happiest minds
The Mindful IT Company
Born Digital . Born Agile

Network transitioning from IPv4 to IPv6 Document

Happiest People . Happiest Customers

Contents

Introduction

As enterprises are moving towards IPv6 it is imaginable that the migration process will happen gradually even within a single enterprise. To facilitate this gradual migration several migration techniques have been defined. This document details about the transitioning techniques and process being followed by Happiest Minds.

- **Scope**

This document will provide guidelines for network transitioning from IPv4 to IPv6.

- **Purpose**

The purpose of this document is to provide an overview of the transitioning techniques and standards.

- **Executive Summary**

As enterprises are moving towards IPv6 it is imaginable that the migration process will happen gradually even within a single enterprise. To facilitate this gradual migration several migration techniques have been defined. The following transition techniques provide a mechanism to gradually migrate from an IPv4 infrastructure to an IPv6 infrastructure:

- Dual Stack - The ability to assign an IPv4 and an IPv6 address to the same network interface
- Tunneling - Static and automatic tunnelling of IPv6 packets over an IPv4 infrastructure
- NAT-PT - Static or Dynamic Network Address Translation - Protocol Translation

IPv4 to IPv6 transition stages

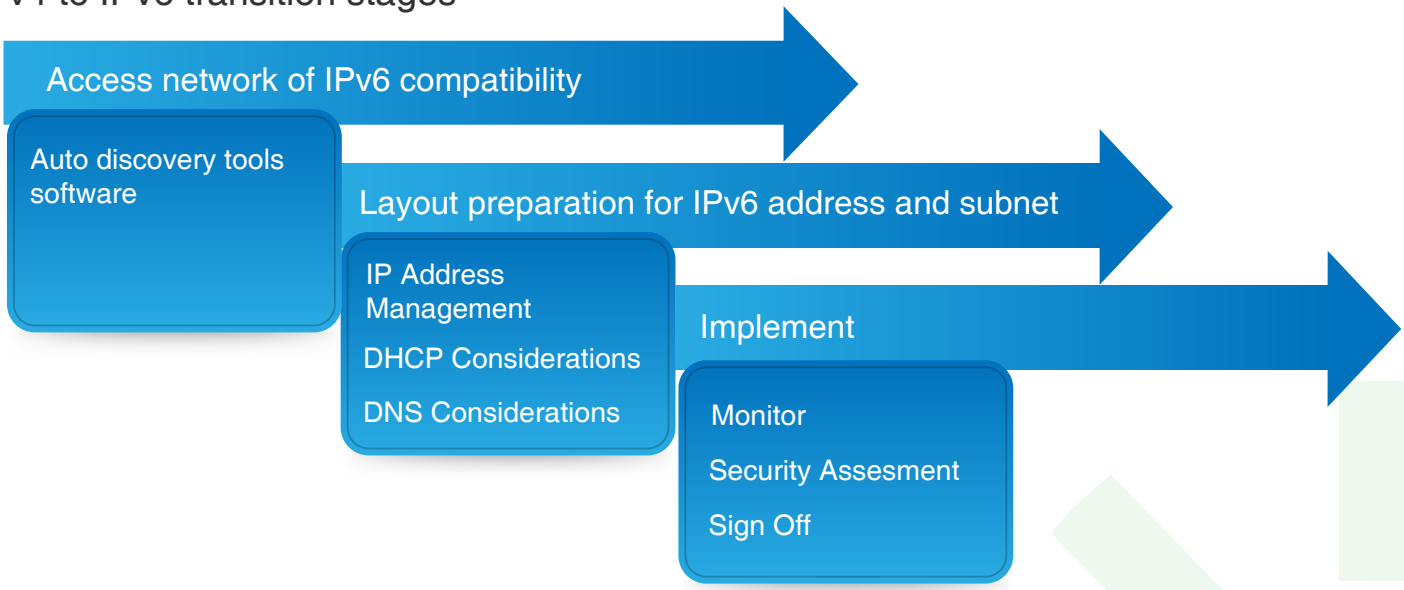

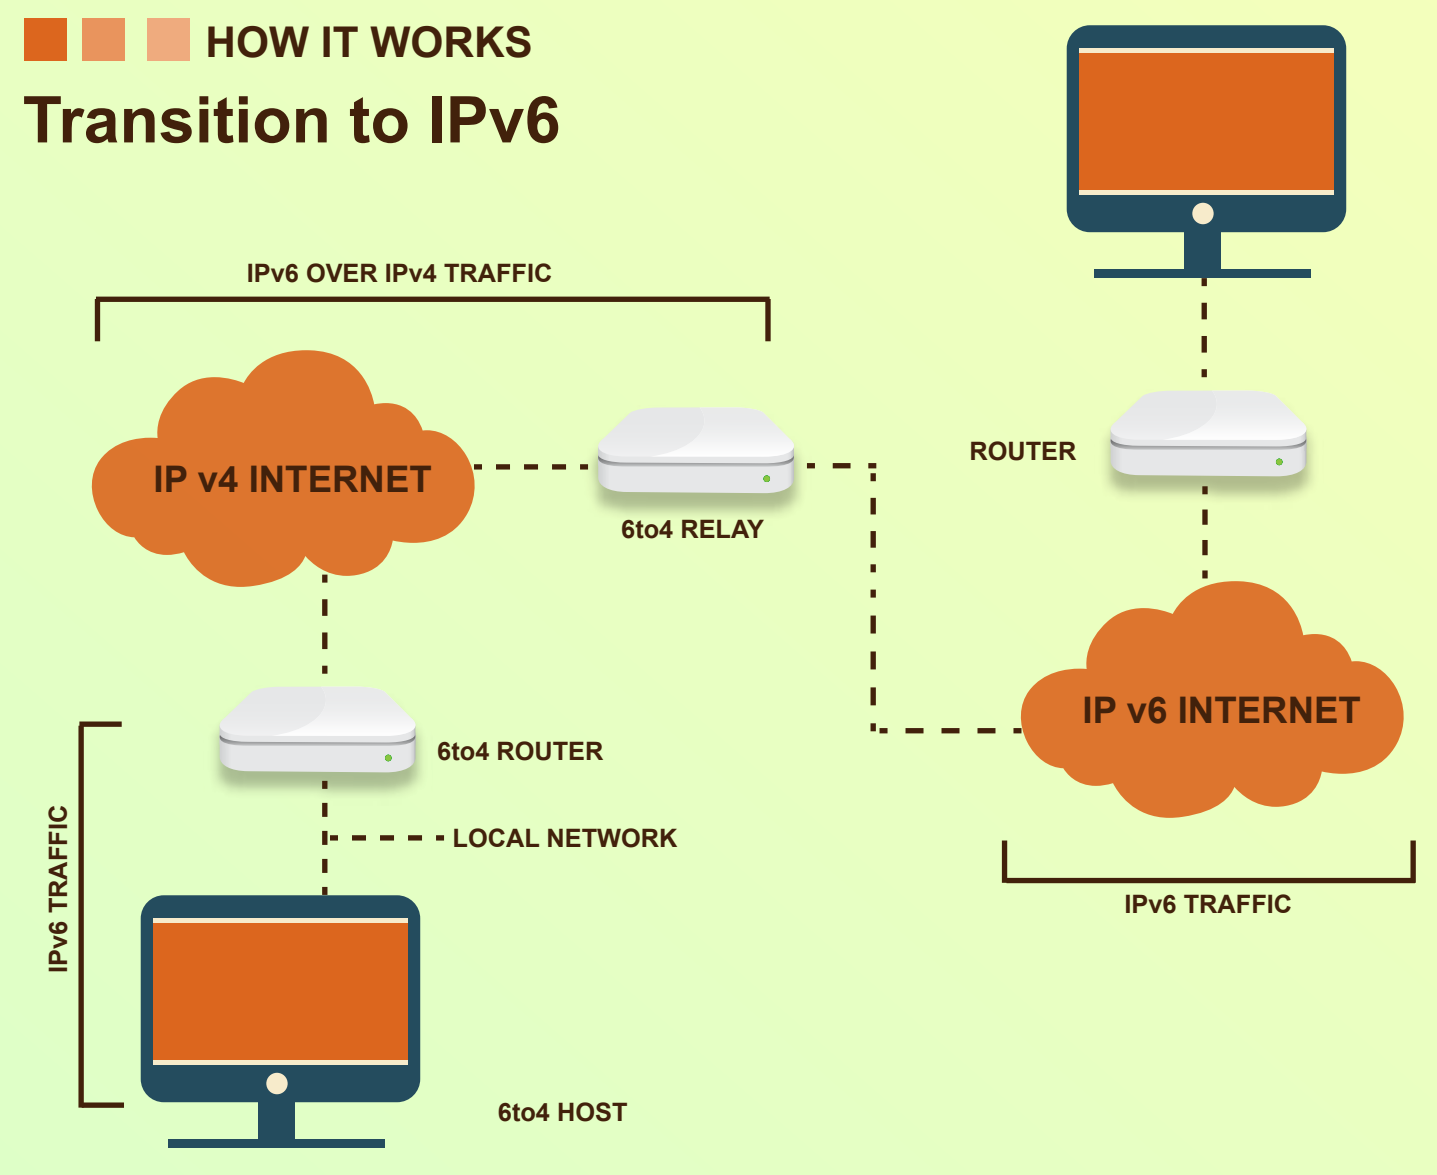

