



IPv6 Optimized Networks and Support for IPv4 Legacy Applications

Yanick Pouffary
IPv6 Forum Fellow

David Green
SRI International



NAv6TF/ARIN XV IPv6 Conference

Orlando, Florida

April 17 – 21, 2005





Agenda

- V6 Dominant Networks
- Approaches to backwards compatibility
 - DSTM
 - Configured Tunnels
 - Application Gateways
 - Translation
- Conclusions
- Q&A



Problem Statement

- IPv6 networks that continue to support IPv4 may not achieve many of the netcentric benefits of IPv6 if backwards compatibility limits IPv6 because of poor transition mechanism architectural choices



IPv6 Dominant Networks

- Network has majority of network traffic in IPv6 format
- Critical mass of applications pushed to IPv6
- IPv4 services maintained for backwards compatibility with legacy applications/systems



Optimizing Networks for IPv6

IPv6 Only Cores

- Some networks will want to quit supporting IPv4 sooner rather than later
 - Supporting 2 protocols has an inherently higher cost
 - IPv4 has limited scalability due to security and stateful configuration models
 - New services like MIPv6, NEMO are IPv6 only – a network application built on this may not be IPv4 compatible at all
 - IPv6 is a better protocol for end-2-end / Peer-to-Peer applications
 - New end-2-end wireless services may be born IPv6 only due to the lack of v4 address space
- If a network turns off IPv4 support in its routers, or is born without IPv4 support in its routers, IPv4 services can still be accessed through transition mechanism at the network edges




Approaches to Backwards Compatibility

- Four mechanisms:
 - DSTM & Tunnel Broker
 - Configured Tunnels
 - Application Gateways
 - Translation
- All mechanisms sit at the edge of a network or as software on a host and act as a gateway to IPv6 domains



DSTM

- Dual Stack Transition Mechanism (DSTM) defined IETF draft-bound-dstm-exp
- Allows “dual stack” node on IPv6 only network to talk to an IPv4 only node on an IPv4 only network by tunneling from a dual stacked node to an IPv4 network
- Uses IPv4-over-IPv6 tunnels to carry IPv4 through IPv6 dominant backbone network to a tunnel endpoint at a DSTM router/translation gateway
- Can be added to a tunnel broker to automate the process and enforce security policy



