



DREN IPv6 Implementation Update

Internet2 Joint Techs, Summer 2008

July 22, 2008

Lincoln, NE

Ron Broersma

DREN Chief Engineer

High Performance Computing Modernization Program

ron@spawar.navy.mil



Background

- WAN provider for the DoD R&D community
- Also serving as DoD's IPv6 pilot implementation of the DoD CIO June '08 mandate
- Deploying IPv6 where possible in a production environment
 - See what works and what's broken
 - See what's missing
 - Share lessons learned



Previously...

- Reported at TIP 2008:
 - Phasing out of DRENV6 testbed, and expansion of native IPv6 peering on production network.
 - DHCPv6 implementations not very interoperable.
 - Still finding lots of IPv6 implementation flaws, and they are getting harder to diagnose.
 - Major impediment to IPv6 deployment is lack of parity between IPv6 and IPv4 in most products.
 - Vendors aren't "eating their own dogfood"

Not eating own dogfood

June 2007		Jan 2008		July 2008	
Organizations	Attributes	Organizations	Attributes	Organizations	Attributes
	WWW		WWW		WWW
	MX		MX		MX
	DNS		DNS		DNS
DoD Organizations					
5	HPCMO	Green	Red	Green	Green
6	SPAWAR	Red	Green	Green	Green
Network Equipment Vendors					
10	Cisco	Red	Red	Red	Red
11	Extreme	Red	Red	Red	Red
12	Force 10	Red	Red	Red	Red
13	Foundry	Red	Red	Red	Red
14	Juniper	Red	Red	Red	Red
15	Eriasson	Red	Red	Red	Green
16	Nortel	Red	Red	Red	Red
17	3Com	Red	Red	Red	Red
Computer and OS Companies					
21	Microsoft	Red	Red	Red	Red
22	Apple	Red	Red	Red	Red
23	Crag	Red	Red	Red	Red
24	HP	Red	Red	Red	Red
25	IBM	Red	Red	Red	Red
26	Sun	Red	Red	Red	Red
27	SGI	Red	Red	Red	Red
Network Security Products					
31	Checkpoint	Red	Red	Red	Red
32	ISS	Red	Red	Red	Red
Networks					
36	DREN	Red	Green	Green	Green
37	AARNET	Red	Red	Red	Red
38	NTT	Red	Red	Red	Red
39	ESNET	Green	Green	Green	Green
40	NISN	Red	Red	Red	Red
41	NLR	Red	Red	Red	Green
42	NREN	Red	Green	Green	Green
43	NYSERNET	Red	Red	Red	Red
44	Internet2	Red	Red	Red	Green
45	Abilene	Red	Red	Red	Red
46	Qwest	Red	Red	Red	Red

- Look in DNS to see if there were AAAA records for www, MX, and DNS.
- Quick sampling of major computer and network companies showed no public facing IPv6.
- Also see Mark Prior's version at http://www.mrp.net/Internet2_IPv6_Survey.html



In other news

- DoD (and agencies under the OMB mandate) met the June 2008 deadline by sending an IPv6 ping across their nets.
 - But didn't leave the networks IPv6-enabled. ☹
- July 7, 2008 - DoD NIC announced that they are now ready to issue IPv6 prefixes to customers.



New initiatives

- Network Management via IPv6
- Migrating the VTC network to IPv6
- Finding the right IPv6-capable IDP



Network Mgmt using IPv6

- Goals
 - Determine if network management can be performed using IPv6.
 - What works? What is missing?
 - Determine if ALL network management can be performed using IPv6.
 - Can we make the Management LAN IPv6-only? If not, what are the showstoppers?
 - Work with vendors to IPv6-enable all management functions on their products.



