DREN IPv6 Implementation Experiences and Challenges

Australian 2006 IPv6 Summit
Dec 2006
Canberra ACT Australia

Ron Broersma
DREN Chief Engineer
High Performance Computing Modernization Program
ron@spawar.navy.mil

Background

- DREN …
  - is DoD's ISP for the RDT&E community
  - also serves as the DoD IPv6 "pilot" network
  - operates 2 IPv6 wide area networks (testbed, production)

Some History

- 2001
  - January - May: DREN builds the DREN IPv6 testbed
- 2003
  - June: DoD CIO sets goal to transition all DoD and Service inter and intra networking by FY '08
  - July: DREN chosen at the DoD IPv6 "pilot"
  - August: HPCMP Director directs HPC Centers to transition to dual-stack infrastructure
- 2004
  - DoD makes plans and organizes. DREN just does it.
- 2005
  - March: DoD IPv6 Transition Plan signed out
  - Services working on their own transition plans

DREN IPv6 philosophy

- Push the "I believe" button, and turn on IPv6 everywhere to see what works (and what doesn't)
- Do it in a production environment
  - can get away with this in an R&D environment, but not on operational networks.
- Go native. (no tunnels)
- Even if the world doesn't convert for years, R&D environments need it now.
- Figure out how to deploy IPv6 to the rest of DoD in the future.

