, service providers are next in the IP addressing food chain behind Regional Internet Registries to suffer the consequences of being unable to meet customer needs for IPv4 address space. Without address space to allocate, service provider offerings which entail address allocation will grind to a deafeningly screeching halt. Though one could argue that if such an offering featured only IPv6 space, some customers would look elsewhere, probably at least for a 6-12 month period while IPv4 completely putters out, offering IPv6 address space enables a service provider to continue with such offerings and to "future proof" its customers supporting both its IPv4 and IPv6 allocations.

Even service providers that do not allocate IP address space must be able to support their customers who choose to utilize their own IPv6 address space. For example, broadband or generally residential service providers typically provide an IP address to the customer premises equipment (CPE) gateway device terminating the service provider link into the residence. Beyond this point with the customer's network, IP addresses may be assigned autonomously using DHCP, static configuration or IPv6 autoconfiguration. While the vast majority of such customers will simply rely on the DHCP server within the CPE gateway, more technically savvy users may choose to experiment. And with Microsoft Windows XP, Windows Vista and Windows 7, Mac OS, and Linux distributions all supporting IPv6 by default, use of IPv6 may be more prevalent than expected though fully unintentional.

This white paper summarizes IPv6 deployment approaches for service providers. We'll discuss various strategies that can be employed by service providers offering residential service, Internet Service Provider (ISP) services and business services which offer value-added transport services. Our companion white papers, *IPv6 Addressing and Management Challenges* (1) and *IPv4-IPv6 Transition and Co-Existence Strategies* (2) provide an overview of IPv6 addressing and a general review of defined IPv4-IPv6 co-existence strategies respectively.

IPv6 Deployment Scope

The first step is to define the scope of your IPv6 deployment: are you looking to offer IPv6 addresses to customers, to support customer-assigned IPv6 addresses for transport using your network or simply upgrade your backbone? Support of IPv6 customer addressing requires not only corresponding IPv6 routing infrastructure support, but also IPv6 support of network management and general operations support systems, customer care systems and training, and effective IPv4/IPv6 address management with corresponding Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) support. We'll focus primarily on routing infrastructure ("data plane") deployment approaches, but we'll also touch on impacts on the additional supporting infrastructure ("management plane").

Introduction to Deployment Approaches

Broadband residential, business, and wireless service providers offer transport and in some cases, multimedia services to residential and/or business customers. The customer premises equipment (CPE), typically a router, cable modem, fiber termination unit, or wireless router device, terminates the service provider access link and is referred to as a customer-edge (*CE*) router. The CE router forwards all outbound customer-initiated IP packets to the service provider network via the provider-edge (*PE*) router to which its connection to the service provider network terminates. The PE router then routes the packets to other customer-facing PE routers or the Internet either directly or via service provider core or backbone routers (also known as *provider* or *P* routers).

Conversely, the PE router routes inbound traffic to the customer site from the service provider network originating from the Internet or other customer sites for business network applications. The service provider "core" network consists of P routers which route packets among themselves and PE routers. The service provider generally provisions the service-provider facing network interface of the CE device with an IP address. For business applications, the CE router is the interface to several networks within the corresponding customer site, whereas for residential applications, the CE device sometimes serves as a DHCP server to provision addresses to a relatively small number of hosts for IP services.

Reference Architecture

Within the service provider network, IP packets can be routed over a variety of underlying services such as multi-protocol label switching (MPLS) ultimately to the intended destination. For destinations accessible via the public Internet, the destination host may be accessible via an IPv4 address (via "IPv4 Internet"), IPv6 address (via "IPv6 Internet") or both. Physically, this is one Internet, though we illustrate it as logically separate to distinguish IP routing by version. Connectivity to one or both of these "Internets" depends on service provider capabilities and route advertisements in terms of supporting routing via IPv4, IPv6 or both.

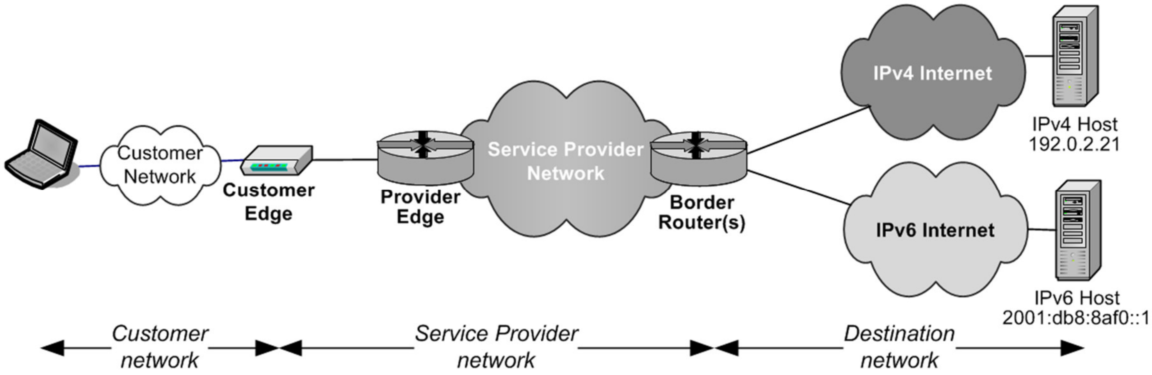

Figure 1: Basic Three Layer Architecture: Customer/Service Provider/Destination

In this paper, we'll consider the deployment of IPv6 along each of the three tiers of this basic architecture:

- Customer network – depends on whether the service provider supports IPv4, IPv6 or both versions for customers.

