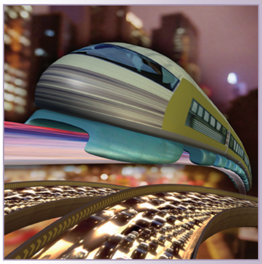

A Scenario-Based Review of IPv6 Transition Tools

Momentum for IPv6 transition is on the rise, and many transition tools and techniques are available to ISPs, enterprise networks, and unmanaged networks.

Michael Mackay, Christopher Edwards, and Martin Dunmore
Lancaster University

Tim Chown
University of Southampton

Graca Carvalho
Cisco Systems

The rationale for introducing IPv6 is beginning to sway many network providers in commercial and research environments.[1] Among IPv6's unique benefits over IPv4 are increased address space, simpler "plug and play," network security, and improved mobility support.[2] These benefits are interrelated. Increased address space lets networks globally address more and new types of devices, which removes the need for network address translation (NAT). In turn, this restores the Internet's original end-to-end principle, allowing host-to-host IPsec and novel services to run into networks previously "hidden" behind NAT boxes. Peer-to-peer applications and ad hoc and mobile networking will all benefit from IPv6 deployment.

Commercially supported IPv6 software, including operating systems and popular applications, is now available from vendors such as Cisco, Juniper, Microsoft, Apple, IBM, and Sun. In Europe, the European Commission sees IPv6 introduction as key to its 2005 e-Europe Action Plan. Several EU Information Society Technologies' (IST) IPv6 projects, including 6NET (www.6net.org) and Euro6IX (www.euro6ix.org), are examining how to deploy and support IPv6 in academic and commercial settings. There is also significant IPv6 work going on in the US and Asia.[3] We can consider the challenge of introducing IPv6 from two angles. First, when introducing a new "native" IPv6 network, we clearly must ensure that all components (network devices, host operating systems, applications, and so on) support the new protocol. Second, when introducing IPv6 to an existing IPv4 infrastructure, we must have transitioning mechanisms that enable the protocol's seamless introduction, minimizing any impact on existing network users.

1089-7801/03/$17.00©2003 IEEE

Glossary

- 3GPP: Third-Generation Partnership Project
- ALG: Application-layer gateways
- ATM: Asynchronous transfer mode
- BIA: Bump in the API
- BIS: Bump in the stack
- DSTM: Dual-stack transition mechanism
- GGSN: Gateway GPRS support node
- GPRS: General Packet Radio Service
- GSM: Global System for Mobile Communication
- ISATAP: Intra-site automatic tunnel addressing protocol
- IST: Information Society Technologies program
- MPLS: Multiprotocol label switching
- NAT: Network address translation
- NAT-PT: Network address translation–protocol translation
- NBMA: Nonbroadcast multiple access
- Ngtrans: Next-generation transition
- NRENs: National research and education networks
- PDP: Packet data protocol
- SIIT: Stateless IP/Internet control message protocol translation
- TEP: Tunnel end point
- TRT: Transport relay translator
- VLAN: Virtual local area network
- WLAN: Wireless local area network

The Internet Engineering Task Force recently established the IPv6 Operations (v6ops) working group (www.ietf.org/html.charters/v6ops-charter.html) to supersede the next-generation transition (ngtrans) working group. The v6ops group will be looking at four major IPv6 deployment scenarios: ISPs, cellular networks, enterprise networks (including universities), and unmanaged networks (such as those found in homes or small offices). Here, we offer an overview of the potential migration scenarios facing three of these groups (excluding cellular networks), and the transitioning mechanisms available to them.

IPv6 Research and 6NET

The IETF has been developing IPv6 protocols for more than six years. The ngtrans working group developed many tools aimed at assisting IPv6's introduction, offering a "transitioning toolbox" for operators migrating to IPv6. More recently, v6ops' work is focusing on the operational aspects of introducing IPv6 services into IPv4 environments, and on deploying green field IPv6 networks (new networks that primarily use IPv6).

In the EU's IST Fifth Framework Programme, two major projects are running trials of various IETF-proposed IPv6 protocols and mechanisms. In Euro6IX, large telecommunications companies and other commercial organizations are studying IPv6 deployment in the context of Internet exchange points. In 6NET, we're deploying an IPv6-only backbone to connect numerous national research and education networks (NRENs) at participating universities throughout Europe. All 6NET participants are potential IPv6 adopters, and they include all major European NRENs and other universities with long-standing IPv6 interests.