The test subject

- Management LAN
 - A network for all management traffic that is separate from the normal data network, usually for reasons of security and reliability.
 - Physical separation uses separate cables/switches, while logical separation might just use a VLAN dedicated to management traffic. Doesn't just rely on ACLs for separation.
 - The management traffic is comprised of SNMP, SYSLOG, netflow, sflow, ssh, RADIUS, etc.
 - Isolated (private) network with no routed/natted access outside.
- One DREN customer, spread across a dozen sites worldwide, uses management LANs at each location, and interconnects them out-of-band from the rest of the network.
 - Try to make it IPv6-only



Subject environment

- Equipment in the network
 - Switches: Foundry (ethernet), Ericsson (ATM)
 - Routers: Juniper, Cisco
 - Firewalls: Netscreen (Juniper)
 - Inter-site VPN: netscreen ns204
 - Various security and other appliances
- Management apps
 - Ironview Network Manager (Foundry)
 - snmp, sflow, syslog
 - Inmon Traffic Sentinel
 - sflow, netflow, snmp
 - Others: mrtg, ssh, RADIUS



IPv6 on Mgmt LAN

- Configuring an IPv6 address on device interfaces
 - Foundry: OK
 - Ericsson ATM switch: failed
 - Cisco: OK
 - Juniper router: OK
 - Juniper Netscreen: OK (5.4 or later)
 - Misc appliances: mostly failed
- Inter-site IPSEC mesh
 - ns204 didn't support IPv6 tunnels
 - Replace all with SSG-5s
 - Tunneled IPv4 and IPv6 traffic in IPv4 IPSEC tunnel.
 - Tested IPv6 traffic in IPv6 tunnel – worked well
 - Moved IPv4 traffic to IPv6 tunnel – also worked well
 - Shut down IPv4 tunnels!
 - But v4 traceroutes never show the v6 hop. ☹



Address plan

- Addressing
 - Wanted something akin to private address space.
 - Used ULA (RFC4193), but without the ugly random “Global ID”.
 - ULA = FC00::/7
 - FD00::/7 implies “locally assigned”
 - FDgg:gggg:gggg:ssss:iiii:iiii:iiii:iiii
 - g : random global ID, s : subnet, i : interface ID
 - First try: g = 0, s = small integer (“site”), i = match host num from v4 address
 - FD00:0:0:1::10:30
 - Problem: network discovery took too long
 - range was 0 to 0xfffff (took weeks)
 - Second try: g = 0, s = small integer, i = hex value of host num from v4 address
 - FD00:0:0:1::A1E
 - Discovery much faster, range now 0 to 0xffff (2 hrs)
 - Third try: s = 0, g = small integer in first byte followed by 0’s
 - FD01::A1E
 - Shorter, less typing, often shorter than old IPv4 address. ☺



Network mgmt apps

- InMon Traffic Sentinel
 - Tried to make it do snmp to a switch using IPv6.
 - Could not configure an IPv6 target address.
 - Feature request for full IPv6 support
 - Delivered in less than 3 months
- InMon and sflow
 - InMon relies on sflow for discovery and autoconfiguration.
 - Tried to make Foundry switch send sflow to an IPv6 target.
 - Could not configure an IPv6 target address.
 - Feature request to Foundry to implement sflow via IPv6 in entire product line
 - Nothing yet

Lesson: We don't have time to discover all these shortcomings serially



Find all showstoppers

- Set up IPv6-only mgmt LAN, with switches configured for IPv6-only.
- Learned that additional things were not implemented:
 - Foundry:
 - FDP (like CDP) can't report neighbor's IPv6 address
 - » Have to wait for LLDP
 - FES class switches – IPv6 MIB not implemented
 - FESX class switches – IPv6 MIB not supported until release 4.1 (very latest)
 - ServerIron – no IPv6 support (but coming soon)
 - Netscreen:
 - Snmp via IPv6 not supported
 - Freeradius
 - No IPv6 support until 2.0 (what's in RHEL 5 or later).



