White Paper

Managing IPv6 Networks

Michael Shevenell,
November 2005

Introduction

For nearly two decades the dominant networking protocol has been the Internet Protocol version 4 (IPv4). However, within the next three to five years the Internet Protocol Version 6 (IPv6) is likely to make significant inroads in the networking space. This will be especially true in certain market segments and locations.

Driving some of the early interest in IPv6 was the perceived exhaustion of address space in IPv4. Although this is still an issue (primarily in Asia), it is really the delivery of ever more complex services in a global network that is driving the current interest. Key factors in this drive include the following:[4]

- Increased Address Space
- Improved Auto Configuration
- Improved Security
- Improved End-to-End Quality of Service (QoS)
- Improved Mobility

A forecast by Price Waterhouse suggests that IPv6 will begin deployment in North America in 2007, and that IPv4 will be phased out by approximately 2011[1]. As mentioned previously, the IPv4 address space exhaustion will continue to be an issue, even using current address limiting techniques. One report predicts that the unallocated IPv4 address pool will be exhausted between April 2015 and June 2020[2].

Major Adopter – U.S. Army

The U.S. Army is seen as one of the major driving forces in the implementation and deployment of IPv6-based network technology. According to one report, the Army has a stated goal to complete the transition to IPv6 by FY 2008[3]. It is not surprising that the Army would require advanced networking technologies to perform its work. Specific factors that require next-generation networks include a large, highly distributed, global workforce with a high need for reliability, mobility and security. To show their commitment to this technology, the Army required that as of October 2003 all assets being developed, procured or acquired shall be IPv6 "capable"[4]. In addition to the Army, it appears clear that the entire U.S. government has begun taking a stronger stance in this technology. Recent mandates by the U.S. Office of Management and Budget require all federal agencies to use IPv6 by June 2008[14].

Major Adopter – Japan

Another of the early adopters of IPv6 is the country of Japan, as well as much of industrialized Asia. This comes for a good reason: Address Space Allocation. Due to the U.S.-centric development of the early Internet, Asia was allocated very little of the total address space. Although the U.S. has less than 5% of the world's population, it has been assigned over 74% of the total address space in IPv4. As early as 2001, the Japanese government invested over $100M (USD) in IPv6 technology with the goal of 2005 for a full implementation[12].

Summary of Standards Work

Agreed-upon standards is the first step in any global communication systems development work. They define the framework, as well as the details of the information exchange. This includes everything from architectures to protocols and encodings.

Protocol Development

The initial phase of protocol specification is essentially complete. From this effort a number of RFCs have been published, starting in 1995 and continuing for about five years. See Appendix A for a listing of the relevant RFCs. In addition, a number of successful interoperability events have taken place.

MIB Development

In contrast to the protocol development described above, there has been fairly little progress in the development of IPv6-compatible MIBs in the IETF. As of September 2005, the following standard MIBs, derived from search of MIB Depot[8], show support for the IPv6 address type:

The total numbers of MIBs currently supporting IPv6 are a very small fraction of the more than 6500 publicly available standard and enterprise MIBs. The preponderance of IPv6 support in MPLS-oriented MIBs is due to these MIBs being developed very recently at a time of high awareness for IPv6.

Table 1: Current Standard MIBs with IPv6 Support

Docs-Cable-Device	Exp-Docs-Cable-Device	IPv6-ICMP (RFC2481)
IPv6 (RFC2485)	OSPFv3 (RF C2328)	MPLS-FTN (RFC381)
IPv6 TCP (RFC2452)	IPv6 UDP (RFC2454)	MLD-v6 (RF C3019)
MPLS-TE (RFC3812)	IPv6-UDP (RFC2454)	
MLPPLS-LDP (RFC3815)	MPLS-LSR (RFC3813)	

IPv6 Addressing

As mentioned previously, one of the major features of IPv6 is the enormous address space it makes available. The address length has been extended from 4 to 16 bytes (128 bits); this provides 3.4 x 1038 unique addresses. To illustrate the magnitude of this space, it is calculated that this amounts to 1015 addresses for each square inch of the Earth's surface.

