Deploying IPv6

Alain Durand • *Sun Microsystems* • *Alain.Durand@sun.com*

Version 4 of the Internet Protocol (IPv4) has been in wide use for the past 25 years, but it has outlasted its effective lifetime and is straining to keep up with the demands of today's Internet. IPv6 is a natural evolution from IPv4 and attempts to address many of the older protocol's shortcomings. IPv6 is designed to keep up with the rapid growth of the Internet. The core specifications have been standardized through the IETF's IPng working group, and the NGtrans working group is studying issues surrounding the challenging task of smoothly transitioning from IPv4 to IPv6.

Benefits of IPv6

The Internet has been facing the depletion of IPv4 address space since the early 90s. The registries (RIPE-NCC, ARIN, APNIC) responsible for global IP address allocation have enacted strict policies to help minimize the problem, but the Internet continues to grow exponentially. Network administrators must increasingly rely on network address translation (NAT) technologies to deploy networks. Managing those NAT devices is a complex and time-consuming task and, by their nature, NATs break the end-to-end principle of the Internet. Some applications, such as IPsec, simply cannot work across a NAT device, and other applications, such as VoIP, require adding special code to the NATs in order to function.

With the deployment of high-speed, always-on connections at home via DSL and cable modems, and the upcoming third-generation wireless devices, the Internet is changing. Users are no longer connecting just computers, but a whole new range of gadgets that will require global IP addresses.

In the face of these issues, IPv6's main driving force and benefit is its large address space. For example, the current IPv6 address allocation policy recommendation is to allocate a 48-bit prefix to every site on the Internet, whether homes, small offices, or large enterprise sites. The 48-bit prefix allows 65,000 subnets within each site, each of which could accommodate a virtually infinite number of hosts.

IPv6 also brings such benefits as stateless autoconfiguration, more efficient mobility management, and integrated IPsec.

Obstacles to Deployment

IPv4 is an extraordinary success story, with about 200 million users around the world. The foundation of such a successful industry cannot be changed overnight. Deployment can only happen incrementally, and extra care should be taken to avoid impacting IPv4 traffic in any way.

A basic concept of the Internet, almost as famous as Moore's law, is Robert Metcalf's law that states that a network's "value" or "power" increases in proportion to the square of the number of nodes on it. A brand new network with no established applications or services cannot attract users, and thus, is not as interesting as a mature, well established network.

The foundations of an industry can't be changed overnight.

Following Metcalf's law, IPv6 currently has low value because IPv4 has a huge population of users, and IPv6 has a much smaller number of users. Network managers will naturally face the question of why to transition to a new network that their peers do not yet support.

Another obstacle to the widespread deployment of IPv6 is the missing infrastructure required for realistic production-level deployment of the protocol. Hardware support, operating systems, middleware, applications, management tools, and trained technical staff are needed in order to complete the picture. Each of these missing parts represents a significant investment in time, resources, and of course, money. For example, high-speed network routers will very likely require some hardware

1089-7801/00/$10.00 ©2001 IEEE

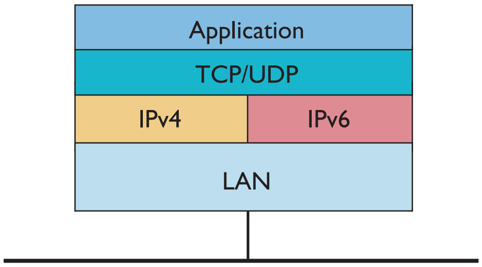

Figure 1. Hybrid stack. Applications running on a hybrid, or dual, stack node use the same code to access both IPv4 and IPv6 networks.

upgrade to switch IPv6 packets at the same speeds they are processing IPv4 packets today. Investments in legacy applications may need to be amortized before upgrades can be adopted. Software vendors might also be reluctant to enhance IPv6 support without a clear picture of the return on their investment.

Tunneling can connect IPv6 islands across an ocean of IPv4 networks.

Although porting an application from IPv4 to IPv6 is usually a straightforward process, hidden issues do surface. For example, third-party applications are rarely compiled on the latest version of the targeted operating system, but on an earlier release that is likely to be in wide use by the customer base. Operating systems usually provide backward compatibility, so there is generally no problem running the new application on a more recent release of the OS. To run on IPv6, however, an application must be compiled on a version of the OS that supports it.

