White Paper

IPv6 as a Catalyst for Service Growth

Multiple System Operators Making the Transition

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408.745.2000
1.888 JUNIPER
www.juniper.net

Part Number: 200229-001 June 2007

Table of Contents

Copyright ©2007, Juniper Networks, Inc.

Executive Summary

Multiple System Operators (MSOs) are beginning to recognize that, given the importance of convergence to their long-term viability, the eventual transition to supporting and using IPv6 will be critical to their continued success. IPv6 helps enable convergence and is quickly becoming a necessity for larger operators. With the numerous transition mechanisms available, and given a reasonable time span, transitioning to IPv6 can be relatively easy and inexpensive. However, all MSOs would be wise to begin planning for IPv6 transition soon. Those who do will have plenty of time to implement those plans, therefore helping ensure success and minimizing potential disruption.

Although this paper mainly focuses on MSOs making the transition to IPv6, it is targeted at anyone interested in the history and development of IPv6 and IPv4-to-IPv6 transition technologies and strategies. The paper begins with a history and overview of IPv6 and then discusses the significance of this technology and the transition mechanisms available for migration to IPv6. The paper then presents an implementation example, examining Comcast's IPv6 transition strategy; fundamental to Comcast's approach is the incremental introduction of IPv6, which allows sufficient time for a learning curve for the operations personnel while minimizing disruption to existing networks, devices, and customer services.

Introduction

Convergence has been a top buzzword in the networking industry for some time now. Its meaning has shifted over time, but today it refers to a multitude of services supported on a common infrastructure. These services include:

- Voice (fixed and mobile)
- Video (on-demand and streaming)
- Conferencing (which may include voice, video, messaging, presentation graphics, audience data, and other information)
- Entertainment (in too numerous forms to list here)
- Wide-area, secure private networks (point-to-point and multi-access)
- Traditional Internet

For cable operators, convergence provides the capability to offer a rich portfolio of services while using a common infrastructure, enabling them to pursue multiple revenue streams, control costs, and strengthen competitive offerings. The financial benefits of not having to build, maintain, and operate separate systems for each service can be significant. For customers, convergence provides the capability to receive a wide range of communications and media services, all from a single provider and through a single connection.

IP: One Protocol to Bind Services

The infrastructure onto which MSOs are converging these current and future services is a router-based IP network. Donald Davies of the British National Physical Laboratory proposed packet switching of voice and other communication data in the mid 1960s, a decade before Vint Cerf and Bob Kahn were credited with creating the Internet Protocol. The rest of the story is now history: the ARPANET, for which IP was originally invented, eventually evolved into the Internet that we know today. Today, applications ranging from the ARPANET's original Telnet to the newest Internet applications all run over IP.

Copyright ©2007, Juniper Networks, Inc.

At the time that IP and the Internet were evolving, many other networks were also being developed, from cable television and telephony networks to Asynchronous Transfer Mode (ATM) and Frame Relay (FR) networks. Similarly, numerous network protocol suites were being developed, from Open Systems Interconnection (OSI) to Netware to the Data Over Cable Service Interface Specification (DOCSIS) family of protocols. So why has IP become the standard on which network services are converging?

To answer this question, we need look no further than the prevalence of the Internet. At the turn of the twenty-first century, most MSOs were still specializing in one or a few services (with service-specific infrastructures), but almost all had IP networks, and although their customers were consumers of some combination of network services—video in almost all cases, plus perhaps voice or commercial services—most customers also had the need for Internet service. This commonality of infrastructure among MSOs and customers has made IP the logical protocol on which to build multiservice networks.

IPv6: More Addresses. Lots More.

The main problem with the traditional version of IP (IPv4) is its address space—that is, the total number of addresses that the protocol can supply. When it was created for the relatively tiny ARPANET research network, IPv4's 32-bit address space seemed enormous, but that was before the advent of e-mail and the Web, search engines and the IP Multimedia Subsystem (IMS), and the proliferation of such applications into homes and businesses worldwide.

