

IP NETWORK

Brocade IPv6 Routing Innovations and Solutions for NetIron and MLX Platforms

Next-generation networks require an overall architecture evolving to accommodate new technologies that support the growing numbers of users, applications, and services. IPv6 is designed to be more robust, dramatically increase the pool of global IP addresses, simplify network administration, resolve security and mobility issues, and improve Quality of Service (QoS). This paper defines IPv6, explores the advantages of the new protocol suites, describes Brocade® IPv6 innovations, and details the Brocade IPv6 solutions and transition strategies for service providers on the Brocade NetIron and MLX platforms.

BROCADE

CONTENTS

introduction.....	3
The Limitations of IPv4.....	3
Insufficient IP Address Space	3
Address Prefix Allocation	3
Complexity of Configuration	3
Data Security.....	4
Quality of Service (QoS).....	4
IPv6 Technical Advantages.....	4
New Header Format.....	4
Much Larger Address Space	4
Multicast and Anycast	5
Stateless Address Configuration.....	5
Built-in Security	5
Better Support for QoS	5
Extensibility	6
Brocade IPv6 Industry Presence and Customers.....	7
Brocade IPv6 Innovation and Solutions.....	8
High-Performance IPv6 Traffic Forwarding.....	8
Carrier Ethernet Services with IPv6 Enabled.....	8
Increased Resilience and Performance with MCT and 64-Port Carrier Trunk.....	8
IPv6 Traffic Monitoring with sFlow	8
IPv6-Ready with 40/100-GbE Technology and Platform	9
Brocade 6PE: IPv6 over MPLS	9
Router Redundancy with VRRPv3	9
Brocade Hardware Supporting IPv6	9
Brocade Software.....	11
Brocade IPv6 Deployment and Transition Strategies.....	12
Tunnel: 6to4 Tunnel and Configured Tunnel.....	13
Dual Stack: IPv6 and IPv4 Coexist.....	14
6PE: IPv6 over MPLS	14
Total IPv6.....	15
Moving Forward with Brocade	15
About Brocade.....	15

INTRODUCTION

The vast majority of networks today, including the Internet, are based on the Internet Protocol version 4 (IPv4). Developed almost 30 years ago, IPv4 is showing its age with 32-bit addressing that cannot accommodate the explosive growth of network-connected devices around the world. It is predicted that the IP address will run out as early as 2011. Moreover, IPv4 was not designed with security in mind, and its security solutions have not kept pace with the expanding requirements of IP-based networks.

Today's Internet is being gradually replaced by tomorrow's Internet. IPv4 is slowly being supplemented by Internet Protocol version 6 (IPv6), a more robust suite of protocols and standards designed to dramatically increase the pool of global IP addresses, simplify network administration, resolve security and mobility issues, and improve Quality of Service (QoS). IPv6 accommodates the new technologies that support the growing numbers of users, applications, and services.

The U.S. government has been pushing IPv6 adoption for many years. In June 2003, the Department of Defense (DoD) issued a memo mandating IPv6 adoption. In August 2005, the U.S. Office of Management and Budget (OMB) issued a similar memo. In late 2006, the Defense Information Technology Standards Registry (DISR) produced its first "Product Profile Document," which classified all networked products into categories, each with specific IPv6 requirements. Companies that do business with the DoD must use this document to develop IPv6 product plans. In October 2010, OMB issued a memo that specified the IPv6 transition timeline and requirements to all its agencies.

Service providers, enterprises, and government agencies around the world have started to examine new IPv6 equipment and explore the implications of the changes to come—many of them have even started deploying IPv6-based networks. Brocade believes in providing our customers and partners with comprehensive information about significant emerging technologies so that they can be fully informed and can partner with Brocade to implement the best solutions for their environments.

This paper defines IPv6, explores the advantages of the new protocol suites, describes Brocade IPv6 innovations, and details Brocade IPv6 solutions and transition strategies for service providers.

