IPv6 deployment on Production Networks

2010 Rocky Mountain IPv6 Summit

Denver, CO

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 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"

Historically

- Late 70's, early 80's
  - many DREN customers were on original ARPAnet and MILnet
  - we already experienced one transition (NCP -> IP)
- 1990's
  - early IPv6 implementations (NRL)
  - Ad hoc experiments, 6bone, etc.
- 2001-2002
  - Became official coordinated project
  - Built and operated dedicated IPv6 testbed for early adopters
- 2003
  - DREN selected as DoD IPv6 "pilot"
  - Started efforts to IPv6-enable all production networks and services
- Today
  - Have been successfully operating IPv6 across production networks and services for years
  - Starting to think in terms of IPv4 as "the old/legacy protocol"

Progress to date

- ✔ WAN – dual stack everywhere, peering (unicast+multicast)
- ✔ LANs – 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
- ✔ Servers, desktops, laptops – 93% dual stack

IPv6 Survey
http://www.mrp.net/IPv6_Survey.html

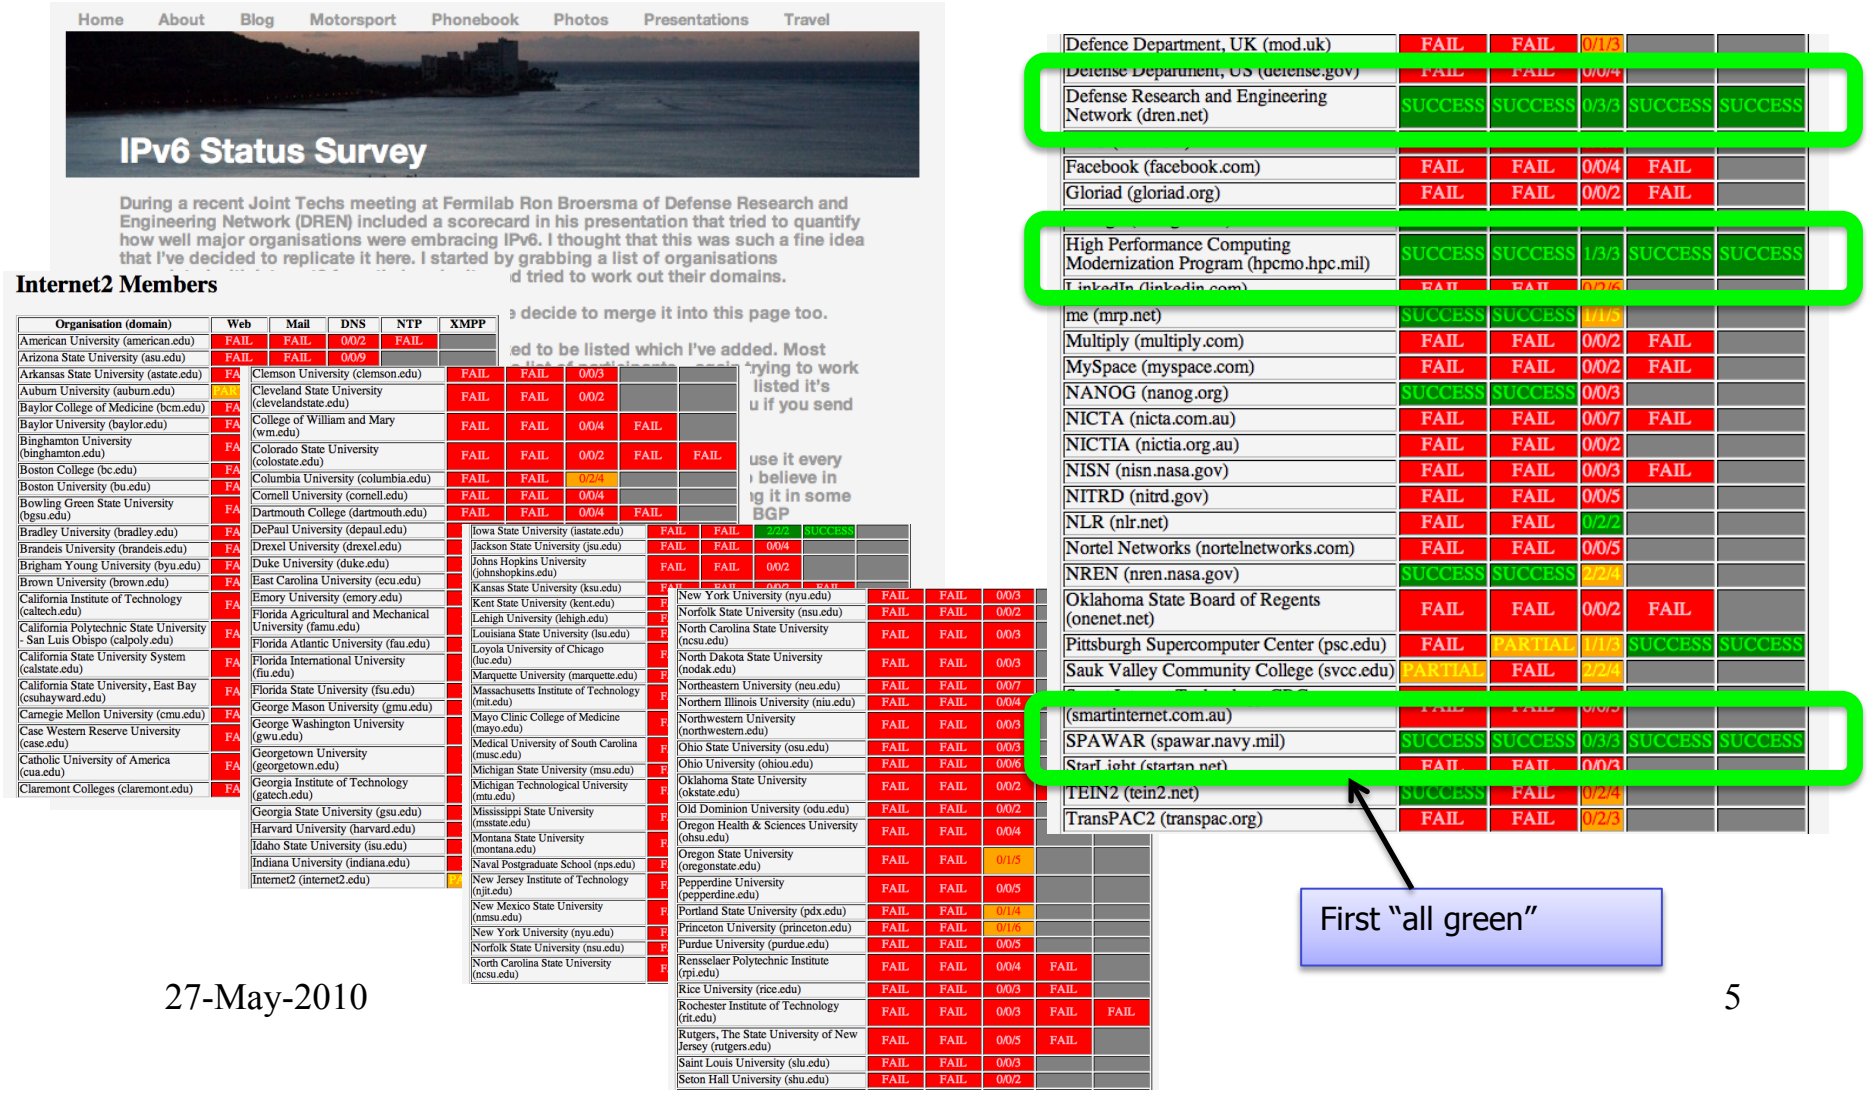

Home About Blog Motorsport Phonebook Photos Presentations Travel

IPv6 Status Survey

During a recent Joint Techs meeting at Fermilab Ron Broersma of Defense Research and Engineering Network (DREN) included a scorecard in his presentation that tried to quantify how well major organisations were embracing IPv6. I thought that this was such a fine idea that I've decided to replicate it here. I started by grabbing a list of organisations d tried to work out their domains.

e decide to merge it into this page too.