- Service provider network – the PEs and “core” of the service provider’s network interconnecting customers and the Internet.
- Destination network – depending on the capabilities of the destination web site, email server, etc. reachability may require one or either IP version.

Deployment Approaches Overview

The service provider generally controls their network and in most cases, the IP address assigned to the service provider-facing interface of the CE device. The customer may independently implement either protocol version, though their ability to connect via a particular version of IP will depend in part on the version supported by the intended destination and on the transport supported by the service provider. The following table summarizes the connectivity options available across this simple three-tier architecture for differing IP versions at each tier.

Table 1: Basic Deployment Options Based on Protocol Versions Supported

Customer network	Service Provider Network	Destination network	IPv6 Deployment Approach*									
			No IPv6	NAT444	Dual Stack	6PE/ 6VPE	Config. Tunnels	6rd	4over6	Dual Stack Lite	NAT64 with DNS64	Full IPv6
IPv4	**IPv4**	**IPv4**	●	●	●							
IPv4	**IPv4**	**IPv6**			●							
IPv6	**IPv4**	**IPv4**			●							
IPv6	**IPv4**	**IPv6**			●	●	●	●				
IPv4	**IPv6**	**IPv4**			●		●		●	●		
IPv4	**IPv6**	**IPv6**			●							
IPv6	**IPv6**	**IPv4**			●						●	
IPv6	**IPv6**	**IPv6**			●		●			●		●

* Note there are numerous “transition technologies” defined in various RFCs, though many techniques have been deemed insecure or unstable. Please consult (2) or (3).

The first four rows of the table illustrate connectivity options for a service provider maintaining an IPv4 network, at least for the time being. For example, the first row of the table highlights an end-to-end IPv4 connection, which reflects today’s dominant transport scenario. To support any other scenario in the following three rows, implementation of some form of IPv6 compatibility is required. In the second and third rows, support of IPv4 customers communicating with IPv6 destinations and vice-versa, requires deployment of dual stack. The fourth row interconnects IPv6 endpoints via the service provider IPv4 network. Various approaches are available to address this scenario including dual stack, end-to-end configured tunnels (e.g., VPNs), 6rd, or if using MPLS, 6PE or 6VPE.

The four rows below the dotted line at the bottom of the table illustrate a service provider implementation of IPv6 and its ability or requirements to support various connection types. Of course service providers may implement multiple technologies at once during a phased IPv6 deployment; e.g., phased in by market or geography.

As is clear from the table, the dual stack option offers support of all combinations, though it requires a comprehensive IPv4/IPv6 address plan and implementation of dual stack on PE or all service provider routers. To state the obvious, note that when configuring dual stack or other approaches serving both versions facing the Internet, border routers must be configured to advertise both IPv4 and IPv6 routes.

Routing Infrastructure Deployment Approaches

This section provides an overview of each of these implementation strategies. We won't address the configured tunnels option as this implies a pre-arranged VPN or tunnel between the customer and destination network hosts to tunnel over the service provider's network, without direct involvement by the service provider. Automated tunnels or "softwires" are tunnels created on the fly and are core components of several deployment approaches as we'll see.

NAT444

Like the "No IPv6" approach, the NAT444 (4) strategy is not an IPv6 implementation strategy. But the NAT444 approach serves as a means of "buying time" in prolonging the lifecycle of IPv4 in order to deploy supplemental address space in the form of IPv6. NAT444 features a large scale IPv4-IPv4 network address translation (LSN NAT44) gateway which allows multiple subscribers to share a common IPv4 address. NAT44 provides the benefit of not requiring replacement of existing IPv4 CPE though at the cost of limiting the number of customer sessions often required by applications such as AJAX and RSS feeds as well as loss of CPE geolocation information (E911) and more (5).

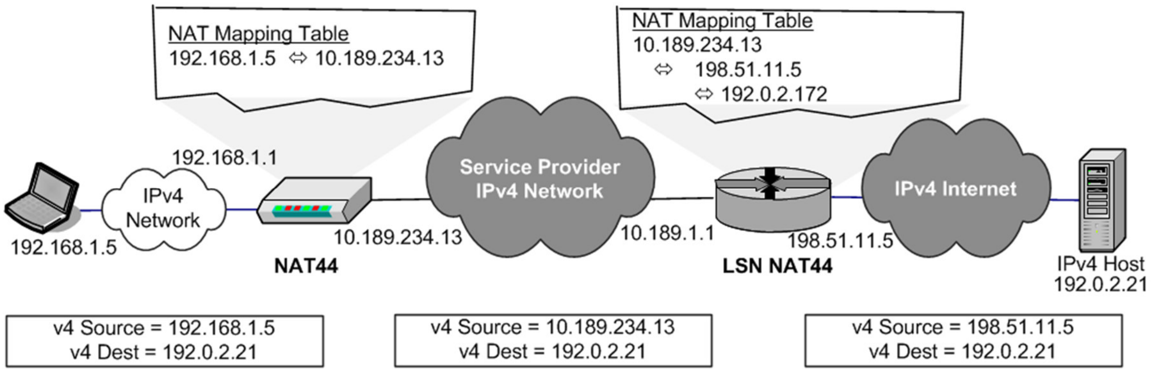

Figure 2: NAT444 Architecture

This architecture shows the use of two NATs in the data path, one within the CPE, translating the home network IPv4 space to service provider-supplied private address space. A second NAT, the LSN NAT44, translates the subscriber's CPE address to a public IPv4 address. Port numbers differentiate different sessions among the same and multiple end customers. The term "NAT444" is illustrative of the dual concatenated IPv4-IPv4 NATs and the resultant use of IP addresses from three IPv4 address spaces (customer private, service provider access, and public Internet).

