Introduction

- **Transition** - the process or a period of changing from one state or condition to another.
- The transition between the IPv4 Internet today and the IPv6 Internet of the future will be a long process during both protocols coexists. A mechanism for ensuring smooth, stepwise and independent changeover to IPv6 services is required. Such a mechanism must help the seamless coexistence of IPv4 and IPv6 nodes during the transition period.
- IETF has created the **NGTrans** (Next Generation Transition) Group to facilitate the smooth transition from IPv4 to IPv6 services. The various transition strategies can be broadly divided into three categories, including dual stack, tunneling and translation mechanisms.
- Complete transition from IPv4 to IPv6 might not be possible because IPv6 is not backward compatible. This results in a situation where either a site is on IPv6 or it is not. It is unlike implementation of other new technologies where the newer one is backward compatible so the older system can still work with the newer version without any additional changes.
- To overcome this short-coming, we have a few technologies that can be used to ensure slow and smooth **transition (***the process or a period of changing from one state or condition to another***)** from IPv4 to IPv6.

5.1 Tunneling or Encapsulation

Tunneling technique is easy to use and flexible. Hosts can communicate with IPv4 hosts using IPv4 or communicate with IPv6 hosts using IPv6. When everything has been upgraded to IPv6, the IPv4 stack can simply be disabled or removed. Whenever you can, deploying dual-stack hosts and routers offers the greatest flexibility in dealing with islands of IPv4-only applications, equipment, and networks. Dual stack is also the basis for other transition mechanisms. Tunnels need dual-stacked endpoints, and translators need dual-stacked gateways. Disadvantages of this technique include the following: you have two separate protocol stacks running, so you need additional CPU power and memory on the host. All the tables are kept twice: one per protocol stack. A DNS resolver running on a dual-stack host must be capable of resolving both IPv4 and IPv6 address types. Generally, all applications running on the dualstack host must be capable of determining whether this host is communicating with an IPv4 or IPv6 peer. In a dual-stack network, you need to have a routing protocol that can deal with both protocols (such as IS-IS) or a routing protocol for the IPv4 network (such as OSPFv2) and another routing protocol for the IPv6 network (such as OSPFv3). If you are using dual-stack techniques, make sure that you have firewalls in place that protect not only your IPv4 network, but also your IPv6 network, and remember that you need separate security concepts and firewall rules for each protocol.

The disadvantages are known from other tunneling techniques used in the past. Additional load is put on the router. The tunnel entry and exit points need time and CPU power for encapsulating and decapsulating packets. They also represent single points of failure. Troubleshooting gets more complex because you might run into hop count or MTU size issues, as well as fragmentation problems. Management of encapsulated traffic (e.g., per-protocol accounting) is also more difficult due to encapsulation. Tunnels also offer points for security attacks

Tunneling – *Enables IPv6 islands or individual nodes to communicate over an IPv4 network*

In a scenario where, different IP versions exist on intermediate path or transit networks, tunneling provides a better solution where user's data can pass through a non-supported IP version.

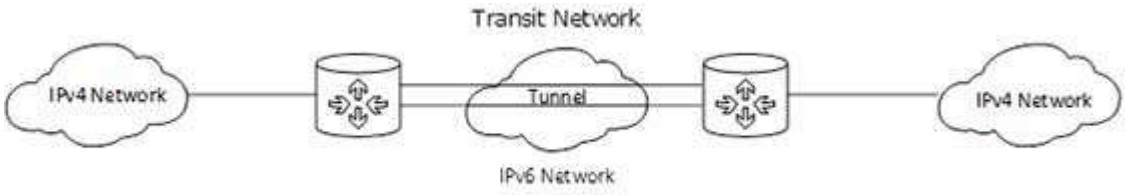

[*Image: Tunneling*]

The above diagram depicts how two remote IPv4 networks can communicate via a Tunnel, where the transit network was on IPv6. Vice versa is also possible where the transit network is on IPv6 and the remote sites that intend to communicate are on IPv4.

Tunneling Configurations

RFC 4213 defines the following tunneling configurations to tunnel IPv6 traffic between IPv6/IPv4 nodes over an IPv4-only infrastructure:

■ **Router-to-router: Figure 11-6** Two IPv6/IPv4 routers connect two IPv6-enabled infrastructures over an IPv4-only infrastructure.

■ **Host-to-router and router-to-host: Figure 11-7** An IPv6/IPv4 host that resides within an IPv4-only infrastructure uses an IPv6-over-IPv4 tunnel to reach an IPv6/IPv4 router.

■ **Host-to-host: Figure 11-8** An IPv6/IPv4 node that resides within an IPv4-only infrastructure uses an IPv6-over-IPv4 tunnel to reach another IPv6/IPv4 node that resides within the same IPv4-only infrastructure.

Here: IPv6/IPv4 node Implements both IPv4 and IPv6 and is assigned both IPv4 and IPv6 addresses. **IPv6 node** Implements IPv6 and can send and receive IPv6 packets. An IPv6 node can be an IPv6-only node or an IPv6/IPv4 node.