THE LIMITATIONS OF IPv4

The current version of Internet Protocol or IP (known as version 4 or IPv4) has not been substantially changed in the past 30 years, a lifespan over which IPv4 has proven to be robust, easily implemented and interoperable, and scalable enough to accommodate the ever-expanding Internet. However, continued exponential growth of Internet-enabled devices and the evolving requirements for secure data transfer over the Internet are outstripping the practical capabilities of IPv4 and exposing its limitations.

Insufficient IP Address Space

With only 32-bit capacity, IPv4 addresses have become relatively scarce. The current IP address space is unable to satisfy the potential huge increase in the number of users and the geographical needs of the Internet expansion, let alone the requirements of emerging applications such as Internet-enabled wireless devices, home and industrial appliances, Internet-connected transportation, integrated telephony services, sensor networks such as RFID, IEEE 802.15.4/6LoWPAN, and distributed computing or gaming.

Address Prefix Allocation

Because of the way IPv4 address prefixes have been and are currently allocated, Internet backbone routers are routinely required to maintain unreasonably large routing tables—as large as 85,000 specified routes and more. The current IPv4 Internet routing infrastructure is a combination of both flat and hierarchical routing.

Complexity of Configuration

Most current IPv4 implementations must either be manually configured or must use a stateful address configuration protocol such as Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a

need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

Data Security

Private communication over a public medium such as the Internet requires encryption services that protect the data that is sent from being viewed or modified in transit. Although an add-on standard now exists for providing security for IPv4 packets (known as Internet Protocol Security, or IPsec), this standard is optional, and proprietary alternatives are commonly used.

Quality of Service (QoS)

While standards for QoS exist for IPv4, no identification of packet flow for QoS handling by routers is present in the IPv4 header. Instead, real-time traffic support relies on the IPv4 Type of Service (ToS) field and the identification of the payload, typically using a UDP or TCP port. However, the IPv4 ToS field has limited functionality, and payload identification using a TCP or UDP port is not possible when the IPv4 packet payload is encrypted.

IPv6 TECHNICAL ADVANTAGES

IPv6 was developed to address the IPv4 limitations described above. It offers a number of enhancements over its predecessor, including auto-configuration, header simplification, extended routing headers, and flow labeling. The most significant improvement is the increase in the number of bits used for addressing, from 32 in IPv4 to 128 in IPv6. The designers of IPv6 correctly envisioned that some day in the future, the number of addressable IP endpoints would exceed the number of addresses supported by IPv4.

New Header Format

IPv6 introduces a more streamlined header format that reduces overhead processing on intermediate routers and speeds throughput. IPv4 and IPv6 headers are not interoperable, and IPv6 is not backward compatible with IPv4. A host or router must use an implementation of both IPv4 and IPv6 (for example, dual stack) in order to recognize and process both header formats. Although IPv6 addresses are four times as large as IPv4 addresses, the new IPv6 header is only twice as large as the IPv4 header.

Much Larger Address Space

IPv6 has 128-bit (16-byte) source and destination IP addresses, which allows, for example, each cell phone or mobile electronic device to be assigned a unique IP address. IPv4 supports 4.3×10^9 (4.3 billion) addresses, which is incapable of furnishing even one address to every living person on the planet. Remember, millions of people have multiple IP-enabled devices. With 128 bits, IPv6 can express over 3.4×10^{38} possible combinations or 5×10^{28} addresses for each of the roughly 6.5 billion people alive today.

Multicast and Anycast

Multimedia applications can take advantage of multicast: the transmission of a single datagram to multiple receivers. Multicast (both on the local link and across routers) is a requirement of IPv6, in contrast to IPv4, in which multicast is optional and rarely deployed across routers.

In addition, IPv6 defines a new broadcasting method termed “anycast.” Like multicast, anycast has groups of nodes that send and receive data packets. However, when a packet is sent to an anycast group, it is delivered to only one of the group members, thereby limiting the data flooding that characterizes IPv4 networks. IPv6 eliminates broadcast packets, allowing greater use of switches instead of routers, flattening networks, and improving performance at the physical level.

