DREN III
IPv6 Lessons Learned

Ron Broersma

DREN Chief Engineer

Background

- DREN III
  - The 4th generation of DREN
  - prime contract awarded to CenturyLink (Qwest)
    - Dec 2012
  - 18 months to build, test, accept network, and migrate all customers from DREN2
    - finished June 2014
- IPv6 requirements were clearly specified in the DREN III acquisition and contract.

Acquiring IPv6 capable products

- Vendors will say that their products support IPv6, or are IPv6-capable.
  - This means nothing.
- Lessons learned over last decade:
  - All products lack IPv4/IPv6 feature parity
  - Vendors aren't "eating their own dogfood"
  - IPv6 bugs and missing features do not get resolved unless the company has a strong corporate commitment to IPv6, or there is airtight contractual language that requires it.

Doesn’t the FAR and USGv6 help?

- The FAR and USGv6 profile are very important, but not sufficient to address the issue.
  - Contacting officers don’t enforce FAR IPv6 policy.
  - The FAR IPv6 policy apparently applies to equipment and systems, not “services”.
- Unless the contractual requirements specification directly addresses IPv6, and the provider is held to those requirements, you will not get fully capable IPv6 products and services.

DREN IPv6 contractual requirements

- DREN III is an IPv6 network, with legacy support for IPv6.
  - Establish the vision
- IPv6 must work as good as or better than IPv4.
  - this is measurable, and enforceable
- Must not deploy anything in the network that does not comply with this requirement.
  - non-compliant components can be rejected
- All network management functions are IPv6-only (no IPv4).
  - no cheating

Actual language from PWS

DREN III RFP (Jan 2011)

"DREN is identified as an IPv6 network with IPv4 legacy support. Therefore, all systems, software, and equipment supporting the DREN network and its services shall handle IPv6 in an equivalent or better way than current IPv4 capabilities, performance, and security. No systems, software, or equipment shall be deployed on the DREN that does not meet this requirement. Additionally, all network management shall be enabled using IPv6."

Results

- During test and acceptance, most testing was done using IPv6, but comparisons to IPv4 were made.
  - IPv6 had equivalent performance
- When the NOC was built, we asked for a listing of all addresses used on all devices.
  - we pointed out every use of IPv4 and asked CTL to remove it, or explain why it was there.
  - all IPv4 addresses were removed from the management interface of all SDPs (site routers).
  - IPv4 addresses removed from the Spirent Test centers.
- RSA SecurID – clients did not support IPv6, but this authentication was used almost everywhere, causing most systems to require IPv4.
  - RSA could not commit to a near term fix, so this product will be replaced.
- InfoVista – could not support netflow capture over IPv6
  - replaced with SevOne.
- ALU Service Aware Manager (SAM) has multiple IPv4 dependencies (replication, SevOne interaction, Oracle)
  - unresolved
- HP iLO – NTP and Syslog are IPv4 only
  - fixed in later release
- ALU 7750 – NTP over IPv6 not supported
  - fixed in newer version of code
- Cisco 2960 switches – no SNMP over IPv6
  - replaced with Juniper switches
- Juniper MX, SRX, EX – TACACS did not work over IPv6
  - fixed in MX, awaiting fixes for others.
- Infoblox grid only worked over IPv4
  - fixed
- Perle IOLAN – PPP only works over IPv4.
  - fixed

Other observations

- DREN III requires that all customers connect with dual-stack (IPv4 + IPv6), run BGP, support jumbo frames, and support 802.1q “tagging”.
- We thought the big problem for some customers was going to be routers and other CPE that didn’t support IPv6.
- Surprisingly, IPv6 was supported in all customer products we interfaced with.
  - even if those customers didn’t care about IPv6, nor had ever tried to make their network support IPv6, nor tried to purchase IPv6-capable products.
- Lesson: mainstream products have basic IPv6 support today.

Observations

- We encountered significant apathy towards IPv6 at many customer locations.
  - Most had no knowledge of the Federal or DoD mandates.
  - Without some incentive, IPv6 does not get turned on.
  - Too many other priorities, and no perceived benefit.
- But there was no resistance to enabling it on the CPE router.
  - a few tried to quote the STIGs or other policies that appear to say IPv6 must be disabled.
- Lesson: we need to provide more incentives.