The 6NET project began in January 2002 and runs until December 2004. 6NET researchers have more than 1,000 human-months of effort dedicated to:

- network services such as IPv6 multicast, mobile IPv6, IP security protocol, quality of service, and domain name servers;
- IPv6 applications;
- network deployment, monitoring, and management; and
- IPv6 transition.

By offering IPv6 connectivity in existing academic environments (and beyond that, into the homes of staff and students), 6NET hopes to encourage the development of innovative services and applications that are both mobile and secure, offering new solutions to old problems.

Transitioning Tool Overview

The IETF's ngtrans activity developed a transitioning toolbox that operators can use to migrate to IPv6. Generally, the tools operate either within a site or between sites communicating across the Internet, and fall into one of two categories:

- Tunneling tools address IPv6-to-IPv6 communications in the initially IPv4-dominated Internet.
- Interoperation tools provide communication between IPv6 and IPv4, through either protocol-translation or dual-stack mechanisms.

The new v6ops activity is reviewing these tools in relation to the basic transitioning scenarios. In the process, some tools might be relegated to historic status, while others will require simplification if they are to move forward as RFCs for deployment.

Dual-Stack Tools

RFC 2893 introduced the dual-stack mechanism, in which the operating system of a host or router — an "IPv4–IPv6 node" — is equipped with both sets of protocol stacks (although in practice, the stacks share many elements).[4] Thus equipped, the node can send and receive both IPv4 and IPv6 packets. This is the simplest way for IPv4 and IPv6 to coexist. Dual-stack networks must support addressing, DNS, and other management mechanisms for both protocols across the network.

The dual-stack transition mechanism (DSTM) lets developers deploy dual-stack hosts in an IPv6-only network using a pool of IPv4 addresses.[5] The IPv4 address pool is dynamically assigned to dual-stack-enabled IPv6 hosts to allow communication with the IPv4 world as needed. DSTM uses a dynamic tunneling interface, which facilitates tunneled IPv4 communication to an explicit tunnel end point (TEP) inside the IPv6-only site; the TEP then forwards the IPv4 packets to the external network.

Tunneling Tools

Tunneling tools aim to simplify IPv6-to-IPv6 communication both within and between sites. Such tools will be very important in IPv6's early deployment period, as network operators can use intrasite tools to test IPv6 before full site migration and intersite tools to obtain connectivity to other IPv6-aware sites across the IPv4 Internet.

6to4. The 6to4 tool defines a mechanism that lets IPv6 sites communicate over the IPv4 Internet without using explicit tunnels (those that a network operator must manually configure).[6] 6to4 effectively treats the IPv4 Internet as a unicast point-to-point link layer, specifying an encapsulation mechanism for transmitting IPv6 packets over the Internet by assigning a unique IPv6 address prefix to any site with at least one globally unique IPv4 address. The mechanism constructs a 48-bit site prefix using the 2002::/16 6to4 prefix and the site's IPv4 address. Among its benefits, this technique introduces no new IPv4 routing table entries and only one new /16-length entry into IPv6's global routing table.

Tunnel broker. The tunnel broker automatically manages IPv6 tunnels and tunnel requests from isolated IPv6 sites on behalf of one or more dedicated servers.[7] It thus removes the management load for network administrators, who otherwise must configure and maintain each tunnel. The tunnel broker works best for isolated IPv6 sites and IPv4 Internet hosts who want to connect to an IPv6 network.

IPv6 over ATM and MPLS. Asynchronous transfer mode (ATM) and multiprotocol label switching (MPLS) offer similar approaches for transitioning to IPv6. Both use an overlay network model, in which the core network elements can handle encapsulated IPv6 packets with no knowledge of IPv6. Only the edge network devices need to be IPv6-aware. However, ATM usage is declining somewhat as network providers deploy the latest high-speed technology. In contrast, providers are currently using MPLS for various traffic-engineering purposes. Methods such as Cisco's IPv6 Provider Edge Router (6PE) encapsulation let providers deploy IPv6 services between edge devices on an MPLS core network. However, deploying MPLS — with all its associated complexity — solely to introduce IPv6 is impractical.

