IPv6 at Virginia Tech

Operational experiences from a large-scale production IPv6 deployment

Carl Harris
Chief Technology Officer
Virginia Tech IT

Timeline

- 1997 – 6Bone experimentation between VT Department of Electrical Engineering and IT division
- 1998 – VT has Early Field Trial IPv6 firmware running on a Cisco router; handful of subnets in the information systems building
  - VT was first U.S. site to do native IPv6 over National Science Foundation's vBNS network.
- 2001 – Microsoft Research releases IPv6 add-on support for Windows XP
- 2003 – Mac OS X 10.3 (Panther) includes full support for IPv6

Timeline

- 2004 – Started executing the *Turn it on and fix whatever breaks* strategy.
  - Parallel IPv4 and IPv6 routers (separate hardware)
  - About 20 campus buildings
- 2006 – Native IPv6 routing on all subnets in VT's primary data center
- 2009 – Google apps via IPv6; search, Gmail, YouTube, etc.
- 2010 –IPv6 running on VT's primary core backbone; parallel routing infrastructure removed

Current Status

- Tens of thousands of network clients on our campus using native IPv6 daily for real applications
  - As it should be, most network users don't know or care – "it just works"
  - Many VT applications are IPv6-enabled
  - Google apps especially significant – virtually all traffic between Virginia Tech and google.com is IPv6
  - Lots of systems administration using SSH over IPv6
    - our large-scale virtualization environment is IPv6-only for management access

Current Status

- Vast majority of hosts are "dual stack"
  - Sufficient IPv4 addresses to meet projected needs, so not yet motivated for IPv6-only deployments
  - Windows, Mac OS X, Linux and most other UNIX derivatives have dual-stack support enabled out-of-the-box
  - More work needed on approaches to allow IPv6-only hosts to talk to IPv4-only services

Current Status

- Native IPv6 connectivity to the Internet at large
  - via Internet2 and National LambdaRail networks
  - our regional networking entity working on peering agreements for native IPv6 with commercial providers

Browser Behavior

- Virtually all shipping browsers will utilize an IPv6 network layer in preference to IPv4, if available.
  - Underlying this behavior are the facilities of the socket API
- Basic idea:
  - If these conditions are met:
    - client host has a global IPv6 address
    - target server (the host name in the URL) has a AAAA resource record in DNS (i.e. the name resolves to an IPv6 address)
  - Then attempt to connect to the target via IPv6
    - fallback to IPv4 on ICMP unreachable or connection timeout

Common Resolvable Issues

- IPv6 “islands”
- Router advertisements from misconfigured hosts
  - a.k.a. “Rogue RAs”
- Unexpected tunneling

IPv6 Islands

- Commonly experienced during the initial rollout of IPv6.
- Easy to omit IPv6 networks from the routing protocol process.
  - If no one is really using IPv6, the problem goes unreported.
- The basic problem is a network with disconnected subgraphs, and is easily resolved
  - just fix the routing configuration
- Because of the behavior of the browser (and more generally TCP-based applications) the reported symptom usually isn't "can't connect" but "slow connection"
- Helpful to do troubleshooting on IPv6-only hosts
  - easy to get fooled by a fully functional IPv4 layer

Rogue RAs

- A misconfigured host can send router advertisements on a link layer network that identify the host as a first-hop router
  - Windows Internet Connection Sharing option
- Same kinds of issues introduced by rogue DHCP servers.
  - broken connectivity
  - inappropriate addressing/routing
- Especially troublesome on large, flat wireless LAN networks
  - larger number of potentially misconfigured hosts and larger impact from a single host

Rogue RAs

- Symptoms
  - slow connections (see also “unexpected tunneling”)
  - no connection
- Mitigation strategies
  - RA priority – assign a non-default priority to legitimate RAs
  - Block inbound RAs and DHCP6 from untrusted ports
  - “RA Guard” feature
    - akin to DHCP Snooping feature
  - Potential solution: Secure Neighbor Discovery (SEND)

Unexpected Tunneling

- Some IPv6 capable hosts will resort to automatic (transparent) 6-to-4 tunneling if no first hop IPv6 router is available
  - in most cases, there's a knob to turn to enable, but Windows has been an exception in certain configurations
  - "automatic" uses IPv4 anycast to locate the "nearest" available 6-to-4 relay
    - Where is that?
- Symptoms:
  - very long round trip times – i.e. IPv6 works, but very slowly
  - host has only one global IPv6 address and it starts with 2002::/16

Unexpected Tunneling

- Mitigation:
  - Don’t put AAAA records for services into DNS until your client networks are fully IPv6 enabled
  - Don’t enable automatic 6-to-4 on client hosts unless you need it
  - Make sure you have a local 6-to-4 relay
    - i.e. know what “nearest” means

Outstanding Issues

- VT's production web load balancing infrastructure is not IPv6 enabled
  - Workarounds with some dedicated solutions
  - Need a significant hardware investment to replace, but current investment still has some time on its lifecycle
- Wireless LAN solutions for IPv6 are "not quite there yet"
  - VT peaks at 9,000 current wireless clients, daily
  - Existing solutions support seamless "roaming" for IPv4 only
- Want/need better network management controls for IPv6 in network hardware
  - e.g. rogue router (RA) suppression

Outstanding Issues

- Still need better tooling for managing and monitoring an IPv6 topology using IPv6.
  - Key to proactive trouble resolution
- Very few network-based security products are IPv6 aware
  - however, ominous “security concerns” for IPv6 are just FUD
  - most host-based approaches admit IPv6 solutions

Larger Issue

- Networking equipment and software vendors slow to roll out IPv6 solutions
  - Feature parody, not feature parity
  - IPv6 support != ping + traceroute
  - Still seeing new products appearing with IPv4-only architectures
  - Seeing substantial IPv6 advances in products designed for China, Japan, and other Asian-Pacific countries where IPv4 address space is extremely limited

Larger Issue

- .edu customers in U.S. cannot alone create enough demand to drive IPv6 technology development
- Some service providers beginning to step up deployment timelines
  - e.g. Comcast
- Need significant IPv6 deployments in Fed networks to help drive industry.
- The time window for "wait and see" strategies is quickly closing.