Today, this means compiling the application on the latest releases of operating systems—often bypassing more than one previous version. Thus, a third-party application vendor who wants to create an IPv6-ready application will have to choose between forcing customers to upgrade their operating systems to a IPv6-compliant version, or supporting two versions of the application: one that runs on IPv6 and IPv4 on the newer OS and one that runs only on IPv4 on the previous release.

Current Deployment Status

Many major vendors announced or introduced IPv6-compatible products during 2000. Sun Microsystems, for example, shipped Solaris 8 with IPv6 support, and Microsoft revealed a “technology preview” that prefigures what will be in the next major release of Windows. Cisco also published its three-phase roadmap for delivering IPv6 services. More information on IPv6, and a more complete list of available implementations, can be found at http://playground.sun.com/ipng.

Some of these implementations are in beta or developer releases, while others are fully supported production-ready products. Vendors have demonstrated the interoperability of these implementations in the Internet Labs (ilabs) of the Networld+Interop venues (in Atlanta, Georgia, and Las Vegas, Nevada) as well as at events such as the annual Connectathon event in San Jose, California, where vendors conduct thoughtful tests of their “prerelease” software.

On the Internet Service Provider front, Japanese ISPs (NTT, IIJ, and others) provided IPv6 trial service in both tunnel mode and on native links during 2000. They are preparing to make full-scale IPv6 service available soon. A new U.S.-based startup company, Zama Network, is building its business plan on an IPv6 backbone, and several other big ISPs in the U.S., Asia-Pacific, and Europe are also moving toward providing IPv6 services.

It is now clear that a strong undercurrent of deployment activity is already taking place in some significant key industries. One such area where IPv6 is expected to take off is the third-generation cellular phone industry. The Mobile Wireless Internet Forum has mandated IPv6 support in its architecture. For its architecture for the multimedia domain, the Third Generation Partnership Program (3GPP) has chosen to use IPv6 exclusively, and the Third Generation Partnership Program 2 (3GPP2) is considering IPv6 for its all-IP architecture.

From IPv4 to IPv6

The IETF NGtrans working group has designed a set of IPv4-to-IPv6 transition tools to address the various needs of different networks. The two most basic building blocks of the toolbox are the *hybrid stack mechanism* and *tunneling*.

Figure 1 illustrates the first of these tools. A hybrid stack host, also known as a dual stack host, implements both IPv4 and IPv6—usually in a single stack in which most of the code is shared by the two protocols. The host “speaks” IPv4 with IPv4 peers, and IPv6 with IPv6 peers. When both options are available, the host will usually choose the IPv6 path, which increases the value and power of the IPv6 network by creating more users.

Tunneling provides a convenient way for an IPv6 island to connect to other IPv6 islands across an ocean of IPv4 networks. Figure 2 shows how an IPv6 packet can be encapsulated within the payload of an IPv4 packet.

In addition to these two methods, the NGtrans working group has developed other transitioning tools as well. These can be grouped into two categories:

- tools that attempt to automate the configuration of tunnels connecting isolated IPv6 networks over the IPv4 Internet;
- tools that attempt to enable communication between IPv6 nodes and IPv4 nodes.

Two tools have emerged so far from the first category: the *tunnel broker* (which aims to connect single hosts or small networks to the IPv6 Internet backbone) and the *6-to-4 mechanism* (illustrated in Figure 3).

The 6-to-4 mechanism automatically derives a 48-bit IPv6 prefix from any global IPv4 address. It then spans stateless tunnels over the IPv4 infrastructure to connect to other 6-to-4 domains. The beauty of this mechanism is that a single global IPv4 address can bring IPv6 connectivity to an entire site. Adding IPv6 support and 6-to-4 functionality to IPv4 NAT boxes presents a very attractive transition scenario. In an enterprise with multiple branch offices, each of which uses IPv4 private addresses and NAT technology, the 6-to-4 strategy could create a virtual IPv6 extranet. It would also reestablish end-to-end IP connectivity and allow the enterprise to use IPsec between all the servers in the different locations.

