Enterprise IPv6 Deployment in the US Government

7th Slo IPv6 Summit
18 October, 2012
Ljubljana, Slovenia

Ron Broersma
DREN Chief Engineer
SPAWAR Network Security Manager
Federal IPv6 Task Force
ron@spawar.navy.mil

Slovenia is #1

Rank	Country	Sample	Green
1	Slovenia	50	36.0% (18)
2	Czech Republic	50	30.0% (15)
3	Brazil	50	30.0% (15)
4	Singapore	50	22.0% (11)
5	United States of America	50	20.0% (10)
6	Netherlands	50	18.0% (9)
7	Indonesia	50	18.0% (9)
8	Germany	50	16.0% (8)

Source: Eric Vyncke, http://www.vyncke.org/ipv6status/

U.S. Federal Mandate

- Signed by U.S. CIO, Sept 28, 2010
  - By Sept 2012, all public content IPv6-enabled
  - By Sept 2014, all internal systems dual-stack
- Previous OMB mandate
  - everything “IPv6 capable” by June 2008
  - Success(?): Everyone did a “ping6”, and then turned if off. ☹
- “Federal IPv6 Task Force”
  - team working to make it happen
  - transition managers assigned in every agency

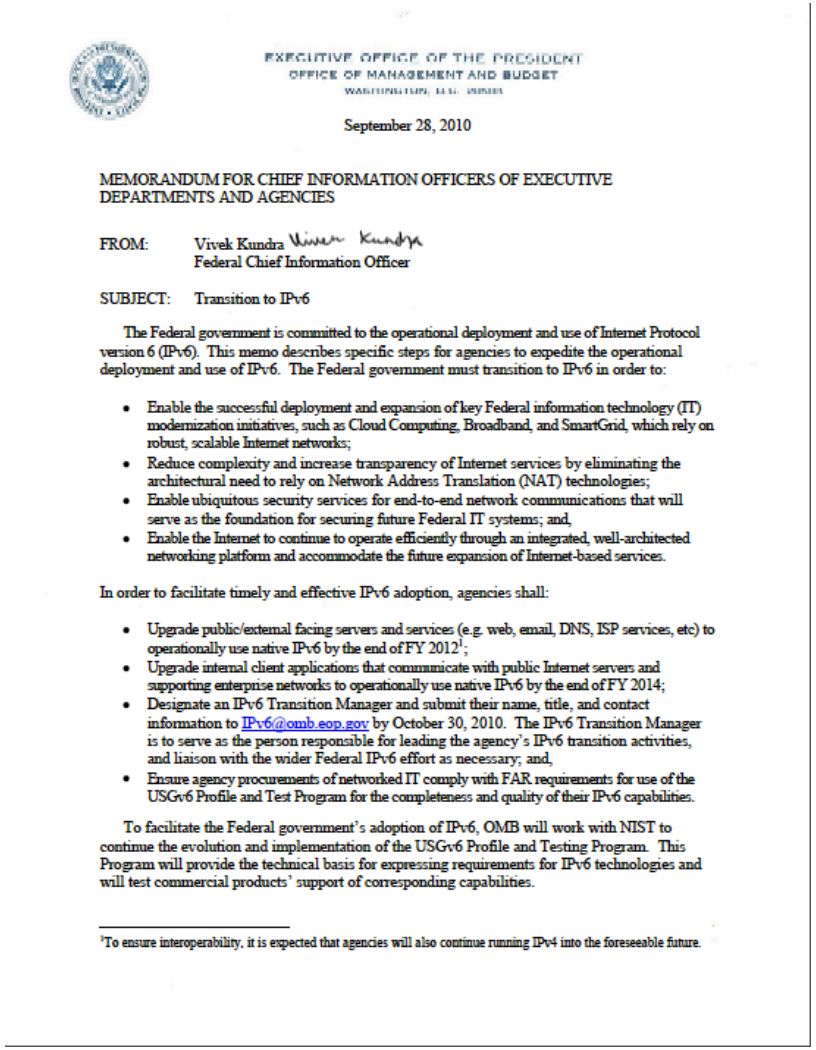

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF MANAGEMENT AND BUDGET
WASHINGTON, D.C. 20503

September 28, 2010

MEMORANDUM FOR CHIEF INFORMATION OFFICERS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: Vivek Kundra
Federal Chief Information Officer