In addition, a number of refinements have been made in the address format to provide native support for anycast, multicast communications, as well as scoped addresses (all defined in RFC2373). Scoped addresses allow the user to define topological span of an address: local, site, regional or global. They would define a domain in which a given address is unique. It should be noted that there is no longer a broadcast address in this protocol; this function is provided by multicast addressing. The 128 bits are allocated as follows:

Table 2: Address Allocation

NUMBER OF BITS	FUNCTION
3	Format Prefix
13	Top Level Aggregation ID
8	Reserved
24	Next Level Aggregation ID
16	Site Level Aggregation ID
64	Interface ID

For Ethernet-based networks, the Interface ID can be simply generated from the MAC address of the interface. This is specified by IEEE EUI-64 which is described in RFC2373.

Given the length of the IPv6 address, it was necessary to devise a shorthand form which would provide a more compact representation. For example, an address like:

3FFE:0C01:0000:0001:0204:27FF:FE0C:91C0

can be expressed in more compact form as:

3FFE:C01::1:204:27FF:FE0C:91C0

where leading zeros are suppressed and blocks of zeros can be eliminated.

Summary of Vendor Support

The success of CA's Network Fault Management (SPECTRUM®) solution and any management system in the IPv6 arena is closely tied to the availability of management interfaces from the IT infrastructure equipment and operating system vendors. A solution can only be deployed at a rate that allows end-to-end management to be maintained.

Protocol Support

Fortunately, the IPv6 protocol is widely supported by a large number of network equipment vendors, including Cisco and Juniper. In addition, most of the major operating systems currently support IPv6 as well. Protocol support has been released in the following operating systems of interest to CA:

- Microsoft Windows: XP, 2000, 2003 Server
- Solaris 8, 9 & 10
- Linux

The wide range of workstation support greatly simplifies the task of deploying CA's Network Fault Management solution.

MIB Support

The current lack of MIB support is one of the significant challenges in deploying any SNMP-based IPv6 management solutions. Although new MIB development efforts are being designed to handle IPv6 addresses, the huge body of IPv4 MIBs, of limited use in their current format in a native IPv6 world, complicates the management picture.

Interoperability Testing

In conjunction with the standards work and vendor implementation, there has been some significant interoperability efforts to verify compatibility. In addition, the IPv6 Forum has developed an "IPv6 Ready Logo" program to identify compatible IPv6 solutions that pass specific self tests and interoperability tests[9]. In order to participate in IPv6 discussions and testing, CA is a member of the IPv6 Forum.

Interoperability – 6bone

One of the largest successful test beds for this new technology is the 6bone. This experimental network of over 1,000 current sites has its roots in the Internet Engineering Task Force (IETF) project previously called Internet Protocol Next Generation (IPng). The initial use of the 6bone was to test standards and implementations. With that phase completed, its current focus is on testing the transition, migration and operational issues[8]. This test bed has been so successful in helping the technology mature that the 6bone is in the process of being phased out by 2005[12].

Interoperability – MoonV6

This effort is led by the North American IPv6 Task Force and is the largest permanently deployed IPv6 network in

the world. It involves the University of New Hampshire (UNH) Interoperability Lab, Internet2, the U.S. Department of Defense, service providers, and regional IPv6 Forum pilots worldwide.

Interoperability – Management

Interoperability testing between IPv6-compatible management systems and devices to date has not received much attention. This may be due in part to the lack of IPv6-compatible MIBs and IPv6-compliant SNMP management systems. CA recommends the formation of an independent interoperability consortium (like the aforementioned UNH Interoperability Lab) to design and execute interoperability tests.

Managing the Whole Transition with CA's Network Fault Management Solution

CA's Network Fault Management (SPECTRUM®) product family is well positioned to handle the transition to IPv6 networks. This transition can be divided into two broad objectives:

- Managing IPv6-compatible networking devices
- Operating the Network Fault Management solution in IPv6 networks

Clearly, managing IPv6 devices is the primary objective.

The second objective is making the family of Network Fault Management applications communicate in a native IPv6 network. This would require all interapplication communication to be done using IPv6. This is a lower priority task because most organizations (including the U.S. Army) are planning long transition periods where both IPv4 and IPv6 technologies will coexist (typically referred to as dual-stack configuration). Therefore, it will be possible for CA's Network Fault Management solution to manage devices using either IPv4 or IPv6, but all interapplication communication can still be done using IPv4.

The solution will also provide continuous management services throughout the entire IPv6 transition process. Possible steps for the process of moving from the current, fully native, IPv4 network to a fully native IPv6 network of the future are outlined in the following table. It is expected that enterprises and service providers will perform incremental network upgrades and have deployments representing each case at some point in time.