Dual Stack

Dual stack implementation requires the configuration of both an IPv4 address and an IPv6 address on each infrastructure (or at least routing) device, and possibly device interface. A dual stack CPE can function effectively regardless of the version(s) supported by the corresponding service provider; connecting end-to-end however does require protocol continuity between the service provider network and the corresponding destination network.

Support of dual stack within the service provider network may be deployed throughout or on the "edge(s)". For example, PE routers facing customers can implement dual stack to enable support of IPv4 and IPv6 customers, while PE routers facing the Internet enable connectivity using the protocol version of the destination. However, without full dual-stack deployment also on core routers interconnecting PEs, tunnels between PEs would be required to transport traffic of the version not implemented. This implementation is illustrated in the next section on IPv6 over MPLS deployments.

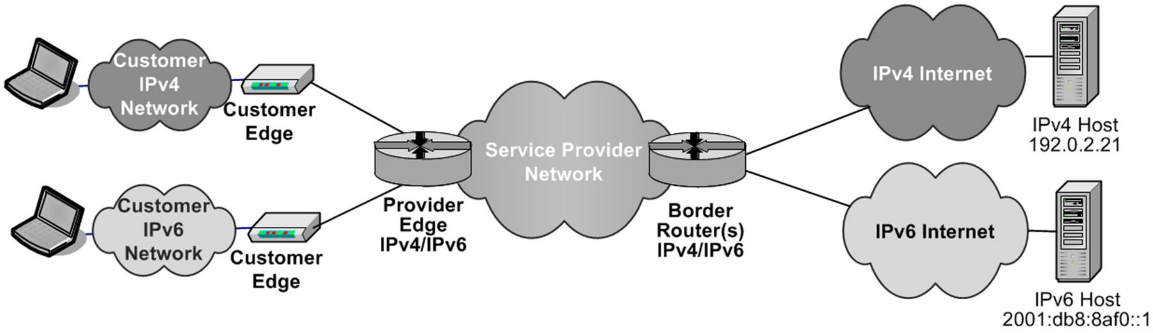

Figure 3: Dual Stack Architecture

IPv6 over MPLS

There are several approaches for implementing IPv6 over MPLS including support of native IPv6, but the most common interim approaches include IPv6 PE (6PE) over MPLS or IPv6 VPN PE (6VPE) over MPLS. The 6PE architecture (6) features dual stack PE routers with IPv4-only core routers; this approach could serve as an intermediate step to full dual stack or IPv6 deployment.

The dual stack PE routers communicate IPv6 network reachability via their respective core-facing IPv4 address(es) via multi-protocol border gateway protocol (MP-BGP) over IPv4. This enables the ingress 6PE router to identify the IPv4 address of the appropriate egress 6PE router and to identify a label switch path (LSP) and associated IPv4 label to enable label-switching through the core IPv4 routers to the egress 6PE router. This technique requires use of both an IPv6 label and an outer IPv4 label but obviates the need to pre-define IPv6 over IPv4 tunnels.

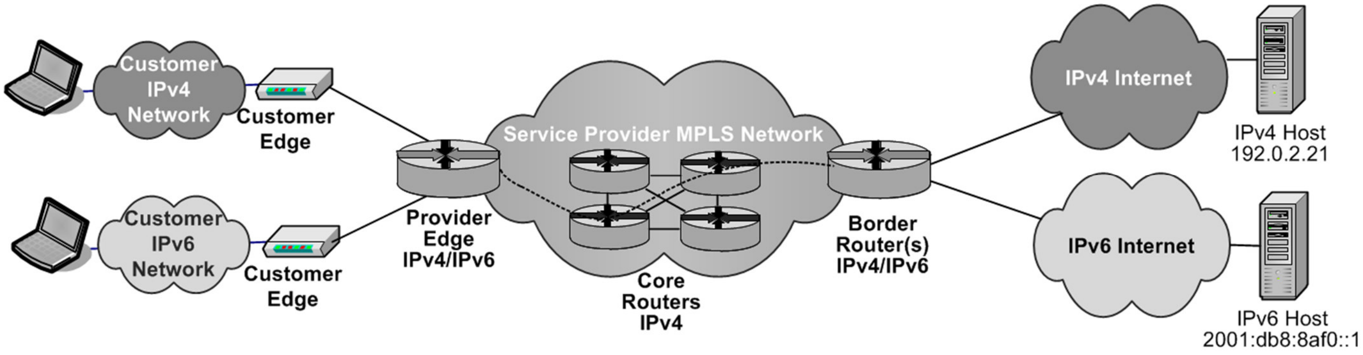

Figure 4: 6PE Architecture

The 6VPE (7) architecture is similar at a high level though IPv6 over IPv4 "VPN tunnels" are utilized to traverse the core, which provides improved privacy over 6PE and support of overlapping address space. A given customer's VPN is associated (provisioned) with one or more VPN Routing and Forwarding (VRF) table entry(ies) in corresponding PE devices. The PE router associates the customer's physical circuit into the network (and possibly layer 2 header information) with the appropriate VPN identified within the VRF table.

Each CE device advertises routes (reachable networks at the corresponding site) to its connected PE router(s). An MPLS label is assigned to the VPN (whether by VPN, CE device, or route), and the label is communicated with the corresponding route during MP-BGP route distribution within the service provider's network among other PE routers serving the customer's network (VPN). As IPv4 packets arrive on a PE interface from a CE router, the PE router determines the VPN from the VRF table, then applies the corresponding label for use in switching the packets to the appropriate PE and ultimately the destination.

