DREN IPv6 Implementation Update

Internet2 Joint Techs, Summer 2009
21 July, 2009
Indianapolis, IN

Ron Broersma
DREN Chief Engineer
High Performance Computing Modernization Program
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)
- These are production networks with 10's of thousands of users and systems.
  - i.e., not just a testbed
- Goals
  - See what works and what's broken
  - See what's missing
  - Share lessons learned

IPv6 deployment mostly complete

- ✔ 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.

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

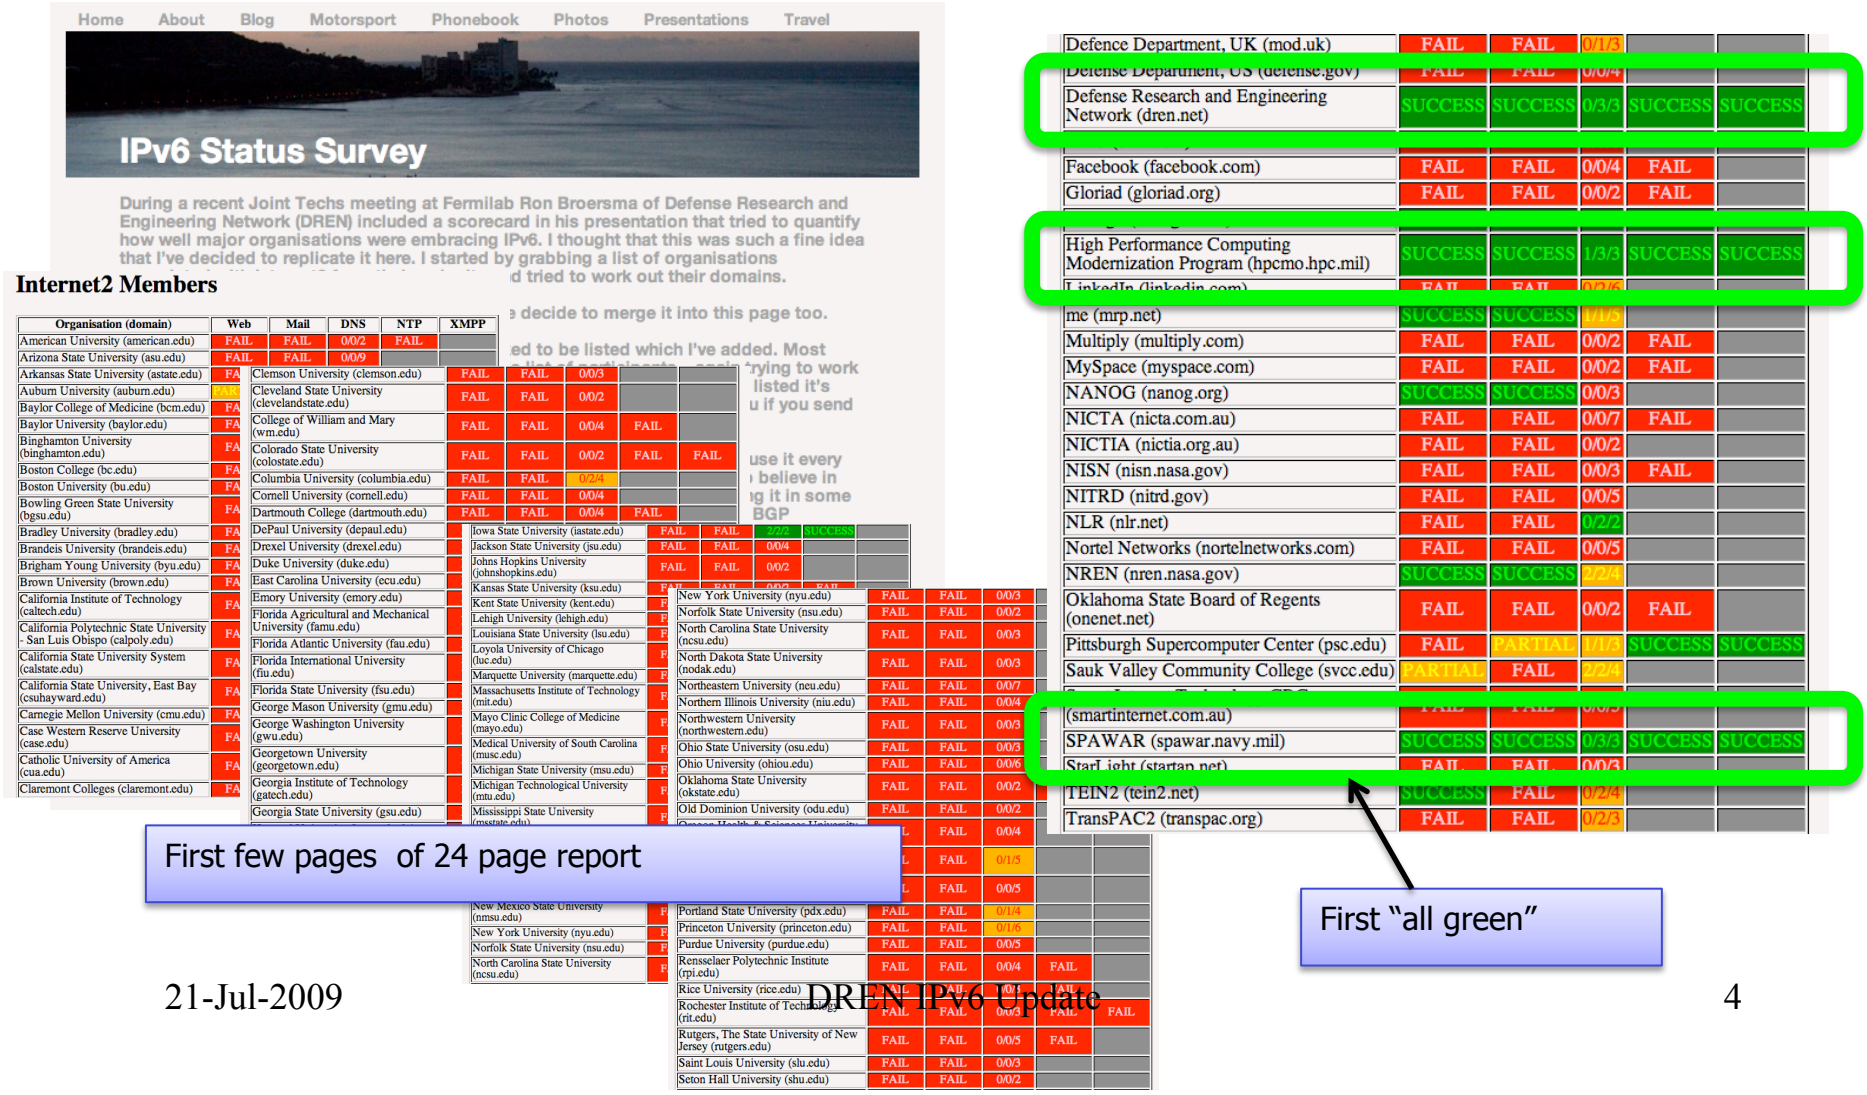

Home About Blog Motorsport Phonebook Photos Presentations Travel

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)					
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)					
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		
		FAIL	0/0/4		
		FAIL	0/1/5		
		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)					
Rochester Institute of Technology (rit.edu)	FAIL	FAIL	0/0/2	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 few pages of 24 page report

First "all green"

Previously discussed…

- Continuing issues
  - Vendors still not eating their own dogfood
  - Big problem is lack of feature parity between v4 and v6 in all products
- Reported at Texas meeting:
  - Trying to make management LAN IPv6-only
  - All DNS zone-xfers via IPv6
  - DNS updates from DHPC via IPv6 (failed)
  - Google via IPv6

Some current initiatives

- Try to move some things to IPv6-only:
  - Network Management
  - VTC network
  - DNS zone xfer and dynamic updates
  - AAA (RADIUS, LDAP, Kerberos)
- Broaden Google support to all DREN customers
- Fixing remote users on VPNs
- Deployment of IPv6-capable IPS
- IPv6 RRs in DNS, automatically
- v6-enable all systems (servers, desktops, etc.)
- Incentives for customers

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

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.

Keeping DNS updated

- Need to get 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

- Scheme #1
  - Cook all the PTRs based on MAC addrs
    - We can find the MAC addrs of all devices on-net
    - We can determine their IPv6 prefix based on their IPv4 address
    - Assume bottom half of address (IID) is EUI-64
    - Generate PTRs for everything, whether doing IPv6 or not.
  - It works, except for all those windows machines that prefer temporary (privacy) addresses
  - Disabling privacy addresses everywhere is a harder problem.

DNS auto-update

- Scheme #2
  - 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 well