Evaluating new products

- Our #1 rule:
  - if we can't get to the company or product website via IPv6, we won't consider such products.
- Why this hard line?
  - we learned the hard way that without strong corporate commitment to IPv6 support, it will take forever to get IPv6 bugs fixed or features added.
  - we learned that the corporate website being IPv6-enabled was a good indicator of corporate commitment to IPv6.
  - this has been tested many times, and it works.
  - in the process, we encourage industry to IPv6-enable their public facing services.
- Examples
- #2 rule:
  - Verify "eating your own dogfood"
  - Test product in production IPv6 network.

IPv6 Today

- Basic IPv6 capabilities are in all mainstream products today.
- Don’t deal with companies that lack strong corporate commitment to IPv6, or aren’t eating their own dogfood.
- Look for opportunities to provide incentives to expand IPv6 deployment, in your customers and in your suppliers.

Even my TV does IPv6

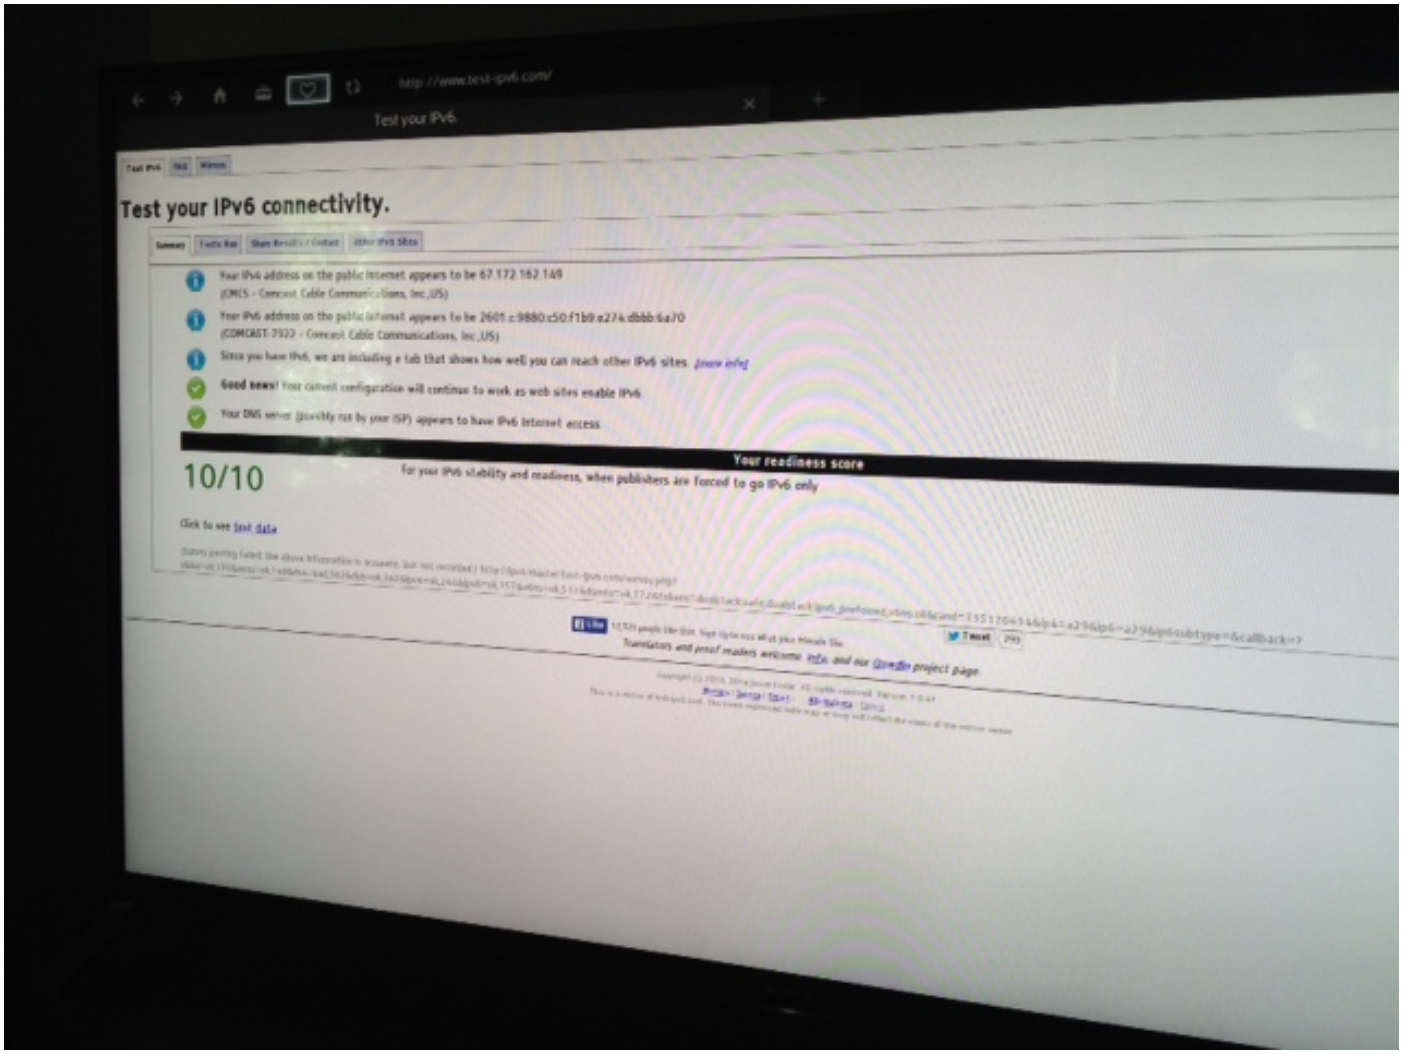

IPv6 in enterprise networks

a note on DHPCv6

IP address assignment to devices

- What is required in many enterprises:
  - control of what IP address gets assigned to a device.
  - stable addressing (always the same IP address)
  - addresses in DNS (PTR, A, AAAA)
  - auto-configuration (the device gets its address and other configuration details from the network).

IPv6 address assignment

- In the beginning – SLAAC
  - autoconfiguration, stable, predictable
  - some effort required to get it all into DNS, but doable.
- Privacy addresses broke everything
  - loss of stability and predictability and control.
- Lets try DHCPv6
  - hopefully has same functionality as DHCP(v4)

DHCPv6

- can't make fixed IPv6 address assignment based on MAC address like in DHCP
  - all you get is a "DUID" (DHCP Unique ID)
- DUID problems:
  - How do you get all DUIDs registered?
  - Many DUIDs are NOT unique!
