SOLUTION SHEET

IPv6

SCALING THE MOBILE INTERNET

Dr. Paul Francis
Chief Scientist, Tahoe Networks

Dr. Francis has nearly 20 years experience in the design of internetworking technologies, during which time he has worked for MITRE, Bellcore, NTT, ACIRI, and Inktomi. His innovations include NAT (Network Address Translation), the multiple addresses method of site multihoming used in IPv6, and multicast shared trees used in PIM sparse mode. Paul has chaired two IETF working groups, and has published numerous RFCs, US and international patents, and research papers. As Chief Scientist, Paul leads Tahoe Networks' technology efforts for IPv6, security, advanced mobility, services and network optimization.

One in a series of technology white papers examining challenges that mobile operators face in building the Mobile Internet.

Problem Statement: The massive number of new users as well as non-phone devices entering service in mobile data networks presents a significant potential address scaling issue for Operators, well beyond the capabilities of IPv4. In addition to the volumes of users and devices, popular, emerging applications such as Peer-to-Peer and Push are poorly supported by certain scaling approaches like Network Address Translation (NAT) and could potentially hamper Operators' ability to roll out new services or differentiate their network offerings. To succeed in building a profitable, service-rich business model, Operators must consider scaling approaches like IPv6 and NAT and determine – from a cost and complexity standpoint – when and where to deploy these protocols in their network.

Tahoe Networks' Value: Tahoe Networks is building a rich set of infrastructure solutions to help Mobile Operators to deploy mobile data services at the Mobile Internet Edge™. Tahoe Networks supports IPv6, NAT, protocol translation (NAT-PT), and automatic and configured tunnels to lower the expense and risk associated with the transition to new scaling approaches.

TAHOE NETWORKS™

IPv6 IN PERSPECTIVE

With some predicting the number of Mobile Internet users exceeding 1.2 billion before the end of the decade, there is concern around the world about the ability for each of these devices to have a unique, always reachable IP address. The deployment and widespread use of IPv6 is popularly understood as a long-term solution for providing globally unique addresses to all of the mobile devices on the planet. Success in deploying IPv6, though, presents significant challenges to Mobile Operators without well-planned and executed strategies.

The basic design work for IPv6 was done in 1991, in the "pre-Web" stage of the Internet. In those days, IP addresses were the only reliable permanent "names" around. Domain names existed, but they weren't uniformly available and were often unreliable. All Internet computers had global IP addresses, and these did not change often. Firewalls, which were only starting to be commercially available, controlled access mainly using IP address. The majority of computers could still "ping" each other (send an IP packet and get a reply). IPv6 was designed to preserve this world.

Today, however, IP plays a much more restricted role. Few users ever type an IP address, using domain names embedded in URLs and email instead. Subscriber authentication is done with user names, not IP addresses, and firewalls increasingly use secure logins and IPsec, not IP addresses, for access control. Indeed for client-server applications like Web, email, and instant messaging, the IP address of the client may be modified in transit by NAT boxes.

Many Internet functions and applications operate today without the permanent, globally unique, end-to-end addresses of 1991. But not all do. Peer-to-peer and push applications are awkward and inefficient without globally unique addresses. Furthermore, management of a large NAT'd network is more difficult than a non-NAT'd network.

Untangling the Complexity, Mitigating Risk

Deploying an all-IPv6 network can be costly. Ultimately it requires modifying applications, new IP stacks in all end-user devices, new software and hardware in routers and switches, new software in network management system such as DHCP servers and SNMP managers, and extensive training for network operators.

All of these changes and upgrades cannot happen overnight. Fortunately they do not have to. There is no IPv4 address exhaustion deadline like there was with the Y2K bug. Rather, pressure will build up over time as IPv4 addresses get harder to obtain and more NAT is used. Rather than try to have IPv6 in place everywhere by a certain date, Operators must selectively choose when and where to apply IPv6 so as to relieve the pressure where it builds up most.