More mgmt apps

- Ironview Network Manager (Foundry)
 - Didn't support IPv6 at all until 3.0 release (10/2007)
 - Now works, but has cosmetic issues
 - Doesn't shorten any of the IPv6 addresses
 - FD01:0:0:0:0:0:0:A1E (should be FD01::A1E)
 - Device discovery works much faster if only the bottom few bits are used for individual device numbers.



VTC project

- Connect VTC equipment to the IP network
 - Normally all done with ISDN
 - Can't put it on the net (policy issues)
 - But, we can do IPv6 "experiments" 😊
 - We'll try to do it "IPv6-only"
- Tandberg – manufacturer of existing VTC units
 - Supports IPv6 very well (we were pleasantly surprised)
 - Not much else to say... It just works.



Other random findings

- JUNOS 8.3 IPv6 statistics

```
>show interfaces ge-0/0/0 statistics detail
```

```
Traffic statistics:
```

Input bytes :	15138812973645	14460136 bps
Output bytes :	3884124913540	6987744 bps
Input packets:	15998657542	2592 pps
Output packets:	9736387437	1422 pps

```
IPv6 transit statistics:
```

Input bytes :	127381430782
Output bytes :	9107850097
Input packets:	209102407
Output packets:	44348632

- Cisco FWSM

- doesn't support IPv6 GRE (IP47) nor IPv6 in IPv4 (IP41)


- MRTG

- No straightforward way to report percentage of traffic that is IPv6.
 - Octet count only in interface MIB, not IP or IPv6 MIBs.




MRTG IPv6 stats

- Split IPv6 from IPv4 traffic onto different interfaces. Then count octets.



IPv4 Traffic



IPv6 Traffic



IPv6 utilization - growth

- My site (SPAWAR)
 - IPv6 > 1% of all traffic at Internet border
 - Was sitting at < 0.1% for a long time
 - Internal percentage of network devices actively using IPv6 increased from 3% to 5% over last 6 months.



END



Backup



IPv6-enablement milestones and scoring (proposed)

1. IPv6 address allocation and address/routing plan developed
2. LAN (wired and wireless) is fully v6-enabled (all routers do v6 on all interfaces) and is connected to the IPv6 Internet (WAN).
 - The implication is that any v6-enabled device can connect anywhere in the LAN and get IPv6 Internet connectivity.
 - (worry about routing implementation, make sure address plan is right, think about security and performance, play with multicast, make sure network staff is trained to support it, etc)
3. Internal infrastructure services (DNS, NTP, DHCP, SMTP, etc) are IPv6-enabled
4. Public facing services (public DNS, MXs, public web site) are IPv6-enabled
5. Network management infrastructure (management LAN, SNMP, SYSLOG, MIBs, management access to switches/routers/etc) is IPv6-enabled
6. Most workstations and servers are v6-enabled
 - (behind this is the support infrastructure, i.e. help desk and other IT support, and a message to sys admins to turn on IPv6 where possible, and servers have AAAA records in DNS, and your help desk tools/scripts for things like host registration and update are upgraded to handle IPv6 addresses)
7. Once critical mass is reached on the client side, remove "A" records for servers (resulting in final incentive and some pain for those that didn't dual-stack their workstations).
 - You don't need to do them all at once, just one at a time as their clients all become dual-stack
8. Migrate some network segments to IPv6-only, with no IPv4 addresses anywhere
 - (this forces one to figure out how to make hosts operate in a v6-only environment, helps the organization figure out what's missing, forces the implementation of some kind of transition mechanism to bridge to the IPv4-only world, plus adds continued incentive to get more stragglers upgraded since they can't even get there by typing the IP address)
9. Deploy advanced features (mobile-ip, v6 multicast, etc.), provide remote IPv6 access (travellers, telecommuters, home, etc) from v4-only environments, cleanup and reduce complexity (readdressing network if necessary),
10. Declare victory
 - you claim a perfect score in public, and are willing to stand up to scrutiny

Scoring: Scale of 0 to 10. 1 point for any milestone that is 95% complete.