Intrasite automatic tunnel addressing protocol. ISATAP is designed to connect isolated IPv6 hosts and routers (nodes) within an IPv4 site.[8] It facilitates incremental IPv6 deployment by treating the site's IPv4 infrastructure as a nonbroadcast multiple access (NBMA) link layer. ISATAP uses a new IPv6 interface identifier format that enables automatic IPv6-in-IPv4 tunneling within the site, whether the site uses global or private IPv4 addresses.[9] The new interface identifier format can be used with both local and global unicast IPv6 prefixes to enable local and global IPv6 routing. ISATAP mechanisms do not impact routing table size, and they require no special IPv4 services.

Teredo. Teredo (formally known as Shipworm) proposes a mechanism that tunnels packets over user datagram protocol (UDP) to bring IPv6 connectivity to IPv6 nodes located behind IPv4 NATs.[10] To run the service, a network needs Teredo servers, which are stateless and manage only a fraction of the traffic between Teredo clients and the Teredo relays that act as IPv6 routers between the service and the native IPv6 Internet. Teredo will likely be used only as a last resort, where IPv4 NATs prevent other mechanisms from working.

Translation Tools

Neither dual-stack nor tunneling mechanisms work for communications between an IPv6-only node and an IPv4-only node. Such communication requires a translation mechanism at either the network, transport, or application layer.

Stateless IP/Internet control message protocol translation. SIIT specifies a key translation algorithm for enabling interoperation between IPv6-only and IPv4-only hosts.[11] In SIIT, temporarily assigned IPv4 addresses are used for IPv4-translated IPv6 addresses. Packets travel through a SIIT translator, which converts the packet headers between IPv4 and IPv6, and translates the header addresses between IPv4 on one side and IPv6 on the other.

Deploying MPLS — with all its associated complexity — solely to introduce IPv6 is impractical.

Network address translation–protocol translation. NAT-PT builds on the common IPv4 NAT device to provide an IPv4–IPv6 translation tool.[12] NAT-PT binds the internal network's IPv6 addresses with the external network's IPv4 addresses to transparently translate packets. As sessions are initiated, NAT-PT uses a pool of IPv4 addresses for dynamic assignment to IPv6 nodes. NAT-PT keeps state information on each session, and thus session packets must pass through the same NAT-PT device. For the actual header translation, NAT-PT relies on SIIT functionality. NAT-PT also supplies a range of application-layer gateways (ALGs), including DNS and FTP, for more complicated protocol translation involving embedded IPv4 addresses.

Bump in the stack/Bump in the API. BIS adopts a unique translation approach, moving it inside the individual hosts rather than translating in a centralized server.[13] All hosts translate between IPv4 and IPv6 internally by adding the necessary segments to their IP stack. BIS is an extreme extension of NAT-PT, in that IPv4 addresses are dynamically allocated to hosts from a pool. The BIA mechanism is similar in spirit to BIS, but it does not translate between IPv4 and IPv6 headers.[14] Instead, it inserts an API translator between the socket API and the host stack's TCP/IP modules, allowing translation to occur without the overhead of translating every packet's header.

Transport relay translator. TRT provides a transport-level translator that relays TCP and UDP connections at the border between IPv4 and IPv6 domains, acting as an intermediary between them.[15] Because no flow state is kept, network operators can run TRT on a single server or a group of servers for scalability. When one host tries to connect to another through TRT, two connections are established: one between the source and the TRT device, and one between the TRT device and the destination, each in its native protocol. Thereafter, TRT transparently relays packets between the source and destination, translating each packet as needed.

Socks64 transport relay. Socks64 is an IPv6–IPv4 gateway mechanism based on Socks (short for "sockets"), an IETF-approved proxy protocol for developing secure communications. Socks64 uses a dedicated Socks server for relaying flows between IPv4 and IPv6 hosts. This method simply extends Socks 5.0 to allow application-level IPv4–IPv6 translation. Basically, the hosts forward IPv6 packets to a Socks server, which translates the flow into IPv4, and vice versa. Socks64 thus offers a useful interoperation tool for sites already using Socks 5.0.

Application-layer gateways. ALGs are typically dual-stack devices that hosts can contact over IPv4 or IPv6, and that can fetch responses over IPv4 or IPv6. ALGs are not direct translation devices; instead, they act as relaying proxies at the application layer. A typical example of an ALG is a Web cache or proxy, or a simple mail-transfer protocol server.

