IPv6 Everywhere:
Living with a Fully IPv6-enabled environment

Australian IPv6 Summit 2010
20 October 2010
Melbourne, VIC Australia

Ron Broersma
DREN Chief Engineer
ron@spawar.navy.mil

Introduction

- Aggressive deployment of IPv6 to DoD's R&E WAN (DREN) and to all campuses of one major customer (SPAWAR)
- May be different than other IPv6 initiatives that you've heard about
  - this is real production stuff, not just a testbed
  - this isn't just an ISP view of the world, or just a campus view, or system or application view, it is ALL of the above
  - the systems and users are autonomous customers, not part of a centrally managed (i.e. active directory) environment
  - this is a heterogeneous environment, not just Windows
    - Win2K, XP, Vista, Win7, Win2K3, Win2K8, Linux, MacOSX, Solaris, HP/UX, BSD, ESX, SCO, etc.
  - this isn't just a few systems, its everything on the network
- Goals
  - Push the envelope with IPv6 deployment and see what's possible
    - See what's missing or broken and work with the vendors to get it fixed
  - Dual stack everywhere, IPv6-only where possible
  - Share lessons learned

DREN: "Defense Research and Engineering Network"

Progress to date

- ✔ WAN – dual stack everywhere, peering (unicast+multicast)
- ✔ LANs, WLAN – all subnets fully support v6, renumber v4
- ✔ Infrastructure services – recursive DNS, NTP, SMTP, XMPP
- ✔ Support services – RADIUS, LDAP, Kerberos
- ✔ Public facing services – authoritative DNS, MX's, www, NTP
- ✔ "Security stack" – firewall, IDS, IPS, etc.
- ✔ Security services – WSUS, McAfee ePO (HBSS)
- ✔ Servers, desktops, laptops – 98% dual stack