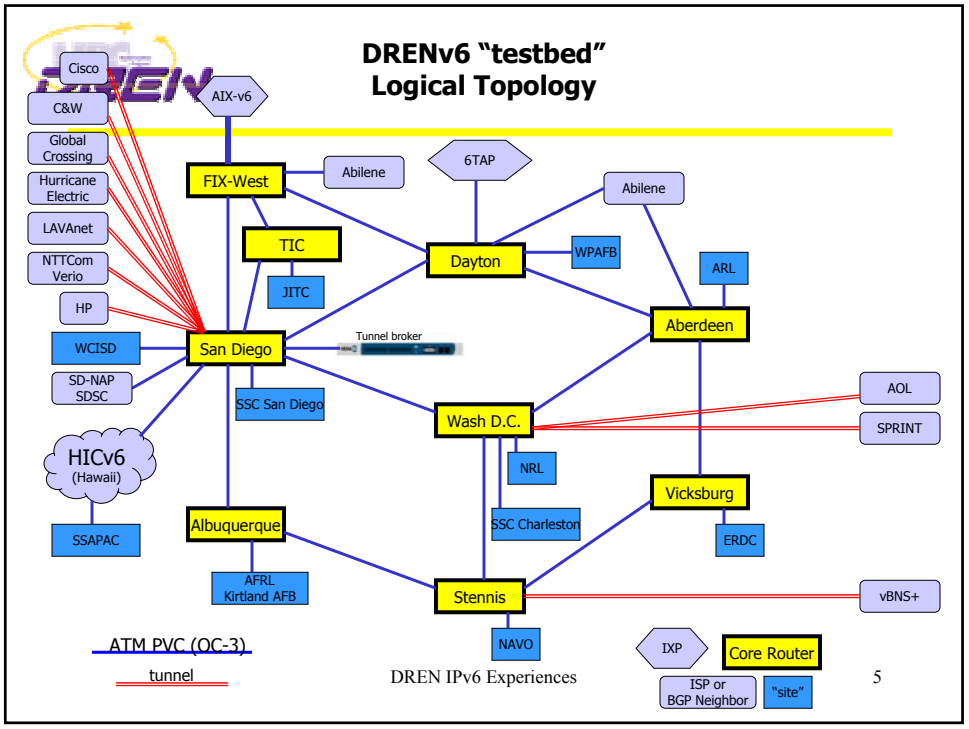

DREN "production" network

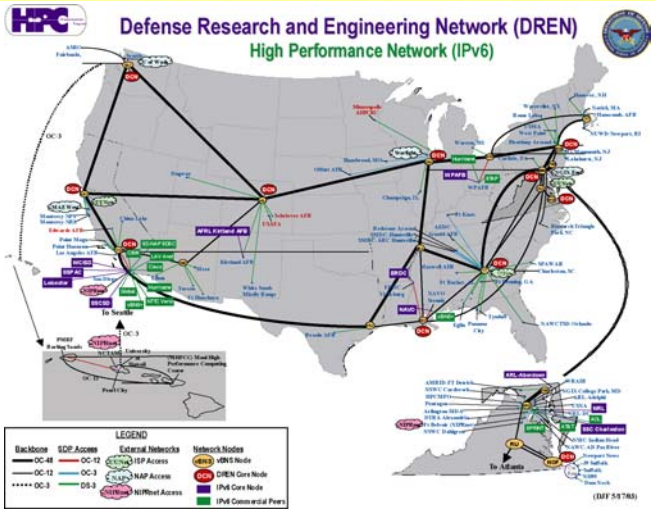

Original Goals

DREN IPv6 Pilot FY2003 Goals:

1. IPv6 enabled WAN infrastructure: border routers (called SDPs), the Verizon-provided backbone, and the Network Operations Center (NOC). Complete
2. Security and Performance as good as existing IPv4-only network. Complete
3. Facilitate IPv6 deployment into HPCMP funded sites' infrastructures. Complete
4. IPv6 enabled:
   - HPCMP funded sites' infrastructures. Mostly complete
   - HPCMP provided applications. Complete
   - COTS applications at HPCMP sites. Ongoing
   - Selected user application candidates. Outside scope
5. Provide equipment feedback, lessons learned, via web and via briefings. Complete

Overall difficulty

- Easy parts
  - Dual-stacking the nets (WANs, LANs)
  - Enabling IPv6 functionality in modern operating systems
  - Establishing basic IPv6 services (DNS, SMTP, NTP)
  - Enabling IPv6 in some commodity services (HTTP)
- A little more challenging
  - Getting the address plan right
  - Operating and debugging a dual stack environment
  - Multicast (but easier than IPv4)
- Hard parts
  - Creating the security infrastructure (firewalls, IDS, proxys, IDP/IPS, VPNs, ACLs)
  - Working around missing or broken functionality
  - DHCP
  - Creating incentives to upgrade and try IPv6
  - Getting the vendors to fix bugs or incorporate necessary features
    - Not enough market pressure, so other activities take priority

DREN sites status

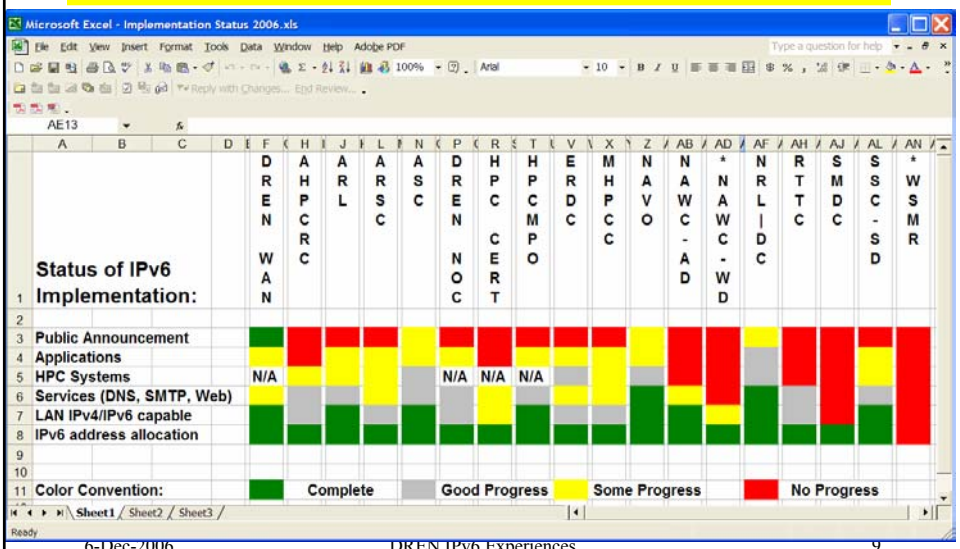

What did it cost?

