

-
- Hierarchical network architecture - The large address space of IPv6 allows addresses to be allocated hierarchically which enables more efficient internet routing by allowing internet routers to achieve greater levels of route table summarization.
 - Enhanced flexibility of multicast management – The larger address space of IPv6 enables many more multicast addresses to be supported, optimizing the way multicast traffic can be handled in the network. Broadcasts which go to all hosts in a network, and the accompanying problem of broadcast storms, are no longer supported with IPv6 and only multicast groups are supported.
 - Simplified configuration options – IPv6 incorporates options for auto-configuration of hosts without DHCP servers to simplify configuration and network attachment. While auto-configuration may not be recommended for larger enterprise networks, it does offer a simple implementation option where these mechanisms are appropriate.
 - Enhanced mobility support – IPv6 supports options which are optimized for handling mobile network devices such as smart phones and wireless devices. Mobile IPv6 allows more efficient routing to devices which are not statically located.
 - Embedded security mechanisms – IPSec is a mandatory integral component of IPv6. Enabling these IPSec capabilities allows for easier implementation of encryption, authentication, and VPNs.

IPv6 migration – General Challenges, Preparation and Planning

IPv6 is not only a network infrastructure feature but it also affects the entire IT infrastructure. It extends from server to clients, from edge to core, from OS to applications. Careful planning and assessment is required to make the migration.

The scope of planning for an IPv6 migration goes beyond the networking group. Network Administrators should be aware that IP addressing is embedded into multiple applications and protocols, including:

- ICMP
- FTP
- SIP/H323
- IM protocols
- DNS
- P2P protocols

Migration to IPv6 must be embraced in an end-to-end basis by all IT departments in an organization, detecting all applications and protocols that require migration and applying the patches needed.

From a communications perspective special attention must be paid to:

- Firewalls, ACLs and IPSs – new ICMP types need to be managed in order for IPv6 to work at all. Filter and chains rewrite might be needed to adopt the new addressing.
- DNS – IPv6 uses new DNS AAAA records that must be created, usually in parallel with existing DNS A records for IPv4.
- DHCP servers – although IPv6 brings new auto configuration modes, DHCP provides a level of control and management which most enterprises will want to retain. For this reason, DHCPv6 is going to be the main address management tool in IPv6.
- Communication servers – SIP and H323 are impacted in their internals by IPv6 addressing so these systems need to be patched/verified for compatibility.
- Management systems – network and systems managers must be checked for compatibility with IPv6.
- Routers – must be able to route IPv6 packets.
- Other network infrastructure – switches and L2 devices must support IPv6 addresses for management and an IPv6 host stack.
- Load balancers – must support IPv6 VIPs and hashing of IPv6 addresses.

As you prepare for actual implementation, IT architects should start with proper address planning and allocation:

- Determine your IPv6 address architecture - Since IPv4 and IPv6 will coexist in your network for the foreseeable future, careful analysis needs to be made as to how the new IPv6 address scheme will overlay your current IPv4 scheme. The flexibility of IPv6 addressing is also a potential pitfall since there will be numerous ways a new IPv6 address architecture could be implemented. Even the IPv6 technical community is sometimes divided on what the best practices actually encompass. An understanding of the growth and future requirements of your network will allow a constructive dialogue with Enterasys to help assess the best approaches to take.
- IPv6 Internet - Apply and register a global address prefix by your Internet service provider (ISP) or Regional Internet Registry (RIR).

Transition mechanisms, or how IPv4 and IPv6 will coexist in your network, are a major consideration for planning any IPv6 migration. The big picture migration strategies include, in order of preference:

- IPv4/IPv6 Dual Stack
- Tunneling. IPv6 at the edge and tunnel over/through existing IPv4 networks to the other IPv6 networks
- Translation between IPv6 only devices and IPv4 only devices

In most of today's enterprises the dual-stack migration is the recommended migration strategy. Tunneling techniques should only be used to address specific design challenges. Address translation is not fully standardized and brings the complexities of NAT back into the picture. While NAT may remain a tool available for transition, it is certainly desirable that it only be used for exceptional cases where other options are not practical. Hence, we will focus on Tunneling and Dual Stack techniques.

Most servers in the enterprise are already IPv6-ready, but a thorough check will avoid issues during the migration. Custom or in-house developed applications require greater attention to ensure IPv6 compatibility. Some checks to perform:

- Can its data structures support a 128 bit IPv6 address?
- Can it resolve addresses against a DNS6 server, e.g. manage AAAA records?