Even in the early 1990s, before the World Wide Web set off an explosion in Internet use, concerns about the continued availability of IPv4 addresses were being raised. Exponential increases in the rate of address allocation were being observed, with the complete depletion of IPv4 addresses predicted by the mid 1990s.

A new version of IP, with a larger address space, was called for. The result was IPv6, which was first standardized in 1995. With a 132-bit address space, the number of available IPv6 addresses is almost inconceivably greater than the number of addresses available in the 32-bit IPv4 space: 4.3 billion total IPv4 addresses versus 3.4 x 1038 total IPv6 addresses. To put the contrast into perspective, if one IPv4 address was one picometer (one trillionth of a meter) long, the entire IPv4 address space would be approximately 4.3 millimeters long—the length of an average-sized ant. If an IPv6 address were one picometer long, the length of the entire IPv6 address space would be approximately 3.4 billion trillion trillion kilometers, or 36 billion light years. The farthest visible object in the universe from the Earth is estimated to be 30 billion light years away, so that is a very long distance—and a lot more IP addresses.

The Sudden Interest in IPv6

Despite the astronomically larger address space, people showed little interest in adopting IPv6 over the ensuing five years after the IPv6 base specification was authored. Network Address Translation (NAT) and private IPv4 addresses were deployed to slow the depletion of IPv4 addresses while IPv6 was being developed, and NAT was so effective in slowing the rate of IPv4 address allocation that in many quarters—and quite commonly in the United States—people expressed doubt about whether transitioning to IPv6 would even be necessary.

Copyright ©2007, Juniper Networks, Inc.

However, the trend toward network convergence has rekindled the interest in IPv6 for three main reasons:

- **Planned services and devices:** The intention to provide a wide array of IP services to a vast number of customers, plus plans for perhaps billions of new network-enabled devices (from mobile phones to home appliances to entertainment systems) will create a demand for IP addresses that IPv4 simply cannot meet.
- **Multiple service profiles to a single location:** Multiservice offerings to a home or office require varying levels of quality and security. In your home, you may be carrying on a telephone conversation while your spouse is conducting a video conference with clients, your daughter is watching a movie, and your son is playing a video game—all over the same IP connection from the same cable operator. If these multiple applications require different levels of quality and security, they cannot work together through a NAT device, and if NAT is taken out of the picture, the conservation of IP addresses also goes away.
- **New markets:** The expansion of service offerings to existing customers by itself represents a big demand for new IP addresses. When developing regions of the world with enormous populations and rapidly expanding economies—China and India being prime examples—are taken into account, IPv4 becomes entirely insufficient. The population of the People's Republic of China alone—some 1.3 billion people—is larger than the number of remaining, unallocated IPv4 addresses.

Plans for new network-enabled devices plus the presence of emerging markets explain why interest in IPv6 first arose in Asia. In the days of the ARPANET, American universities and companies participating in that U.S. research network reserved huge blocks of IPv4 addresses. There was nothing malicious in this; at the time no one anticipated the scope of the modern Internet or the resulting worldwide demand for IP addresses. Nevertheless, the result of this early development was that, while North America was awash in IPv4 addresses, other parts of the world—particularly Asia—found it difficult to acquire sufficient addresses for even their current applications, much less for the growth predicted in Internet infrastructure, new users, and new devices.

Japan became the first country to set a determined direction toward IPv6 transition. Influenced by the country's extensive consumer electronics industry, the Japanese government set transition goals and timelines through the e-Japan Initiative. South Korea and Taiwan soon followed Japan for the same reason: economies dependent upon vibrant consumer electronics industries requiring an extensive, readily available supply of IP addresses to continue innovating. The People's Republic of China began the China Next Generation Internet (CNGI) IPv6 project, partly driven by its own electronic industries, but also in recognition of the inability of IPv4 to supply enough addresses to its huge, increasingly middle-class population.