Stateless Address Configuration

To simplify host configuration, IPv6 supports both stateful address configuration (with DHCP) and stateless address configuration (autoconfiguration without DHCP). With stateless address configuration, IPv6 hosts can configure themselves automatically. In this scenario, when first connected to a routed IPv6 network, a host sends a link-local multicast request for its configuration parameters. An IPv6 router on the network hears this request and responds appropriately with an advertisement packet containing the address. If stateless configuration is not suitable, a host can still use stateful configuration or be configured manually, just as with IPv4 networks.

Built-in Security

Support for IPsec security framework is an IPv6 mandatory requirement. This requirement provides a standards-based solution for network security needs and promotes interoperability between different IPv6 implementations.

Better Support for QoS

New fields in the IPv6 header define how traffic is handled and identified. Traffic identification using a Flow Label field in the IPv6 header allows routers to identify and provide special handling for packets belonging to a flow, that is, a series of packets between a source and destination. Because the traffic is identified in the IPv6 header, support for QoS can be achieved even when the packet payload is encrypted via IPsec.

Extensibility

IPv6 can easily be extended for new features by adding extension headers after the IPv6 header. Unlike options in the IPv4 header, which can support only 40 bytes of options, the size of IPv6 extension headers is constrained only by the size of the IPv6 packet.

Table 1 IPv4 vs. IPv6

Feature	IPv4	IPv6	IPv6 Advantage
Address space	4 billion addresses	2^{128} addresses	79 octillion (quadrilliard) times the address space
Configuration	Manual/DHCP	Plug-and-play: with or without DHCP	Reduces OpEx, avoids errors
Use of broadcast/multicast	Uses both broadcast and multicast	Relies on different forms of multicast, no broadcast	Better efficiency
Anycast support	Unofficial, not part of the original concept	Explicit support of anycast addressing	Allows for newer applications in mobility, data centers, and so on
Network reconfiguration	Labor-intensive	By design, facilitates the renumbering of hosts and routers	Reduces OpEx, facilitates migration
QoS support	ToS/DiffServ	Flow Classes, Flow Labels	More granular control of QoS
Security	IPsec for data packet protection, originally designed for IPv6	IPsec plays a central role in protecting data and control packets	Unified framework for security, better suited for secure computing environments
Mobility	Mobile IPv4	Mobile IPv6 provides enhancements: fast handovers, route optimization, and hierarchical mobility	Better efficiency and scalability, better suited for new age applications such as 3G wireless and beyond.

BROCADE IPV6 INDUSTRY PRESENCE AND CUSTOMERS

Brocade has actively participated in the IPv6 definition, standardization, and architecture and continues to lead and innovate with IPv6 features and solutions. Brocade engineers have been and continue to be engaged with leading IPv6 organizations and groups:

- Was an early member of **IPv6 Forum**
(<http://www.ipv6forum.com/navbar/members/generalmembers.htm>)
- Participates in the industry's largest IPv6 implementation and interoperability test bed: **Moonv6 Project**
(<http://moonv6.sr.unh.edu/>)
- Involved in the **IPv6 Ready Logo Program**
(<http://cf.v6pc.ip/>)
- Attended **US IPv6 Summit**
(<http://www.usipv6.com/>)
- Certified by **UNH IOL IPv6**
(<http://www.iol.unh.edu/>)

Brocade has a long history of implementing IPv6 in leading products. It offers one of the industry's most complete sets of IPv6 unicast, multicast, and transition protocols implementation. The Brocade NetIron and MLX family of products includes hardware-based, wire-speed IPv6 switching and routing products, software upgradeable systems for IPv6 switching and routing, and software-based IPv6 management feature upgrades.