The US Department of Defense (DoD) has published a set of recommendations to qualify applications and devices as IPv6-Ready as part of "[DoD IPv6 Standard Profiles For IPv6 Capable Products.](#)"

This document extends the discussion above and clarifies the set of requirements that an application or network device must fulfill in order to interoperate in IPv6 networks.

There are public lists with IPv6 ready software, like the one published by [Wikipedia](#).

Dual Stack Operation

The dual stack approach requires that every router and host have both an IPv4 and IPv6 address. Hosts should have IPv4 and IPv6 protocol stacks for accessing the network using either mechanism and routers must be able to simultaneously route both IPv4 and IPv6 packets. Over time, when all services are reachable via IPv6 and all hosts support IPv6, one can shut down the IPv4 network access and routing mechanisms. It is anticipated, however, that this will require several years in any of today's enterprises. Because dual stack allows the most seamless means of introducing IPv6, it is the preferred strategy for initiating an IPv6 migration.

The key to effectively implementing a dual stack IPv4/IPv6 infrastructure is DNS. IPv6 DNS is obviously different than the DNS address record used for IPv4. DNS standards have been extended to accommodate IPv6. DNS records for IPv4 nodes are called A records and IPv6 nodes use AAAA (Quad A) records. The DNS infrastructure must be adapted to handle a dual IPv4 and IPv6 environment.

Transition criteria summary:

- Existing IPv4 hosts can be upgraded with dual stack IPv6 independently of the upgrade of other hosts or routers. Once an IPv6 network and routing infrastructure is available, the IPv6 stack can be turned on.
- New hosts using only IPv6 can be added at any time without dependencies on other hosts or routing infrastructure.
- Existing IPv4 hosts with IPv6 installed can continue to use their IPv4 address and do not need additional addresses.

Typically, core routers and switches are the first components that are migrated to a dual-stack configuration. Servers follow along with DHCPv6 and DNSv6 infrastructure which will enable IPv6 to operate on the network. Afterwards clients are migrated. In parallel, firewall configurations must be adjusted as well. Eventually, additional migration techniques for the internet service provider are required to take advantage of an IPv6 infrastructure.

A host will use IPv4 or IPv6 protocols depending on a series of factors:

- IPv4 address will be used if the destination address used by the application is an IPv4 address
- IPv6 address will be used if the destination address used by the application is an IPv6 address
- An IPv6 packet encapsulated inside an IPv4 packet will be used if the destination address used by the application is an IPv6 address with an embedded IPv4 address

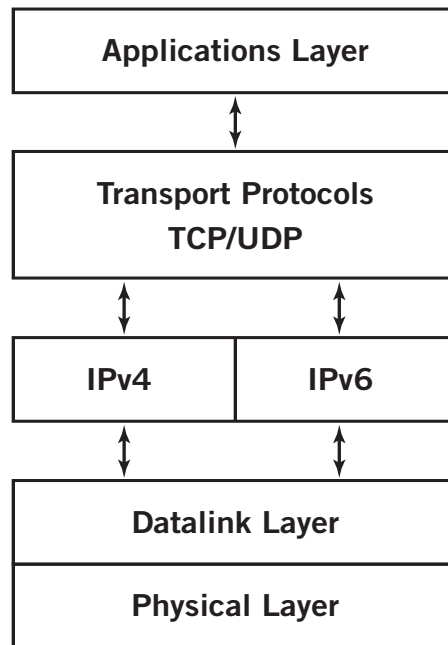


Figure 1: Dual Stack Operation

The Dual Stack approach makes it possible for IPv4 devices to continue to operate while IPv6 addressing is deployed in the network. The IP address assignment in the dual stack nodes occurs by using IPv4 mechanisms, like DHCP, or IPv6 mechanisms, such as stateless address auto-configuration (SLAAC) or DHCP6. Enterasys recommends the use of DHCPv6 due to its flexibility, control, and customization capabilities.

The ability to interoperate with different modes makes the dual stack approach flexible and simple to deploy in networks that are connected with both IPv4 and IPv6 nodes. Enterasys core [routers](#) can handle the routing and addressing load of running two stacks at full wire-rate, without impacting the performance of the network.

Hosts and routers in a dual stack environment can use tunneling mechanisms to route IPv6 traffic over existing IPv4 networks to ensure compatibility with existing IPv4 core networks that, for whatever reason, cannot be migrated at the same time.

The dual stack protocol implementation in an operating system is a fundamental IPv4-to-IPv6 transition technology. Dual-stack hosts are described in RFC 4213. Fortunately, while IPv6 network deployment has been slow in coming, the vendors of the major operating systems have been implementing a dual stack capability for some time, so many of your hosts are today IPv6 capable already.