ed to be listed which I've added. Most rying to work listed it's u if you send

use it every believe in g it in some BGP

Internet2 Members

Organisation (domain)	Web	Mail	DNS	NTP	XMPP
American University (american.edu)	FAIL	FAIL	0/0/2	FAIL	
Arizona State University (asu.edu)	FAIL	FAIL	0/0/9		
Arkansas State University (astate.edu)	FA				
Auburn University (auburn.edu)	FAI				
Baylor College of Medicine (bcm.edu)	FA				
Baylor University (baylor.edu)	FA				
Binghamton University (binghamton.edu)	FA				
Boston College (bc.edu)	FA				
Boston University (bu.edu)	FA				
Bowling Green State University (bgsu.edu)	FA				
Bradley University (bradley.edu)	FA				
Brandeis University (brandeis.edu)	FA				
Brigham Young University (byu.edu)	FA				
Brown University (brown.edu)	FA				
California Institute of Technology (caltech.edu)	FA				
California Polytechnic State University - San Luis Obispo (calpoly.edu)	FA				
California State University System (calstate.edu)	FA				
California State University, East Bay (csuhayward.edu)	FA				
Carnegie Mellon University (cmu.edu)	FA				
Case Western Reserve University (case.edu)	FA				
Catholic University of America (cua.edu)	FA				
Claremont Colleges (claremont.edu)	FA				
Clemson University (clemson.edu)	FAIL	FAIL	0/0/3		
Cleveland State University (clevelandstate.edu)	FAIL	FAIL	0/0/2		
College of William and Mary (wm.edu)	FAIL	FAIL	0/0/4	FAIL	
Colorado State University (colostate.edu)	FAIL	FAIL	0/0/2	FAIL	FAIL
Columbia University (columbia.edu)	FAIL	FAIL	0/2/4		
Cornell University (cornell.edu)	FAIL	FAIL	0/0/4		
Dartmouth College (dartmouth.edu)	FAIL	FAIL	0/0/4	FAIL	
DePaul University (depaul.edu)					
Drexel University (drexel.edu)					
Duke University (duke.edu)					
East Carolina University (ecu.edu)					
Emory University (emory.edu)					
Florida Agricultural and Mechanical University (famu.edu)					
Florida Atlantic University (fau.edu)					
Florida International University (fiu.edu)					
Florida State University (fsu.edu)					
George Mason University (gmu.edu)					
George Washington University (gwu.edu)					
Georgetown University (georgetown.edu)					
Georgia Institute of Technology (gatech.edu)					
Georgia State University (gsu.edu)					
Harvard University (harvard.edu)					
Idaho State University (isu.edu)					
Indiana University (indiana.edu)					
Internet2 (internet2.edu)					
Iowa State University (iastate.edu)	FAIL	FAIL	2/2/2	SUCCESS	
Jackson State University (jsu.edu)	FAIL	FAIL	0/0/4		
Johns Hopkins University (johnshopkins.edu)	FAIL	FAIL	0/0/2		
Kansas State University (ksu.edu)	FAIL	FAIL	0/0/2	FAIL	
Kent State University (kent.edu)					
Lehigh University (lehigh.edu)					
Louisiana State University (lsu.edu)					
Loyola University of Chicago (luc.edu)					
Marquette University (marquette.edu)					
Massachusetts Institute of Technology (mit.edu)					
Mayo Clinic College of Medicine (mayo.edu)					
Medical University of South Carolina (musc.edu)					
Michigan State University (msu.edu)					