DSTM Server
or DHCPv6 with DSTM enhancement

**IPv6 Dominant
Global Backbone**

**IPv4 Only
Net**

IPv4 Only
Host

Dual Stack Host +
DSTM Client Software

**IPv6 Only
Net**

IPv6 Only
Router

DSTM
Router





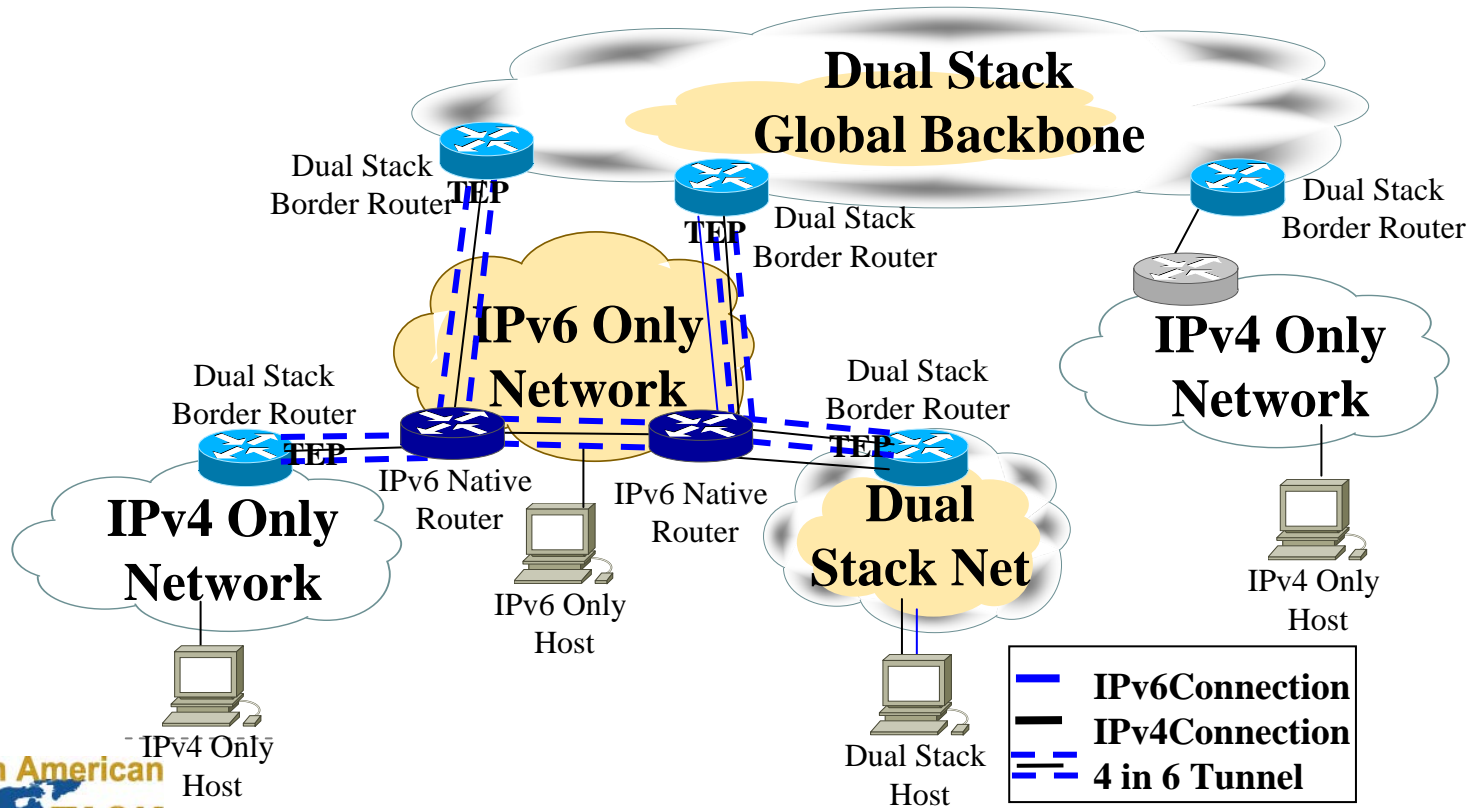
Mechanism Review: DSTM

- Performance
 - Tunnel at DSTM host on IPv6 domain allows for high scalability of IPv6-end TEP
 - Pays tunneling overhead of at least 40 bytes per IPv4 packet due to V4 in V6 overhead
- Scalability
 - Since interoperability is at the edge, this is highly scalable on V6 network
 - Native IPv6 address prefix allows global routing scalability
- Security
 - Good, but tunnel brokered DSTM adds even better security via Authorization, Authentication, Accounting (AAA) to determine who can set up the tunnel
 - Can use native IPv6 IPSEC only up to IPv4 tunnel TEP
 - Can use native IPv4 IPSEC
- Cost/Complexity
 - Can be offered as an enterprise service with tunnel broker & router in one box
 - Requires client software on hosts, but this can be distributed via DSTM tunnel broker
 - To reduce cost/complexity can use a DHCPv6 server, tunnel broker, or static assigned IPv4 addresses in place of DSTM address server on IPv6 domain
 - Administrators must set up DSTM, then tunnels can be set up and torn down automatically as needed



Configured IPv4 in IPv6 Tunnels

- Tunneling encapsulated IPv4 in IPv6 packets (IPv6 next header ID = 4)
- Most modern dual-stacked routers can act as tunnel endpoint (TEP)
- Hand configured mechanism – useful for static networks
- Could be incorporated into a tunnel brokering mechanism to automate setup