The Present: IPv4-only Network

As shown in the table, the situation in the first row is characterized by a fully IPv4 network. Actually, this is not an accurate reflection of the present, since a significant amount of IPv6-capable equipment is being deployed at the current time.

Case 1: IPv4 at the Edges and IPv6 in the Core

In this case, the network upgrade happens to the "middle" of the network—possibly a transport core service. The simplest architecture for this situation is the implementation of IPv6 tunnels to encapsulate the IPv4 traffic. In this case, Network Fault Management of the infrastructure continues to operate as it does currently. A key addition is the discovery and monitoring of the IPv6 tunnels carrying the IPv4 traffic. Existing Network Fault Management Modules (MMs) and Applications will function in this environment.

Table 3: Possible Transition Phases to Fully-Managed Native IPv6

CASE	NMS SNMP TRANSPORT	INTERMEDIATE NET TRANSPORT	DEVICE SNMP TRANSPORT	SPECTRUM APPLICATION	TRANSLATION MECHANISM
Now all v4	V4	V4	V4	V4	None Required
Case 1 v6 Core	V4	V6	V4	V4	Tunnel
Case 2 v6 Devices	V4	V4	V6	V4	NAT PT or Tunnel
Case 3 v6 NMS & Net	V4/V6	V6	V4	V4	Dual Stack NAT PT
Case 4 v6 NMS & Device	V6	V4	V6	V4	Tunnel
Future Native v6	V6	V6	V6	V6	None Required

Case 2: IPv6 Devices and IPv4 Transport

In this case, the user has begun to deploy IPv6 on only devices that must be managed by SNMP over IPv6. A number of vendors (Cisco included[13]) already support this capability to manage IPv4 MIBs using an IPv6 transport for SNMP. A key part of this solution is the use of a Network Address Translator-Protocol Translator (NAT-PT) to convert IPv4 management requests to IPv6 transport. Refer to the chart on page 7 for more information on NAT-PT translator devices. Network Fault MMs will be enhanced to support the new device capabilities as standard and enterprise MIBs are developed.

An alternate approach to managing IPv6 devices with an IPv4 management transport is to use proxy management techniques similar to the Ping MIB (RFC2925). The NMS would communicate with the device supporting the Ping MIB (the proxy agent) using IPv4. The proxy agent would then communicate to the end systems via IPv6. The management system can then monitor the health of these end systems without requiring an IPv6 transport. As an example, Juniper routers support the Ping MIB and could be used in the proxy management role. Cisco does not currently support this functionality in Service Assurance Agent or the RTTMON MIB.

Case 3: Managing IPv4 Devices in an IPv6 Network

In this case, a significant part of the network has been upgraded to IPv6—with the exception of some legacy devices which still support only an IPv4 transport. One solution to this is to deploy CA's Network Fault Management solution in a dual-stack configuration. The IPv4 stack will be used to manage the legacy devices while the IPv6 stack will be used to manage the newer IPv6 devices. Another solution to this situation is to use the NAT-PT to provide protocol translation services, allowing the Network Fault Management solution's purely IPv6 management engine to communicate with IPv4 devices. In either case, the existing MMs may be used to manage the legacy devices.

Case 4: IPv6 at the Edges and IPv4 in the Core

In this case, the edges of the network have been upgraded to IPv6 and some interior part of the network remains IPv4. Like case 1, the preferred solution for this situation is the implementation of tunnels. In this scenario, an IPv4 tunnel is set up to carry IPv6 traffic. A key aspect of this solution is native IPv6 communication from the management solution to the devices being managed. This is accomplished by extending the Device Communications Manager (DCM) to issue SNMP requests over an IPv6 transport. In addition, the Network Fault Management solution will manage the IPv4 tunnel which encapsulates the IPv6 traffic.

The Future: Fully Native IPv6 Network

This case represents some point in the future where a fully native IPv6 infrastructure has been deployed. All components have transitioned to IPv6, and there are the necessary MIBs to provide complete IPv6 management. It remains to be seen when and if the networks of the world will completely transition to IPv6.

Summary

The table on the following page presents a summary of CA's Network Fault Management solutions for each case in the transition matrix.