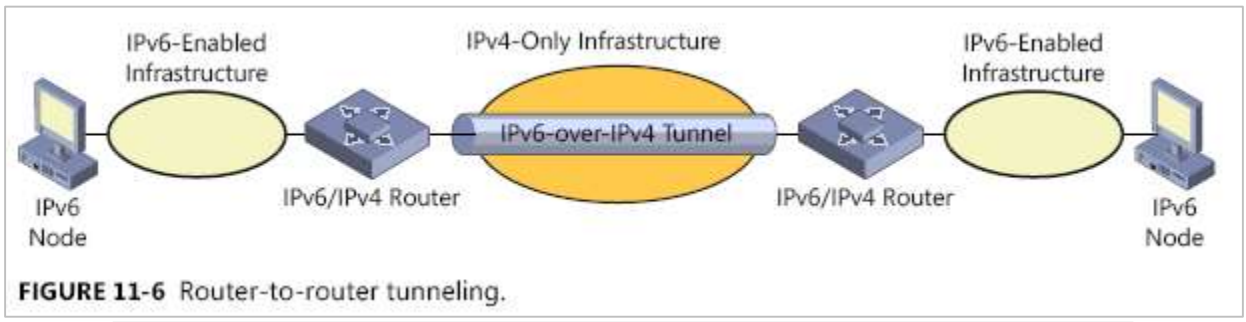

FIGURE 11-6 Router-to-router tunneling.

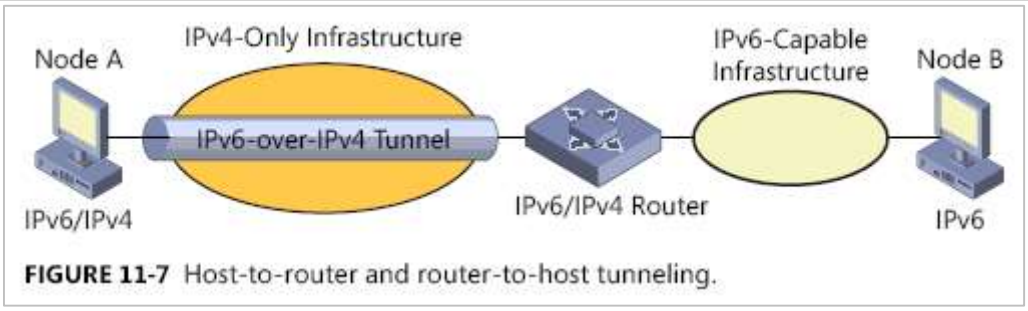

FIGURE 11-7 Host-to-router and router-to-host tunneling.

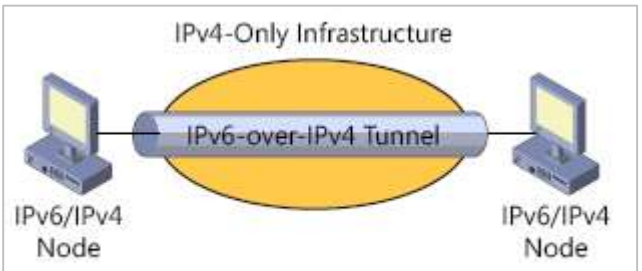

FIGURE 11-8 Host-to-host tunneling.

Types of Tunnels

1. Automatic Tunnels
2. Configured Tunnels:

5.1.1 Automatic Tunneling: *IPv4 Compatible IPv6, IPv6 Over IPv4 (6Over4), IPv6 to IPv4 (6to4)*

Automatic tunnels require IPv4–compatible addresses. Automatic tunnels can be used to connect IPv6 nodes when IPv6 routers are not available. These tunnels can originate either on a dual host or on a dual router by configuring an automatic tunneling network interface. The tunnels always terminate on the dual host. These tunnels work by dynamically determining the destination IPv4 address, the endpoint of the tunnel, by extracting the address from the IPv4–compatible destination address.

(i) IPv4 Compatible IPv6: *IPv6 traffic to be carried across an IPv4 infrastructure*

Description - RFC 4213, 2893:

- Automatic tunneling allows IPv6 traffic to be carried across an IPv4 infrastructure without the need for tunnel destination pre-configuration (P2MP)
- v4 address: 123.234.20.1
- v6 address: 0:0:0:0:0:123:234:20:1 or ::123:234:20:1

Benefits:

- Simple to deploy – no pre-configuration required
- Can use BOOTP, DHCP, RARP or manual configuration to obtain IPv4 address

- Allows transport of IPv6 packets over an IPv4 network using v4 tunnel endpoints
- Available on most platforms such as Cisco IOS and Microsoft XP

Issues:

- Requires a globally unique IPv4 address (No NAT tunnel endpoint allowed)
- Tunnel must not send IPv4 packets to: broadcast, multicast, loopback addresses
- 0:0:0:0:0:0/96 static route is required for automatic tunneling

Deployment Applications:

- A potentially cost-effective method of obtaining IPv6 connectivity
- Not recommended – currently being deprecated

