

IPv6: A catalyst for service growth

Flourishing Internet services, devices rekindle interest in IPv6

By Jeff Doyle, JUNOS Product Manager, Juniper Networks

Convergence has been one of the top buzzwords in the networking industry for some years now.

What that one small term means is that a multitude of services—voice (fixed and mobile); video (on-demand and streaming); conferencing (which might 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); and traditional Internet—can all be supported on a common infrastructure.

A rich portfolio of services means multiple revenue streams and a stronger competitive stance for MSOs and other service providers. The financial benefits of not having to build, maintain, and operate separate systems for each service can be significant. For the customer, receiving a wide range of communications and media services, all from a single provider through a single connection, has tremendous appeal.

IP: One protocol to bind them

The infrastructure on which all of these services are being converged 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, Bob Kahn and Jon Postel dreamed up IP. You know the rest of the story: The ARPANET, for which IP was invented, evolved into the Internet we know today. And applications from the ARPANET's original Telnet to the newest Internet applications all run over IP.

But there have been a great number of networks over that span, from telephony and cable television networks, to ATM and Frame Relay; as well as a host of net-

work protocol suites—from OSI to Netware to SNA. Why has IP become the standard on which network services are converging?

The answer is the Internet itself. At the turn of the century, most service providers were still specializing in one or a few services, with service-specific infrastructures, but almost all had an IP network. And while customers were consumers of some combination of network services—voice in almost all cases, video in the case of cable MSOs, plus perhaps private line or VPN or

With a 132-bit address space, the number of available IPv6 addresses is almost inconceivably larger than the addresses available within the 32-bit IPv4 space

early Transparent LAN services—most had Internet service. So this commonality of infrastructure among service providers and customers has driven IP as the protocol on which to build multi-service networks.

IPv6: More addresses. Lots more.

The 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 was thought to be enormous. But that was before the advent of e-mail and the Web, search engines and IMS (IP Multimedia subsystem), and the proliferation of such applications into homes and businesses worldwide.

Even in the early 1990s, before the World Wide Web caused an explosion in Internet usage, concerns about the continued availability of IPv4 addresses were being raised. Exponential increases in the rate of address allocation were 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, first standardized in 1995. With a 132-bit address space, the number of available IPv6 addresses is almost inconceivably larger than the addresses available within the 32-bit IPv4 space: 4.3 billion total IPv4 addresses versus 3.4×10^{38} total IPv6 addresses. Putting 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. That's a very long distance; that's a great many IP addresses.

Why the sudden interest in IPv6?

Despite the astronomically larger address space, little interest in adopting IPv6 was shown over the ensuing five years after its standardization. Network Address

E-mail: jdoyle@juniper.net

Translation (NAT) and private IPv4 addresses were deployed to slow the depletion of IPv4 addresses while IPv6 was being developed. NAT was so effective in slowing the rate of IPv4 address allocation that in many quarters—and quite commonly in the United States—doubt was expressed whether transitioning to IPv6 would ever be necessary.

But the trend to network convergence has rekindled the need for IPv6, for three 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 cannot meet.

- **Multiple service profiles to a single location:** Multi-service offerings to a home or office will require varying levels of quality and security. In your home, you might be having 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 service provider. 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 is also removed.

- **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 took huge blocks of IPv4 addresses. There

was nothing malicious in this; at the time no one anticipated the scope of the modern Internet, and the resulting worldwide demand for IP addresses. Nevertheless, the result was that while North America was awash in IPv4 addresses, other parts of the world—particularly Asia—found it difficult to acquire sufficient addresses even for 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 transi-

- An acknowledgement of the serious efforts taking place in Asia: American operators are feeling the need to stay competitive with their Asian counterparts.

- The aggressive IPv6 transition plans of several branches of the federal government: The government is a huge customer of IP services, and service providers understand that they must support IPv6 if they want to keep or gain federal agencies as customers.

- Growing multi-service plans: As service providers plan multiple service offer-

Device	Addresses per device: 2005 HSD only	Addresses per device: 2006+ Triple play
Cable modem (CM)	1 <small>(private only)</small>	1
Home computer/router	1	1
Voice adapter <small>(eMTA)</small>	0	1-2
Set-top box (STB)	0	2
Total IP addresses <small>(average 2.5 STBs per household)</small>	1-2	8-9

Figure 1: Effect of triple play address requirements in the Comcast network.

tion. 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 electronics industries, but also in recognition of the inability of IPv4 to supply enough addresses to its huge, increasingly middle-class population.