Due to the gradual migration to fully native IPv6 networks, it is expected that this process will take several steps as outlined below:

- Management of Dual Stack devices via v4 transport and v4 MIBs (current)
- Management of Dual Stack devices via v4 transport and selected v6 MIBs
- Management of v6 Devices using translators / dual stacks and v4 MIBs
- Management of v6 Devices using translators / dual stacks and v6 MIBs
- Management of v6 Devices using native v6 Transport and v6 MIBs

In addition to the broad objectives mentioned above, there are several meaningful short-term steps which enable useful network management in an IPv6 environment. These steps would include:

- Support for BGP Traps with IPv6 Information
  - bgpEstablishedNotification: Discover new BGP Peers and Services
  - bgpBackwardTransNotification: Detect status of BGP Peers and Services
- Support for Discovery via IPv6 MIB (RFC2465)
- Support for Discovery via TCPv6 MIB (RFC2452)
- Support for Discovery via ISISv6 MIB

Table 4: Possible Transition Phases to Managing IPv6 Networks					
CASE	NMS v4 STACK	NMS v6 STACK	MANAGE TUNNELS	NEW v6 MMS	INTER-APPLICATION V6 COMMUNICATIONS
All v4	Y	N	N	N	N
Case 1: v4 Edge v6 Core	N	N	Y	N	N
Case 2: v6 Edge	N	N	N	Y	N
Case 3: v4 Legacy	Y	Y	N	N	N
Case 4: v4 Core	N	Y	Y	Y	N
All v6	N	Y	N	Y	Y

Recommendations

There will be a slow migration to IPv6 over the next several years. Therefore it makes sense for the solution to be developed and deployed in an incremental manner. One of the simplest approaches to this migration is to deliver some sort of gateway or proxy outside of CA's Network Fault Management solution to perform the conversion. This will allow the management platform to continue managing IPv6 devices (with IPv4-compatible MIBs). An overview of this architecture is shown below.

Figure 1: Possible Proxy Solution for Managing IPv6 Networks

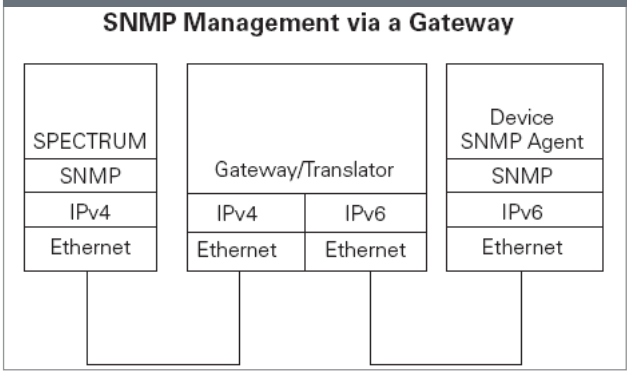

The IPv6 Forum has certified a number of products in the "IPv6 Ready Logo" as Proxies, Gateways or translators[10]. See Appendix B for a complete listing. In addition, Cisco has outlined several approaches to the IPv6 to IPv4 translation problem. They are outlined in Appendix C.

As the deployment of IPv6 progresses, it is expected that some networks will become exclusively IPv6 (over time). In these networks, it is recommended that a native IPv6 communication capability be implemented in CA's Network Fault Management solution. Essentially, this requires an IPv6-compatible Device Communications Manager (DCMv6).

An alternative approach to going through translators (like NAT-PT) is to use distributed management techniques to manage IPv6 devices from an IPv4 management system. This approach utilizes the IETF RFC2925 Ping and Traceroute MIBs to invoke Distributed Management commands (DistMan) to determine the health of IPv6 components. Currently, Juniper is the major router vendor that supports this functionality. An IPv6 test could be set up using IPv4 commands to the Ping MIB resident in a Juniper device. When this test is completed, the results could be read by the Network Fault Management platform and used to reflect the status of the device in question. The advantage of this approach (currently only available for customers who operate Juniper networks) is that it does not require the purchase or configuration of

additional components to accomplish translation. This is simpler than the NAT approach, but does not offer the ability to manage devices directly via SNMP. It is likely that each approach would have use in transition networks. A summary of the two approaches is shown in the following table:

Table 5: NAT-PT versus DISTMAN

	NAT-PT	DISTMAN
Require Additional Equipment	No	No
Vendor Support	Cisco	Juniper
Configuration Complexity	More Complex	Less Complex
Flexibility	High	Low
Supports Ping/Traceroute	Yes	Yes
Supports SNMP	Yes	No

Planning

The success of any deployment is based largely on adequate planning and understanding. It is strongly recommended that service providers and enterprises adopt a migration strategy which defines extended periods of co-existence for IPv4 and IPv6.