Although U.S. interest in IPv6 was lukewarm (at best) until as recently as two years ago, most of the larger MSOs and other providers are now acknowledging that migration to IPv6 is in their near future. Some are in the early exploratory stages, many already have firm transition plans in place, and a few are actively implementing IPv6.

In the United States, the change of attitude toward IPv6 can be attributed to three factors:

- **Acknowledgment of the serious efforts taking place in Asia:** American operators are feeling the need to stay competitive with the Asian market.
- **Aggressive IPv6 transition plans within several branches of the U.S. federal government:** The U.S. government is a huge customer of IP services, and MSOs understand that they must support IPv6 if they want to keep or gain federal agencies as customers.
- **Growing multiservice plans:** As MSOs plan multiple service offerings, they are seeing that IPv4—even private 10/8 address space—will not support projected addressing requirements.

Copyright ©2007, Juniper Networks, Inc.

For example, Comcast predicts that by 2010 the number of IP addresses required to serve their customers—taking into consideration both ongoing traditional services and planned new services—will exceed 80 million. As Figure 1 shows, beginning in 2007 most of the pressure on their address requirements will come from new triple-play services: high-speed Internet, cable television, and telephone. Yet in 2005, even ongoing addressing demands of traditional high-speed data (HSD) service offerings surpassed 16.8 million addresses—the entirety of the 10/8 private IPv4 space.

Device	Addresses per Device: 2005 HSD Only	Addresses per Device: 2006+ Triple Play
Cable Modem (CM)	1 (private only)	1
Home Computer/ Router	1	1
Voice Adapter (eMTA)	0	1-2
Set-top Box (STB)	0	2
Total IP Addresses (Avg. 2.5 STBs per Household)	1-2	8-9

Figure 1: Effect of Triple Play Address Requirements in the Comcast Network

The good news is that, although MSOs are wise to begin planning for IPv6 transition sooner rather than later, they still have plenty of time to carry out these plans. Furthermore, the IPv6 transition toolbox already has several mature transition technologies that fit a wide range of strategies, which means that the transition to IPv6 need not be difficult or expensive.

The IPv6 Transition Toolbox

IPv6 transition technologies available today can be grouped into three general categories:

- **Dual-stacked interfaces:** Dual-stacking enables a single interface on a host, router, server, or other networked device to support both IPv4 and IPv6 simultaneously.
- **Tunneling***:* A tunnel is a means for encapsulating an IPv6 packet behind an IPv4 header, or vice versa. This approach is used for connecting, for example, IPv6 devices or sites across IPv4 networks or, later in the transition cycle, IPv4 devices or sites across an IPv6 network.
- **Translators:** These devices allow an IPv4-only device to speak to an IPv6-only device.

Dual stacking of all devices in the network is the simplest means of transitioning. If all or most interfaces are “bilingual,” the transition is controlled through the domain name system (DNS). In other words, a dual-stacked device queries the name of a device to which it must speak, and if the DNS returns an IPv4 address, the device sends IPv4 packets. By the same token, if the DNS returns an IPv6 address, the device sends IPv6 packets. Most modern operating systems and most modern routers support dual stacking.

However, several factors can make dual-stacking an impractical or incomplete solution in some situations:

- Some older routers or interfaces in the network do not support dual stacking, or do not support IPv6 at an acceptable performance level.
- Some devices in the network support IPv6 only, rather than being dual-stack capable. For example, IP-capable mobile phones are planned that will support IPv6 only, but users will still need to be able to browse IPv4 Websites and send messages to IPv4 destinations.
- IPv6 may be used in some networks or applications specifically because there are not enough IPv4 addresses to serve all participating interfaces.

Copyright ©2007, Juniper Networks, Inc.

Tunnels come into play in circumstances where not all network interfaces are dual-stack capable, where IPv6 routing and forwarding performance on some devices is insufficient, or where some parts of the network do not support IPv6 at all. Tunnels can be classified in two categories: configured tunnels and automatic tunnels.