Mechanism Rev Mechanism Review: Configured Tunnels

- Performance
 - Static
 - No auto-reconfiguration – only manual
 - Pays tunneling overhead of at least 40 bytes per IPv4 packet due to V4 in V6 overhead
- Scalability
 - Administrators must set up each tunnel endpoint manually and this makes the labor requirements difficult to scale to large numbers of tunnels
- Security
 - Can use IPSEC
- Cost/Complexity
 - Labor intensive to set up
 - Labor intensive to re-establish in order to reconfigure networks, administrator must re-enter all tunnels TEPs if network topology or addressing changes
 - Can add a tunnel broker to automate part of tunnel setup and maintenance



Application Gateway

- A dual stacked application server that sits on a dual-stacked network and runs as a gateway system between v4 & v6 networks
- Ideal for client-server architectures where clients exchange information without P2P connections (Examples: E-mail, Weblogs)
- A legacy v4 client program can connect to the server and send/receive information, a v6 client program on an IPv6-only network can connect to the same server
- Where proxy servers exist in an architecture, they can be dual stacked and turned into an ALG. (Example: SOCKS HTTP caching proxy)
- Application gateways are application specific, not general translators





Mechanism Review

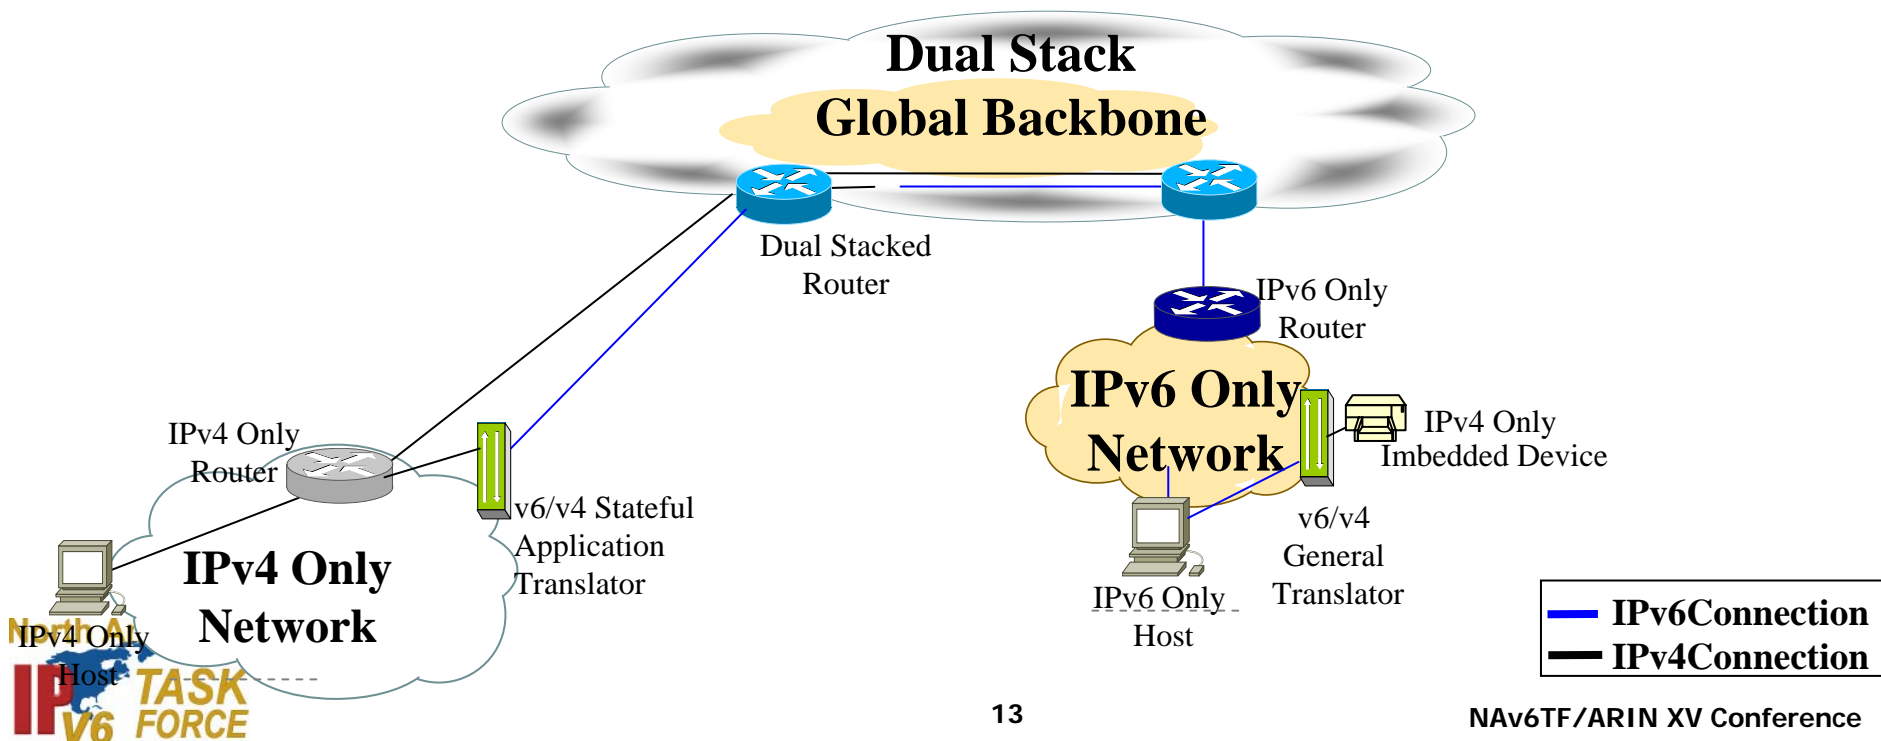
Application Gateway

- Performance
 - Depends heavily on design: Proxy may be a performance bottleneck, simple application server like e-mail can handle high performance
 - Proxy as an core enterprise service often requires great memory and processor resources – deploy on network edge
 - Often applications have client-server gateways/fusion points at operation centers and these servers can make natural ALG deployment points
- Scalability
 - If proxy interoperability is at the network edges, this is highly scalable on V6 network
 - Separate application gateway needed for each possible application
- Security
 - Good, Can be excellent as gateway or proxy site can be firewalled on both v4 and v6 sides