Protocol Development

Cisco recommends the following steps in planning the migration of network management applications[5]:

- IPv6 stack on network management station (NMS)
- IPv6 stack on network devices
- NMS applications running over an IPv6 stack
- SNMP over an IPv6 transport
- IPv6 address family support on public and private MIB when required

This process is critical because it is expected that both IPv4 and IPv6 will need to be managed simultaneously for many years.

Additional migration guidelines are presented by Waters Creek Consulting[6].

1. Upgrade DNS server(s) to handle IPv6 addresses
2. Introduce dual stack systems that support IPv4 and IPv6
3. Add IPv6 addresses to DNS records of those systems
4. Rely on tunnels to connect IPv6 islands separated by IPv4 networks
5. Remove support for IPv4 from systems
6. Rely on header translation to reach remaining IPv4-only systems

References

1. Price Waterhouse Coopers, Technology Forecast: 2002:2004, Volume 2: Emerging Patterns of Internet Computing, Oct 2002, pp. 541-542.

2. IPv4 Address Space Report, 26 April 2004, http://bgp.potaroo.net/ipv4/.

3. ASD(NII) Memorandum, "Subject: Internet Protocol Version 6 (IPv6)", 9 June 2003.

4. Shipp, John H. III, "IPv6 Army Transition Planning: NAIPv6 Summit", 17 June 2004. http://usipv6.unixprogram.com/North_American_IPv6_Summit_2004/Thursday/PDFs/John_Shipp.pdf

5. Cisco IPv6 Solutions Calendar Year 2004 and Beyond, http://www.cisco.com/en/US/tech/tk872/technologies_white_paper09186a00802219bc.shtml

6. Thomas, Stephen, Waters Creek Consulting, "Example Migration Plan", http://www.waterscreek.com/ipv6/tsld051.htm

7. MibDepot, http://mibdepot.com/index.shtml

8. 6bone, "What is the 6bone?", http://6bone.net/about_6bone.html

9. IPv6 Forum, "IPv6 Ready Logo Program", http://www.ipv6ready.org/about_policy_for_use.html

10. IPv6 Forum, "Latest Approved Application List 2004/12/8", http://www.ipv6ready.org/logo_db/approved_list.php

11. Cisco, "IPv6 Deployment Strategies", http://www.cisco.com/en/US/tech/tk872/technologies_white_paper09186a00800c9907.shtml#1075253

12. Japan Inc. "IPv6 Asia's Agent of Change", July 2003, http://www.japaninc.net/article.php?articleID=1137

13. Cisco, "Cisco IOS Software Release Specifics for IPv6 Features", http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d65ed.html

14. Perera, David, "OMB: IPv6 by June 2008", http://www.fcw.com/article89432-06-29-05-Web.html

Appendix

Appendix A - IPv6 RFCs

This appendix contains a listing of some of the RFCs which make up the protocol standards of IPv6.

IETF RFC NUMBER	RFC DESCRIPTION
RFC1881	IPv6 Address Allocation Management
RFC1887	Architecture for IPv6 Unicast Address Allocation
RFC1924	A Compact Representation of IPv6 Addresses
RFC2080	RIPng for IPv6
RFC2185	Routing Aspects of IPv6 Transition
RFC2460	Internet Protocol Version 6
RFC2373	IP Version 6 Addressing Architecture
RFC2452	TCP over IPv6 MIB
RFC2454	UCP over IPv6 MIB
RFC2463	ICMP for IPv6
RFC2465	IPv6 MIB
RFC2466	ICMP over IPv6 MIB
RFC3596	DNS Extensions for IPv6
RFC2893	Transition Mechanisms for IPv6 Hosts and Routers
RFC2461	Neighbor Discovery for IPv6
RFC2462	IPv6 Stateless Address Autoconfiguration
RFC2374	An IPv6 Provider-based Unicast Address Format
RFC2375	IPv6 Multicast Address Assignments
RFC2529	Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC2710	Multicast Listener Discovery
RFC2740	OSPF for IPv6
RFC2925	Remote Ping, Traceroute and NS Lookup MIB
RFC3056	Connection of IPv6 Domains via IPv4 Clouds

Appendix B - Proxies, Gateways and Translators

This section includes a snapshot of the IPv6 Ready Logo approved translators.