(ii) IPv6 Over IPv4 (6OVER4): *IPv6 hosts to become fully functional v6 hosts without direct v6 connectivity on a IPv4 physical link*

Description – RFCs 4213, 2893, 2529:

- Allows isolated IPv6 hosts to become fully functional v6 hosts without direct v6 connectivity on a IPv4 physical link
- Often used in conjunction with configured tunneling: embedding the nodes IPv4 address into an IPv6 address 123.234.20.1 becomes 0:0:0:0:0:0:123:234:20:1
- Site local using IPv4 multicast as virtual link layer

Benefits:

- Simple to deploy
- Allows transport of IPv6 packets over an IPv4 network
- Available on most platforms
- Existing standard (RFC 2373 and RFC 2529)

Issues:

- Not Scalable
- 0:0:0:0:0:0/96 static route is required
- IPv6 multicast is implemented over IPv4 multicast

Deployment Applications:

- Used rarely

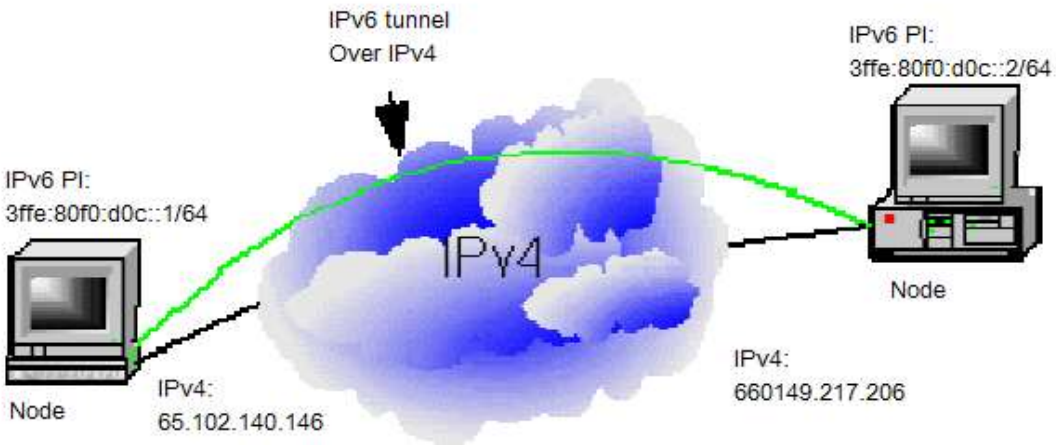

Fig. Tunnel Example

Fig. Encapsulation Example

(iii) Pv6 to IPv4 (6to4): *Encapsulation of IPv6 addresses automatically into IPv4 address*

Description – RFC 3056, 3068:

- Stateless automatic tunneling
- Encapsulation of IPv6 addresses automatically into IPv4 address (P2MP)
- IANA defined 2002:: /16 (improvement over use of IPv4 addresses)
- 123.234.20.1 becomes 2002:7bea:1401:: /48 (6to4 32-bit IPv4 GW address)

Benefits:

- Extremely easy for IPv6 "islands" located in IPv4 network to communicate
- Provides for 65,536 /64 networks with 2^{64} nodes/network
- Creates a globally unique /48 IPv6 prefix for use within the AS
- Good for IPv6 domains/sites with no IPv6 support

Issues:

- Requires one globally unique IPv4 address (No NAT) and 6to4 relay router
- Scales well for sites, NOT for individual hosts
- Number of problems remain for communication between an isolated IPv6 network and the IPv6 Internet
- Tunnel must not send IPv4 packets to: Broadcast, multicast, loopback addresses

Deployment Applications:

- 6Bone, switch.ch – gateway to other IPv6 clouds, APANA Melbourne, Australia
- Designed for site-to-site and site to existing IPv6 network connectivity
- Site border router must have at least one globally-unique IPv4 address
- Uses IPv4 embedded address

Example:

Reserved 6to4 TLA-ID:	**2002::/16**
IPv4 address:	**138.14.85.210 = 8a0e:55d2**
Resulting 6to4 prefix:	**2002:8a0e:55d2::/48**

- Router advertises 6to4 prefix to hosts via RAs
- Embedded IPv4 address allows discovery of tunnel endpoints

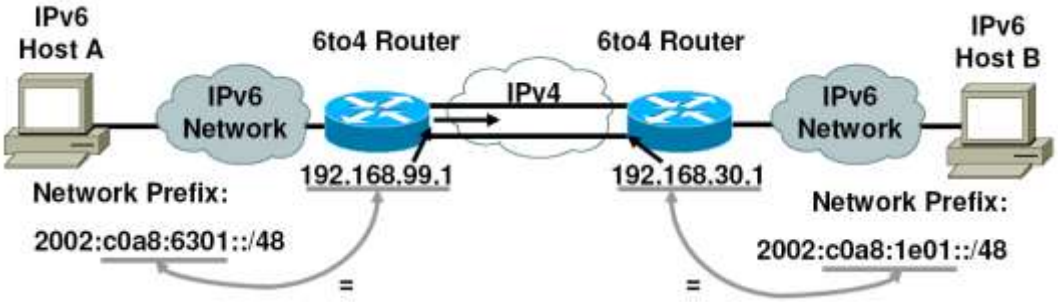

6to4:

– Is an automatic tunnel method

– Gives a prefix to the attached IPv6 network

5.1.2 Configured or Manual or Static or Point-to-point Tunneling

A configured tunnel requires manual configuration of the local and remote tunnel endpoints. In a configured tunnel, the IPv4 addresses of the remote tunnel endpoint are not embedded or encoded in the next-hop IPv6 address for the destination IPv6 address.