- Regarding site purchases to enable IPv6 transition
  - Hardware: remarkably few purchases were necessary
    - Only 2 of the Pilot sites had to buy a router, which was scheduled to be replaced anyway (@ $25,000 and $8,200 each)
    - Some sites with Cisco routers chose to replace their Supervisor Blade with newer ones that did IPv6 in hardware rather than software (avg =$25,000). At a site with lesser bandwidth requirements, the old Blades would have been OK as is
    - It was fairly typical to expand the memory on routers, at a cost of $500 to $2,000 per router
    - No site had to replace any computers for the IPv6 transition
  - System Software:
    - Upgrades from Windows 2000 to Windows XP operating system and from Windows 2000 server to 2003 server were common
    - All other necessary software upgrades were covered under maintenance contracts at no additional cost

Planet DREN

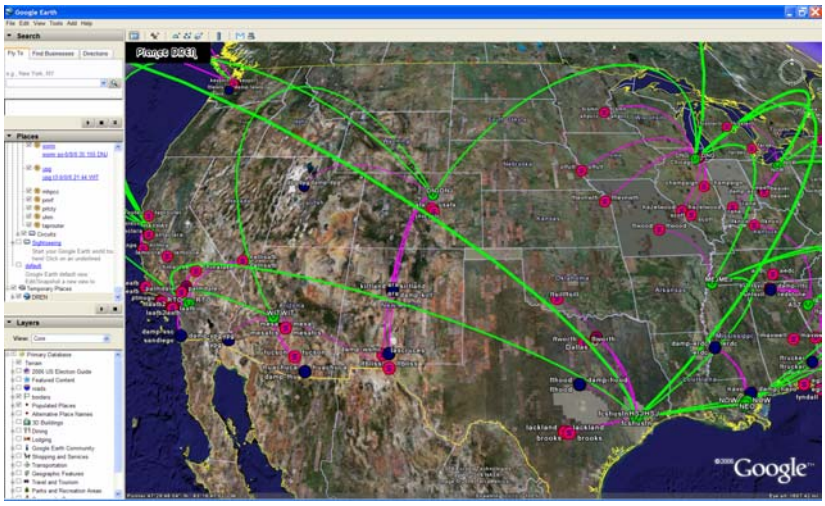

Performance measurements (production environment)

- Monitoring TCP performance between some high-end sites.
  - Using nuttcp, 9K MTU, Linux 2.4.26-web100 kernel
- Observations
  - RTT nearly identical between v4 and v6
  - TCP jumbo between ARL and ASC fails.
  - One or more paths demonstrated near line rate performance for both v4 and v6
  - In some cases, v4 appeared more robust. Reasons unknown.
- See http://www.wcisd.hpc.mil/~phil/ipv6