Michigan Technological University (mtu.edu)					
Mississippi State University (msstate.edu)					
Montana State University (montana.edu)					
Naval Postgraduate School (nps.edu)					
New Jersey Institute of Technology (njit.edu)					
New Mexico State University (nmsu.edu)					
New York University (nyu.edu)					
Norfolk State University (nsu.edu)					
North Carolina State University (ncsu.edu)					
New York University (nyu.edu)	FAIL	FAIL	0/0/3		
Norfolk State University (nsu.edu)	FAIL	FAIL	0/0/2		
North Carolina State University (ncsu.edu)	FAIL	FAIL	0/0/3		
North Dakota State University (nodak.edu)	FAIL	FAIL	0/0/3		
Northeastern University (neu.edu)	FAIL	FAIL	0/0/7		
Northern Illinois University (niu.edu)	FAIL	FAIL	0/0/4		
Northwestern University (northwestern.edu)	FAIL	FAIL	0/0/3		
Ohio State University (osu.edu)	FAIL	FAIL	0/0/3		
Ohio University (ohiou.edu)	FAIL	FAIL	0/0/6		
Oklahoma State University (okstate.edu)	FAIL	FAIL	0/0/2		
Old Dominion University (odu.edu)	FAIL	FAIL	0/0/2		
Oregon Health & Sciences University (ohsu.edu)	FAIL	FAIL	0/0/4		
Oregon State University (oregonstate.edu)	FAIL	FAIL	0/1/5		
Pepperdine University (pepperdine.edu)	FAIL	FAIL	0/0/5		
Portland State University (pdx.edu)	FAIL	FAIL	0/1/4		
Princeton University (princeton.edu)	FAIL	FAIL	0/1/6		
Purdue University (purdue.edu)	FAIL	FAIL	0/0/5		
Rensselaer Polytechnic Institute (rpi.edu)	FAIL	FAIL	0/0/4	FAIL	
Rice University (rice.edu)	FAIL	FAIL	0/0/3	FAIL	
Rochester Institute of Technology (rit.edu)	FAIL	FAIL	0/0/3	FAIL	FAIL
Rutgers, The State University of New Jersey (rutgers.edu)	FAIL	FAIL	0/0/5	FAIL	
Saint Louis University (slu.edu)	FAIL	FAIL	0/0/3		
Seton Hall University (shu.edu)	FAIL	FAIL	0/0/2		
Defence Department, UK (mod.uk)	FAIL	FAIL	0/1/3		
Defense Department, US (defense.gov)	FAIL	FAIL	0/0/4		
Defense Research and Engineering Network (dren.net)	SUCCESS	SUCCESS	0/3/3	SUCCESS	SUCCESS
Facebook (facebook.com)	FAIL	FAIL	0/0/4	FAIL	
Gloriad (gloriad.org)	FAIL	FAIL	0/0/2	FAIL	
High Performance Computing Modernization Program (hpcmo.hpc.mil)	SUCCESS	SUCCESS	1/3/3	SUCCESS	SUCCESS
LinkedIn (linkedin.com)	FAIL	FAIL	0/2/6		
me (mrp.net)	SUCCESS	SUCCESS	1/1/5		
Multiply (multiply.com)	FAIL	FAIL	0/0/2	FAIL	
MySpace (myspace.com)	FAIL	FAIL	0/0/2	FAIL	
NANOG (nanog.org)	SUCCESS	SUCCESS	0/0/3		
NICTA (nicta.com.au)	FAIL	FAIL	0/0/7	FAIL	
NICTIA (nictia.org.au)	FAIL	FAIL	0/0/2		
NISN (nisn.nasa.gov)	FAIL	FAIL	0/0/3	FAIL	
NITRD (nitrd.gov)	FAIL	FAIL	0/0/5		
NLR (nlr.net)	FAIL	FAIL	0/2/2		
Nortel Networks (nortelnetworks.com)	FAIL	FAIL	0/0/5		
NREN (nren.nasa.gov)	SUCCESS	SUCCESS	2/2/4		
Oklahoma State Board of Regents (onenet.net)	FAIL	FAIL	0/0/2	FAIL	
Pittsburgh Supercomputer Center (psc.edu)	FAIL	PARTIAL	1/1/3	SUCCESS	SUCCESS
Sauk Valley Community College (svcc.edu)	PARTIAL	FAIL	2/2/4		
(smartinternet.com.au)	FAIL	FAIL	0/0/5		
SPAWAR (spawar.navy.mil)	SUCCESS	SUCCESS	0/3/3	SUCCESS	SUCCESS
StarLight (startap.net)	FAIL	FAIL	0/0/3		
TEIN2 (tein2.net)	SUCCESS	FAIL	0/2/4		
TransPAC2 (transpac.org)	FAIL	FAIL	0/2/3		