Dual Stack

Dual stack is the core technique which the other transition techniques are builds upon. As stated earlier Dual Stack is the technique to allow an IPv4 and IPv6 address to be defined on the same network interface.

Tunnelling

Static Tunnels, Point to Point Mechanisms

The advantage of doing static tunnels is that dynamic routing protocols can be run over the tunnel. There are 2 variants involved here for static tunnels;

- GRE - Default tunnel mode
- IPv6IP - Less overhead, no CLNS (IS-IS) transport Configuring static tunnels is a pretty straight forward process and the difference between configuring a GRE or a IPv6 IP tunnel is marginal.

(config) ipv6 unicast-routing
(config) interface INTERFACE-ID
(config-if) ipv6 address ADDRRESS

Next
(config) interface tunnel NUMBER
(config-if) tunnel source INTERFACE/IPv4-ADDRESS
(config-if) tunnel destination IPv4-ADDRESS
(config-if) ipv6 address IPv6-ADDRESS
(config-if) tunnel mode ipv6ip <-- tunnel IPv6 over
IPv4. GRE tunnel mode is the default

Automatic tunnels, Point to Multipoint Mechanisms

- 6to4 site to site tunnel
- ISATAP host to host & host to router tunnel protocol (Within a Site)

The advantage of automatic tunneling is that only one tunnel is needed (together with BGP or static routing) in order to establish a full mesh of connectivity. IPv6 prefix 2002: is reserved for the use of automatic tunneling. The IPv4 tunnel address is converted to Hex and pre-ended with 2002:. For instance with a tunnel IPv4 address of 1.1.1.1 the corresponding IPv6 subnet is 2002:0101:0101::/48.