COMPANY	PRODUCT	FUNCTION
iBIT Technologies Inc.	Forsix-1000R	Host Box(IPv6/IPv4 translator based on RTOS)
Alpha Networks	Home GateWay	Service Gateway for next-generation IPv4/IPv6 services transition platform
Yokogawa Electric Corporation	TTB3010	IPv6/IPv4 Translator Series
Panasonic Communications Co.	KTR5	Host box (IPv6/IPv4 Protocol Translator)
ZyXEL	ZyWall-200	IPv4/v6 Dual Stack Security Gateway in Router mode
NEC Corporation	MX6350-PG-M	IPv6/IPv4 Translator and Mobile IPv6 CN proxy

In addition, several versions of Cisco Routers have earned the IPv6 Ready Logo. See the table below:

ROUTER MODEL	IOS REVISION
Cisco 12000 series	Cisco IOS 12.0(26)S
Catalyst 6500/Cisco 7600 series	Cisco IOS 12.2(17a)SX1
Cisco 7300/7200/7500 series	Cisco IOS 12.2(14)S

Appendix C - Various Translation Approaches

Cisco has outlined several approaches to this problem[11].

TRANSLATION MECHANISM	PRIMARY USE	BENEFITS	LIMITATIONS	REQUIREMENTS
NAT-PT	IPv6-only hosts to IPv4-only hosts.	No dual stack. To be supported in IPv6 for Cisco IOS software Phase II.	No end-to-end IPSec. Dedicated server is single point of failure.	Dedicated server. DNS with support for IPv6.
TCP-UDP Relay	Translation between IPv6 and IPv4 on dedicated server.	Freeware. No changes to Cisco IOS software.	No end-to-end IPSec. Dedicated server is single point of failure.	Dedicated server. DNS with support for IPv6.
BIS	IPv4-only hosts communicating with IPv6-only hosts.	End-system implementation.	All stacks must be updated.	Updated IPv4 protocol stack.
DSTM	Dual-stack hosts (but with IPv6 address only).	Temporary IPv4 address allocated from pool.	No current support in Cisco IOS software.	Dedicated server to provide a temporary global IPv4 address.
SOCKS-Based IPv6/IPv4 Gateway	IPv6-only hosts to IPv4-only hosts.	Freeware. No changes to Cisco IOS software.	Additional software in the router.	Client and gateway software in the host and router.

Appendix D - Comparison of Deployment Strategies

Cisco has outlined and compared several deployment options[11].

DEPLOYMENT STRATEGY	KEY USER/PRIMARY USE	BENEFITS	LIMITATIONS	REQUIREMENTS
IPv6 over IPv4 Tunnels	Service provider wanting to offer initial IPv6 service. Enterprise wanting to interconnect IPv6 domains or link to remote IPv6 networks.	Can demonstrate demand for IPv6 for minimal investment. Easy to implement over existing IPv4 infrastructures. Low cost, low risk.	Complex management and diagnostics due to the independence of the tunnel and link topologies.	Access to IPv4 through dual-stack router with IPv4 and IPv6 addresses. Access to IPv6 DNS.
IPv6 over Dedicated Data Links	Service provider WANs or metropolitan area networks (MANs) deploying ATM, Frame Relay, or dWDM.	Can provide end-to-end IPv6 with no impact on the IPv4 traffic and revenue.	Lack of IPv6-specific hardware acceleration and support for IPv6 network management in currently deployed hardware.	Access to the WAN through dual-stack router with IPv4 and IPv6 addresses. Access to IPv6 DNS.
IPv6 over MPLS Backbones	Mobile or greenfield service providers, or current regional service providers deploying MPLS.	Integrates IPv6 over MPLS, thus no hardware or software upgrades required to the core.	Implementation required to run MPLS. High management overhead.	Minimum changes to the customer edge (CE) or provider edge (PE) routers, depending on the technique.
IPv6 using Dual-Stack Backbones	Small enterprise networks.	Easy to implement for small campus networks with a mixture of IPv4 and IPv6 applications.	Complex dual management of routing protocols. Major upgrade for large networks.	All routers are dual-stack with IPv4 and IPv6 addresses. Access to IPv6 DNS. Enough memory for both IPv4 and IPv6 routing tables.

About the Author

Michael Shevenell is a Development Architect for CA's Enterprise Systems Management business unit.

MP28914