IPv4/IPv6 Performance Comparison

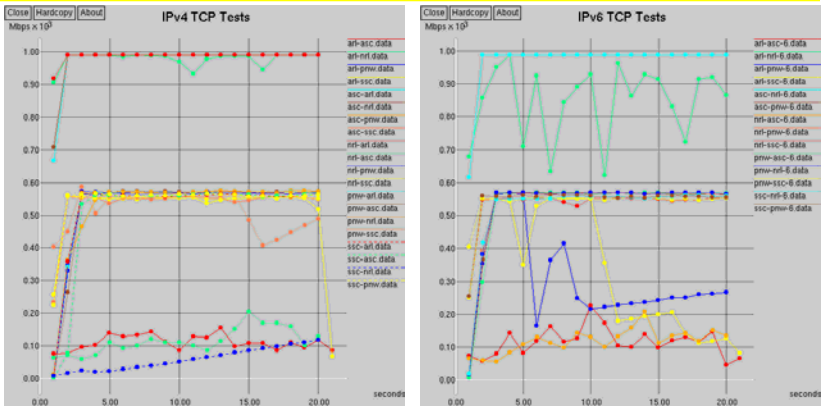

The above graphs show TCP throughput second by second for the 20 second tests for IPv4 and IPv6. Colors may not be the same between the windows because some IPv6 tests are missing (due to filter problems). The first second or two are usually TCP slow start followed by equilibrium.

The 1 Gbps and OC12 line rate tests stand out. Also clear from these graphs is the greater stability or robustness of IPv4 over IPv6 on some paths. The reason(s) for this are TBD. It could be from the Linux IPv6 implementation, or from hardware along the path.

IPv6 Security Review

- Independent security review performed by SAIC for DREN
  - Publicly available
- Some of the conclusions:
  - protocol is no less secure than v4
  - multicast is still spoofable
  - mobility is scary
  - ND – spoofable, but no exploits found yet
  - Windows – ack's things twice in all v6 TCP streams???
  - router renumbering – can spoof – possible DoS
  - landv6 attack works, but doesn't crash machine

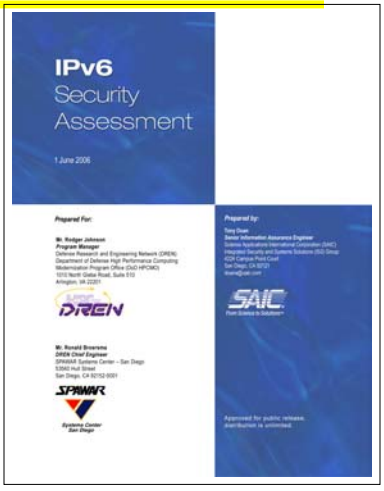

Firewall workaround

- 2003
  - Very few firewalls available
  - Limited functionality
  - Some sites used their own software-based firewalls
- 2005
  - Several vendors now offering IPv6-capable firewalls
  - Mixed results in using them
    - Small sites with limited requirements may be OK
    - HPC Centers need to spread traffic across several boxes because no one box can meet all our requirements