First "all green"

IPv6 deployment

- 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
- Don't buy from vendors unless they support IPv6
  - check out their web site to see if it is IPv6-enabled
  - beg for "feature parity"
- Don't be afraid to "break some glass"
- If you haven't started yet, you're already behind

Today's talk

- Review issues we've run in to over the last year or so
- Share lessons learned and solutions
- Point out some unresolved issues

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
  - Need to modify code to ask for IPv6
  - Perl modules need to be made IP version agnostic

A note on Freeradius 2

- Freeradius 2 supports IPv6
- For RedHat, 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

DNS zone transfers

- Can we force all zone xfers over v6?
- Need to change the config on the slaves

```
zone "nosc.mil" {
        type slave;
        file "slaves/nosc";
        masters {
                128.49.4.20;
        };
};
```

Change this to the IPv6 address of the master server

```
zone "nosc.mil" {
        type slave;
        file "slaves/nosc";
        masters {
                2001:480:10:4::20;
        };
};
```

Like this

DNS zone transfers

- Need to update ACLs on the master server

```
Allow-transfer {
        198.253.48.7;
        2001:480:10:1048::7;
}
```

← Add the IPv6 address of the slave server

- Notifications:
  - The master sends a "notify" to all slaves listed by "NS" records.
  - Dual stack slaves get notified at both the IPv4 and IPv6 addresses.
  - The IPv4 notifications now fail

```
named[15666]: zone nosc.mil/IN: refused notify from non-master: 128.49.4.20#51718
```

  - Adding the IPv4 address back in the master list on the slave server seems to be OK, and the zone xfers go via IPv6
  - Also change all instances of "also-notify" to IPv6 address

Oracle

- Oracle Applications Server fails when running on a Solaris machine that has IPv6 enabled:

Part Number B32217-05
Oracle Application Server Release Notes
10g Release 3 (10.1.3.1.0) for Solaris Operating System (x86) and Solaris Operating System (x86-64)

2.1.2 **IPv6 Not Supported**
This release of Oracle Application Server is not certified to run on machines that are configured with IPv6. You have to install and run this release of Oracle Application Server on machines that are configured with IPv4.

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 is supposed to work, and be more reliable, but every time we enable IPv6, all mounts start failing.

Unresolved

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
```

Fixing the VPN problem

- Travelers and telecommuters use client VPNs to connect to the corporate Intranet securely
  - Like Cisco IPSEC VPN or Juniper SSL VPN
- Only tunnels IPv4 traffic (today)
- IPv6 traffic, if supported at all, goes outside this tunnel, and is blocked by the site firewall.
  - Seriously impacts performance for IPv6-enabled remote users.
  - They disable IPv6 to fix it (bad scenario)
- Solution:
  - Deploy ISATAP to Intranet. Works well!
  - But MACs don't have ISATAP client support.
    - Bug report filed with Apple
    - Already reported: original Bug ID# 4550554

Wrong tunnel metrics

- RFC 3484 specifies preference for choice of source address

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
2002::/16	30	2
::/96	20	3
::ffff:0:0/96	10	4

- Windows ends up with same metric for native and tunneled routes.
  - Systems often choose the wrong one to use, and end up tunneling when they have native IPv6 available.
- Workaround (for ISATAP):
  - Block ISATAP RA's to/from Native IPv6 subnets.

black-hole issue

- Found that connections were failing repeatedly over certain paths
  - Large packets dropped in transit
- After analysis:
  - Juniper router not sending the ICMP6 “too big” when packet larger than egress interface MTU
    - Broke path MTU discovery
  - Broken in most versions of JunOS, when one side is MPLS.
- Fix:
  - Upgrade to JunOS 9.3R3 or later

Expanding internal IPv6 adoption

- Jan 2009 – only 5% of our systems (servers, desktops, laptops, etc.) were doing IPv6
  - Double from the year before
- Today: A major internal campaign has us now at 93%.
  - A totally volunteer and optional effort
  - We had to provide encouragement and incentives for over 500 independent projects and systems administrators

Creating incentives

- We don't centrally control most customer devices (desktops, laptops, servers, printers, etc.)
  - Have to look for mechanisms to get these users to turn on IPv6 and use it
  - Modern Operating systems (Windows 7, MAC OSX, Linux) get IPv6-enabled automatically, but (for example) XP users need to "ipv6 install".
- For some servers, when their client base is mostly IPv6-enabled, we remove "A" record from DNS for that server.
  - The rest of the clients migrate quickly
  - Customers in environments lacking IPv6 generate local demand for fully IPv6 support

More incentives

- On some servers we just block IPv4 to specific services (HTTP, HTTPS), to encourage clients to migrate
  - Doesn't work well on servers with java applets, because then the java apps can't connect
    - See java problem reported earlier

Making progress visible within organizations – another incentive

Code	IPv6 Count	Non IPv6 Count	Total Count	IPv6(%)
23000	29	1	30	96.7%
40000	7	1	8	87.5%
41000	309	31	340	90.9%
42000	38	1	39	97.4%
43000	47	6	53	88.7%
53000	1444	98	1542	93.6%
55000	959	93	1052	91.2%
56000	804	61	865	92.9%
71000	582	33	615	94.6%
72000	459	36	495	92.7%
83000	19	1	20	95%
84000	9	0	9	100%
H0000	24	2	26	92.3%
H4000	73	1	74	98.6%
H5000	105	3	108	97.2%
TOTAL:	4908	368	5276	93%

Percentage of systems doing IPv6

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
- Older versions of MS Outlook
  - We tell users to upgrade to MS Office 2007
- Printers, and various odd devices
  - Too hard right now
  - For HP printers we are replacing the jetdirect cards with new ones that support IPv6

Google via IPv6

- We registered for Google AAAA's