As of today, a number of OS's support dual-stack operation:

- Windows XP SP1 and newer
- Windows Server 2003
- Windows Vista
- Windows Server 2008
- Windows 7

- Mac OS X 10.2 (Jaguar) and newer
- Linux kernel 2.4 and newer
- FreeBSD 4.5 and newer
- A more comprehensive list can be found [here](#)

Modern hybrid dual stack implementations of IPv4 and IPv6 allow programmers to write networking code that works transparently on either IPv4 or IPv6. The software uses hybrid sockets designed to accept both IPv4 and IPv6 packets. The bottom line is that all the major host and server operating systems are already capable of supporting applications using both IPv4 and IPv6.

Tunneling

Because IPv6 will have to interoperate with existing IPv4 infrastructure, tunneling provides a way to use an existing IPv4 routing infrastructure to carry IPv6 traffic. This could be applied where there are “islands” of IPv6 only nodes, perhaps for research or specialized applications, in advance of having a fully IPv6 network routing infrastructure, which needs to be interconnected. Tunneling IPv6 packets over the IPv4 infrastructure can be done by encapsulating IPv6 packets inside IPv4 packets.

The IPv6 header contains the address of the final destination and the IPv4 header contains the address of the tunnel endpoint.

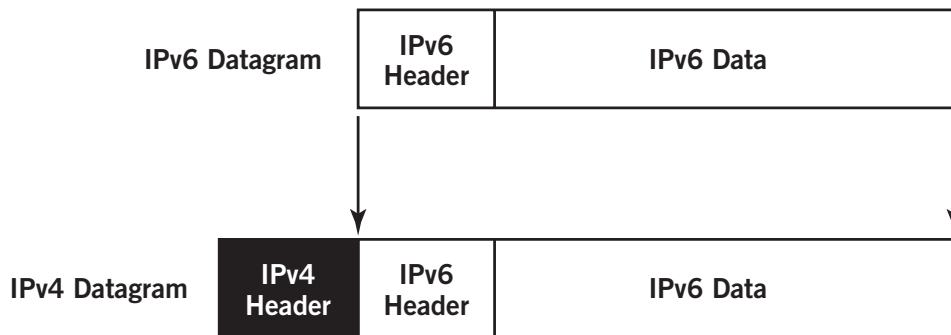


Figure 2: 6in4 encapsulation

Configured Tunneling of IPv6 over IPv4

While there are mechanisms defined for automatic tunnel set up, to ensure control, Enterasys recommends manual configuration of tunnels given the relative simplicity of the tunnel configuration. This also makes the most sense since it is expected that tunneling will only be an interim mechanism used in a network. Because the IP address of the router performing the tunneling is not the same as the IP address of the final destination, manual configuration of the tunnel endpoint is required. This is referred to as configured tunnels. Which packets are routed and which are tunneled is decided by the routing stack at the tunnel starting point.

Many tunnels deployed in current IPv6 networks are realized in this manner. It has been proven that this is a simple and efficient way to carry IPv6 packets across an IPv4 infrastructure where the tunnel acts as one hop in the IPv6 infrastructure. As the network infrastructure becomes capable of handling IPv6, tunnels can be taken down and native IPv6 enabled between the endpoints as part of a ubiquitous dual stack network environment.

Translation

Many protocols and applications must be adapted to work with the new IPv6 addressing. For this reason, address translation between IPv4 and IPv6 can be very complex and caution should be taken when implementing address translation as a transition tool. However, there may be circumstances where neither dual stack or tunneling are able to be used. Since the inception of IPv6, several methods have been developed for address translation from IPv6 to IPv4 and vice versa:

- NAT-PT (moved to historical status, no longer recommended)
- NAPT-PT (moved to historical status, no longer recommended)
- SIIT (RFC2765, soon to be obsoleted by draft-ietf-behave-v6v4-xlate-23)
- NAT64

The subject is still an active topic of work in the standards body and has been debated extensively. The conclusion is that in the enterprise market there is no optimal v4 to v6 translation strategy that ensures clean communication from an IPv6 host to an IPv4 host and vice versa, but some address translation mechanisms will be developed to provide for an interoperability mechanism of last resort.

The main problem is not the address translation itself but all the address references found in the protocols' internals and the applications. Some applications will have to be rewritten just to adapt the new address length into their data structures. As was the case with IPv4 NAT, IPv4/IPv6 address translation may require a whole new set of Application Layer Gateways (ALGs) for the existing protocols, not to mention the application data sets and structures.