Configured tunnels are used in situations where permanent IPv6 connectivity is required between sites. They may be IP-in-IP, IP Security (IPSec), or MPLS-based tunnels. MSOs already operating an MPLS backbone have a head start on transition; not only can they interconnect IPv6 sites with native MPLS tunneling (6PE), but they can offer their customers Layer 3 IPv6 virtual private networks (6VPE), Layer 2 VPNs, and virtual private LAN service (VPLS) as options for customer IPv6 connectivity.

Automatic tunnels are used when temporary connectivity between sites or devices is required, and they include mechanisms for signaling the information necessary for setting up and tearing down the tunnels. The automatic tunneling mechanisms that have found general acceptance in IPv6 networks are:

- **Tunnel brokers:** Proprietary implementations of tunnel brokers, in which a server provides tunnel information for site-to-site connectivity, have been around for quite a few years now.
- **6to4:** Like tunnel brokers, 6to4 is used for site-to-site connectivity. However, instead of a server, 6to4 encapsulates the IPv4 tunnel endpoint information inside the IPv6 source and destination addresses.
- **Intra-Site Automatic Tunnel Addressing Protocol (ISATAP):** As the name implies, ISATAP is used within a single site to connect individual devices. Normally it is used to connect an IPv6 device to an IPv6 router across an IPv4 site, but it can also be used to connect two IPv6 hosts within a site.
- **Teredo:** Most automatic tunneling mechanisms do not work through a NAT device (although a few tunnel brokers do). Teredo gets around this problem by encapsulating IPv6 packets in User Datagram Protocol (UDP)/IPv4.
- **Dual Stack Transition Mechanism (DSTM):** DSTM is useful in the later stages of IPv6 transition, allowing IPv4 packets to be tunneled over an IPv6 site. It can also be useful in situations where dual stacking is practical but IPv4 addresses are scarce because it has the capability to assign an IPv4 address to a dual-stacked host for only the amount of time the address is needed.

When IPv6 devices exist in the network—either specialized devices such as the advanced mobile phones mentioned previously or devices for which IPv4 addresses are not available for dual stacking—and those devices must speak to IPv4 devices, translators are required. Although many translator mechanisms have been proposed, the only solution that has gained wide acceptance is Network Address Translation with Protocol Translation (NAT-PT).

Copyright ©2007, Juniper Networks, Inc.

Controlling the Cost of Transition

One factor contributing to early skepticism about IPv6 was the perceived cost of transition. Many saw transition as a short-term project in which most software and hardware in their network would have to be upgraded to IPv6 capability, but the reality is that, worldwide, transition plans span years. Objectives just to attain full IPv6 capability range from 2008 to 2012. Taking a long-range view of transition means that capital expenditure can be minimized. As network devices and software are upgraded as a regular part of the network lifecycle, MSOs can help ensure that IPv6 capability is migrated to the network by including IPv6 support in their selection criteria. In fact, as more and more vendors support IPv6 (and keep in mind that most already do), it will actually be difficult to find network products that do not support IPv6.

Nevertheless, there is a cost to transition. The cost will be primarily in operational and human expenditures: training, inventory, planning, and implementation. Yet even here the costs are acceptable. IPv6 routing is not much different from IPv4 routing, IPv6 DNS records look almost the same as IPv4 DNS records, and IPv6 address design and management are in fact much easier than with IPv4, because there are no complexities such as variable-length subnet masking.

With the abundance of transition mechanisms available, and given a reasonable time span, transitioning to IPv6 can be relatively easy and inexpensive. All MSOs need to do is decide how they want to approach it.

Comcast: A Transition Observed

Comcast is addressing its projected address requirements by planning for IPv6 now. The company's initial objective is to deploy IPv6 for the management and operation of customer-premise devices such as CMs, STBs, and multimedia terminal adaptors (MTAs). These first phases of transition allow Comcast to control and contain the scope of its IPv6 deployment. As worldwide IPv6 migration progresses, Comcast will be positioned to begin offering IPv6 to provide global services to its customers.