And where American interest in IPv6 was lukewarm at best as recently as two years ago, most of the larger telcos, ISPs, and MSOs are now acknowledging that IPv6 is in their future. Some are in the early exploratory stages, while many already have firm transition plans in place. A few are actively implementing IPv6. The change of attitude toward IPv6 in the U.S. can be attributed to three factors:

ings, they are seeing that IPv4—even private 10/8 address space—will not support projected addressing requirements.

For example, Comcast’s use of IP addresses for its high-speed data service exceeded 16.8 million addresses—the entirety of the 10/8 private IPv4 space—in 2005. Its current customer base exceeds 20 million, including just over 9 million HSD subscribers, and subscribers are projected to continue to grow more or less linearly. But the real address demand will come from Comcast’s planned new services, as Figure 1 shows. In addition to an address to control the cable modem and perhaps a home router or computer, one or two addresses will be required for VoIP and two for a set-top box. Considering an average of 2.5 STBs per household, a single triple play (HSD, VoIP, and video) household will require as many as nine IP addresses. So, 20 million triple play customers represents a need for more than 180 million addresses.

The good news is that although you are wise to begin planning for the IPv6 transition soon, there is still plenty of time

to carry out those plans; there is a grab-bag of mature transition technologies to fit a wide range of strategies; and the transition to IPv6 need not be expensive.

The transition toolbox

Transition technologies fall into three general categories:

- **Dual-stacked interfaces:** This allows a single interface on a host, router, server, or other networked device to support both IPv4 and IPv6 simultaneously.

- **Tunnels:** A tunnel is a method for encapsulating an IPv6 packet behind an IPv4 header, or vice versa. They are 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 all devices in the network is the simplest means of transitioning. If all or most interfaces are “bilingual,” you control your transition through DNS (domain name server): A dual-stacked device queries the name of a device to which it must speak; if DNS returns an IPv4 address, the device sends IPv4 packets; if DNS returns an IPv6 address, the device sends IPv6 packets.

Most modern operating systems and most modern routers support dual stacking. There are several factors, however, that 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 are IPv6-only, rather than dual-stack-capable. For example, IP-capable mobile phones are planned that will support IPv6 only, but the users will still need to be able to browse IPv4 Web sites and send messages to IPv4 destinations.


- IPv6 might be used in some networks or applications specifically because there are not enough IPv4 addresses to serve all participating interfaces.

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 might be IP-in-IP, IPSec, or MPLS-based tunnels. Service providers already operat-




VistaMAX

Broadband Wireless




WiMAX Ethernet Switch

Power over Ethernet
Power up to Eight VistaMAX Base Stations



VistaMAX OBR3500 Base Station

Integrated GPS
Point to Multipoint




Network Management System

Flexible Subscriber Bandwidth Management


All outdoor Base Stations and Subscriber Stations

VistaMAX launching at 3.5 GHz with future models planned for other bands including 5 GHz, 1.9 GHz



Subscriber Home / Office


Efficiently Scalable



VistaMAX OSR3500 Subscriber Station

Power over Ethernet
VCom Power Connect (audible alignment beeper)

VistaMAX Products will be submitted for certification testing by WiMAX Forum™



Principal Member

www.vcom.com

Advanced Broadband Products

Tel: (888) 292-VCOM / (306) 955-7075 Fax: (306) 955-9919 Email: sales@vcom.com

ing an MPLS backbone have a leg up on the transition; not only can they interconnect IPv6 sites with native MPLS tunneling (6PE), 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.

vices. Normally it is used for connecting an IPv6 device to an IPv6 router across an IPv4 site, but it can also be used for connecting two IPv6 hosts within the 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 en-

Controlling the cost of the transition

A factor contributing to early skepticism about IPv6 was the perceived cost of the transition; many saw the 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. Target dates just to attain full IPv6 capability range from 2008 to 2012. Taking a long-range view of the transition means that capital expenditure can be minimized; as network devices and software are upgraded as a regular part of the network lifecycle, you can insure that IPv6 capability is migrated into your network by including in your selection criteria IPv6 support. 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 associated with the 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. 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 an abundance of transition mechanisms available and given a reasonable time span, transitioning to IPv6 can be relatively inexpensive. All you need to do is decide how you want to approach it.