```
$ dig www.google.com aaaa

;; ANSWER SECTION:
www.google.com.                 152199  IN      CNAME   www.l.google.com.
www.l.google.com.       90      IN      AAAA    2001:4860:0:2001::68
```

- Nobody noticed (good), until…
- When one site's web proxy broke IPv6, then people really noticed, and it got fixed quickly
  - Canary in a coal-mine effect

Google 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

Utilization comparison

IPv4 traffic

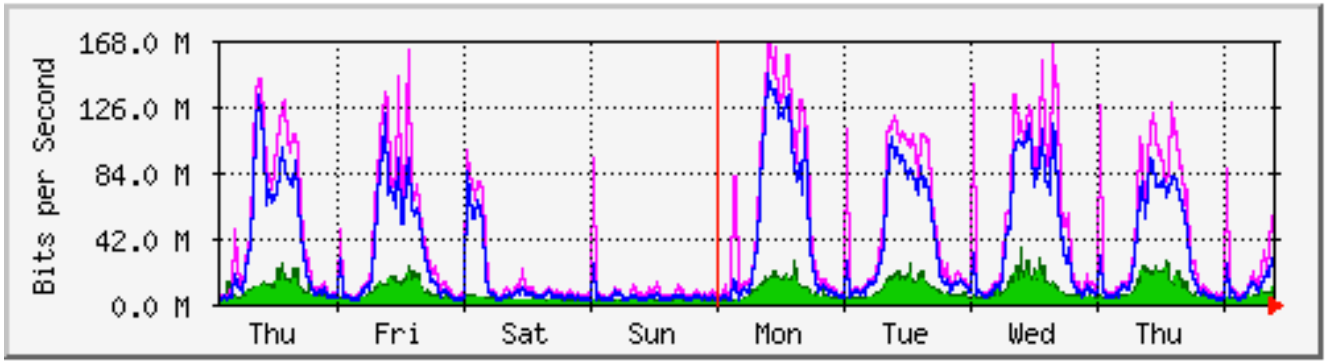

IPv6 traffic

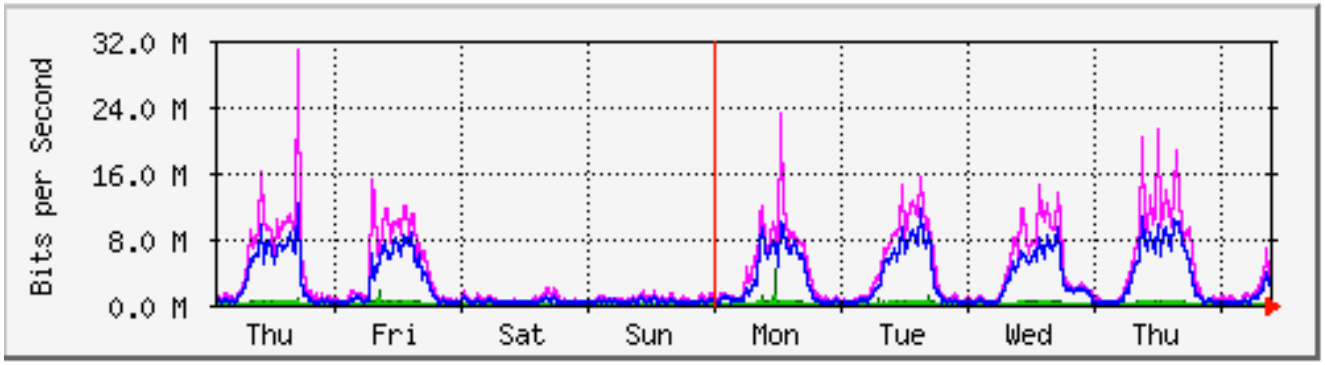

Almost 10% of traffic is IPv6

More challenges

- Maintaining all the new IPv6 addresses in DNS
- Large groups of systems that are under "configuration control", and can't be modified.
- Sys admins that are too busy with other priorities.
- Rogue 6to4 relays sending RAs
  - Windows systems with ICS enabled.
- Symantec Endpoint Protection (SEP) breaks IPv6
- Broken external DNS servers prevent some of our clients from running IPv6
- Blackberry Enterprise Services (BES) on IPv6-enabled Windows server will crash.

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)

DHCPv6 in MacOSX

- No DHCPv6 client in Mac OSX
- If you report the bug to Apple, they will say
  - "Already reported under original Bug ID# 3598535"
- Status from Apple:
  - "nothing new to report" and "engineering is still investigating"
  - "duplicate bug reports do raise visibility"
- and…
  - "is IPv6 support a big deal?"
  - "as of now there hasn't been a lot of movement towards IPv6 networks"
  - "is there a timeframe when you need this support?"
  - "does it work in XP?"
- Recommendation: everyone report this bug to Apple
  - (and ask for ISATAP too, while you're at it)

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 – we've been running it for a year
  - Yes, there are some additional complexities, and optimizations required, like garbage collection of temporary and privacy addresses.
- Lingering problems with IPv6 objects in the IP-MIB and IPV6-MIB
  - Many routers do not properly support RFC 4293 (version independent IP-MIB)

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.
- How To: (next slide)

Disabling privacy addresses

- Windows XP