Over the years, many of our key customers have deployed IPv6 in their environments using Brocade products and solutions. Hurricane Electric, for example, successfully deployed the largest IPv6 backbone in the world as of Q1 2008 (shown in Figure 1), running on Brocade NetIron® XMR Series routers. To learn more about the Brocade solution for Hurricane Electric, see the Brocade press release at:

http://newsroom.brocade.com/easyir/customrel.do?easyirid=74A6E71C169DEDA9&version=live&prid=730390&releaseisp=custom_184

Other key customers include:

- Highwinds
- Netnod (Europe)
- Several U.S. Federal government customers
- SPAWAR, US
- Several leading Japanese service providers

Brocade provides a strong roadmap to support both enterprise and service provider customers in their IPv6 deployments, and we continue to be a leading player in helping governments and industries worldwide in their transition to and support for IPv6.

BROCADE IPv6 INNOVATION AND SOLUTIONS

Brocade has implemented IPv6 based on industry open standards to guarantee interoperability. In addition to offering similar features as IPv4, Brocade continues to develop new features, technologies, and capabilities—all of which bring innovation to the IPv6 world.

High-Performance IPv6 Traffic Forwarding

The Brocade NetIron and MLX products deliver IPv6 traffic and services at wire speed without oversubscription. Both the Brocade MLX Series and Brocade NetIron XMR Series platforms are based on a distributed, high-performance, ultra-low latency, and network-processor-based architecture. This provides data forwarding performance that is unparalleled in the industry.

Table 2. IPv6 Routing Performance

Packet Size (bytes)	76	128	256	512	1024	1518
Input Packet Rate Per Interface (pps)	13,020,833	8,445,945	4,528,985	2,349,624	1,197,318	812,743
Total Input Packet Rate (pps)	104,166,662	67,567,560	36,231,878	18,796,992	9,578,544	6,501,944
Packet Loss	0	0	0	0	0	0

Table 3. IPv6 Routing Latency

Packet Size (bytes)	76	128	256	512	1024	1518
Minimum Latency (µs)	6.90	7.20	7.90	8.70	10.90	13.10
Average Latency (µs)	11.06	11.15	11.60	12.65	14.44	16.76
Maximum Latency (µs)	16.40	11.80	12.30	13.30	15.40	30.70

Carrier Ethernet Services with IPv6 Enabled

As an early and active member of the Metro Ethernet Forum (MEF), Brocade has incorporated requirements from MEF specifications in its new generation of products. The Brocade MLX and NetIron XMR product lines are certified for both MEF 9 and MEF 14. All the Carrier Ethernet services—such as VPLS and VLL—are IPv6 ready.

Increased Resilience and Performance with MCT and 64-Port Carrier Trunk

With Brocade Multi-Chassis Trunking (MCT), customers can achieve enhanced system availability through redundant systems, loop management without the use of Spanning Tree Protocol, full system bandwidth high availability, rapid link-failure recovery, and link aggregation to an IEEE 802.3ad-capable edge device. Carrier trunks allow 64x10G link aggregation groups (LAGs) to be deployed using an advanced load-sharing algorithm, which operates on IPv6 headers in addition to IPv4, Layer 2, MPLS Layer 2 VPN, and MPLS Layer 3 VPN headers.

IPv6 Traffic Monitoring with sFlow

Brocade understood the power of sFlow early and was one of the first to implement it. Brocade switches and routers support sFlow version 5. With extensive experience in building high-density routers and switches with a large number of high-speed interfaces, Brocade has pushed the envelope with a scalable and mature implementation of sFlow technology. IPv6 is natively supported with sFlow technology on both Brocade XMR and MLX/MLXe product lines.

IPv6-Ready with 40/100-GbE Technology and Platform

The Brocade next-generation MLXe core router integrates multiport 100-GbE, multi-Terabit trunks and comprehensive 100-GbE packet processors to ensure wire-speed performance. The router provides one platform, one software, and five times the throughput, with features that significantly reduce network complexity and ready-to-deliver IPv6 services from the service provider core through to the data center network.

Brocade 6PE: IPv6 over MPLS

Most service providers around the world today run MPLS-based IP backbones. For this reason, Brocade developed an IPv6 Provider Edge (6PE) router over MPLS to help service providers integrate IPv6 into their MPLS network. Brocade 6PE-enabled backbones allow IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires no backbone infrastructure upgrades and no reconfiguration of core routers and represents a very cost-effective strategy for IPv6 deployment for service providers.