Evaluating Available Transitioning Tools

Developers created each of the transitioning tools discussed earlier with a particular role in mind. It is essential therefore that each mechanism can be deemed as "fit for purpose" — that is, fit for a given role when applied in a particular network scenario. To help assess this, we defined the following criteria.

Scalability. Perhaps the most important consideration is how a given mechanism will scale. For example, NAT-PT can handle a few connections quite well, but — like IPv4 NAT — as the number of devices initiating connections grows, so too do the processing and state-maintenance loads, which can cause service degradation or failure.

Security. Developers should present a mechanism's security implications in its Internet draft or RFC. Some generic attack types, such as IP spoofing, IP-in-IP tunnels being allowed through firewalls, or relayed denial of service attacks, might be relevant to many mechanisms.

Performance. A transition mechanism's performance is quite closely tied with its scalability, but performance impact can be direct or indirect. Traffic encapsulation has a direct computational impact, for example, and it also affects packet sizes. Indirect effects include IPv4 performance degradation when a device also acts as an encapsulating end point for IPv6 traffic. Also, transitioning mechanisms (particularly translators) can often act as single points of failure, especially if stateful in nature (that is, for each session, there is a single device maintaining its state).

Functionality. In certain transitioning scenarios, some of IPv6's "new features" cannot be exploited. For example, SIIT cannot translate IPv4 options or IPv6 extension headers. Also, several mechanisms have difficulty translating multicast addresses (without using some form of bridge or gateway). There are other IPv4 "knock-on" implications when dual-stack services are deployed on a QoS-enabled backbone. For example, will IPv6 support impact the network's ability to support IPv4-based QoS requirements? And what of IPv6-based QoS requirements?

Host and router requirements. Some mechanisms place requirements on the host, such as the need for a specific configuration within the network layer. With ISATAP, for example, a host must be explicitly configured to use ISATAP tunneling, rather than, say, autoconfiguring an IPv6 address from observed IPv6 router advertisements.

IPv4 and IPv6 address requirements. Different mechanisms place different requirements on available IPv4 (and IPv6) addresses. DSTM, for example, calls on a pool of IPv4 addresses for the IPv4-in-IPv6 tunneling. In contrast, 6to4 requires only one global IPv4 address.

Application requirements. Applications might need to know which transition mechanism is being used. With ALGs, for example, users or system administrators must configure the application to use the given ALG (although the application could also discover the ALG automatically). For applications that embed IP addresses in the payload, translation techniques such as NAT-PT can break the application unless a specific ALG is added.

Ease of use. Transition tool configuration should be hidden from the application's end user; if IPv6 is successfully deployed, end users are unlikely to notice the change.

Ease of management. This property refers to both the deployment effort required and the effort for ongoing management of the transitioning network. Explicit tunnels will obviously have a greater management burden than automatically configured tunnels, such as 6to4 and tunnel brokers. Network managers also face some indirect implications, such as the potential for growing numbers of IP-in-IP tunnels to pass through firewalls, and the increasing use of end-to-end IPsec (also through firewalls). Technique interactions are an additional concern. For example, ISATAP and 6to4 can be used together for external and internal connectivity, but other methods do not integrate so well.

Scenario A: Internet Service Providers

If IPv6 is to be viable over the long-term, ISPs must be early adopters. There are two main classes of ISP service. The first provides IP services to enterprise networks. In this case, the network will transport IPv6 through its ISP to the wider IPv6-enabled Internet, and the site will manage the IPv6 deployment internally. The second class offers connectivity to single dial-up or broadband customers, who will expect the ISP to provide their connectivity and will want the IPv6 functionality to be as plug-and-play as possible.

Transitioning Tool Options

If an ISP decides to introduce native IPv6 into its core, it might install a new set of dedicated IPv6-only routers, perhaps creating a separate IPv6 infrastructure (which is less risky). However, assuming that vendors offer stable, hardware-enabled code, it's likely to be more cost-effective for ISPs to upgrade their core routers to dual-stack and route both IPv4 and IPv6 packets over the same hardware and links.