```
ipv6 -p gpu UseTemporaryAddresses no
```

- Windows 2003

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

- Windows Vista

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

- Windows 2008, Windows 7

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

nga.mil

- Query to resolve www.extranet.nga.mil with AAAA returns RCODE 3, "no such name" (NXDOMAIN).
  - Windows XP will never do the "A" query
- If the name exists, even if no RRs for it, it should not return NXDomain.

```
09:29:59.312403 IP newiview.17577 > ins1.sd.domain:  17998+ [1au] AAAA? www.extranet.nga.mil. (49)
09:29:59.392933 IP ins1.sd.domain > newiview.17577:  17998 NXDomain 0/0/1 (49)

09:30:15.731744 IP newiview.11851 > ins1.sd.domain:  35028+ [1au] A? www.extranet.nga.mil. (49)
09:30:15.895239 IP ins1.sd.domain > newiview.11851:  35028 1/2/1 A 164.214.10.84 (105)
```

- Due to faulty behavior of Cisco CSS load balancer doing DNS functions
- Windows XP machines that are IPv6-enabled can't get to web site.
- See 4.2 of RFC 4074.

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

Windows patching

- We upgraded to Windows Software Update Service (WSUS) 3.0
  - supports IPv6
- All of our Windows patching now happens over IPv6

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

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.

IPv6 on Mgmt LAN