Utilization comparison

IPv4 traffic

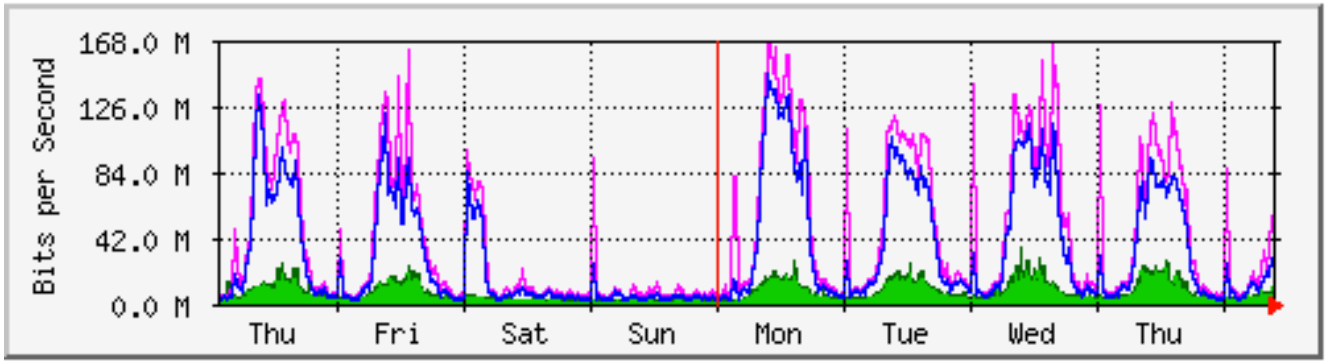

IPv6 traffic

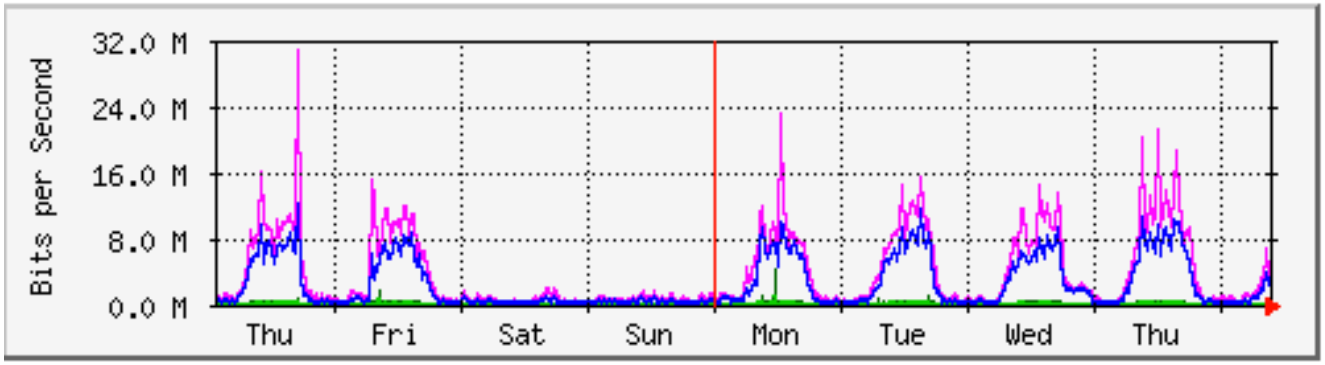

Almost 10% of traffic is IPv6

Lessons Learned

- Its not really that hard, and doesn't have to be very expensive
- But you need to make it a corporate culture, that permeates all levels of the organization
- Don't wait until it's a crisis, just roll it out gradually as part of normal tech refresh or other upgrades
- If you haven't started yet, you're already behind
- Training must be simplified
- Work from the outside inward
- Don't be afraid to "break some glass"
  - things get fixed quicker that way
- Don't buy from vendors unless they support IPv6
  - beg for "feature parity"
  - check out their web site to see if it is IPv6-enabled…

Eating your own dogfood

- Many IPv6 proponents (vendors) not eating their own dogfood.
- Example: sponsors for this Summit
- http://www.mrp.net/IPv6_Survey.html
- What's wrong with this picture?
- But this is MUCH better than previously

Australian IPv6 Summit 2010 Sponsors/Supporters

Organisation (domain)	Web	Mail	DNS	NTP	XMPP
AUDA (auda.org.au)	FAIL	FAIL	0/3 0/3		
Australian Computer Society (acs.org.au)	FAIL	FAIL	0/0 0/3		
Australian Industry Group (aigroup.asn.au)	FAIL	FAIL (M)	0/3 0/3		
Australian Information Industry Association (aiia.com.au)	FAIL	FAIL	0/0 0/3		
Blue Coat (bluecoat.com)	FAIL	FAIL (PP)	2/2 2/4	FAIL	C:FAIL
Bluecat Networks (bluecatnetworks.com)	FAIL	FAIL (M)	0/2 0/2		
Cisco Systems (cisco.com)	www.ipv6	FAIL	0/2 0/2	FAIL	FAIL
Communications Alliance (commsalliance.com.au)	FAIL	FAIL	0/0 0/3		
eintellego (eintellego.net)	SUCCESS	SUCCESS	0/0 0/2		
Engineers Australia (engineersaustralia.org.au)	FAIL	FAIL (M)	0/0 0/2		
HP (hp.com)	FAIL (N)	FAIL	0/6 0/6		
ICANN (icann.org)	SUCCESS	SUCCESS	0/1 3/5	FAIL	
Internet Society (isoc.org)	SUCCESS	FAIL	0/0 6/6		
Internet Society of Australia (isoc-au.org.au)	SUCCESS	SUCCESS	0/0 3/4		
Internode (internode.com.au)	FAIL	FAIL	0/0 4/4		
IPv6 Forum (ipv6forum.com)	SUCCESS	SUCCESS	0/0 2/2		
IPv6 Forum Australia (ipv6forum.org.au)	SUCCESS	FAIL	0/0 3/4		
IPv6 Now (ipv6now.com.au)	SUCCESS	SUCCESS	1/2 4/6		
Juniper Networks (juniper.net)	FAIL (A)	FAIL (P)	0/3 0/5		
Mach Technology (mach.com.au)	FAIL	FAIL	0/0 0/4		
Multimedia Victoria (mmv.vic.gov.au)	FAIL	FAIL	0/0 0/3		
NICTIA (nictia.org.au)	FAIL	FAIL	0/0 0/2		
Sophos (sophos.com)	FAIL (A)	FAIL	0/2 0/2	FAIL	
Vocus (vocus.com.au)	SUCCESS	SUCCESS	3/3 3/3	Stratum 1	SUCCESS