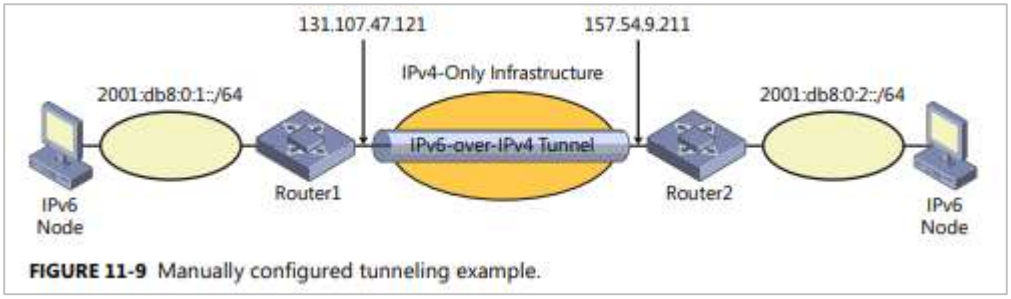

FIGURE 11-9 Manually configured tunneling example.

These tunnels will carry IPv6 traffic encapsulated within IPv4 packets, requiring that both tunnel endpoints must be dual stack. The encapsulation of IPv6 traffic within an IPv4 packet is as depicted below:

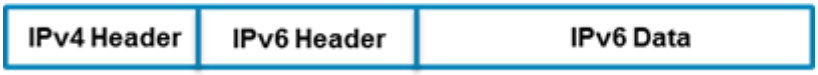

Because the tunnel options have to be manually set up, configured tunnels are not as flexible as the other automatic tunnel mechanisms; however, these tunnels are stable and straightforward to troubleshoot.

You can use configured tunnels to do the following:

- **Connect IPv6 islands on an intranet** For example, if you have different portions of your intranet that are not contiguous but are IPv6-capable, you can connect them with configured tunnels. This method is often used when service provider WAN solutions do not support IPv6 but you require end-to-end IPv6 connectivity.
- **Connect to the IPv6 Internet across the IPv4 Internet** Rather than using an ISP with native IPv6 connectivity, you can also connect to the IPv6 Internet through an ISP that offers tunneled connectivity to the IPv6 Internet. In this case, you configure an edge router of your organization with a configured tunnel to the ISP's router.

Description – RFC 4213, 2893:

- Tunneling allows IPv6 traffic to be carried across an IPv4 network
- Tunnel destination address is specified in the tunnel source configuration creating a P2P topology
- The tunnel acts as 1 hop for a IPv6 packet whereas an IPv4 encapsulation packet may take many hops

Benefits:

- Simple to deploy
- Allows transport of IPv6 packets over an IPv4 network
- Available on most platforms

Issues:

- Must be manually configured
- Potential (unknown) issues with delay and latency through the tunnel
- Additional CPU load for encapsulation/de-capsulation
- Single Point of Failure (tunnel endpoint)

Deployment Applications:

- Cost-effective method of obtaining IPv6 connectivity
- NTT commercial product/service, LONG Network (experimental)

5.2 Dual Stack

IPv6 was delivered with migration techniques to cover every conceivable IPv4 upgrade case, but many were ultimately rejected by the technology community, and today we are left with a small set of practical approaches. One technique, called dual stack, involves running IPv4 and IPv6 at the same time. End nodes and routers/switches run both protocols, and if IPv6 communication is possible that is the preferred protocol.

Description:

_ Deployment and utilization of IPv4 and IPv6 concurrently
_ v4 address: 123.234.20.1
_ v6 address: 2001:3f0:4c02:2:320:a6ff:fe4c:350f

Benefits:

_ Can be deployed on hosts, routers, on same interface as IPv4
_ Deals with most address selection and DNS resolution issues
_ Allows hosts to continue to reach IPv4 resources, while also adding IPv6 functionality
_ Simple to deploy. Allows backwards compatibility (IPv4 support)
_ Available on most platforms. Easy to use, flexible.