6rd (IPv6 Rapid Deployment)

RFC 5569 (8) defines "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)," a technique to enable a service provider to provision IPv6 addresses to end customers for IPv6 connectivity while maintaining an IPv4 infrastructure. RFC 5969 (9) specifies the 6rd protocol. This method calls for softwire tunneling of customer IPv6 traffic from the customer premises to an IPv6 destination via modified 6to4 technique. The modification entails use of the service provider's IPv6 prefix (/32) in lieu of the 6to4 prefix, 2001::/16.

Like 6to4, the next 32 bits of the 6rd IPv6 prefix consists of the IPv4 address of the 6to4 gateway, in this case the customer premises broadband router. Hence a 6to4 prefix is defined as:

2001:{32-bit IPv4 address}::/48, while the 6rd prefix is

{32-bit service provider IPv6 prefix}:{32-bit IPv4 address}::/64.

This enables the service provider to provision a /64 to each customer, which comprises a single IPv6 subnet. Thus, a service provider with an RIR allocated IPv6 block 2001:db8::/32 would provision a customer gateway device with IPv4 address 192.0.2.130 with a 6rd subnet address of 2001:db8:c000:282::/64 as shown in Figure 5.

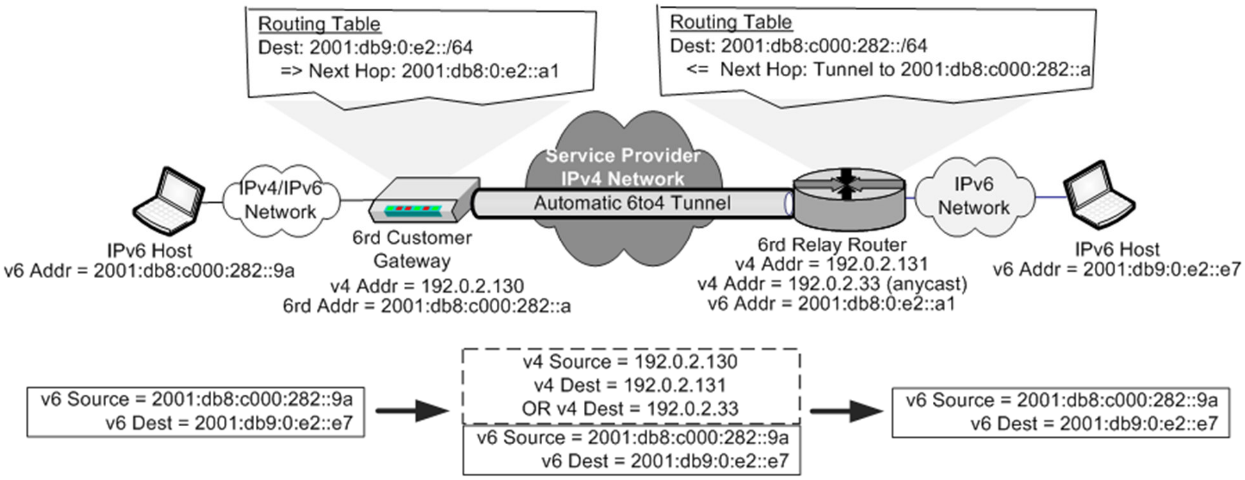
Routing Table
Dest: 2001:db9:0:e2::/64
=> Next Hop: 2001:db8:0:e2::a1
Routing Table
Dest: 2001:db8:c000:282::/64
<= Next Hop: Tunnel to 2001:db8:c000:282::a
Service Provider
IPv4 Network
Automatic 6to4 Tunnel
IPv4/IPv6
Network
IPv6
Network
IPv6 Host
v6 Addr = 2001:db8:c000:282::9a
6rd Customer
Gateway
v4 Addr = 192.0.2.130
6rd Addr = 2001:db8:c000:282::a
6rd Relay Router
v4 Addr = 192.0.2.131
v4 Addr = 192.0.2.33 (anycast)
v6 Addr = 2001:db8:0:e2::a1
IPv6 Host
v6 Addr = 2001:db9:0:e2::e7
v6 Source = 2001:db8:c000:282::9a
v6 Dest = 2001:db9:0:e2::e7
v4 Source = 192.0.2.130
v4 Dest = 192.0.2.131
OR v4 Dest = 192.0.2.33
v6 Source = 2001:db8:c000:282::9a
v6 Dest = 2001:db9:0:e2::e7
v6 Source = 2001:db8:c000:282::9a
v6 Dest = 2001:db9:0:e2::e7

Figure 5: 6rd Deployment Example (3)

A device within the residence requiring an IPv6 address would assign an address from this subnet. For example, in Figure 5 a PC is assigned IPv6 address 2001:db8:c000:282::9a. The 6rd customer gateway tunnels native IPv6 packets over IPv4 to a 6rd gateway (relay router). The other address-related change between 6rd and 6to4 is that the 6to4 anycast address is fixed (192.88.99.1), while the 6rd anycast address is defined by the service provider themselves from its own address space. Each customer router must be provisioned with the 6rd relay agent or anycast address(es).

The 6rd relay router terminates the IPv4 tunnel, then routes the IPv6 packet natively to its destination. The use of the service provider's prefix enables 6rd-reachable destinations to be advertised along with the service provider's native IPv6 traffic.