6to4 Tunnels

The configuration of a 6to4 tunnel is needed when one IPv6 site has to be connected with another IPv6 site through an IPv4 infrastructure. The principle is simple as the mechanism to use is to create a tunnel interface. The core part is to calculate the two 48 bit IPv6 subnets to be used between the two sites. As we saw earlier these subnets are calculated from the IPv4 IP address on both site of the point to point link. After these 2 subnets have been calculated on the sites, a static route to the remote 48 bit subnet pointing to the remote site's 2002:: address is all what is needed to connect these two IPv6 sites together. At each site the 48 bit subnet can be divided into several 64 bi
t prefixes for a hierarchical addressing scheme .
(config) ipv6 unicast-routing
(config) interface INTERFACE-ID
(config-if) ipv6 address ADDRRESS
(config) interface tunnel NUMBER
(config-if) tunnel source INTERFACE/IPv4-ADDRESS <-- no tunnel destination, we are multipoint
(config-if) ipv6 address 2002::[EXAMPLE]
(config-if) tunnel mode ipv6ip 6to4
As stated earlier a static route must be configured to the next-hop address to each endpoint. Static routes can be omitted by using (e)BGP to peer using the 6to4 addresses between the two endpoints. By redistributing on each site the IPv6 IGP routes into BGP full IPv6 connectivity can be obtained between the sites.