Issues:

_ May require 2 routing tables & routing processes
_ Additional CPU, memory
_ IPv6 network security requirements are same as IPv4 networks today – don't overlook enforcing security on parallel protocol

Application Deployment:

_ Easy way to deploy, flexible, bring up 2nd protocol in parallel with first

Implementing Dual-Stack

The term dual-stack normally refers to a complete duplication of all levels in the protocol stack from applications to the network layer. An example of complete duplication is the OSI and TCP/IP protocols that run on the same system. However, in the context of IPv6 transition, dual-stack means a protocol stack that contains both IPv4 and IPv6. The remainder of the stack is identical. Consequently, the same transport protocols, TCP, UDP, and so on, can run over both IPv4 and IPv6. Also, the same applications can run over both IPv4 and IPv6.

The following figure illustrates dual-stack protocols through the OSI layers. In the dual-stack method, subsets of both hosts and routers are upgraded to support IPv6, in addition to IPv4. The dual-stack approach ensures that the upgraded nodes can always interoperate with IPv4-only nodes by using IPv4.

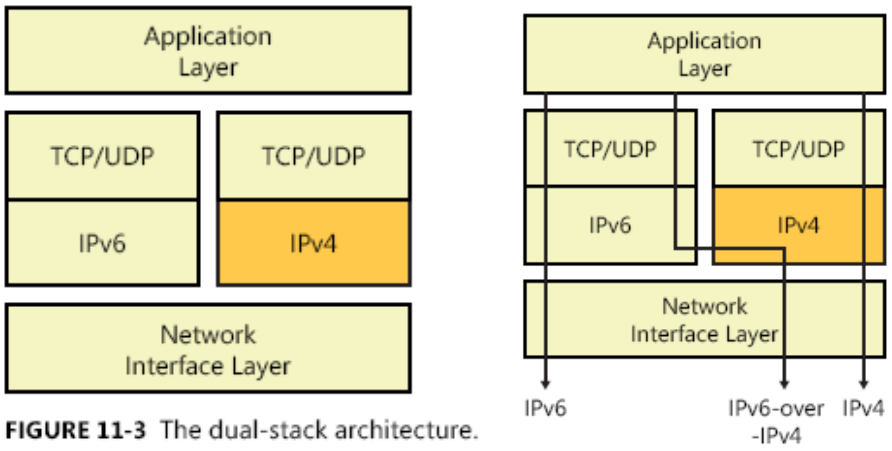

FIGURE 11-3 The dual-stack architecture.

Fig. Dual Stack

IPv4/IPv6 Dual-Stack Mechanism

As the word means, dual-stack mechanisms include two protocol stacks that operate in parallel and allow network nodes to communicate either via IPv4 or IPv6. They can be implemented in both end system and network node. In end systems, they enable both IPv4 and IPv6 applications to operate at the same time. The Dual-stack capabilities of network nodes support the transport of both IPv4 and IPv6 packets.

Dual Stack Routers – IPv6 with IPv4 network

A router can be installed with both IPv4 and IPv6 addresses configured on its interfaces pointing to the network of relevant IP scheme.

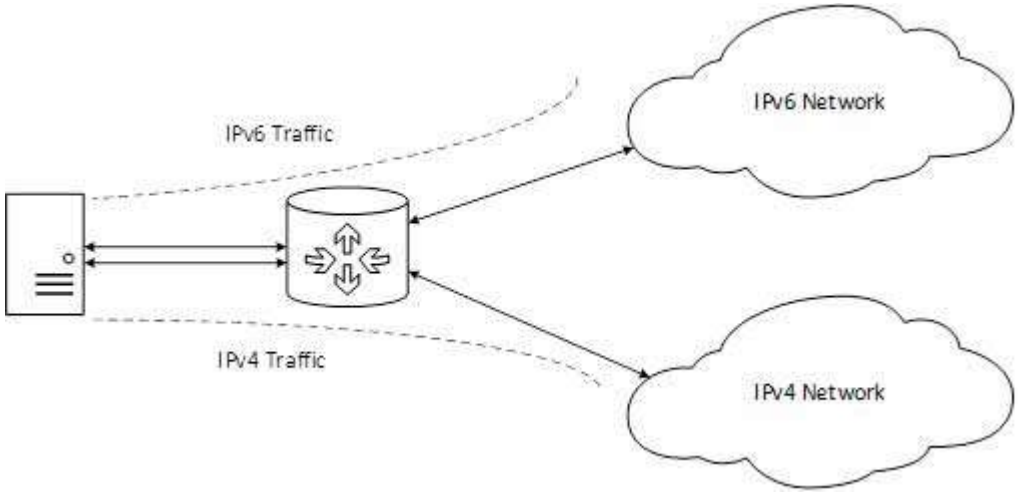

[*Image: Dual Stack Router*]

In the above diagram, a server having IPv4 as well as IPv6 address configured for it can now speak with all the hosts on both the IPv4 as well as the IPv6 networks with the help of a Dual Stack Router. The Dual Stack Router, can communicate with both the networks. It provides a medium for the hosts to access a server without changing their respective IP versions.