4over6

The 4over6 approach, specified in RFC 5747 (10), is an automated tunneling (i.e., softwire) approach for interconnecting IPv4 subscribers to IPv4 destinations via an IPv6 network. As a converse to the 6PE approach, 4over6 features an IPv6 core of P routers which route native IPv6 packets and IPv4 packets tunneled over IPv6 among the PE routers. A subscriber with IPv4 address space can communicate with an IPv4 destination using this technique.

Each CE router provides routing updates to its connected PE router. The PE routers would use MP-BGP to communicate 4over6 routes and PE routers route traffic accordingly. Considering Figure 6, The subscriber CE on the left side of the diagram advertises reachability to the 198.51.100.0/24 network while the destination CE (not shown in the figure explicitly) advertises reachability to 192.0.2.0/24. The respective PE routers use MP-BGP to communicate this reachability. When a packet arrives at the PE router on the left of the diagram, the route to the destination is identified and the packet is encapsulated with an IPv6 header for routing via the IPv6 core (P) routers. Upon receipt of the IPv6 packet at the egress PE router, the PE router decapsulates the packet (removes the IPv6 header) and routes the original IPv4 packet to its destination via the serving CE router.

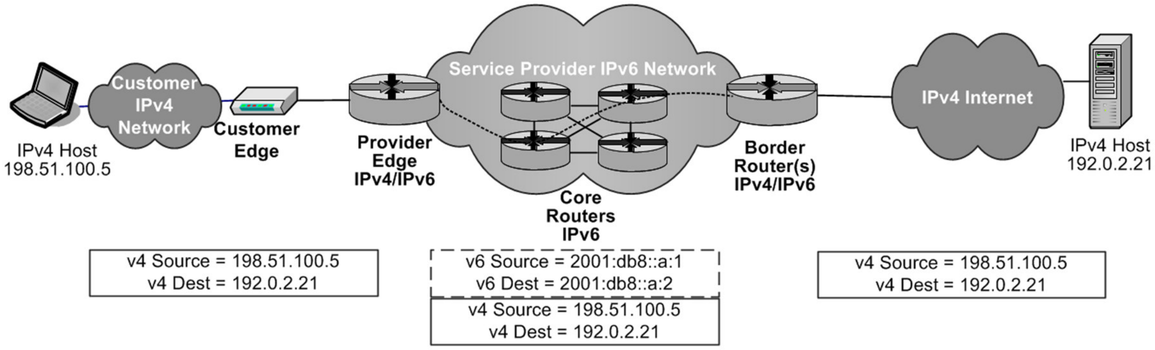

Figure 6: 4over6 Example

The current 4over6 architecture supports only a single Autonomous System (AS) number, so support of multiple customer private networks is limited, though support of multiple AS numbers is an area of future study.

Dual-Stack Lite

Dual-stack lite (11) is a technology that enables a service provider to deploy IPv6 within their network, while facilitating long-term support and efficient utilization of IPv4 addresses assigned to customer network devices. Service providers typically assign an IP address to a customer router or gateway which interfaces directly to the broadband access network. The customer gateway performs DHCP server functions in assigning IP addresses to IP devices in the home network. It is expected that such home network devices will support only IPv4 for quite some time.

The components comprising a dual-stack lite implementation include the following:

- Basic Bridging BroadBand (B4) element - bridges the IPv4 home network with an IPv6 network; the B4 function may reside on the customer gateway device or within the service provider network.
- Softwire IPv4-in-IPv6 tunnel – tunnels IPv4 traffic between the B4 and the AFTR over IPv6.
- Address Family Translation Router (AFTR) - terminates the IPv4-in-IPv6 softwire tunnel with the B4 element and also performs IPv4-IPv4 network address translation (NAT) functionality.

Figure 7 illustrates the inter-relationship of these three components within an end-to-end IP connection. Starting on the left of the figure, the IPv4 host obtains an IPv4 address, 10.1.0.2, from the DHCP server function of the customer gateway. Let's say this IPv4 host desires to connect to a website, which has been resolved to IP address 192.0.2.21. The IPv4 host formulates an IP packet with source address 10.1.0.2 and source port of 1000 for example, and destination address 192.0.2.21 port 80. The host transmits this packet to its default route, the customer gateway.

The customer gateway in this example includes the B4 element, which sets up the softwire IPv4-in-IPv6 tunnel if it is not already established. The customer gateway has been assigned an IPv6 address on its WAN port (facing the service provider network) and it is over this connection that the tunnel is established. The customer gateway has also been configured with the AFTR IPv6 address manually or via DHCPv6. As shown in Figure 7, the B4 element encapsulates the original IPv4 packet with an IPv6 header and transmits it to the AFTR.

The AFTR terminates the tunnel and removes the IPv6 header. The AFTR then performs an IPv4-IPv4 NAT function. This is required to translate the original packet's private (RFC 1918) IPv4 source address into a public IPv4 address. Thus, the service provider must provision a pool of public IPv4 addresses which can be used as source IP addresses on packets destined for an IPv4 destination as in this case. This pooling enables the service provider to more efficiently utilize the increasingly scare public IPv4 address space. The AFTR also generally performs port translation as well and must track this mapping for each NAT operation in order to properly map IPv4 addresses and port numbers bi-directionally.