Teredo tunnelling

Instead of using routers to tunnel packets, Teredo tunneling has the hosts perform the tunneling. This requires the hosts to be configured with double stacks. It is mostly put to use to move about packets through an IPv4 address translation device.

ISATAP Tunnels

ISATAP tunnels are needed when point to multipoint connections must be made within one site. What ISATAP does is use an improvised EUI-64 mechanism, resulting in the link-local address (FE80:) to be taken into account. So ISATAP is used to connect two dispersed IPv6 islands in a single site. As the whole mechanism is based on the tunnel, this concept is not the hardest to understand. The important difference here is in how the IPv6 endpoints addresses are calculated. The rest of the concept remains the same as with the 6to4 tunneling mechanism. As the ISATAP host prefix is calculated automatically the configuration is even more straightforward than the 6to4 method.

(config) ipv6 unicast-routing
(config) interface INTERFACE-ID
(config-if) ipv6 address ADDRRESS
(config) interface tunnel NUMBER
(config-if) tunnel source INTERFACE/IPv4-ADDRESS <-- no tunnel destination, we are multipoint
(config-if) ipv6 address 100:100::/64 eui-64
(config-if) tunnel mode ipv6ip isatap

The modified EUI-64 mechanism calculates the host prefix/identifier as follows; the first 32 bits are always 0000:5EFE and the last32 bits are the IPv4 address in Hex. The first 64 bit network prefix can be anything and is not bound to any rules. With ISATAP the same routing considerations applies as for the 4to6 mechanism and will not be repeated here.

IPv6 translation

IPv6 translation schemes implement some form of packet-header translations between the IPv6 and IPv4 addresses. The goal is to translate packets with IPv6 addresses to those with IPv4 addresses, so that IPv6-only hosts can talk to the IPv4-only Internet. This may sound not very complex but in some cases it can be. For example, some server load balancers (SLBs) are capable of load balancing and translating packets with IPv6 addresses to IPv4 packets. This competency may turn out to be advantageous in data center deployments in that the existing IPv4 infrastructure can remain unchanged, and only the SLB requires modification. This kind of technique requires other mechanisms (such as tunneling or dual-stack techniques) to get the IPv6 packets to the SLB. However, as a means for accessing the IPv4 Internet, the translation technique can actually be quite complex. Originally, NAT-Protocol Translation (NAT-PT: RFC 2766) was proposed for this cause, but it has consequently been considered non-practical to deploy (due to the need for application layer gateways ALGs), the management of IPv6 Domain Name System (DNS) requests, and the need for session initiation ordering for connection state establishment). In addition, IPv6 translation approaches break the Internet end-to-end connectivity model (a common trait of all NAT implementations), and cannot handle fragmentation in the core.

Benefits of IPv6 transitioning

More Efficient Routing: IPv6 reduces the size of routing tables and makes routing more efficient. IPv6 lets ISPs aggregate the prefixes of their customers' networks into a single prefix and announce this one prefix to the IPv6 Internet. On top of that, in IPv6 networks, the process of fragmentation is taken care of by the sourcing device, rather than the router, using some sort of a protocol to discover the path's maximum transmission unit (MTU).

More Efficient Packet Processing: IPv6's simplified packet header makes packet processing way more effective. Compared to IPv4, IPv6 has no IP-level checksum, so it does not need to be recalculated at every hop of the router. Not using the IP-level checksum was possible because most link-layer technologies already contain it and other error-control capabilities. In addition, most of the transport layers, that handle an end-to-end connectivity, have a checksum that calls for enabling the detection of errors to a great extent.