SUBJECT: Transition to IPv6

The Federal government is committed to the operational deployment and use of Internet Protocol version 6 (IPv6). This memo describes specific steps for agencies to expedite the operational deployment and use of IPv6. The Federal government must transition to IPv6 in order to:

- Enable the successful deployment and expansion of key Federal information technology (IT) modernization initiatives, such as Cloud Computing, Broadband, and SmartGrid, which rely on robust, scalable Internet networks;
- Reduce complexity and increase transparency of Internet services by eliminating the architectural need to rely on Network Address Translation (NAT) technologies;
- Enable ubiquitous security services for end-to-end network communications that will serve as the foundation for securing future Federal IT systems; and,
- Enable the Internet to continue to operate efficiently through an integrated, well-architected networking platform and accommodate the future expansion of Internet-based services.

In order to facilitate timely and effective IPv6 adoption, agencies shall:

- Upgrade public/external facing servers and services (e.g. web, email, DNS, ISP services, etc) to operationally use native IPv6 by the end of FY 2012[1];
- Upgrade internal client applications that communicate with public Internet servers and supporting enterprise networks to operationally use native IPv6 by the end of FY 2014;
- Designate an IPv6 Transition Manager and submit their name, title, and contact information to IPv6@omb.eop.gov by October 30, 2010. The IPv6 Transition Manager is to serve as the person responsible for leading the agency's IPv6 transition activities, and liaison with the wider Federal IPv6 effort as necessary; and,
- Ensure agency procurements of networked IT comply with FAR requirements for use of the USGv6 Profile and Test Program for the completeness and quality of their IPv6 capabilities.

To facilitate the Federal government's adoption of IPv6, OMB will work with NIST to continue the evolution and implementation of the USGv6 Profile and Testing Program. This Program will provide the technical basis for expressing requirements for IPv6 technologies and will test commercial products' support of corresponding capabilities.

[1]To ensure interoperability, it is expected that agencies will also continue running IPv4 into the foreseeable future.