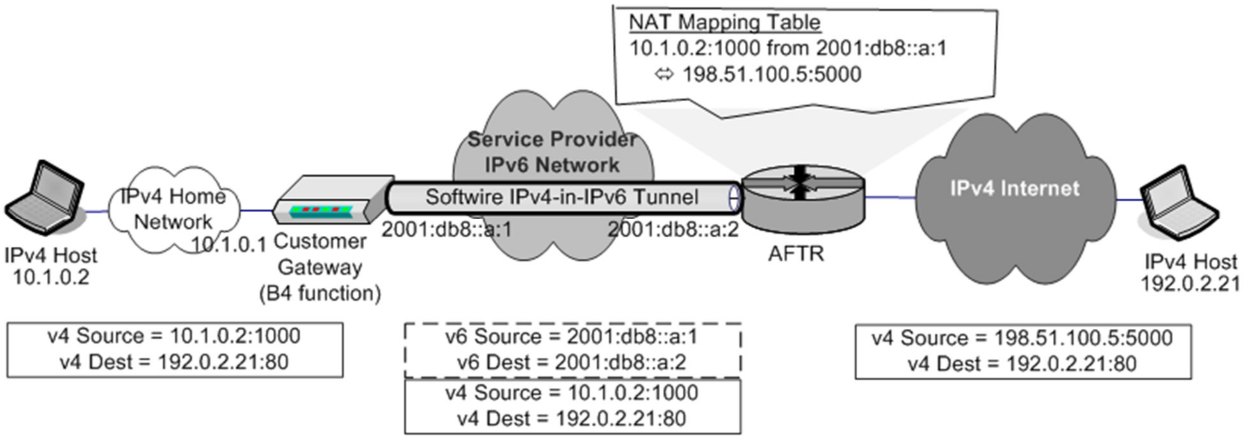

Figure 7: Dual Stack Lite Architecture (3)

In Figure 7, the AFTR has mapped the customer's source IPv4 address and port, 10.1.0.2:1000 to 198.51.100.5:5000. Since customers generally utilize private address space where overlaps may occur, the NAT mapping table also tracks the tunnel over which the packet originated. The packet ultimately transmitted to the destination host includes this mapped IPv4 address and port, 198.51.100.5:5000. Return packets destined for this address/port are mapped to [destination] address 10.1.0.2:1000 and tunneled to 2001:db8::a:1.

Customers deploying native IPv6 or dual stack hosts can have respective IPv6 addresses provided by DHCPv6 functionality implemented in the customer gateway or via autoconfiguration. IPv6 packets transmitted over the home network to the customer gateway would not utilize the softwire tunnel, but instead be routed natively over the service provider IPv6 access network.

NAT64 and DNS64

The NAT64 (12) solution translates IPv6 packets into IPv4 packets. This enables an IPv6 CPE via a service provider IPv6 network to interface to IPv4 destinations. Hence this approach is appropriate for IPv6 greenfield networks or fully migrated IPv6 service provider networks requiring IPv4 destination access (which will likely be required for a number of years to come!).

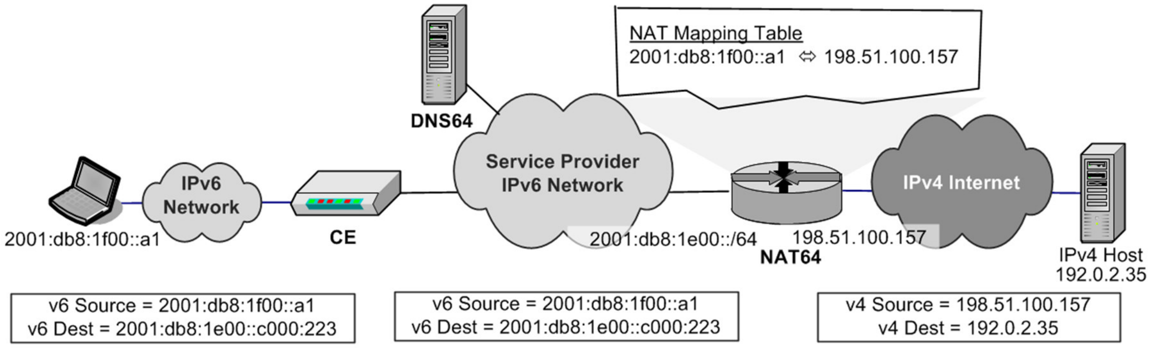
NAT Mapping Table
2001:db8:1f00::a1 ⇔ 198.51.100.157
DNS64
Service Provider IPv6 Network
IPv6 Network
IPv4 Internet
2001:db8:1f00::a1
CE
2001:db8:1e00::/64
198.51.100.157
NAT64
IPv4 Host
192.0.2.35
v6 Source = 2001:db8:1f00::a1
v6 Dest = 2001:db8:1e00::c000:223
v6 Source = 2001:db8:1f00::a1
v6 Dest = 2001:db8:1e00::c000:223
v4 Source = 198.51.100.157
v4 Dest = 192.0.2.35

Figure 8: NAT64 Architecture

Key to this strategy is the DNS64 component, which is a special recursive DNS server in that it processes queries for AAAA records normally and passes through valid responses for IPv6 addresses, but it additionally issues A record queries for failed AAAA responses in an attempt to identify an IPv4 destination address in the absence of an IPv6 address. If a valid A resource record set is received by the DNS64 server, it formulates a response to the resolver for the initial AAAA query comprised of the IPv6 network address of the NAT64 gateway, concatenated with the returned IPv4 address from the A record response.