Once an ISP makes IPv6 connections and peerings to other IPv6-enabled ISPs, it can introduce dual-stack services for connecting customers, including (leased-line) enterprise customers and individual dial-up users. In the absence of dual-stack, an ISP could use manual or automatic tunneling, and perhaps dedicate one or more routers as open 6to4 relays so its customers can reach other ISPs' IPv6 sites through 6to4 tunneling. (These relays are open to abuse, however.)

For networks that are not IPv6-enabled, ISPs could offer transition support services through such mechanisms as a tunnel broker or a 6to4 router and relay. For customers with IPv6-only sites, ISPs might offer a NAT-PT translation device.

ISP Case Study: 3GPP

In the Third Generation Partnership Project (www.3gpp.org), which is standardizing, the principle transitioning components are the user equipment (a mobile handset device) the gateway GPRS support node (GGSN), and the IP multimedia subsystem. At some point during migration, the 3G service provider must deploy IPv6-enabled versions of each. These components — referred to as packet data protocol (PDP) contexts — can be linked via IPv4, IPv6, or point-to-point protocol. User equipment might have multiple PDP contexts of different types open to many GGSNs at a time. However, all PDP contexts in the multimedia subsystem must be IPv6, as mandated within 3GPP release 5.0.

The ISP environment is interesting because the operators' migration approaches will define, quite strictly, the extent of IPv6 services that their customers receive.

In this case, several transitioning scenarios are possible, such as connecting an IPv6 PDP over the IPv4 Internet and allowing IPv6 PDPs to connect to IPv4 end points. To do this, the 3G networks need site-scope tunneling and interoperation mechanisms at the network's edge (beyond the IP's multimedia subsystem) to operate on IPv6 sessions leaving the site. Initially, IPv6-over-IPv4 sessions might require either a tunneling mechanism, such as 6to4, or configured tunnel deployment if end-to-end IPv6 is unavailable. 3GPP is still debating which interoperation mechanism to recommend for IPv6-to-IPv4 communication over 3G networks. However, the clear choice is either NAT-PT or DSTM.

The 6NET Context

The 6NET project has deployed an IPv6-only network backbone spanning some 15 NREN presence points. The backbone is connected to each national IPv6 pilot or service network, and also has native IPv6 peerings with external networks such as Euro6IX via the UK6X exchange, and with Abilene (http://abilene.internet2.edu) via SURFnet (www.surfnet.nl/en/). The Géant network (www.dante.net/geant/), which interconnects all European NRENs, is currently IPv4-only, but developers are upgrading its backbone routers to support dual-stack operation; the network is expected to offer initial native IPv6 pilot services by mid-2003. In recent IPv6 trials, Géant set an Internet2 IPv6 land speed record.[16]

NRENs will next have to contend with how to offer IPv6 connectivity to university end sites. Germany's Deutsches Forschungsnetz (DFN; www.dfn.de) research network has deployed 6WiN, a separate IPv6 infrastructure, while France's network for technology, teaching, and research (Renater; www.renater.fr) supports IPv6 encapsulated in ATM and plans to offer a dual-stack core in the next generation. In Greece and the Czech Republic, networks have encapsulated IPv6 in MPLS using 6PE. In the UK, Janet (www.ja.net) currently offers interested universities connectivity via a tunneled service in advance of dual-stack operation, which will be launched later this year. We can expect all major European NRENs and their interconnecting network to support IPv6 in dual-stack networks by 2004. However, progressive adoption by end universities and regional metropolitan area networks (MANs) will mean that some tunneling is still required.

Summary

From a transitioning perspective, the ISP environment is interesting because the operators' migration approaches will define, quite strictly, the extent of IPv6 services that their customers receive. As such, the ISPs' (scalable) migration decisions have direct knock-on effects for customers. In the future, customers might require ISPs to offer value-added IPv6 services that not only have performance-based restrictions, but security and mobility considerations, as well (leveraging IPv6's end-to-end IPsec and mobile IPv6 features).

Scenario B: Enterprise Networks

Enterprise networks are likely to be early IPv6 adopters because many vital enterprise network concerns (such as mobility and security) are IPv6 drivers. A typical enterprise network might serve a large university campus or business. Such single- or multisite networks might contain a combination of intranets and extranets, and wireless, mobile, or other subnets.