Router Redundancy with VRRPv3

Virtual Redundancy Router Protocol (VRRP) is an election protocol that provides redundancy to routers within a VLAN. VRRP allows alternate router paths for a host without changing the IP address or MAC address by which the host knows its gateway. Both the Brocade MLX and NetIron XMR Series routers support the standard VRRPv3, which supports both IPv4 and IPv6 environments.

BROCADE HARDWARE SUPPORTING IPv6

The industry-leading Brocade NetIron CES and CER products are IPv6-ready edge routers. They provide:

- Brocade IPv6 compact switch solutions
- Standard-based IPv6 management edge features
- Wire-speed IPv6 routing and forwarding in hardware
- Redundant power supply
- Low power consumption for greener use and reduced TCO
- 1-1.5U compact form factor and reduced TCO
- Copper, fiber, and hybrid port choices
- 10-GbE uplink option for high-speed connectivity to aggregate and core



Both Brocade XMR and MLX/MLXe product lines are IPv6-ready aggregation and core routers.



- Brocade IPv6 modular solution with port density and diversity
- Standard-based IPv6 software feature sets
- Wire-speed IPv6 routing and forwarding in hardware
- Up to 1536-GbE ports or 256 x 10-GbE ports with wire-speed performance
- Copper, fiber, and hybrid port choices
- 10/100/1000 and 10-GbE modules to mix and match
- Virtual output queuing for advanced buffering
- Redundant management and fabric modules
- Redundant 1+1 power supplies
- Reduced power consumption for a greener solution and reduced TCO

BROCADE SOFTWARE

Brocade IronWare® OS for NetIron products provides the intelligence behind high-performance switches, routers, and application delivery platforms from Brocade. Leveraging over 12 years of experience in powering global networks, Brocade IronWare OS has demonstrated ongoing innovation to ensure nonstop network operations. The features for IPv6 are summarized in Table 4.

Table 4. Brocade IronWare IPv6 Features

Category	Feature
Essentials	IPv6 (RFC2460)
	ICMPv6 (RFC2463)
	Neighbor Discovery (RFC2461)
	Stateless autoconfiguration
	ECMP
	Link aggregation
	VRRPv3/VRRP-E for IPv6
	BFD for OSPFv3 and IS-IS
	DNS v6
	DHCPv6 relay
Routing	Static routes
	RIPng
	IS-IS, MT-ISIS
	OSPFv3
	MP-BGP
	PBR
	IPv6 prefix list
Multicast	MLD v1 and v2
	MLD access group
	PIM-SM, PIM-DM, PIM-SSM
	MSDP, DVMRP
	Static mRoutes
	MP-BGP for IPv6 multicast
Management and Applications	ICMPv6
	IPv6 MIB
	Stateless DHCP
	DHCPv6 prefix delegation
	sFlow
	Radius AAA

Category	Feature
	Ping, Telnet, Traceroute, FTP, SSH, TCPdump
	Management VFR for IPv6
Security	IPv6 standard ACL
	IPv6 extended ACL
	ACL accounting
QoS	IPv6 QoS
Options and Miscellaneous	Optional headers
	Path MTU discovery
Transition	Configured and automatic tunnels
	6to4
	IPv6 over GRE
	6PE
Performance	Hardware routing for unicast and multicast
	Hardware ACL handling
	Wire-speed IPv6 packet forwarding

Key IPv6 capabilities of the Brocade NetIron product family include:

- Dual-stack router with concurrent IPv4 and IPv6 operations
- Hardware-based IPv6 multicast, 6-to-4 tunneling, and IPv6 Access Control Lists (ACLs)
- Up to 240K IPv6 routes in hardware FIB
- Carrier trunks allow 64x10G LAG to be deployed using an advanced load-sharing algorithm that operates on IPv6 headers.
- Routing protocols supported: Static, RIPng, OSPFv3, ISIS, and BGPv6
- Up to 10 Million IPv6 routes in BGP+ RIB with 2000 IPv6 BGP+ neighbors
- Up to 64K IPv6 OSPFv3 routes
- BFD support for OSPFv3, IS-IS, and Multi-Topology IS-IS
- Multicast support: PIM-SM, PIM-SSM, MLD v1/v2, and 4K multicast IPv6 entries
- VRRP-v3 and VRRP-E with VRID reuse