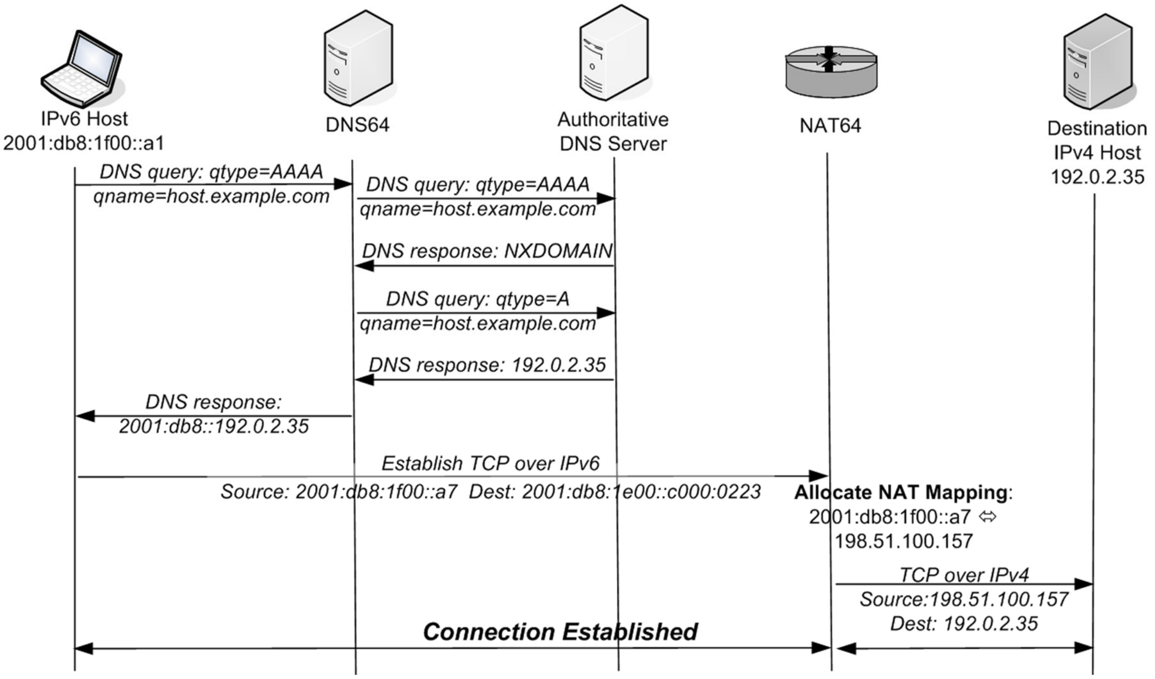

Figure 9: DNS64 Resolution Process

While IPv4 NAT (NAT44) devices have been implemented in IP networks for over 20 years, NAT64 is a relatively new breed and is yet unproven. However, the NAT44 experience increases the likelihood of successful NAT64 implementations.

Comparison of Deployment Approaches

The following table summarizes the relative characteristics of each deployment approach.

Table 2: High Level Comparison of Deployment Approaches

Basic Criteria	IPv6 Deployment Approach									
	No IPv6	NAT444	Dual Stack	6PE/ 6VPE	Config. Tunnels	6rd	4over6	Dual Stack Lite	NAT64 with DNS64	Full IPv6
Business or Residential	N/A	Both	Both	Bus	Bus	Res	Bus	Res	Both	Both
Provides IPv6 support	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IPv4-IPv6 Co-existence	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Implementation complexity	None	High	High	Mod	Mod	Low	Low	Mod	High	High
Requires CPE changes	No	No	No	No	No	Yes	No	Yes	No	Yes
Requires new SP equipment	No	Yes	No	Yes	No	Yes	No	Yes	Yes	No
Incremental operations & troubleshooting complexity	Low	Mod	Low	Mod	Low	Low	Low	Mod	Mod	Mod
Supports overlapping IPv4	N/A	Yes	No	6VPE	No	N/A	No	Yes	No	N/A

Addressing and DNS Considerations

Whichever strategy(ies) you choose, appropriate IPv6 address allocation and DNS configuration is crucial to the success of IPv6 deployment. Unless you're planning a greenfield deployment, IPv6 address space will need to be managed and allocated in conjunction with the current and future deployed IPv4 address space. For example, with a dual-stack deployment, tracking of IPv4 and IPv6 addresses not only at the subnet level but down to the dual-stack device interface is critical for accurate IP address inventory and management.

Allocating IPv6 assigned from your RIR requires careful planning. As a general guideline, the difference in prefix length from that of the block you received from your RIR and the address space size you intend to allocate to your customers will dictate what you have to work with in terms of address hierarchy. For example, if you received a /32 from your RIR and plan to allocate /64s to subscribers, you'll have 32 bits with which the address space can be allocated hierarchically, e.g. by region (e.g., /36s), city (e.g., /44s), service node (e.g., /52s), and PE (e.g., /56s) etc. This example allocation hierarchy allows up to 16 region allocations from which 256 city allocations may be made respectively, from each of which 256 service node allocations may be made, from which 16 PE allocations may be made, each of which can support 256 /64s (customers). Of course this is just one example and you may have more or fewer levels of differing sizes.

If you are allocating larger blocks to customers, you will have fewer bits to work with. If you are allocating non-uniform address blocks to customers (e.g., /48s to businesses, /64s to SOHO customers), a more sophisticated allocation and tracking mechanism should be considered in terms of allocating space in a sparse, best-fit or a random manner. Please consult (1; 3; 13) for details and examples on sparse, best-fit and random allocation techniques.