4 years ago…

- Some of the issues I brought up when I spoke here in 2006:
  - Windows 2000 issues
  - Lack of IPv6 support in firewalls
    - beta code only
    - we had to run dual firewalls
  - Juniper mirroring problems (only IPv4 supported)
  - DHCPv6 support lacking (XP, MacOS)
  - lack of feature parity – all vendor products
  - … and many bugs
- Most issues now resolved, except for DHCPv6 support and feature parity

Keeping DNS updated

- Need to get all PTRs and some AAAA's in DNS for all devices doing IPv6
- Manual editing of zone files?
  - Much more painful than IPv4
  - How do you know when some device starts doing IPv6 and gets a SLAAC address?
- DHCPv6?
  - Use DHCPv6 to provide addresses, and use dynamic DNS update
  - Problem: too many clients do not yet support DHCPv6 (Windows XP, MAC OSX, others)

DNS auto-update

- Basic scheme
  - Use SNMP to poll the routers
    - Grab the ARP cache and the ND table
  - For all MAC addresses in the ND table with global unicast addresses matching the site IPv6 prefix:
    - Find the corresponding IPv4 address from the ARP cache
    - Find the FQDN for the IPv4 address in DNS (PTR lookup)
    - Build a PTR record for the IPv6 address, using FQDN from IPv4 address
    - Push to DNS dynamically
  - Works very well
  - Yes, there are some additional complexities, and optimizations required, like garbage collection of temporary and privacy addresses.
  - Hoping to release tool tool as open-source.
- Lingering problems with IPv6 objects in the IP-MIB and IPV6-MIB
  - We really need all routers supporting RFC 4293 (version independent IP-MIB)

AAA services

- RADIUS
  - Needed to upgrade servers to freeradius 2.0 to support IPv6
- Kerberos, LDAP servers
  - Just works, as expected
- LDAP client issue
  - Could not make some perl and PHP based apps connect to LDAP via IPv6
  - Perl module Net::LDAP has no IPv6 support until 0.35
    - Latest RHEL only has 0.33
  - Still need to modify code to ask for IPv6
  - Perl modules need to be made IP version agnostic

IPv6 support in perl is poor, with no near term resolution

Google over IPv6

- Google has “opt-in” program to get everything over IPv6
- Feb 3, 2009 – added all of SPAWAR
- July 28, 2009 – DREN and ALL customers added
- Any DREN user that is IPv6-enabled will get to Google services over IPv6
  - Faster (over non-congested links)
    - DREN private peering with Google is IPv6-only
  - Helps to quickly identify IPv6 connectivity problems
- As incentive, we block IPv4 to Google

Lack of IPv6 support

- vmware ESX 3.x
  - Supported in 4.0, but disabled by default
- Windows 2000
  - We tell users to upgrade to a newer OS
- Printers
  - Too hard right now
  - For HP printers we are replacing the jetdirect cards with new ones that support IPv6
- Various appliances

More challenges

- Older versions of MS Outlook
  - We tell users to upgrade to MS Office 2007 or later
- Large groups of systems that are under "configuration control", and can't be modified, even to enable IPv6.
- Rogue 6to4 relays sending RAs
  - Windows systems with ICS enabled.
  - Workaround: set router priority to "high". Fix: RA Guard
- Symantec Endpoint Protection (SEP) breaks IPv6
- Blackberry Enterprise Services (BES) on IPv6-enabled Windows server will crash.
- Oracle lack of IPv6 support

java

- We noticed that java apps never use IPv6
  - Even when operating on properly configured dual stack systems, and talking to IPv6-enabled servers.
- Java system property `java.net.preferIPv6Addresses` is set to "false" by default
- Fix: Add this to your java options:

```
-Djava.net.preferIPv6Addresses=true
```