Crucial to Comcast's future direction—and indeed to any MSO's transition—are the new Data over Cable System Interface Specifications (DOCSIS) 3.0, which provide for both IPv6 and, through channel bonding, sharply increased bandwidth capabilities. Comcast's transition plan calls for dual stacking of its operational network: the Comcast Regional Access Networks (CRANs), the backbone connecting the CRANs, and the Cable Modem Termination System (CMTS) (Figure 2).

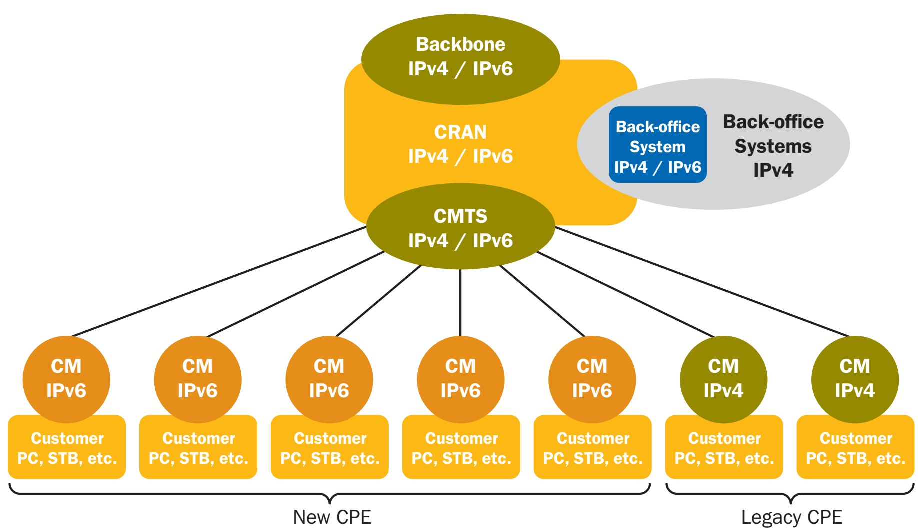

Figure 2: IPv6 Migration in the Comcast Network

Copyright ©2007, Juniper Networks, Inc.

Any new CMs, STBs, and MTAs at the customer premises will be IPv6-only. Back-office systems do not need to communicate with the customer premise equipment (CPE) and will for the most part remain IPv4 for the near term (although some back-office equipment must be modified to manage IPv6 addressing data). Dual stacking of the operational network means that existing IPv4-based customer devices will be unaffected, although they too will migrate to IPv6 as legacy customer equipment is gradually phased out.

Fundamental to the success of Comcast's transition plan is the incremental introduction of IPv6. Dual stacking is implemented in the backbone first (and has already been completed), followed by the CRANs and back-office systems, and finally the CMTS. Only then will the network be ready to push IPv6 to the customer premise. The incremental nature of this strategy allows sufficient time for a learning curve for the operations personnel and also minimizes disruption to existing networks, devices, and customer services.

Conclusion: Catalyzing Service Growth with IPv6

IP is here to stay, and the systems and services converging on IP are putting strain on the available IPv4 addresses that was unforeseen as recently as a few years ago. Cable operators like Comcast not only understand that IPv6 is essential to both the continued growth of their existing services and the rollout of new services, but they also recognize that planning for the introduction of IPv6 into their networks must begin now.

But there is still more to the story. Development of new IP applications has recently been stymied by the scarcity of IPv4 addresses and by the presence of NAT as a stopgap to address depletion. As IPv6 becomes more commonplace and NAT is removed from the picture, an environment will emerge that is far more conducive to the development of innovative new applications. As recently as three years ago, the need for a "killer app" to get IPv6 rolling was commonly cited. Now, forward-thinking MSOs understand that IPv6 must come first; the killer applications will quickly follow.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

Copyright 2007 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS and JUNOSe are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.