Perhaps the most important factor for defining the site's migration approach is how many global IPv4 addresses it has. Sites with numerous addresses will likely adopt a dual-stack approach, as there will be less contention for resources. Sites with limited IPv4 address space might choose an IPv6-only internal network infrastructure. In this instance, to access external IPv4 sites, choosing DSTM rather than NAT-PT avoids the translation requirement, but requires configuration in the client hosts. Where both protocols coexist, IPv6 could run in parallel with IPv4+NAT (which uses private IPv4 addresses). Dual-stack deployment requires careful consideration of protocol use: applications could try to use both protocols (depending on the address records returned by DNS lookups) and might or might not fall back to one protocol when the other is unavailable.

Transitioning Tool Options

In most cases, sites will start by making small portions of their internal networks IPv6-capable, so they can test various tools and configurations. Depending on a site's performance and scalability requirements, providers can then extend the use of the technologies as they prove stable. ISATAP supports such an approach to introduction, letting developers overlay the IPv6 network on the existing IPv4 infrastructure. A site might also deploy an internal tunnel broker service for isolated (or remote) hosts wishing to use IPv6. A site can expect to receive a /48 prefix IPv6 address space from their ISP, and site developers can divide the space among various departments (perhaps with a /56 prefix), according to allocation policies.

Using IPv6 stateless autoconfiguration in the early stages will help lower an operator's administrative burden. The site can deploy IPv6 DNS services on existing name servers, using the Berkeley Internet Name Domain (BIND) DNS software, and then decide whether IPv6 devices will use the same or a separate namespace (lancs.ac.uk or ipv6.lancs.ac.uk, for example). An external IPv6 connection will be required early on; if the ISP cannot supply a native IPv6 link, the site could use 6to4 or configured tunnels to an IPv6 network. The organization firewall will need to pass Protocol 41 (identifying IPv6 in IPv4 packets) for IPv6-in-IPv4 tunnels, and operators should also increase security with additional IPv6-aware filters.

Enterprise sites should be able to adopt IPv6-aware applications as soon as they're available. In the dual-stack model, the same applications should be able to serve both IPv4 and IPv6. Once operators have established IPv6 services and critical applications, they can then think about deploying IPv6-only nodes and IPv6-only links.

As IPv6 becomes more commonplace, IPv6-only nodes will need to interoperate with IPv4-only servers (and even "legacy" devices, such as IPv4-only printers) in the host organization or on the Internet. Interoperation is available through various mechanisms, including translators, such as NAT-PT and SIIT; relays, such as TRT or Socks64; and dual-stack mechanisms, such as DSTM and various protocol-specific ALGs. The choice will depend on several factors, including access to IPv4 address space and other resources, the existing network infrastructure, cost, and the specific application domains.

The 6NET Context

Most European universities have enough IPv4 address space to use static, global addressing on their internal networks. For them, it makes sense to introduce IPv6 in dual-stack mode, while retaining the IPv4 infrastructure and using a technique such as ISATAP for internal IPv6 connectivity. Alternatively, they could use separate IPv6 routing infrastructures to inject IPv6 subnet router advertisements into existing IPv4 virtual LANs, enabling IPv6 connectivity on a 1:1 subnet mapping with the existing IPv4 infrastructure. This technique can be used in advance of robust IPv6 support in internal router or Layer 3 switching equipment.

In many Eastern European NRENs, however, IPv4 address space is not so widely available, and private addressing is more commonplace. It's likely that we'll see earlier adoption of an IPv6-only infrastructure in such sites. Elsewhere, IPv6-only networks are most likely to appear first in wireless LAN (WLAN) deployments to support personal digital assistants and other mobile devices.

Summary

Although it's impossible to cover all cases here, we can make some generalizations that make the task slightly more manageable. Large sites must emphasize scalability and performance when considering interoperation. It's likely that devices such as NAT-PT and DSTM will run into scalability issues, both in terms of available resources and required performance. Other important factors include security provisioning and mobility support. A combination of the available tools might be required to provide the necessary support.

When corporate intranets, public WLAN hotspots, and other networks use internal RFC-1918 addressing without supporting devices providing NAT-PT or DSTM, the Teredo method alone offers a realistic chance for IPv6 connectivity.