Privacy addresses

- See RFC 4941
- Windows systems do this by default (and we don't like it!)
- Breaks many things in our environment
  - Forensics
  - Stable DNS entries
  - Automated management tools
- Could fix with DHCPv6, but client not available in important OS's
  - Windows XP, Mac OSX
- Would be nice if RA's could say "don't do this"
- So we have to visit every Windows machine to disable this.
  - Breaks the "plug and play" goal of IPv6 for clients.

```
netsh interface ipv6 set privacy state=disabled store=persistent
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```

Mac OSX 10.6 (Snow Leopard)

- After upgrade to Snow Leopard, web browsing and other apps no longer seemed to prefer IPv6 over IPv4.
- Behavior is that only the first DNS answer to any query is accepted, and the others are dropped.
  - if you get the A before the AAAA, the AAAA will get dropped
- In 10.6, mDNSResponder is now used for all unicast DNS queries, not just for multicast as was the case in earlier releases.
- mDNSResponder will query for "A" and "AAAA", but will immediately stop listening after the first reply.
  - the application never receives the other responses
- References:
  - http://support.apple.com/kb/HT3789
  - http://openradar.appspot.com/7333104

java on Mac OS X

- java defaults to IPv4 instead of IPv6
  - reported earlier
- You can change the behavior by setting a preference
  - -Djava.net.preferIPv6Addresses=true
- This preference setting has no effect in Mac OS X
  - can't override the bad default
- Reference:
  - http://openradar.appspot.com/7100919

Mac OS X and IPv6 printers

- You can't configure an IPv6 address for a printer
- It has to find the printer using Bonjour, or you have to specify a DNS name.
  - an explicit IPv6 address will not work.
  - Apple says: "this is expected behavior"
- Reference:
  - http://openradar.appspot.com/7100507

NetApp Storage Appliance

- We've been waiting a long time for IPv6 support
- Delivered in 7.3.1 (Jan '09) but very buggy
- 7.3.1_P2 (Jul '09) was supposed to work, and be more reliable, but every time we enabled IPv6, all mounts started failing.
- 7.3.3 (May '10) now works, and is enabled
  - All NFS now over IPv6 (Linux clients, Solaris 10 clients)
    - not supported in Solaris 8
  - CIFS now over IPv6 from modern Windows systems
    - IPv6 file sharing not supported in Windows XP
- Possible latent bug – interface flap loses v6 route permanently until interface reset

DNSSEC and IPv6 frags

- DNSSEC over IPv6 problems
  - 13 roots, 8 reachable via IPv6
  - Of the 8, 2 truncate so 6 will fragment on long responses (like DNSSEC answers)