- Cost/Complexity
 - Cost can be cheap for simple dual stacked networked servers, especially applications that already employ gateways between domains or between units
 - High cost of complex application gateways was one of the drivers of moving from the NCP model to E2E IPv4 in 1980s. Avoid using this mechanism as the general solution for all interoperability



Translators

- Many Mechanisms:
 - SIIT, NAT-PT, BIA, BIS (Stateless)
 - TRT, Socks (Stateful)
- Strip IPv4 header off datagram and replace with IPv6 header (and vice-versa)
- Need translators modified for IPv6 native routing prefix for IPv6 dominant networks
- Translation breaks end-to-end model and will break IPv6 end-to-end IPSEC
- Best reserved for adding IPv6 capability to imbedded devices that cannot use ANY other transition mechanism. Ex: Legacy webcams, sensors, printers, storage arrays





Mechanism Review

Translators

- Performance
 - Single point of failure?
 - Interoperability of last resort
- Scalability
 - Service must be located on local IPv6 domain and rely on IPv4 for global reach
 - Multiple IPv6 sources can share a single IPv4 address
- Security
 - Blocks most IPSEC – BAD!
 - Single point of failure
- Cost/Complexity



Architecture For IPv6 Dominance

- Mechanisms in order of preference:
 - DSTM (With or without Tunnel Broker)
 - Configured Tunnels
 - Application Gateways (Use for Client-Server Apps)
 - Translation (Last resort – Breaks E2E Model and IPSEC)



Conclusions

- Without optimizing networks for IPv6 we lose many of IPv6's benefits
- IPv6-only networks can still service traffic from dual stacked hosts through the use of transition mechanisms
- High IPv4 operations and maintenance (O&M) cost may be a large factor in pushing to IPv6 dominance & IPv6 only networks
- Transition mechanisms for IPv6 dominant networks are at a lower state of maturity than current transition mechanisms for v4 dominant networks
 - A focused engineering effort can quickly prepare them for deployment to create v6 optimized networks