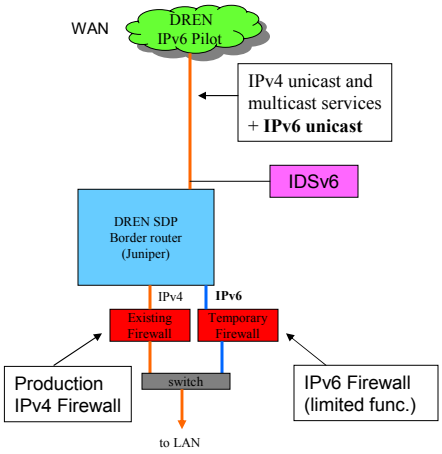

Security layer workaround

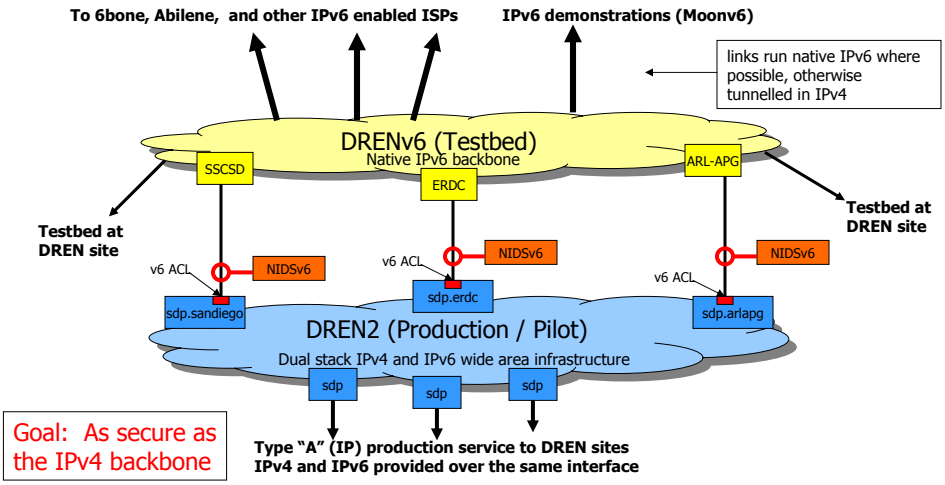

IPv6 Multicast

- Focus: get DREN backbones fully ipv6-multicast enabled.
- Status (work in progress)
  - Testbed – fully operational
    - PIMv2, MLDv2, SSM, ASM, static RP, Embedded-RP
    - Peering with m6bone
  - Production – operational
    - routers all upgraded to JunOS 7.2 or later
    - PIMv2, MLDv2, SSM, ASM, some Embedded-RP
  - Beacon – operational (dbeacon)
    - ASM and SSM, using Embedded-RP group address
  - Test environment
    - Linux 2.6.11, Linux 2.4, Solaris 10
    - Cisco (testbed), Juniper (DREN production), Juniper (site), Foundry BI (site)
      - simulating cross-domain interaction

Test Environment

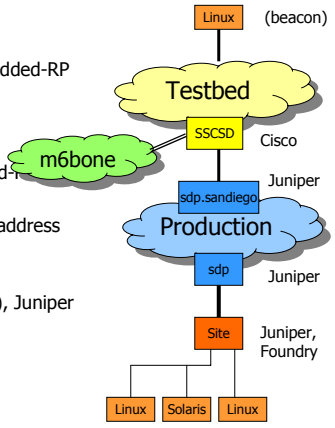

Other IPv6 Multicast activities

- Native IPv6 multicast peering turned up to Abilene at 3 exchanges (NGIX-E, Starlight, Ames)
  - Previously we were only peering with m6bone via tunnel
- Multicast tests to NYSERNET via Abilene
  - SSM works (if not behind a Juniper)
  - ASM works, using Embedded-RP.
- Way cool test…
  - ASM from NYSERNET to UNINETT (Norway) via Abilene, DREN, m6bone, using DREN's IPv6 RP in San Diego (embedded-RP)

```
[cookiemonster:~/Misc/Source/ssmping-0.8.1] owens% ./asmping ff7e:124:2001:480:f000:: ssmping.uninett.no
ssmping joined (S,G) = (2001:700:1:7:211:d8ff:fe8f:1f9b,ff7e:124:2001:480:f000::4321:1234)
pinging S from 2001:468:901:1:203:93ff:fed6:dfcc unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=1 dist=13 time=131.824 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=2 dist=13 time=126.898 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=3 dist=13 time=136.707 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=4 dist=13 time=131.818 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=5 dist=13 time=126.892 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=6 dist=13 time=126.814 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=7 dist=13 time=131.623 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=8 dist=13 time=126.680 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=9 dist=13 time=126.898 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=10 dist=13 time=131.525 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=11 dist=13 time=126.704 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=12 dist=13 time=126.472 ms
  multicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=12 dist=8 time=299.035 ms
  unicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=13 dist=13 time=131.481 ms
  multicast from 2001:700:1:7:211:d8ff:fe8f:1f9b, seq=13 dist=8 time=232.013 ms
```