Directed Data Flows: IPv6 supports multicast more than broadcast. Multicast lets bandwidth-intensive flows (like multimedia streams) to be sent across to multiple destinations simultaneously, saving the network bandwidth to the greatest extent. Hosts with no interest of any sort no longer must process packets of broadcasts. What's more, the IPv6 header has a new field, named Flow Label that can identify packets belonging to the same flow.

Simplified Network Configuration: Address auto-configuration (address assignment) is built in to IPv6. A router will send the prefix of the local link in its router commercials. A host, on the other hand is capable of generating its independent IP address by appending its link-layer address, by converting the same into an Extended Universal Identifier (EUI) 64-bit format, to the 64 bits of the local link prefix.

Support for New Services: By eliminating Network Address Translation (NAT), true end-to-end connectivity at the IP layer is hence retained by enabling new and more valuable services. The advantage with Peer-to-peer network solutions is that they are way easier to develop and maintain, and services such as VoIP and Quality of Service (QoS) become more robust.

Security: IPsec, which provides privacy, authenticity and data compatibility, is wired into in IPv6. Due to its ability to carry forward malware, IPv4 ICMP packets are most often blocked by firewalls that are corporate centric, but ICMPv6, the execution of the Cyber Controlled Message Protocol for IPv6, could be allowed because IPSec can be implemented in the ICMPv6 packets as well.

Which technique to choose

There is no straightforward answer to this question and difficulty compounds when we consider the vast list of possible techniques suggested in previous sections. Network operators and administrators choose to experiment with a few techniques and the IETF V6OPS workgroup is working towards providing guidance in the choice of these techniques for different usage scenarios. Further, since multiple transition techniques are defined, it is likely that multiple techniques can be used within a local network and hence the network administrator has to consider issues arising out of combinations of techniques. For example, DSTM and ISATAP have opposite functions and should not be used on the same host. The network administrator would require comprehensive tools to configure and manage a network where multiple transition techniques interact with each other.

Things to consider before transition

- Routing protocol
- DNS
- Network management

Offerings

Transitioning from IPv4 to IPv6 requires detailed analysis of the existing infrastructure. We as Happiest Minds believe in providing end-to-end solution and services. The services that we offer that distinguish in the market:

Consultation and planning services:

- Understanding of the requirement
- Detailing of the scope with respect to the existing infrastructure
- Highlighting the key areas of issue
- Validation or Preparation of the hardware inventory for IPv6 enabled NICs
- Asset or infrastructure proposal
- Selection of the transition technique to be used

Transition services:

- Implementation
- Function validation
- Security validation

Support services:

- Operational support

Summary

This document gives an overview of the technologies involved in the IPv4 to IPv6 transitioning. It also lists all the services being offered.

Authors

Kapil Chaturvedi

Kapil Chaturvedi has 11 years of IT experience, including 3 years in Data center planning and implementation and 2 years in wireless planning and implementation. He has worked in multivendor environment including Cisco, Aruba, HP, Nortel, F5, Watchguard, Fortigate and others. He has also worked in Open source technologies and providing solutions based upon them. Key past project includes: Data center planning for one of the client at Malaysia having multiple sites geographically distributed across the globe, NOC planning and new project implementation for Standard Chartered bank, Wireless implementation for Aruba Networks Inc.

Happiest Minds

Happiest Minds, the Mindful IT Company, applies agile methodologies to enable digital transformation for enterprises and technology providers by delivering seamless customer experience, business efficiency and actionable insights. We leverage a spectrum of disruptive technologies such as: Big Data Analytics, AI & Cognitive Computing, Internet of Things, Cloud, Security, SDN-NFV, RPA, Blockchain, etc. Positioned as "Born Digital . Born Agile", our capabilities spans across product engineering, digital business solutions, infrastructure management and security services. We deliver these services across industry sectors such as retail, consumer packaged goods, edutech, e-commerce, banking, insurance, hi-tech, engineering R&D, manufacturing, automotive and travel/transportation/hospitality.

Headquartered in Bangalore, India; Happiest Minds has operations in USA, UK, The Netherlands, Australia and Middle East.

Copyright Information

Business Contact: business@happiestminds.com

Visit us: www.happiestminds.com

Follow us on