Status

- NIST IPv6 Deployment Monitor
  http://fedv6-deployment.antd.nist.gov/

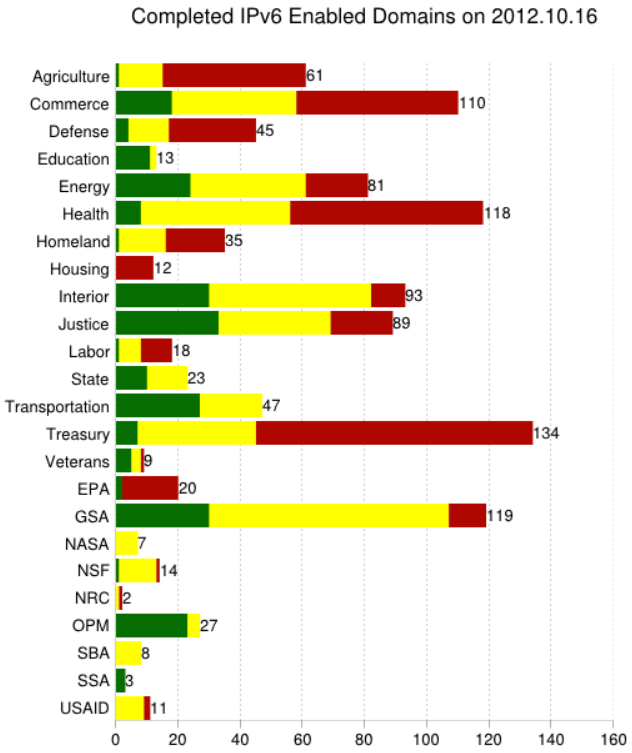

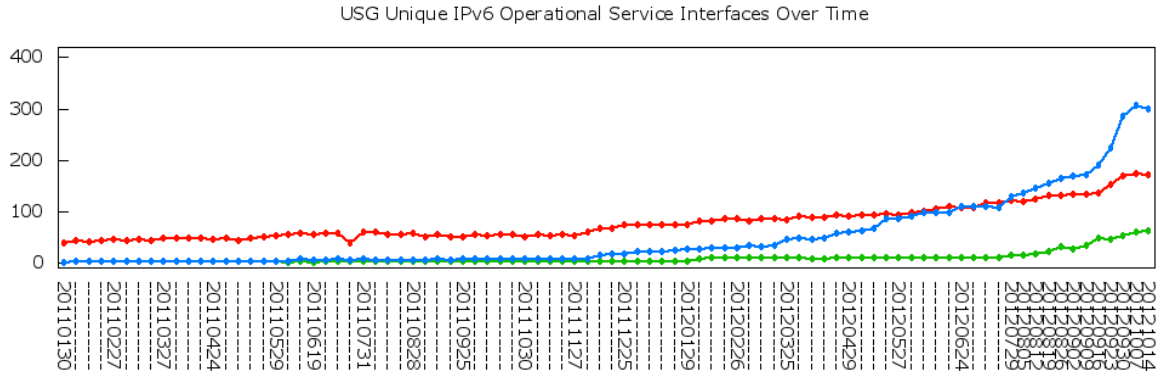

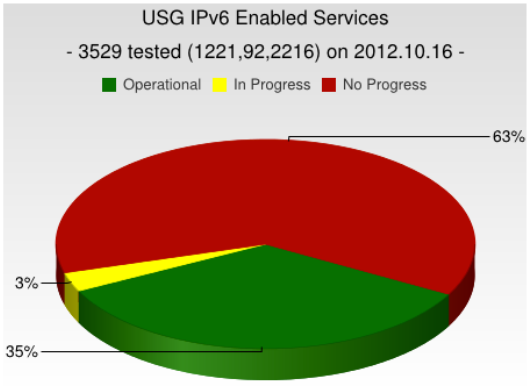

Observations and Questions

- Why did much of the change came right before the deadline?
- If these metrics show only 35% completion, does this indicate a failure to meet the goal?
- After the Sept 2012 deadline, what incentive is there to…
  - leave things turned on
  - continue making progress on the other 65%

Success?

- Yes, this was a big success:
  - A significant increase in demand signal from the U.S. Government to industry, to deliver IPv6 services
    - much harder to ignore us, or give low priority to our requirements
    - explodes the myth that “nobody is asking for IPv6”
  - A huge increase in IPv6 awareness in the Government agencies
    - people holding workshops, getting training, working with their providers, etc.

Success?

- A lot of public Government content is becoming IPv6-enabled, as part of the World goal to IPv6-enable the entire public Internet
  - being the solution, rather than the problem
  - setting an example and paving the way for the rest of the public sector
- This hopefully incentivizes other countries to do something similar

Looking forward

- What is the incentive to keep the pressure on after the deadline?
  - We plan to not allow .gov domains to be renewed if that organization has not met the mandates for IPv6 (and maybe DNSSEC as well).
- Other Governments and organizations should consider similar incentives

Keys to success

- Clear simple achievable vision and mandate, with deadlines, from the top (CIO)
- Responsibility, accountability and authority established and managed at the executive level
- Public reporting of status along the way, both internally and externally
- Bring in experts that have IPv6 operational experience, if you don't have it organically in your organization.
  - (there are few experts available; check with the industry to ensure who you bring in can provide what is needed)
- Early (and consistent) interaction with service and technology providers, to communicate requirements.
  - and be willing to switch providers to acquire IPv6 support
- Dual-stack support from ISP(s)

Challenges Experienced

- Issues (for 2012 mandate)
  - Certain ISPs cannot deliver IPv6 support in time
  - TIC, MTIPS not ready for IPv6
  - Some existing security products lack IPv6 support
  - CDNs weren't ready (in the beginning)
  - Large bureaucracies move very slowly, and many have outsourced their IT expertise
  - Transition planning is happening without IPv6 operational experience.
    - impacts things like addressing plans
  - Contracts for "Managed Services" in legacy status cannot be changed without huge cost and schedule impact
  - Guidance and oversight from departments to subordinate agencies lacked IPv6 operational experience from an enterprise perspective

Challenges Experienced

- Issues (2014)
  - Certain larger enterprises are having difficulty in scoping the 2014 objective effort
  - 2014 guidance and oversight from departments to subordinate agencies lacked IPv6 operational experience from an enterprise perspective
  - Certain Carriers' MPLS networks providing WAN managed services to federal agencies will not support IPv6 in 2013/2014
  - Certain enterprises have not established IT Asset configuration control re IPv6
  - Information security engineers do not have the IPv6 knowledge to support the creation of the required 2014 objective architecture
  - Certain departments did not inform their bureaus of the FAR acquisition criteria or Enterprise Architecture
  - FAR acquisition criteria does not contain enforcement clauses