Comcast: A transition observed

Comcast is answering its projected address requirements shown in Figure 1 by planning for IPv6 now. Its initial objective is to deploy IPv6 for the management and operation of customer premise devices such as cable modems (CMs), set-top boxes (STBs), and multimedia terminal adapters (MTAs). These first phases of the transition allow Comcast to control and contain the scope of the IPv6 deployment. As worldwide IPv6 migration progresses, Comcast will be positioned to begin offering IPv6 to provide global services to its customers.

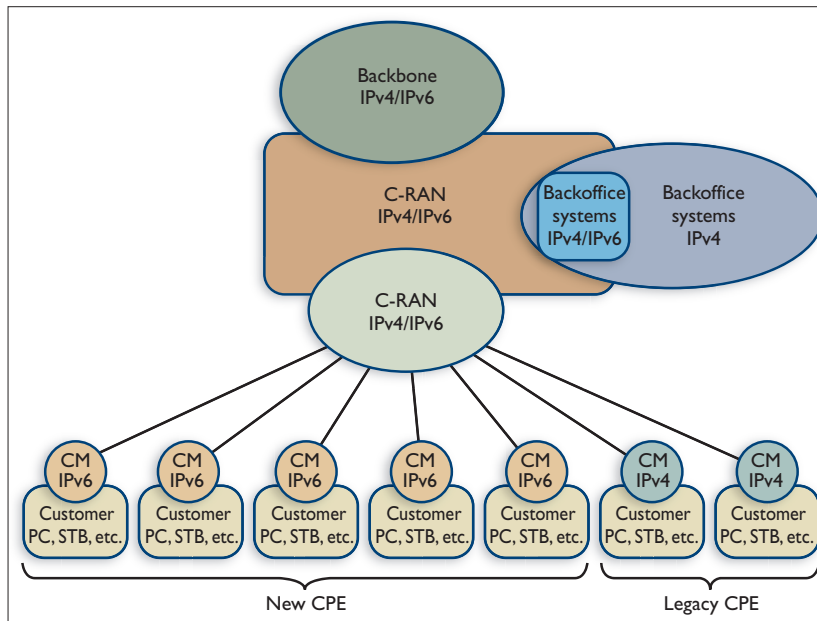


Figure 2: IPv6 migration in the Comcast network.

Automatic tunnels are used when temporary connectivity between sites or devices is required, and 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.

- 6to4: Like tunnel brokers, 6to4 is used for site-to-site connectivity; but 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 for connecting individual de-

capsulating IPv6 packets in UDP/IPv4.

- Dual Stack Transition Mechanism (DSTM): DSTM is useful in the later stages of a 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, by assigning an IPv4 address to a dual-stacked host for only the amount of time the address is needed.

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

Key to Comcast's future directions—and indeed to any MSO's transition—are the new DOCSIS 3.0 specifications, which provide for both IPv6 and, through channel bonding, sharply increased bandwidth capabilities (see "Driving DOCSIS 3.0," *CED*, April 1, 2005).

Comcast's transition plan calls for dual-stacking its operational network: The Comcast Regional Access Networks (CRANs), the backbone connecting the CRANs, and the CMTS (see Figure 2). Any new CMs, STBs, and MTAs at the customer premises will be IPv6-only. Most back office systems do not need to communicate with the customer premise equipment, and will for the most part remain IPv4 for the near term (although a subset of the 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 naturally 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 is the network ready to push IPv6 to the customer premise. The incremental nature of this strategy allows sufficient time for a learning curve among the operations personnel, and also minimizes disruption to existing networks, devices, and customer services.

IPv6: Catalyzing service growth

IP is certainly here to stay, and the systems and services converging on IP are putting strains on the available IPv4 ad-

resses that were unforeseen just a few years ago. Service providers like Comcast not only understand that IPv6 is essential both to the continued growth of existing services and to the rollout of new services, but also that planning for the introduction of IPv6 into their networks must begin now.

But there is more: Development of new IP applications is becoming stymied by the dearth of IPv4 addresses and by the presence of NAT as a stop-gap to address depletion. As IPv6 becomes more commonplace and NAT is moved out of the picture, we will find an environment 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 service providers understand that IPv6 must come first—the killer applications will quickly follow. ■

Coming December 6 and 7...



It's Year Three for this proven and exclusive annual online tradeshow & conference.

- Visit the booths of leading broadband vendors
- Listen to daily keynote addresses
- Benefit from relevant and timely conference sessions organized into a number of high interest tracks.



Do it all from your PC's Web browser, without leaving your office.

Additional details coming soon.

CED

Broadband Marketers
Contact Todd Marquez, Associate Publisher, at todd.marquez@advantagemedia.com for exhibiting and sponsorship information.