BROCADE IPv6 DEPLOYMENT AND TRANSITION STRATEGIES ON NETIRON AND MLX PLATFORMS

IPv4 networks dominate in the market today and will continue to dominate in the future. However, the deployment of IPv6 networks will not happen all at once, and it is expected that the IPv6 networks and IPv4 networks will coexist for quite a while. Therefore, a key part of any IPv6 solution is its ability to integrate into and coexist with existing IPv4 networks.

To get from IPv4 to IPv6, a number of different transition methods have been proposed over the years. Some of these methods have been implemented in products, and some have been deployed in networks. Because many transition technologies are still in motion, confusion can result when service providers and enterprises try to select a transition technology for their environments. Given that time is running out, and large portions of the Internet are not IPv6 ready, a significant challenge confronts the service provider community.

To address this challenge, Brocade has developed a set of dual-protocol networking solutions designed with embedded support for the high-speed switching and routing of IPv4 and IPv6 traffic, as well as IPv6 transition mechanisms, which have become popular in the industry. These solutions are based on known, well-understood, and standardized technologies. These solutions are targeted at service provider backbones when the amount of IPv6 traffic dictates a transition strategy.

Tunnel: 6to4 Tunnel and Configured Tunnel

Automatic 6to4 tunneling refers to a technique in which the tunnel endpoints are automatically determined by the routing infrastructure. Tunnel endpoints are determined by using well-known IPv4 addresses on the remote side and embedding IPv4 address information within IPv6 addresses on the local side. A 6to4 tunnel allows isolated IPv6 domains to be connected over an IPv4 network to remote IPv6 networks, each of which has at least one connection to the shared IPv4 network.

A 6to4 tunnel can be configured on a border router in an isolated IPv6 network, which creates a tunnel on a per-packet basis to a border router in another IPv6 network over an IPv4 infrastructure. The border router at each end of a 6to4 tunnel must support both the IPv4 and IPv6 protocol stacks.

A manually configured tunnel is equivalent to a permanent link between two IPv6 domains over an IPv4 network. The primary use is for stable connections that require regular, secure communication between two edge routers. An IPv6 address is manually configured on a tunnel interface, and manually configured IPv4 addresses are assigned to the tunnel source and destination. The border router at each end of a 6to4 tunnel must support both the IPv4 and IPv6 protocol stacks.

Tunneling requires only the border routers to support both IPv4 and IPv6. Therefore, for service providers in the early stage of IPv6 deployment, tunneling facilitates the implementation of the internetworking between IPv6 islands and the incremental deployment of IPv6 without upgrading the entire network. In this way, the implementation scope of IPv6 is gradually enlarged.