A note on Akamai

- When we first tried to IPv6-enable some large public web sites, there were two major showstoppers
  - existing load balancers that didn't support IPv6
  - content hosted at Akamai, which wasn't IPv6-ready
- Good news:
  - after major efforts on the part of Akamai, many of our Akamai-hosted properties are now being IPv6-enabled.
    - need to "opt-in", but there are no additional charges (for public sector)
    - new customers will get dual-stack right away!
- Bad news:
  - non public-sector has to pay extra to get IPv6 support
  - public sector may have to start paying extra after Sept.
  - from an IPv6-only environment, you must use a dual-stack recursive DNS server, because their internal DNS is not IPv6-enabled

War Stories

- “Don’t ask us for what we can’t deliver”
- “I tried this on my home computer, so I know it is good for the enterprise”
- “Security manager says that I need to enumerate all hosts by scanning subnet”
- an Intelligence agency story
- “don’t listen to this guy”

Addressing Plans

- Common mistakes
  - Doing other than /64 for subnets
    - Didn't read RFC 4291 nor 5375
  - Thinking that the addressing plan has to be perfect the first time
    - because you "believe" you can't afford to re-address
  - Choosing allocations for sites based on size of site
    - because /48 for all sites is too wasteful
  - Justification "upwards", instead of pre-allocation "downwards"
  - Host-centric allocation instead of subnet-centric

Addressing Plans

- Without sufficient operational experience with IPv6 deployment, you WILL get it wrong at first.
  - usually takes the 3rd time to get it right
- Planners are hindered by IPv4-thinking
  - being conservative with address space
  - thinking “hosts” instead of “subnets”

Making the paradigm shift

- You may be un-qualified to develop an IPv6 addressing plan if you think:
  - /64 for subnets is wasteful
  - /64 for point-to-point links is wasteful
  - /48 for small sites is wasteful

Updates, Observations, and other News…

World IPv6 Launch

- Since the SPAWAR enterprise network (AS 22) is 100% dual-stack, how would network utilization (traffic inbound from the Internet) be impacted by an event like this?
- Previously (5 min averages, daytime):

  1% in 2009 before Google whitelisting
  2.5% after Google whitelisting
  just under 10% when Youtube was added (late Jan 2010)
  around 15% after World IPv6 Day (June 2011)

- After World IPv6 Launch

  around 20% average during the day

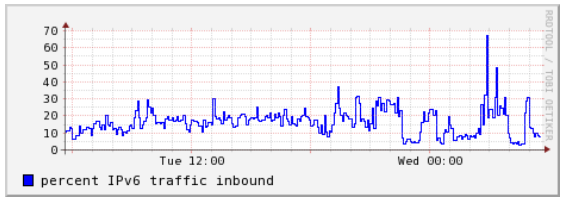

- Another view: overall daily average of traffic:

  Before: range (workdays) : 11-14%
  After: 14-18%

Top Enterprise Deployment Challenges

- Lack of IPv6/IPv4 feature parity
  - taking way too long to get there
- Vendors not eating own dogfood
  - but this is starting to change
- Rogue RAs due to Windows ICS
  - set router priority to “high” as workaround
- Privacy Addresses (RFC4941) break address stability
  - no easy way to centrally disable
- Lack of DHCPv6 client support in older OS’s
- Network Management over IPv6 not quite there
- Operational Complexity with dual-stack

Configuring addresses: Did we break it along the way?

- Enterprise requirement: stable, deterministic addresses, dynamically assigned, working in a heterogeneous environment.
  - "plug 'n play", centrally managed
- SLAAC
  - not perfect, if you were hoping do things the DHCPv4 way, but works quite well.
    - except for those pesky "Rogue RAs"
  - about the only choice when so many devices don't have DHCPv6 client support
  - FAST!!

Configuring addresses

- Privacy extensions (RFC 4941) make SLAAC less useful for enterprise environments.
  - privacy/temporary addresses, enabled by default in Windows, and now appearing in other major OS's.
  - we lost stability and predictability
  - we have to monitor and log all address usage, and build new correlation and search tools
  - if we need to disable privacy addresses, you have to manually configure that on the hosts
    - and this breaks "plug 'n play"
    - no mechanism for the network to disable this behavior in the clients
      - but I wish there was
- So lets try DHCPv6…