- Extracting MAC addresses from DUIDs
  - DUID-LLT and DUID-LL have embedded MACs
  - ISC DHCPv6 has hooks to extract MAC address and make it work like DHCP(v4)
  - But Solaris always sets the last 48 bits to zero.
  - And there is other strangeness from various devices, so doesn't always work.
- Local hacks to ISC DHCPv6
  - Extract MAC address from DUID-EN (for HP Printers)
  - If we still can't find MAC address in our database, extract MAC address from IPv6 link local address of source, sent by the DHCPv6 relay.

DHCPv6 packet (via relay)

13:00:30 dhcpd: Relay-forward message from 2001:480:10:4::1 port 547, link address 2001:480:10:104::1, peer address fe80::f6ce:46ff:fe49:3176
13:00:30 dhcpd: DUID = 0:2:0:0:0:b:f4:ce:46:49:31:76

HP Printer uses DUID-EN, but sends the MAC address anyway

12:54:09 dhcpd: Relay-forward message from 2001:480:10:3::15 port 547, link address 2001:480:10:64::1, peer address fe80::226:b9ff:fe7c:a356
12:54:09 dhcpd: DUID = 0:1:0:1:1b:d:44:eb:0:26:b9:7c:a3:56

MAC Address: 00:26:b9:7c:a3:56

Long term resolution

- RFC 6939
  - "Client Link-Layer Address option in DHCPv6"
  - DHCPv6 Relay-Forward messages will provide the link-layer (MAC) address to the server.
  - We will see this implemented in some routers by middle of next year.
  - Then we can finally reliably assign IPv6 addresses based on registered MAC address, like we did in IPv4.

DHCPv6 dynamic DNS updates

- Not possible to make it work like in IPv4
  - "hostname" option not implemented in IPv6
  - instead, "client fqdn" option (RFC 4704)
    - but that isn't implemented widely
- Adding "ddns-hostname" in each host declaration makes it work.
- But now the ddns updates clash with the IPv4 ddns updates
  - uses a DHCID RR to detect clashes (RFC 4703)
    - IPv6 DHCID based on DUID, but IPv4 DHCID is not.
    - So, DHCIDs are different, and will conflict.
  - Fix: set "update-conflict-detection false".

End

Contact me:
ron@dren.hpc.mil