Overview – Migrating to IPv6 with Enterasys

The preferred approach for IPv6 migration with Enterasys products is the dual-stack approach. Dual-stack approach assumes hosts and routers can use a dual-stack IPv4/IPv6. The dual-stack migration can be augmented by selective tunneling techniques for certain network designs.

Enterasys Switches and Routers

Enterasys L2 and L3 switches provide a full range of support for IPv6. L2 switches provide for dual stack management and provide IPv6 traffic classification policies in order to segment IPv6 traffic onto a VLAN or block it altogether. Enterasys core routers support manually configured tunnels using different tunneling technologies - 4in4, 6in4, 4in6, and 6in6. Additionally, a fully featured set of routing protocols are supported.

IPv6 Feature	A4	B5	C5	G-Series	K-Series	S-Series
IPv6 Host Stack	✓*	✓	✓	✓	✓	✓
Dual IPv6/IPv4 Management	✓*	✓	✓	✓	✓	✓
IPv6 L2 Policy	✓	✓	✓	✓	✓	✓
OSPFv3			✓	✓	✓	✓
Configured Tunnels			✓	✓	✓	✓
Static Routing					✓	✓
RIPng					✓	✓
ISIS					✓	✓
BGP					✓	✓
Multicast PIM-SM/SSM					✓	✓
NAT / LSNAT / Web Cache Balancing						✓

*Follow-on release

Figure 3: Enterasys IPv6 capabilities

Beyond IPv6 Switching and Routing

Other network appliances must be verified for their compatibility with IPv6. IPv6 is supported as a management protocol for its family of management tools for Enterasys Network Management Suite (NMS), including Network Access Control (NAC), Enterasys Intrusion Prevention System (IPS), and Security Information and Event Manager (SIEM). Any tool working above L3 but using an IP address for any kind of processing is sensitive to the protocol change, particularly security appliances must be ready to identify new addresses and attacks. Enterasys identified these challenges and has been developing improvements in its security and management portfolio to help the enterprise in the IPv4 to IPv6 migration.

Network Management

Enterasys NMS supports IPv6 address modeling and IPv6 policy support for the switches that support it. Enterasys NAC is able to detect, track and control IPv6 systems over IPv4/IPv6.

Security

Enterasys IPS supports full IPv6 traffic and inspection. This includes whitelists, blacklists, application filters, event filters, active response (i.e. connection sniping), payload signatures, protocol analysis, and probe detection. The sensor appliances can be managed over an IPv6 network. Nearly all signatures are application layer signatures so they work for either IPv4 or IPv6.

Enterasys Security Information and Event Manager (SIEM) supports several IPv4 to IPv6 migration strategies. These include mixed IPv4 and IPv6, IPv6 tunneled over IPv4 and pure IPv6 implementations. The Enterasys SIEM supports each of these implementations while maintaining all of its monitoring capabilities of an IPv4 network. The appliance can be managed in an IPv6 network.

Flow records, event records, the SIEM user interface and reports support IPv6 fields. Event and Flow Records now use modified fields for normalized IPv6 addresses. Flow processor supports the parsing of IPv6 addresses from packets and NetFlow records. Rule sets have also been added for IPv6 traffic. All user interfaces support the use of IPv6 parameters in searches, sorting and queries.

Conclusion

The migration of an existing IPv4 infrastructure to IPv6 will be one of the most demanding challenges facing IT organizations in the years to come. This is not because of the inherent complexities of the migration, but due to the universal reach of IP and dependency of today's enterprises on the operation of the network. This reach makes the migration a cross departmental undertaking. Anywhere in the enterprise can be a system or device affected by the migration. Proper analysis and careful planning will minimize the impact to business operations.

Enterasys products provide features companies need to ensure a smooth migration from IPv4 to IPv6. Enterasys can also provide the professional services needed to guide an organization through the migration, ensuring the highest level of business continuity. The combination of an IPv6 capable product portfolio and the expertise on how to implement a seamless transition to IPv6 makes Enterasys a preferred partner for customers embarking on this new initiative.

Contact Us

For more information, call Enterasys Networks toll free at **1-877-801-7082**, or +1-978-684-1000 and visit us on the Web at enterasys.com



© 2011 Enterasys Networks, Inc. All rights reserved. Enterasys Networks reserves the right to change specifications without notice. Please contact your representative to confirm current specifications. Please visit <http://www.enterasys.com/company/trademarks.aspx> for trademark information.