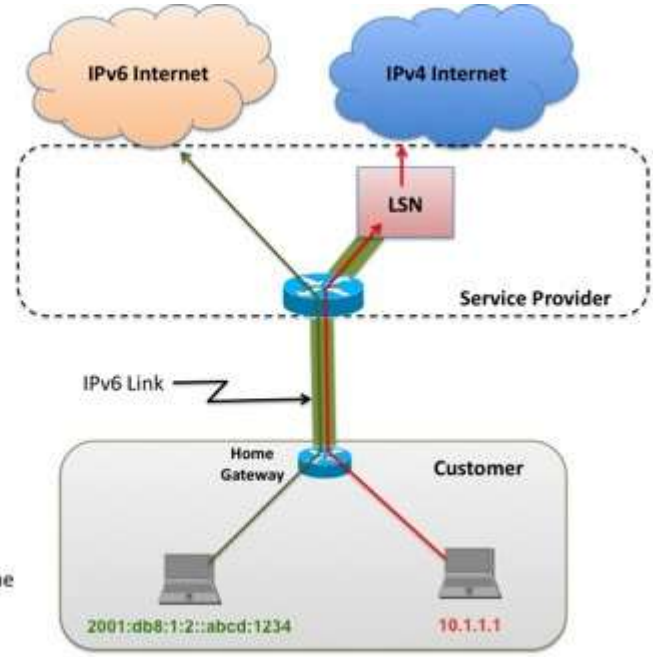

Figure 2
IPv6 packets are routed normally, while IPv4 packets are routed to the LSN.

Fig. Dual Stack

Dual Stack technique is easy to use and flexible. Hosts can communicate with IPv4 hosts using IPv4 or communicate with IPv6 hosts using IPv6. When everything has been upgraded to IPv6, the IPv4 stack can simply be disabled or removed. Whenever you can, deploying dual-stack hosts and routers offers the greatest flexibility in dealing with islands of IPv4-only applications, equipment, and networks. Dual stack is also the basis for other transition mechanisms. Tunnels need dual-stacked endpoints, and translators need dual-stacked gateways. Disadvantages of this technique include the following: you have two separate protocol stacks running, so you need additional CPU power and memory on the host. All the tables are kept twice: one per protocol stack. A DNS resolver running on a dual-stack host must be capable of resolving both IPv4 and IPv6 address types. Generally, all applications running on the dualstack host must be capable of determining whether this host is communicating with an IPv4 or IPv6 peer. In a dual-stack network, you need to have a routing protocol that can deal with both protocols (such as IS-IS) or a routing protocol for the IPv4 network (such as OSPFv2) and another routing protocol for the IPv6 network (such as OSPFv3). If you are using dual-stack techniques, make sure that you have firewalls in place that protect not only your IPv4 network, but also your IPv6 network, and remember that you need separate security concepts and firewall rules for each protocol.

5.3 Translation: *enable IPv4-only devices to speak to IPv6-only devices, which is not possible with any of the tunneling methods.*

- The concept of address translation is also not a new concept to most network engineers; this is because Network Address Translation (NAT) is implemented between different IPv4 networks in almost every residential household. The concept behind this type of NAT and the newer technologies that support address translation between IPv4 and IPv6 networks is similar. IPv6 translation technologies differ from IPv6 tunneling technologies; this is because the translation technologies enable IPv4-only devices to speak to IPv6-only devices, which is not possible with any of the tunneling methods.
- However, IPv4/IPv6 translation and IPv4-only translation cause a certain amount of complexity. What happens when an IPv6-only device is attempting to communicate with a device on the public IPv4 Internet and only an IPv4 DNS record (A) exists? In these situations, a secondary technology is required to step in and provide additional services for the connection to work.
- The first method to be introduced to provide IPv6 translation services was Network Address Translation - Protocol Translation (NAT-PT). NAT-PT defined a mechanism to not only translate between IPv4 to IPv6 addresses but also a built-in ability to provide protocol translation services for Internet Control Message Protocol (ICMP), File Transfer Protocol (FTP), and Domain Name System (DNS). The component that was responsible for these translation services is called the application layer gateway (ALG).

IPv4/IPv6 Translation Mechanism

The basic function of translation in IPv4/IPv6 transition is to translate IP packets. Several translation mechanisms are based on the SIIT (Stateless IP/ICMP Translation algorithm) algorithm [16]. The SIIT algorithm is used as a basis of the (i) BIS (Bump In the Stack) and (ii) NAT-PT (Network Address Translation-Protocol Translation) mechanisms,

NAT Protocol Translation – IPv6 Network to IPv4 host

This is another important method of transition to IPv6 by means of a NAT-PT (Network Address Translation – Protocol Translation) enabled device. With the help of a NAT-PT device, actual can take place happens between IPv4 and IPv6 packets and vice versa. See the diagram below:

A host with IPv4 address sends a request to an IPv6 enabled server on Internet that does not understand IPv4 address. In this scenario, the NAT-PT device can help them communicate. When the IPv4 host sends a request packet to the IPv6 server, the NAT-PT device/router strips down the IPv4 packet, removes IPv4 header, and adds IPv6 header and passes it through the Internet. When a response from the IPv6 server comes for the IPv4 host, the router does vice versa.