Configuring addresses

- Lets try DHCPv6…
  - soon those Windows XP and other machines with no DHCPv6 client will go away, we hope.
  - When you enable DHCPv6, clients can now get an address that you assign centrally
    - we get stability and predictability back
  - But, unless you disable the “A” bit in the RA prefix announcement, the clients still get SLAAC addresses, and privacy addresses
    - and seem to prefer those addresses for sourcing traffic, rather than the DHCPv6 address
    - but if you disable it, then hosts without DHCPv6 clients are dead.
  - And worse, there’s now this DUID thing
    - you can’t control assignment based on MAC address any more
    - to use DUID, you have to get the DUID from the clients somehow
      - back to manual processes
    - and all your cloned devices (very common in an enterprise) all have the same DUID, unless you manually reset it

Configuring addresses

- What's the solution?
  - Microsoft says to eliminate all non-Windows machines, and use Active Directory, and then set up your GPO to disable privacy addresses
  - Some suggest "learn to live with privacy addresses"
  - Others suggest "Take it to the IETF"
  - Hack your DHCPv6 server to pull MAC addresses from the DUID (mostly works, but risky)
  - Wait for the dhcpv6-relay to pass along the source MAC address
  - Other suggestions are welcome.

Playing with IPv6-only environments

Management over IPv6 in some products

Previously (June '2011):

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP
Cisco										
Brocade										
Juniper										

Now:

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP	IPv6 MTU	No v4
Cisco[3]								6				
Brocade[1]											9	
Juniper								5				
ALU								4				
A10							8	7				

1. Can't reboot using SNMP over IPv6
2. .
3. 15.2(2)TR
4. 10.0R6 (Nov 2012)
5. 12.3R1 Nov 2012 (beta in August)
6. ASR1K:3.7S (July 2012)
7. 3.0 release, 2012Q4
8. No plans
9. fix planned for Apr 2013

IPv6-only bug (recently fixed)

- when disabling IPv4 on Brocade FESX switches, they start responding to all ip-subnet-broadcasts, and start ARPing (from 0.0.0.0), and other strange behaviors.

```
11:27:14.103150 00:0c:db:6b:73:c0 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103152 00:12:f2:32:62:80 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103184 00:0c:db:9d:43:00 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103186 00:12:f2:32:63:c0 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103238 00:12:f2:02:83:40 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103275 00:0c:db:6f:7b:40 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103308 00:12:f2:32:56:80 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103343 00:12:f2:32:5a:40 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103385 00:0c:db:69:a8:00 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103424 00:0c:db:c8:61:80 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.103457 00:12:f2:32:56:c0 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.104042 00:0c:db:9d:3b:00 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.104076 00:12:f2:8d:41:40 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.104205 00:12:f2:32:82:00 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.105807 00:12:f2:33:0e:00 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.105840 00:12:f2:32:ad:c0 > Broadcast, ethertype 802.1Q (0x8100), length 64: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.106956 00:0c:db:69:c8:a0 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.109253 00:0c:db:6f:a9:80 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
11:27:14.109291 00:0c:db:c8:87:c0 > Broadcast, ethertype 802.1Q (0x8100), length 60: vlan 2, p 3, ethertype ARP, arp who-has 128.49.9.249 tell 0.0.0.0
```

Other IPv6-only tests

- Test environment:
  - network with ONLY IPv6 turned on (no IPv4 configuration or routing)
    - “A” bit enabled (SLAAC)
    - “M” and “O” enabled (for DHCPv6)
  - Many operating systems connected, to see how they behave
    - Windows7, MacOSX, Linux (multiple distributions), FreeBSD
    - iPhone, iPad, Android
- Anything without a dhcpv6-client won’t get DNS addresses
  - Windows XP, MacOSX before Lion, Android

IPv6-only

- Observation (Lion):
  - You can browse OK with Safari, but Chrome and Firefox hang when trying to browse to IPv6-only web sites
    - happy-eyeballs not working
  - tcpdump shows it ARPing for Internet addresses
  - … because there is a default-route-to-interface installed in the routing table
  - … because it assigns IPv4 link-local (RFC 3927) and implements “ARP for everything” (paragraph 2.6.2)
  - … so it “thinks” it has full IPv4-internet reachability (unlike IPv6 behavior)
- Most other OS’s exhibit similar behavior
- Need to fix happy-eyeballs
- workaround: actually assign IPv4 addresses (like maybe from 100.64/10 space) with default IPv4 route, but have router respond to everything as net/host-unreachable.
  - or just disable IPv4 on the OS (Lion has a knob to do this).

END

Contact me:
ron@spawar.navy.mil