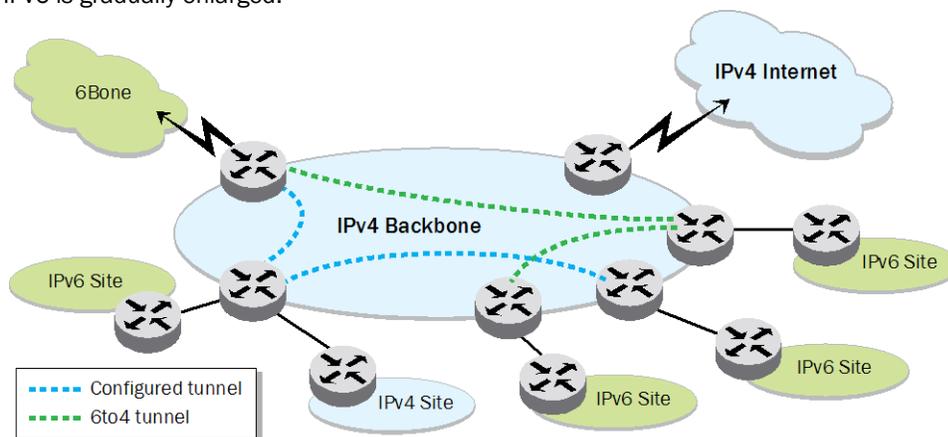


Figure 1. 6to4 Tunnel and Configured Tunnel

DUAL STACK: IPv6 AND IPv4 COEXIST

Specified in the RFC4213 standard, dual stack refers to the technique for providing message interworking between terminal/network nodes and IPv4/IPv6 nodes by implementing both IPv4 and IPv6 protocol stacks on terminal devices and network nodes. Routers supporting IPv4/IPv6 dual stack enable the network to act as two parallel logical networks. With good interoperability, the dual-stack model is the most direct method of making IPv6 and IPv4 node compatible. Moreover, this approach is easy to understand and deploy.

Brocade provides ASIC-based, wire-speed IPv6 and IPv4 packet forwarding in all its products. Since the dual-stack model is easy to understand and straightforward to implement, service providers can use a dual-stack approach with Brocade high-performance products to achieve a smooth transition to IPv6.

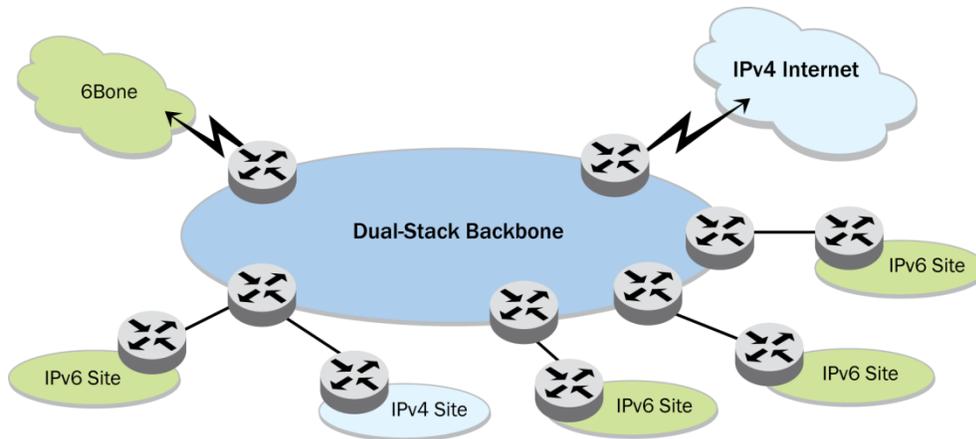


Figure 2. IPv4 and IPv6 Dual Stack

6PE: IPv6 over MPLS

For service providers running an MPLS backbone network, Brocade 6PE (IPv6 over MPLS) allows IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires no backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on labels rather than on the IP header itself, which provides a cost-effective strategy for IPv6 deployment.

Additionally, the inherent Virtual Private Network (VPN) and Traffic Engineering (TE) services available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over infrastructure that supports IPv4 VPNs and MPLS-TE.