TAHOE NETWORKS™

IPv6 DEPLOYMENT CHOICES

Specifically, Operators must choose:

1. Which applications and services to run over IPv6.
2. Which infrastructure components must transition to IPv6 to support these services.
3. Which transition and translation tools to use for the transition.

Applications and Services Drive Deployment Choices

Client-server applications like Web, email, WAP, and streaming media do not require IPv6. These applications have worked for years through NAT and can safely continue to do so for the time being. While emerging push and peer-to-peer applications can operate through NAT, they do so less reliably and less efficiently than with IPv6. Operators should focus their early IPv6 investment on these applications, both because they operate better over IPv6, and because they are new applications and so fewer legacy issues to content with.

Upgrade Only End-user Devices, Edge Routers, and Application Servers

To support push and peer-to-peer services, end-user devices must be upgraded to dual stack, with client-server applications running over the IPv4 stack, and push and peer-to-peer applications running over the IPv6 stack. Push and Peer-to-Peer servers must also run IPv6. Within the infrastructure, only the edge routers and potentially the DNS servers need be upgraded to dual stack. The edge routers must be upgraded because they assign IPv6 prefixes to end devices, and forward IPv6 packets. The edge router management systems can continue to run over IPv4, but must of course be able to manage the edge router IPv6 functionality. During these upgrades, the Operator must utilize the appropriate transition tools, choosing among the seven tunneling mechanisms, six different IPv4-related IPv6 address types, and six translation mechanisms defined by the IETF.

Other routers do not need to be upgraded because the edge routers can tunnel IPv6 over IPv4, for instance using the "6to4" automatic tunneling mechanism. Infrastructure systems such as authorization, authentication, and billing do not need to be upgraded because these rely on identifiers other than IP addresses, for instance MS-ISDN numbers or Network Access Identifiers (NAI). Contrary to common belief, it is neither necessary nor appropriate to provision or "ship" most user devices with embedded IPv6 addresses.

Address and Protocol Translation

Even if a given Operator deploys IPv6 for a given application, it cannot predict how quickly the rest of the world will adopt IPv6. Inevitably the Operator will have to do NAT-PT (NAT-Protocol Translation) at its edge routers, both to translate its users' IPv4 packets into IPv6 and vice versa. And until both the Operator and the rest of the world move completely to IPv6, the Operator will continue to do NAT at its edge routers.

TAHOE NETWORKS™

IPv6 SOLUTIONS

New Capabilities, New Threats

Full end-to-end addressability enables new classes of applications like Push and Peer-to-Peer, but it also exposes new threats. Every new Peer-to-Peer application creates new opportunities for spam, hacking, and DoS (Denial-of-Service). The Operator and users must have full control over what traffic is transmitted over-the-air to the device. The edge router must be able to filter traffic on a per-user, per-application, and volume basis. There must be an interface to the edge router that allows cooperating applications and management systems to punch holes in the edge router for specific "flows". Where appropriate, the user or user device must be able to reject any flow request from outside. The edge router must be able to rate-limit the number of such requests. All of these capabilities must be in place before any Peer-to-peer or Push applications are turned on.

Deploying IPv6 at the Mobile Internet Edge™

The optimal place in the mobile network for Operators to introduce IPv6 with the greatest level of flexibility, control, and cost efficiency is at the Mobile Internet Edge (MIE). The MIE is the intersection of the radio network and the IP network, including the Internet, Intranet, and other private networks (e.g., "walled garden").

Tahoe Networks is building a new category of data networking infrastructure solutions for the MIE that helps Mobile Operators lower the expense and risk associated with scaling their networks through supporting targeted and controlled deployment of IPv6. The Tahoe Networks' solution provides unsurpassed scalability and reliability, reusable building blocks to rapidly rollout new services, and advanced accounting services that leverage the unique characteristics of the mobile network to support subscriber identity and preferences.