  - Yes, there are some additional complexities, and optimizations required, like garbage collection of those temporary addresses.
  - Plan to release tool at some point, after more testing and cleanup.
- Lingering problems with IPv6 objects in the IP-MIB and IPV6-MIB
  - We really need all routers supporting 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

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

speaking of DHCPv6…

- 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 as of this morning:
  - "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 now!
  - (and ask for ISATAP too, while you're at it)

New 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 33%, increasing about 1% per week at the moment.
  - A totally volunteer and optional effort
  - Some Departments are over 80%.
  - Still going too slow
- Goal: 95% by Oct 1, 2009.

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 (Vista, 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
- Latest idea:
  - Block IPv4 access to www.google.com.
  - Plan to start phasing that in next week, site by site.

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

Strong recommendations

- Need to "eat our own dog-food"
  - Vendors selling "IPv6-capable" products
  - We in the I2 community who are advocates
- Need to NOT be afraid to "break some glass"
  - Attempt more deployment to production networks
  - Do more "challenges" and "hour of IPv6-only"
- File bug reports and enhancement requests
- IPv6-enable all public facing services

Things that are breaking us

- DoD DMZ effort
  - New DNS proxys don't support IPv6
  - In fact, none of it supports IPv6
- Efforts to consolidate Navy public web sites to DISA hosting centers
  - We will lose IPv6 for our public web sites, which is unacceptable.

"Our national interest is at stake in this issue"

A national strategic approach

- IPv6-enable all public facing services
  - The stuff outside your firewall
- It is fairly easy to do (low hanging fruit)
  - You don't have to worry about your internal network just yet.
  - Eliminates prerequisite of IPv6-enabling your whole security stack
- Benefit to IPv6 community is huge
  - External parties generally only care about your public facing services, not your internal stuff.
  - IPv6-only clients can then get to you native, without address translators (NATs, CGN, IVI, etc.).
- DoD or OMB should incentivize something like this
  - Simple concept, fairly easy to do, substantial benefit on many fronts.
  - Update (May 2009): Federal CIO IPv6 Planning Guidance
    - "Quick Wins", 5.1.2 "External Facing Servers" by 2010

END