IPv6 Multicast Beacon

DREN

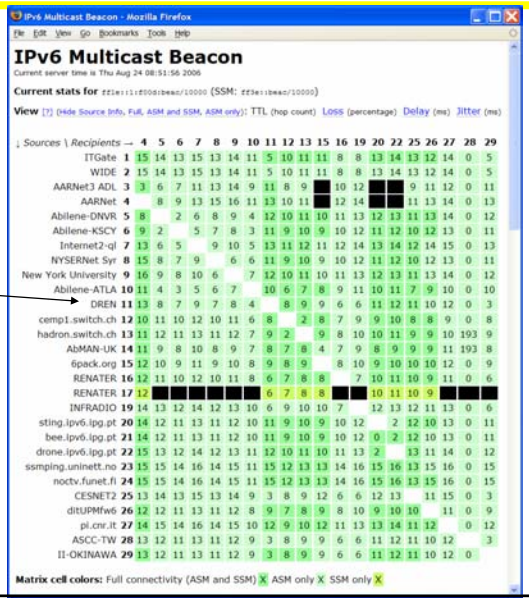

IPv6 Multicast

- Issues
  - Juniper – MLDv2 implementation fundamentally incompatible with modern Linux implementations.
    - Linux uses "exclude mode" for SSM, Juniper only implements "include mode".
    - A fix is "not yet on the product roadmap" ☹
  - Foundry – lacked MLDv2 or Embedded-RP, but good news…
    - BigIron MG8 – Version 2.3 Beta (testing in San Diego)
      - IPv6 PIM-Sparse
      - PIM-SSM
      - MLDv2
      - Embedded-RP
      - And DHCP-relay agent!
    - BigIron (Jetcore)
      - MLDv2 and Embedded-RP in 8.1 (slipping)
  - no MLDv2 in WinXP, broken in old Linux, Solaris.
- Working on…
  - ViPr IPv6 implementation
  - Pressuring the vendors to implement needed features

Some Lessons Learned

- There is no immediate "win" in transitioning to IPv6. The payoff must be viewed as long-term.
- Incentives are needed to encourage near term transition and to make transition a priority.
  - "If you build it, they won't necessarily come"
- Many security components are still not mature nor widely available. Security takes extra thought and effort.
- 1 + 1 > 2
  - managing 2 IP networks (IPv4, IPv6) can be more than double the design complexity due to new interactions.
  - Making topologies congruent can minimize such impact.

Example Re-addressing scheme

- Re-address the network for consistency between protocols
  - IPv4 – move all subnets to /24 or larger
  - Align VLAN number with 3rd octet of IPv4 address
  - Align IPv6 "subnet number" with the above

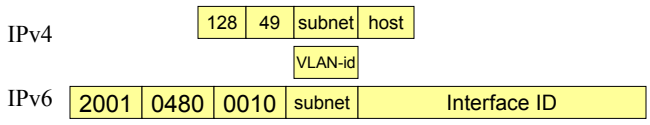

- Benefits
  - Reduction in complexity
  - Easier for operations staff, once re-addressing is complete
- Note
  - Assumes you have enough IPv4 address space to change it as well.

Current DREN IPv6 Pilot Activities: DHCPv6/DNS

- Goal – implement a DHCPv6 environment, similar to how some sites use it in v4.
  - common practice: DHCP (v4) assigns addresses, and performs DNS-update for A and PTR records. DNS master only has to trust DHCP server, not every client.
- Challenge: finding mature and complete DHCPv6 implementation.
- Evaluating
  - ISC (popular dhcp reference implementation), IPv4 only.
  - dhcpv6-linux, incomplete, last update over 2 years ago.
  - dhcpv6 (sourceforge), incomplete but works.
  - Lucent, limited testing appears to work.
  - Dibbler (Poland), in Gentoo Linux distribution.
  - Microsoft Windows Vista©, there but haven't tested yet.
- Issues:
  - no DHCP client in Microsoft Windows XP.
  - some routers don't implement DHCP-relay yet.
  - uncertainty and debate on interactions between stateless and stateful (DHCP) autoconfig.