- Configuring an IPv6 address on mgmt 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 0xffffff (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
    - Delivered in 9 months

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
    - ServerIron – no IPv6 support except in very latest release
  - Juniper Netscreen:
    - SNMP via IPv6 not supported
  - Freeradius
    - No IPv6 support until 2.0

INM issues

- 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.

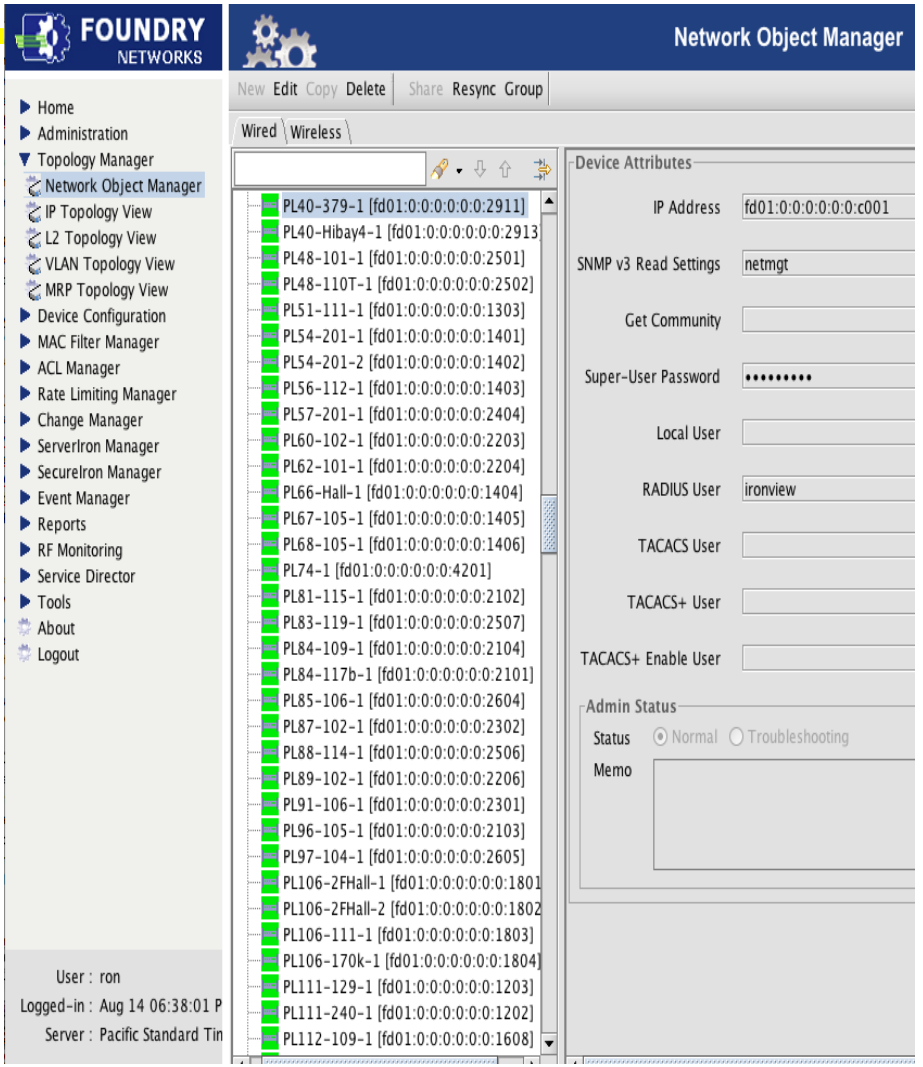

Syslog over IPv6

- syslogd (what's in most Linux distributions, including all RHEL releases) does not support IPv6.
- New "rsyslog" is a replacement
  - IPv6 support
  - Compatible with syslog config file (/etc/syslog.conf)
  - Fedora moved to this in Fedora 8
- Splunk doesn't care
- Converting all the clients out there to use IPv6 destination:
  - Foundry: OK
  - Cisco: OK
  - Netscreen Firewalls and VPNs: OK
  - Aruba (wireless): NO
  - Ascend (dialup): NO
  - Bluecoat (proxy): NO
  - New microwave radio: NO

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 DREN 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

Summary

- Biggest issue right now is lack of feature parity in most products
  - IPv6 features < IPv4 features
- Highest priority for all organizations is to IPv6-enable all public-facing services
- Dual stack works well today as a transition mechanism
- Still much work to do before IPv4 can be turned off anywhere