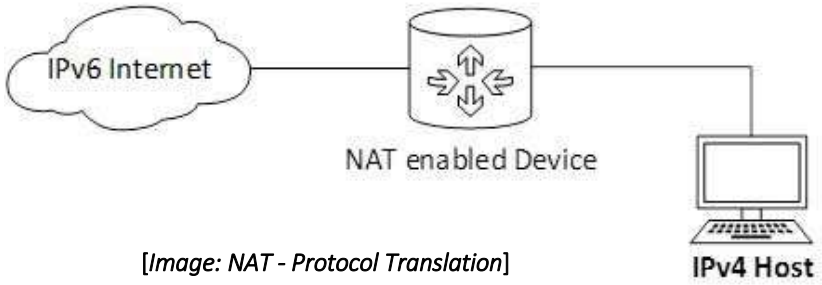

[*Image: NAT - Protocol Translation*]

The concept of address translation is also not a new concept to most network engineers; this is because Network Address Translation (NAT) is implemented between different IPv4 networks in almost every residential household. The concept behind this type of NAT and the newer technologies that support address translation between IPv4 and IPv6 networks is similar. IPv6 translation technologies differ from IPv6 tunnelling technologies; this is because the translation technologies enable IPv4-only devices to speak to IPv6-only devices, which is not possible with any of the tunnelling methods.

However, IPv4/IPv6 translation and IPv4-only translation entail a certain amount of complexity. What happens when an IPv6-only device is attempting to communicate with a device on the public IPv4 Internet and only an IPv4 DNS record (A) exists? In these situations, a secondary technology is required to step in and provide additional services for the connection to work.

The first method to be introduced to provide IPv6 translation services was Network Address Translation - Protocol Translation (NAT-PT). NAT-PT defined a mechanism to not only translate between IPv4 to IPv6 addresses but also a built-in ability to provide protocol translation services for Internet Control Message Protocol (ICMP), File Transfer Protocol (FTP), and Domain Name System (DNS). The component that was responsible for these translation services is called the application layer gateway (ALG).

The ALG piece of the NAT-PT method raised a number of issues. With additional testing and real-life experience, a new method was introduced that separated the address translation functionality and the application layer translation functionalities: NAT64 and DNS64.

DNS64 can synthesize IPv6 address resource records (AAAA) from IPv4 resource records (A); it does this by encoding the returned IPv4 address into a IPv6 address format.

Translation should be used only if no other technique is possible and should be viewed as a temporary solution until one of the other techniques can be implemented. The disadvantages are that it does not support the advanced features of IPv6, such as end-to-end security. It poses limitations on the design topology because replies have to come through the same NAT router from which they were sent. The NAT router is a single point of failure, and flexible routing mechanisms cannot be used. All applications that have IP addresses in the payload of the packets will stumble. The advantage of this method is that it allows IPv6 hosts to communicate directly with IPv4 hosts and vice versa. For the reasons mentioned previously, NAT as described in RFC 2765 and RFC 2766 is going to be moved to experimental.

5.4 Migration Strategies for Telcos and ISPs

Network administrators will soon be forced to migrate to IPv6. This is due to not only technology moving forward into a new standard, but the fact that the number of IPv4addresses is nearly exhausted. It is only a matter of time before the change is mandatory to maintain Internet connectivity.

However, this is not a common migration. There are scans to do, information to gather, and plans to make, among other tasks. In addition, for the best result, an organization needs to go through these tasks together. There will be unexpected problems, and a long list of concerns from everyone involved, but experience shows that most of the tasks will fall to the network administrator and the IT team. This document seeks to examine and assist in solving the top 5 concerns for migrating to IPv6 of network administrators.

1. **Strategy One: Do Nothing** – IPv4 only network
2. **Strategy Two: Extend life of IPv4 network**
   a. Force customer to use NAT

b. Acquire IPv4 address from another organization – IPv4 subnet trading

3. **IPv4 / IPv6 Coexistence / Transition techniques**
   - **a.** Dual stack network
   - **b.** 6rd (Rapid Deploy)
   - **c.** Large Scale NAT (LSN)