DNS configuration will drive end user traffic to IPv4 or IPv6 addressable destinations within your network (name space). Resolution of destination addresses is under the scope of the respective domain name administrator. Return of responses for IPv6 addresses (AAAA query types) and not IPv4 (A query type) will indicate to the querying host that this destination is reachable only via IPv6, in which case IPv6 service support will enable connectivity.

Some IPv4-IPv6 transition technologies have direct requirements from DNS, such as NAT64/DNS64. If you allocate reverse zones to customers along with address space, provisioning of the correct (accurate!) ip6.arpa zone and corresponding NS/glue records in your DNS servers is required. This should follow a process similar to that for IPv4 allocations, but with obvious syntactical differences with hexadecimal vs. decimal domain labels.

The Management Plane

Supporting the chosen IPv6 deployment strategy within the data plane as we've discussed in this paper, management and back-office systems must be analyzed with respect to requirements and capabilities for IPv6 support on the management plane. This includes not only upgrading support systems to enable IPv6 user interface support, but in training operations and customer care users in how to interpret IPv6 addresses and how they map to the network. Sales must be able to sell and the operations and customer care teams must be able to fulfill orders and support customers.

Thus systems not only for IP address management (IPAM) to handle addressing and DNS but also for network monitoring and management, trouble ticketing, order placement and fulfillment, billing, customer care and support applications must all be capable of at least tracking (data entry) if not supporting (transport) IPv6 addressing. This is a broad topic in its own right and perhaps a topic for a future paper.

Summary

With RIRs now holding a finite amount of dwindling IPv4 address space, service providers must plan for and implement IPv6. Managing the IPv4/IPv6 address space resource requires an advanced IPAM tool that intelligently allocates, assigns and tracks the address blocks and individual IP addresses deployed on integrated IPv4/IPv6 networks.

Utilizing an IPAM solution such as IPControl™ from BT Diamond IP which provides a logical, intuitive user interface, with reports and alerts on the current and projected address utilization, service providers can quickly and cost-effectively plan and implement their deployment of IPv6 address space, while effectively managing the ever-changing IP address needs that they face every day. IPControl was the first IPv4/IPv6 IPAM solution and is a highly scalable solution for managing both IPv4 and IPv6 address space together down to the device level along with associated DHCP and DNS configurations, providing a solid management foundation for IPv4 networks, as well as mixed IPv4-IPv6 and IPv6-only networks.

For more information, please consult the references below or www.btdiamondip.com and feel free to contact us directly at 1-800-390-6295 in the U.S., 1-610-423-4770 worldwide or email diamondip-sales@usc-bt.com.

References for Further Reading

1. **Rooney, Tim.** *IPv6 Addressing and Management Challenges.* Santa Clara, CA : BT INS, Inc., March, 2008.

2. —. *IPv4-to-IPv6 Transition and Co-Existence Strategies.* Santa Clara, CA : BT INS, Inc., March, 2008.

3. **Rooney, Timothy.** *IP Address Management Principles and Practice.* Hoboken, NJ : John Wiley & Sons, Inc., 2011.

4. **Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., Ashida, H.** *NAT444.* s.l. : IETF, Jaunary, 2011. draft-shirasaki-nat444-03.txt.

5. **Donley, C., Howard, L., Kuarsingh, V., Chandrasekaran, A., Ganti, V.** *Assessing the Impact of NAT444 on Network Applications.* s.l. : IETF, October, 2010. draft-donley-nat444-impacts-01.txt.

6. **De Clercq, J., Ooms, D., Prevost, S., Le Faucheur, F.** *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE).* s.l. : IETF, February, 2007. RFC 4798.

7. **De Clercq, J., Ooms, D., Carugi, M., Le Faucheur, F.** *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN.* s.l. : IETF, September, 2006. RFC 4659.

8. **Despres, R.** *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd).* s.l. : IETF, January 2010. RFC 5569.

9. **W. Townsley, O. Troan.** *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification.* s.l. : IETF, August, 2010. RFC 5969.

10. **Wu., J., Cui, Y., Li, X., Xu, M., Metz, C.** *4over6 Transit Solution Using IP Encapsulation and MP-BGP Extensions.* s.l. : IETF, March, 2010. RFC 5747.

11. **Durand, A., Editor.** *Dual-stack lite Broadband Deployments Post IPv4 Exhaustion.* s.l. : IETF, February, 2010. draft-ietf-software-dual-stack-lite-03.txt.

12. **Bagnulo, M., Matthews, P., van Beijum, I.** *NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers.* s.l. : IETF, March, 2009. draft-bagnulo-behave-nat64-03.txt.

13. **Rooney, Timothy.** *Introduction to IP Address Management.* Hoboken, NJ : John Wiley & Sons, Inc., 2010.

14. **Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., Palet, J.** *ISP IPv6 Deployment Scenarios in Broadband Access Networks.* s.l. : IETF, Januay, 2007. RFC 4779.

About BT Diamond IP

BT Diamond IP is a leading provider of software and appliance products that help customers effectively manage their complex IP networks. Our award-winning IPControl™ solutions help businesses more efficiently manage IP address space across mid-to-very large enterprise and service provider networks. IPControl is the most comprehensive IPAM solution available today, enabling customers to dramatically improve their IT operational efficiency and service levels by automating IP address management and DNS/DHCP server configuration across their networks. With extensive experience in IP address management, approximately 200 million IP addresses under management, BT Diamond IP is well positioned to help our customers increase performance, visibility and control over their IP networks.

For additional information, please visit www.btdiamondip.com or contact BT Diamond IP at 1-800-390-6295 in the U.S., 1-610-423-4770 worldwide.

IPControl is a trademark of BT INS, Inc.