Scenario C: Unmanaged Networks

The v6ops group is drawing up detailed transitioning scenarios for "unmanaged" networks, comprising small- and medium-sized enterprises (SMEs) and home and mobile users. Users in this group might connect to their ISPs over various mediums, including broadband, dial-up, LAN, or more recently, WLAN. Clearly, deployment in this scenario could prove challenging (providing native IPv6 DSL, for example, is neither trivial nor cheap). However, the mass market might hold the most revenue potential (in novel IPv6 services into the home, for example, or new gaming features). Also, in this scenario, IPv6's acceptance might rest on mass-market penetration.

DSL routers with 6to4 functionality already exist, but to be most effective they need 6to4 relay support from ISPs. Also, the IETF has yet to finalize 6to4's IPv6 multicast support (which might be very useful in the home for broadcast and distribution applications).

This category also includes mobile users — people connecting to dual-stack devices in wireless public hotspots or using the Global System for Mobile Communication (GSM) or GPRS dial-up via a mobile phone, for example — but their usage patterns are similar to home users connecting via single PCs.

Transitioning Tool Options

Typically, this environment consists of one network segment with several hosts using IPv4 private addresses and a few global IPv4 addresses assigned to a gateway device that connects the network to an ISP. Most SME-type organizations don't use internal routers, while others use a router to connect internal network segments. In the home environment, the scenario is typically more simplistic.

Given this, the most significant components to consider here are the gateway device and the ISP access network, as they largely determine how the network connects to the Internet. Also, gateway devices might operate NAT devices or firewalls to give the network private address space and offer some protection from the Internet. In this case, operators must upgrade or replace gateway devices with others that are capable of dual-stack operation. This in turn requires support both for interoperation, which is likely to be NAT-PT (that is, an upgrade from basic NAT), and a tunneling mechanism such as 6to4 or a tunnel broker client. Any other internal routers also need a dual-stack upgrade. Once the gateway supports IPv6, operators can upgrade the network's clients and servers to dual stack.

With IPv6 support achieved, operators can test and install IPv6-aware applications and services and deploy IPv6-only hosts. Finally, they can remove private IPv4 addresses from the dual-stack systems, so that IPv6 is used for all internal communications.

The 6NET Context

In the academic context, it's important to support IPv6 access for mobile, roaming users, whether they're at home or traveling (dialing an ISP through GSM, for example, or connecting from another university's public WLAN hotspot). In such cases, the local ISP (the home, dial-up, or WLAN provider) might not support IPv6 connectivity. Here, the tunnel broker mechanism seems the most appropriate solution.

Summary

IPv6 connectivity options depend on the user's ISP support. An interesting question is how the connecting host can detect supporting mechanisms. Does the WLAN hotspot, for example, have an ISATAP router? Various mechanisms must be able to work together. If no native IPv6 service exists, operators can try falling back on methods such as ISATAP; if no ISATAP router exists, they can try a tunnel broker. If RFC 1918 addressing is used locally, the tunnel-broker option is out, and Teredo becomes the tool of last resort.

Future Work

Transitioning to IPv6 provides many interesting and challenging hurdles. Significant work has been carried out under the IETF umbrella toward standardizing tools that are applicable to different transition scenarios. As outlined here, we're beginning to get a firm understanding of what these different scenarios will be and what they will require, as well as which techniques are most useful to meet the requirements. The next step — deploying and evaluating these techniques in different environments to prove their suitability — is now underway, predominantly in Europe and Asia.

In its effort to drive IPv6 deployment in Europe, 6NET plans to produce transition cookbooks for end sites and ISPs at the end of this year and again in December 2004. These cookbooks will detail scenarios and include configuration examples and usage reports. The first versions are now available for download at www.6net.org.