  - we were losing all IPv6 fragments, so DNSSEC was failing, or falling back to TCP
- IPv6 filtering broken in Linux <2.6.20
  - Latest RHEL is at 2.6.18
  - if you use ip6tables, frags get dropped
    - turn off ip6tables if using DNSSEC on Linux <2.6.20

SPAWAR mgmt LAN update

- Goal:
  - migrate management LAN to IPv6 where possible
    - all devices get IPv6 address
    - all management services use IPv6 transport
  - eventually turn off IPv4, where possible
- Triage
  - ignore devices that were eventually going away anyway (ATM switches, dialup modems, …)
  - use tech refresh to get IPv6 support on various devices
    - old firewalls and vpn devices replaced
    - replaced many older switches
  - Pushed vendors for firmware updates on others

mgmt LAN

- Foundry (Brocade) enhancements requested and delivered
  - snmp (v3) over IPv6 transport
  - DNS over IPv6 transport
  - sflow records over IPv6 transport
  - sflow agent-ID set to IPv6 address
  - RADIUS over IPv6
  - Unified IP MIB support
  - set router priority in RAs
  - … and many bug fixes along the way
- Not yet delivered
  - RA Guard
- Can't disable IPv4 just yet
  - Lack full IPv6 support in FDP and LLDP across all platforms
  - some end-of-support switches still need to be replaced ($$$)

mgmt LAN

- Other devices with IPv6 mgmt support (some only partial support)
  - Spectracom NTP servers
  - Symmetricom NTP servers
  - Netscreen devices (ScreenOS)
  - TippingPoint IPS
  - new APC UPS units ***
- No support
  - Google Search Appliance
  - Aruba WLAN controller and APs
  - Cisco 3000-series VPN servers

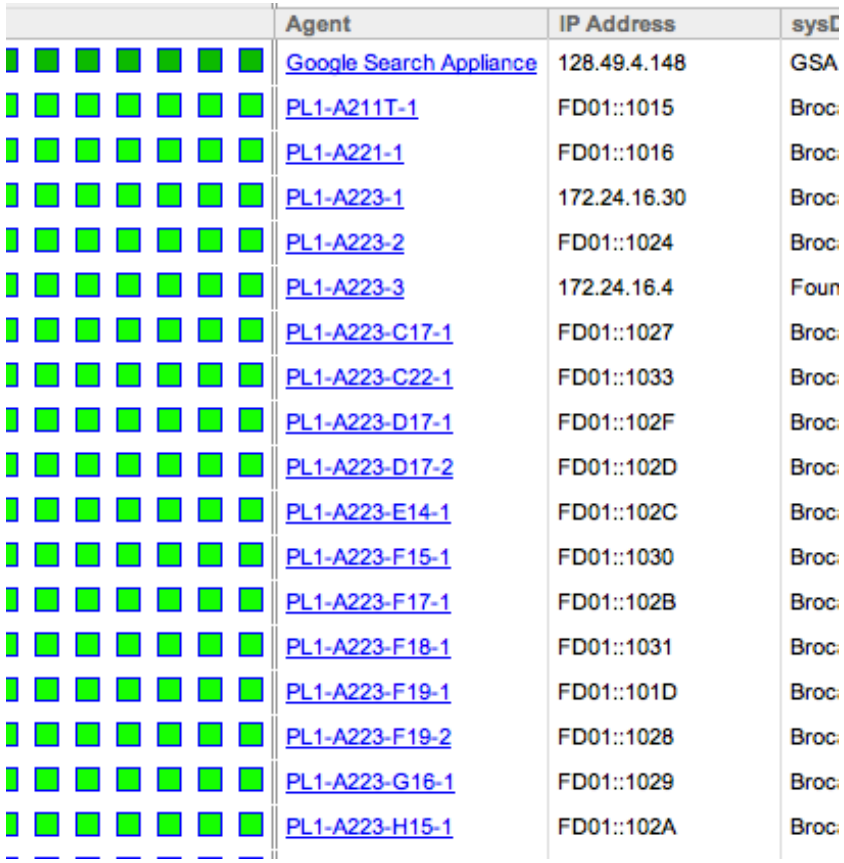

Agent	IP Address	sysD
Google Search Appliance	128.49.4.148	GSA
PL1-A211T-1	FD01::1015	Broc
PL1-A221-1	FD01::1016	Broc
PL1-A223-1	172.24.16.30	Broc
PL1-A223-2	FD01::1024	Broc
PL1-A223-3	172.24.16.4	Foun
PL1-A223-C17-1	FD01::1027	Broc
PL1-A223-C22-1	FD01::1033	Broc
PL1-A223-D17-1	FD01::102F	Broc
PL1-A223-D17-2	FD01::102D	Broc
PL1-A223-E14-1	FD01::102C	Broc
PL1-A223-F15-1	FD01::1030	Broc
PL1-A223-F17-1	FD01::102B	Broc
PL1-A223-F18-1	FD01::1031	Broc
PL1-A223-F19-1	FD01::101D	Broc
PL1-A223-F19-2	FD01::1028	Broc
PL1-A223-G16-1	FD01::1029	Broc
PL1-A223-H15-1	FD01::102A	Broc

A note on Freeradius 2

- Freeradius 2 supports IPv6
- For RHEL 5, there's a separate RPM named "freeradius2"
  - delete "freeradius" and install "freeradius2"
- Documentation and discussion would lead you to believe that it can't do IPv4 and IPv6 at the same time
  - see notes in radiusd.conf
  - see discussion on various web forums
- Actually, all you need to do is add another "listen" clause…

Freeradius 2 example