The NGtrans tools in the second category include *translation*. This simple, natural approach can be performed at the IP, transport, or application layers.

- At the IP layer, the basic translation mechanism is known as Simple IP/ICMP Translator (SIIT), on top of which the Network Address Translator-Protocol Translator (NAT-PT) was built.
- At the transport layer, the SOCKS mechanism has been updated to allow IPv6/IPv4 relaying. The TCP-Relay methodology proposes to terminate TCP or UDP connections at the boundary of an IPv6 domain and relay them to the IPv4 domain.
- At the application layer, proxies can be run very well on dual-stack machines.

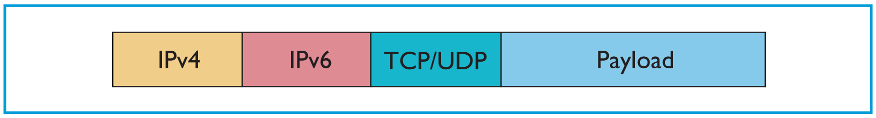

Figure 2. IPv6 in IPv4 tunnel. An IPv4 packet carries an IPv6 packet within its payload.

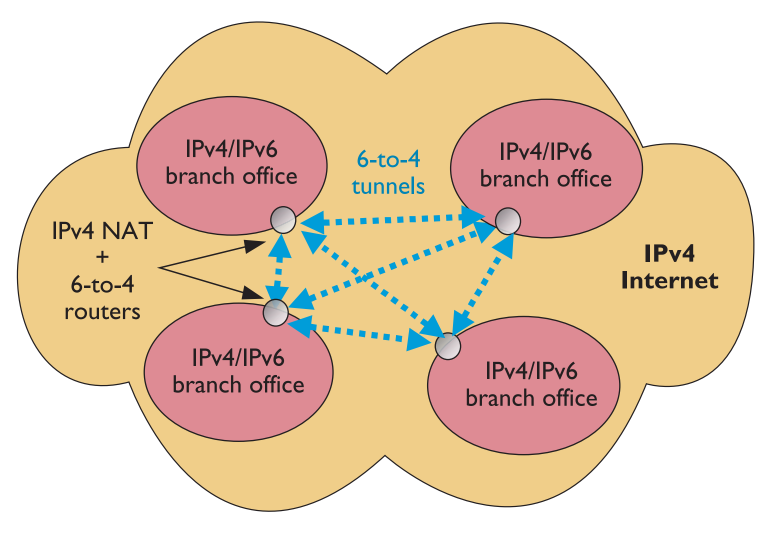

Figure 3. Building an extranet with "6-to-4." The mechanism reestablishes end-to-end IP connectivity disrupted by NATs.

Another approach to the communication between an IPv6 and an IPv4 node, the Dual Stack Transition Mechanism (DSTM), temporally allocates an IPv4 address to the IPv6 node.

Among all these mechanisms, 6to4 and application proxies will probably play the most important roles in the early phases of transitioning. DSTM is likely to be important in the later phase of the efforts. More information on these tools and on NGtrans activities is available at http://www.6bone.net/ngtrans.

What Next?

Many people ask when IPv6 will be widely deployed. This is difficult to answer with any precision because IPv6 will probably be adopted in waves by different business sectors and in different geographic locations.

With the growing availability of IPv6-ready operating systems and routers, ISPs are looking at offering IPv6 services. As mentioned earlier, Asia appears to be the first place IPv6 will take off. (A recent IPv6 forum summit in Osaka, Japan, gathered more than 600 people.) Now that a set of transition tools has been defined and implemented, the next important step is to get applications ported to IPv6. Sectors like the 3G wireless industry are likely to see serious IPv6 deployment in the years to come. The computer gaming and consumer embedded control industries are other areas to expect activities.

With all of the recent activity surrounding IPv6, it's no longer a question of "if," and less a question of "when," but rather "how" IPv6 will be deployed.

Alain Durand is a staff engineer at Sun Microsystems. He is co-chair of the IETF NGtrans working group and a member of the IETF IPv6 directorate and the IPv6 Forum technical directorate.