References

1. S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," Internet Eng. Task Force RFC 2460, Dec. 1998; www.ietf.org/rfc/rfc2460.txt.
2. D. Lee et al., "The Next Generation of the Internet: Aspects of the Internet Protocol Version 6," *IEEE Network*, vol. 12, no. 1, Jan./Feb. 1998, pp. 28-33.
3. H. Esaki, A. Kato, and J.Murai, "R&D Activities and Testbed Operation in WIDE Project," *Proc. IEEE 2003 Symp. Applications and the Internet (SAINT) Workshops*, IEEE CS Press, 2003, pp.172-177.
4. R. Gilligan and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers," Internet Eng. Task Force RFC 2893, Aug. 2000; www.ietf.org/rfc/rfc2893.txt.
5. J. Bound et al., "Dual Stack Transition Mechanism (DSTM)," Internet draft, IETF, July, 2002; work in progress.
6. B. Carpenter and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds," Internet Eng. Task Force RFC 3056, Feb. 2001; www.ietf.org/rfc/rfc3056.txt.
7. A. Durand et al., "IPv6 Tunnel Broker," IETF RFC 3053, Jan. 2001; www.ietf.org/rfc/rfc3053.txt.
8. F. Templin and T. Gleeson, "ISATAP Transition Scenario for Enterprise / Managed Networks," Internet draft, Internet Eng. Task Force, June 2002; work in progress.
9. Y. Rekhter et al., "Address Allocation for Private Internets," IETF RFC 1918, Feb. 1996; www.ietf.org/rfc/rfc1918.txt.
10. C. Huitema, "Teredo: Tunneling IPv6 over UDP through NATs," Internet draft, Internet Eng. Task Force, Sept. 2002; work in progress.
11. E. Nordmark, "Stateless IP/ICMP Translation Algorithm (SIIT)," Internet Eng. Task Force RFC 2765, Feb. 2000; www.ietf.org/rfc/ rfc2765.txt.
12. G. Tsirtsis and P. Srisuresh, "Network Address Translation — Protocol Translation (NAT-PT)," Internet Eng. Task Force RFC 2766, Feb. 2000; www.ietf.org/rfc/rfc2766.txt.
13. K. Tsuchiya, H. Higuchi, and Y. Atarashi, "Dual-Stack Hosts Using the 'Bump-In-the-Stack' (BIS) Technique," Internet Eng. Task Force RFC 2767, Feb. 2000; www.ietf.org/rfc/rfc2767.txt.
14. S. Lee et al., "Dual-Stack Hosts Using 'Bump in the API' (BIA)," Internet Eng. Task Force RFC 3338, Oct. 2002; www.ietf.org/rfc/rfc3338.txt.
15. J. Hagino and K. Yamamoto, "An IPv6-to-IPv4 Transport Relay Translator," Internet Eng. Task Force RFC 3142, June 2001; www.ietf.org/rfc/rfc3142.txt.
16. T. Chown, "IPv6 Initiatives within the European National Research and Education Networks (NRENs)," *Proc. IEEE 2003 Symp. Applications and the Internet (SAINT) Workshops*, IEEE CS Press, 2003, pp.149-152.

Michael Mackay is a PhD candidate at Lancaster University. His primary research is in IPv6 transition, and he is currently developing a site-transitioning framework. He is also interested in IPv6 QoS, mobility, and network management. He received a BSc degree in computer science from Lancaster University. Contact him at m.mackay@comp.lancs.ac.uk.

Christopher Edwards is a lecturer in the Computing Department at Lancaster University, where he teaches courses in telecommunications, networking, and operating systems. His research interests include IPv6 transitioning and support for mobility and QoS. He received a PhD from Lancaster University. Contact him at ce@comp.lancs.ac.uk.

Martin Dunmore is a research associate at Lancaster University, where he works on the IST 6NET project. His research interests include IPv6, QoS, traffic engineering, wireless and mobile computing, ad hoc networking, and intelligent network systems. He received a BSc in computer science and a PhD in distributed computing from Lancaster University. Contact him at m.dunmore@comp.lancs.ac.uk.

Tim Chown lectures in the Department of Electronics and Computer Science at the University of Southampton. His primary research interests are in IPv6 and pervasive computing, and his active projects include 6NET and Euro6IX. He is also interested in wireless networking, security, messaging systems, IP multicast, and multimedia and peer-to-peer applications. He received a PhD in computer science from University of Southampton. Contact him at tjc@ecs.soton.ac.uk.

Graca Carvalho is a consulting engineer at Cisco Systems working for the Europe, Middle East, and Africa (EMEA) market. Since 1999, she has been responsible for EMEA's Advanced Internet Initiatives, associated with R&D initiatives in Europe, which included strengthening relationships between Cisco and the major European NRENs. She has been involved in several European projects including Wineglass and LION, and is one of the principle coordinators of the 6NET project. Contact her at gcarvalh@cisco.com.