```
listen {
        type = auth
        ipaddr = *
        port = 0
        clients = clients-ipv4
}


# Listen on the IPv6 address too
listen {
        type = auth
        ipv6addr = ::
        port = 0
        clients = clients-ipv6
}
```

clients config file for all your IPv4 clients

IPv6 clients config file

Managing the UPSs

- None of the manageable UPS devices supported IPv6
- APC Network Management 2 card now has IPv6 support
  - IPv6-ready Phase-2/Gold Logo
- We're upgrading all APC UPS devices

New approach to training

- Training approach is more pragmatic
  - No more "everything you wanted to know about IPv6"
  - Instead, "turn on IPv6 in 5 easy steps"
    - including templates for emails that you need to send
- Pre-configure IPv6 on all WAN customer interfaces
- Lay out some best practices
  - In very strong terms: "Read my lips".
  - Mostly addressing guidelines.
    - forget about being conservative like in IPv4
    - subnets are /64
      - yes, even the point-to-point links
    - don't encode v4 subnet values into bottom 64 bits
    - no NAT
  - avoid tunneling where possible (go native)

Soapbox

- Enabling IPv6 throughout your environment needs to be a cultural thing.
  - Get everyone involved and on-board
  - Include it as part of tech refresh.
- It may seem overwhelming in the beginning, but its really not that hard to get started.
- Don't be afraid to break some glass
- Very important that we focus on making our public facing services dual-stack as soon as possible.
  - otherwise we'll be in translator-hell
  - eventually some clients won't be able to reach you
- IPv6 is an "unfunded mandate", and everyone needs to do their part.
- Need v4/v6 feature parity in products
- Avoid vendors that don't have a good IPv6 story

END

IPv6 in the Enterprise
What is needed for deployment

Ron Broersma

DREN, SPAWAR

General

- Feature Parity in mainstream vendor products
  - IPv6 needs to be "as good as" IPv4
    - equivalent functionality
    - equivalent performance
      - think ASICs
- Better vendor QA of products
  - they aren't "eating their own dogfood"
  - QA suites are not mature, or don't exist
  - Things would get fixed a lot faster if they had to feel the pain
- IPv6 Internet needs to be as robust at the IPv4 Internet
  - get off tunnels
  - kill the black holes
  - fix PMTUD everywhere

IPv6 on public facing services

- Just to enable IPv6 on www, there are often many showstoppers
  - Akamai doesn't provide IPv6 support
  - Co-Lo or hosting facility provides no IPv6 connectivity
  - Existing load balancers don't yet support IPv6
  - Network engineers have no influence with the IT or marketing staff that runs their web site
  - Won't consider simple alternatives to get started, like a v6v4 proxy.

OS Vendors

- Microsoft
  - We need to phase out XP
    - IPv6 not on by default, no DHCPv6, no DNS over IPv6, easy to become a rogue router sending RAs
  - Life is better with Windows 7 and 2K8, but…
    - privacy/temporary/randomized addresses are often incompatible with enterprise requirements
    - we need a knob somewhere (besides AD) where we can disable this
      - a new bit in the RAs?
- Apple
  - Need serious attention to MacOSX support (has been degrading over last few years)
    - fix brokenness in 10.6 (mDNSresponder, 6to4 preference, etc.)
    - support DHCPv6
    - ISATAP support would help transition in some enterprises
  - Apple needs to dual-stack their own network and eat their own dogfood, and get whitelisted, so they feel the pain that the rest of us feel.
  - Need IPv6 in iOS now

Other pieces

- Need to be able to do DHCPv6 instead of SLAAC
  - requires broad client support, which just isn't there today
- Need RA-Guard standardized and implemented in all switch products
  - still just an I-D
- VPN products need to support dual-stack

Just do it

- Corporate culture
  - make IPv6 permeate the IT culture
  - get buy-in from CIO/CEO on down
  - have a local champion/evangelist
  - include IPv6 in every IT initiative, especially all tech-refresh
- Training
  - simplify it, easy steps to get started
- Triage, and ordering of steps
  - addressing, then connectivity via ISP, then testbed, then training, then public facing services, then security perimeter, then internal networks, then systems and apps, then….