Other News

- IPv6 Firewall interim solution
  - 10 ISG-2000's provided to DREN by Juniper, to run in parallel with existing IPv4 firewalls until integrated solution is available.
- IPv6 Firewall software update
  - Juniper/Netscreen – 5.4 recently released
    - DREN part of beta program – did extensive testing and reporting
    - Merges all previous code trains, production support for IPv4, IPv6, Multicast.
    - Issues:
      - No transparent mode for IPv6
      - No IPv6 multicast
      - No OSPFv3, No BGP
      - Only supported in a single product
- DHCPv6
  - No reference implementation from ISC
  - DHCPv6 relay not implemented in some routers.
    - Support recently added by Foundry, based on our feature requests.

Other News

- S/DREN-2
  - Will be fully IPv6 enabled
- NTP
  - IPv6-enabled time servers now commercially available
  - Testing and deploying
    - Spectracom
    - Symmetricom

IPv6 capability in products

- These are necessary but not sufficient to show functional equivalence to IPv4:
  - Standards activities (IETF, DISR), theoretical analysis of standards (NSA), test equipment (Agilent, Ixia, Spirent), JITC generic test plans and approved product lists, and test beds (DRENv6, MoonV6).
- These are sufficient but not conclusive to show equivalence:
  - Extended use in real networks to expose and fix remaining errors (Internet2, DREN IPv6 pilot, still more would be nice).
- To really determine IPv6 support for your needs, query the vendor for specific features that matter to you. Be careful in evaluating their response. Try not to let your expectations dictate the results you find, or you will overlook/misinterpret results that contradict those expectations.

*It is **crucial** that IPv6 products have **functionality equivalent** to IPv4 products!*

Challenges

- Keeping security policies consistent
  - ACLs
  - Firewall policies
- Adversaries now have a new entry vector
  - Don't allow IPv6 path to be a new weakest link
- Diagnosing network problems
  - Especially if the routing topology isn't congruent
  - Confusion over which protocol is broken, and what protocol is being tested using diagnostic tools.
- Trying to outlaw NAT
  - Some think that it brings important features (i.e. "security").
- Fighting the pressure to disable IPv6 in Vista
  - Uncertainty in whether it is "safe", from a security perspective.
  - We need to make sure this doesn't happen

Examples of things that are broken or missing

- Netscreen firewall
  - Finally have IPv6 in mainline code, but…
    - Only in one of the hardware products (ISG-2000)
    - Still missing OSPFv3, BGP, IPv6 multicast, transparent-mode, GRE, …
- Juniper
  - Port-mirroring doesn't support IPv6 except in very high-end devices.
  - IPSEC for IPv6 only recently added
- Red Hat
  - RHEL4 feels slow with IPv6 load, due to kernel bug. Not officially fixed until RHEL5
- Mozilla Thunderbird
  - LDAP fails if IPv6 is enabled. A long term problem.
  - Emergence of Vista added pressure to achieve a fix.

Examples of things that are broken or missing

- Many products that are critical to security infrastructure are not IPv6-enabled
  - Bluecoat cache/proxy
  - Netscreen IDP
  - Tipping-Point IPS
  - Many VPN products
    - Both SSL VPNs and IPSEC VPNs
  - Netscreen Security Manager
    - Can't manage IPv6-enabled products
  - Vulnerability assessment and forensics tools from most vendors

Situation Today

- We've been successfully using IPv6 in a production environment, with many dual-stack systems and services, for at least 3 years.
  - Modern operating systems just work, out of the box (MacOSX, Vista, Solaris, etc)
- Most urgent needs from our perspective:
  - Need parity with IPv4 in all implementations
  - Enabling IPv6 must NOT break things
  - Need to make security stacks fully IPv6 capable
    - Firewalls, IDS, proxies, IDP/IPS, ACLs
  - Need more incentives to do IPv6 (generate demand)
- Basic layer 3 (IP routing) implementations are mature
  - ISPs and WANs should be IPv6-enabled now.
    - What about SOHO modems/routers?