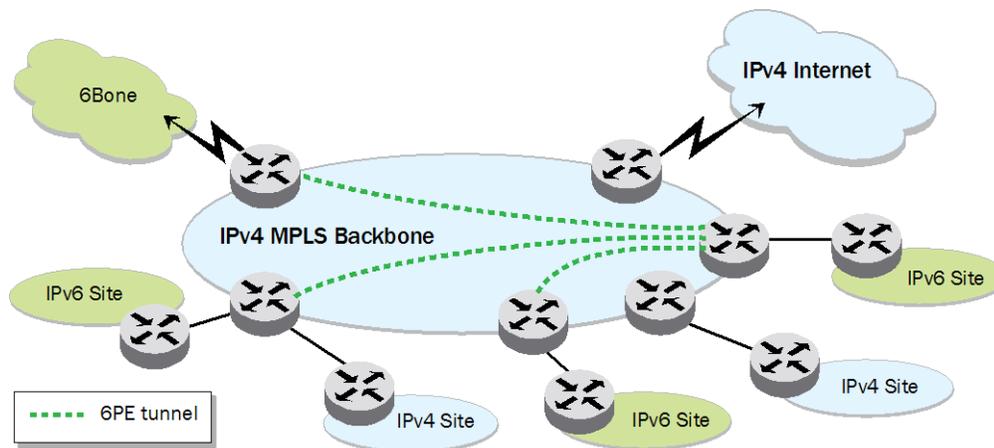


Figure 3. IPv6 over MPLS Backbone

Total IPv6

The end goal of an all-IPv6 Internet infrastructure, with IPv4 used sparingly if at all, is probably decades away. All sections of the network would support IPv6 natively—a true end-to-end IPv6 Internet. With this deployment, all devices would be globally reachable because they all have unique addresses. This true IPv6 network will enable new revenue-yielding services for the service provider, and the ultimate driving force will be a “killer” application that runs on an IPv6 network.

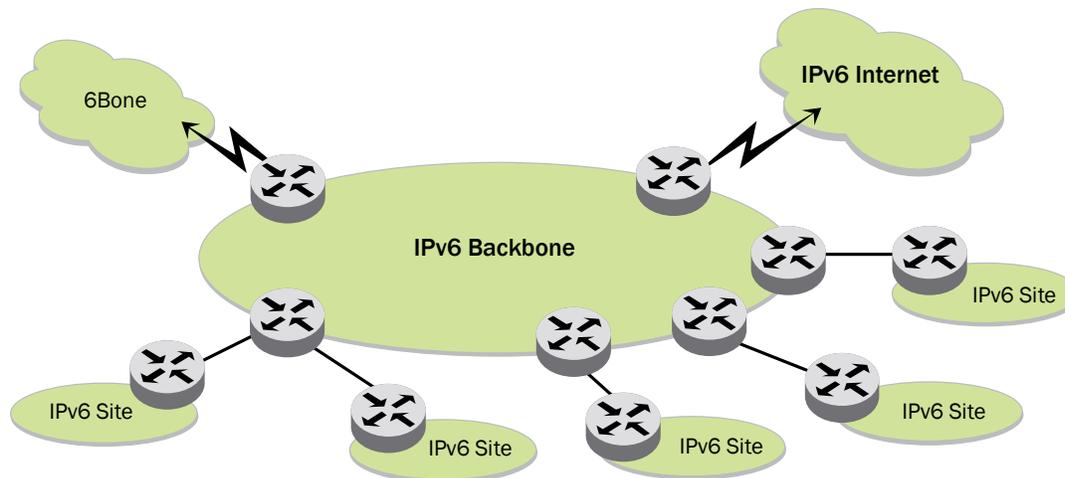


Figure 4. End-to-End IPv6 Network

MOVING FORWARD WITH BROCADE

The wait is over—IPv6 is here. Still, the transition to IPv6 will be a step-by-step process. Brocade understands that both IPv4 and IPv6 versions will coexist for the foreseeable future, and we are committed to delivering IPv6 technology to help our customers and partners with this transition. Brocade is providing mature and highly interoperable solutions with IPv6 capabilities, which are described in this paper, and Brocade will continue to innovate, enhance, and deliver additional IPv6 features across its product offerings. Enterprises and service providers seeking a smooth transition to IPv6 can rely on Brocade’s strong commitment toward IPv6.

ABOUT BROCADE

Offering an industry-leading family of Ethernet, storage, and converged networking solutions, Brocade helps organizations achieve their most critical business objectives through unmatched simplicity, non-stop networking, application optimization, and investment protection. To ensure a complete solution, Brocade partners with world-class IT companies and provides a full range of education, support, and professional services offerings. Learn more about IPv6 and Brocade products and solutions at www.brocade.com.

© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 08/11 GA-TB-369-02

Brocade, the B-wing symbol, BigIron, DCFM, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, Turbolron, and Wingspan are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, Extraordinary Networks, MyBrocade, VCS, and VDX are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.