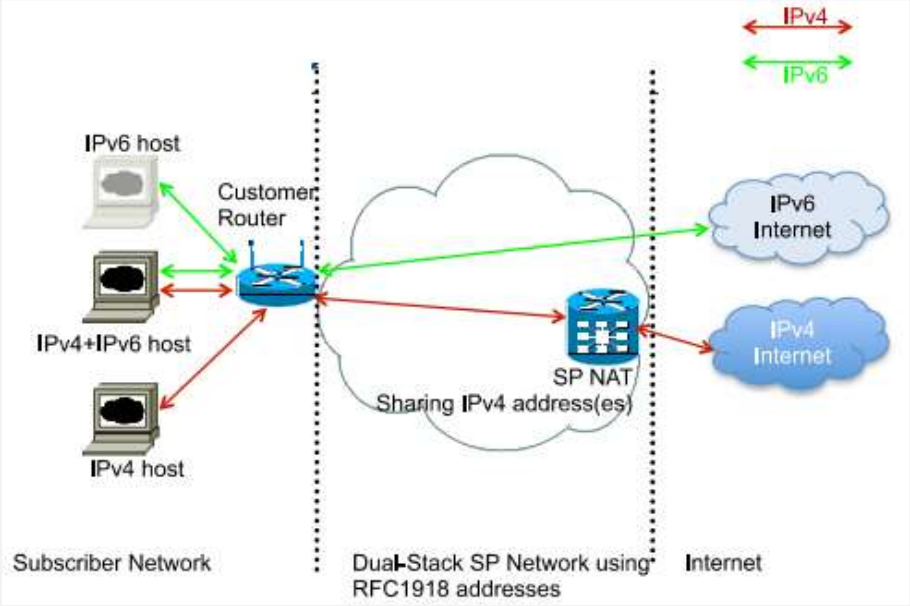

Fig. Dual Stack

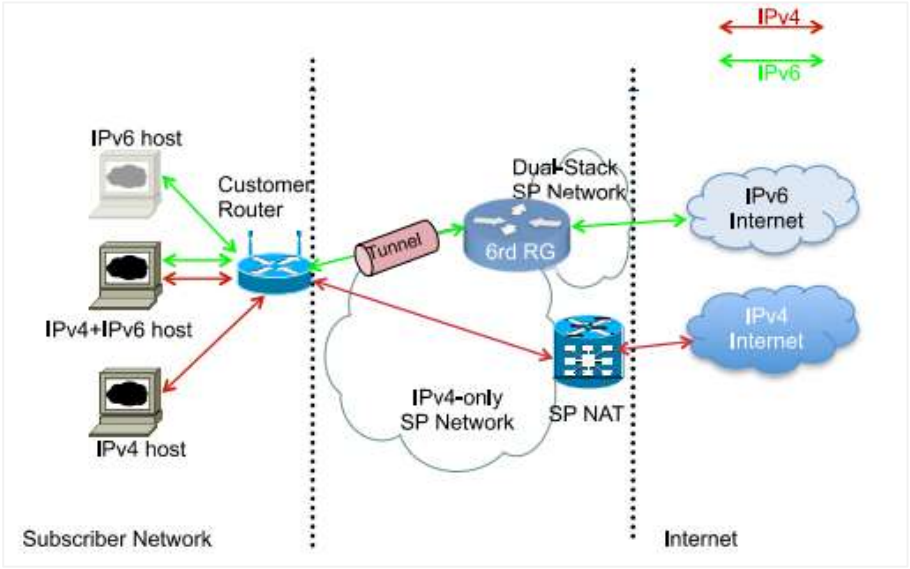

Fig. 6rd

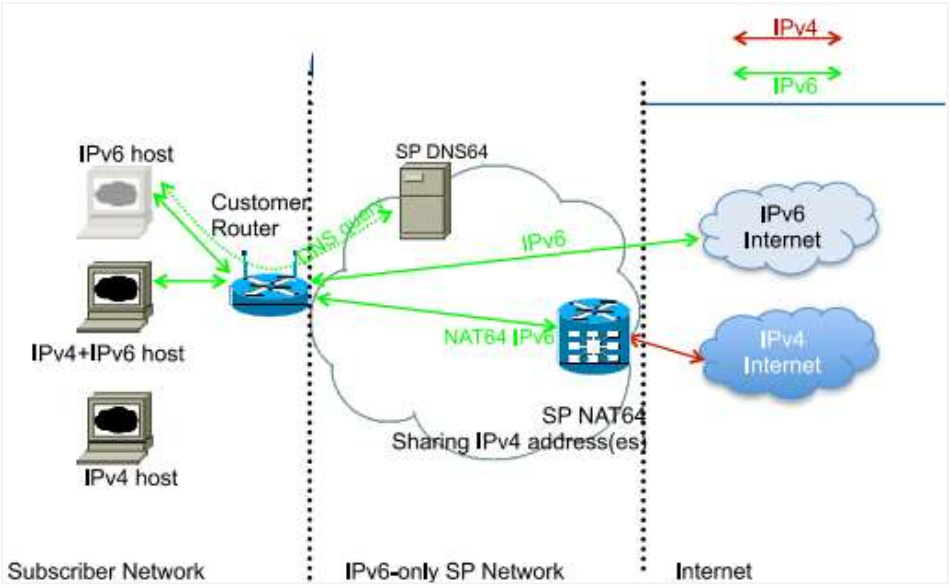

Fig. NAT 64

Strategize

- Don't plan IPv6 the same way you did IPv4. Take a step back and figure out how you would deploy the network if you had absolutely unlimited IPs
- Think through your addressing scheme carefully. Make sure it is scalable and well understood by administrators
- Scrub your technical research to eliminate outdated information. Do not assume the most popular search engine hit is the most recent way IPv6 is done
- Configure your networks in stages out to the customer/business
- Carefully consider how you will deploy IPv6 in the last mile (firewalls, DNS, DHCP, etc.)
- Plan your address space properly and logically in order to keep the routing table small
- Get your own address space and ASN so you can be provider independent