IPv6: Enterprise Addressing Plans

Prepared for the IPv6 Business Information Exchange

Dr. Peter J. Welcher,
Chesapeake NetCraftsmen

About the Speaker

- **Dr. Pete Welcher**
  - **Cisco CCIE #1773, CCSI #94014, CCIP**
  - **Specialities: Large Network Design, Multicast, QoS, MPLS, Wireless, Large-Scale Routing & Switching, High Availability, Management of Networks**
  - **Customers include large enterprises, federal agencies, hospitals, universities, cell phone provider**
  - **Taught many of the Cisco router/switch courses**
  - **Reviewer for many Cisco Press books, book proposals**
  - **Designed and reviewed revisions to the Cisco DESGN and ARCH courses**
  - **Presented lab session on MPLS VPN Configuration at Networkers 2005-2007; presenting on BGP at Cisco Live 2008-2009**
- **Over 150 articles plus blogs at http://www.netcraftsmen.net**

Agenda

- Quick Review of IPv6 Addressing
- Quick Review of Some IPv4 Addressing Techniques
- Choosing Prefix(es)
- Addressing Plan
  - IPv4 Mapping to IPv6
  - New Plan
- Summary, Q&A, References

IPv6 Addressing Review

- Defined in RFC 4291
- 128 bit addresses, written as 8 blocks of hexadecimal digits separated by colons (:)
- Leading zeros in any of the 8 blocks omitted
- Any <u>one</u> block of consecutive zeros can be replaced by double colon (::)

IPv6 Addressing Review – 2

Type of IPv6 address	Prefix
Global unicast addresses	**2000::/3**
Site-local unicast addresses **(deprecated)**	FEC0::/10
Unique local unicast addresses **(replaces site-local)**	FC00::/8 (managed), FD00::/8 (random)
Link-local unicast addresses	**FE80::/10**
IPv4-mapped IPv6 address (96 0's then the hex form of the IPv4 address)	::pqrs:tuvw, also written as ::M.N.P.Q
IPv4 compatible IPv6 address **(somewhat deprecated)** (80 0's then FFFF then the IPv4 address)	0:0:0:0:0:**FFFF**:IPv4 address, also written as ::FFFF:M.N.P.Q
ISATAP tunnel	FE80:0000:0000:0000:0000:5EFE:IPv4 address
Unspecified address	0:0:0:0:0:0:0:0 or 0::0 or ::/128
Loopback	0:0:0:0:0:0:0:1 or ::1

IPv6 Addressing Review

- **Format of an IPv6 address (in hexadecimal)**

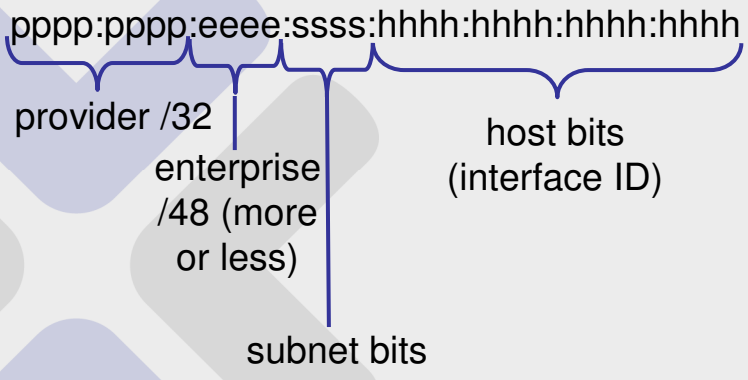

Quirks of IPv6 Addressing

- Get used to huge subnets
  - 64 bits of host
  - EUI-64 mildly wasteful (perception of "plenty" leads to wastage?!)
  - Known vs. random end station addresses
  - Known vs. random router addresses

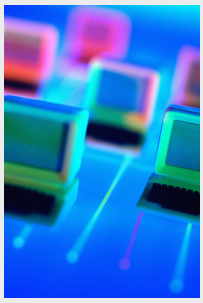

IPv6 Addressing -- Original concepts:

- IPv6 globally summarizable routing
  - Minimize global routing table growth
  - Stay within performance regime of current / likely router hardware (compared to driving costs way up)
- /32 assigned to provider
- /48 or smaller block assigned by provider (***)
  - That provides 16 bits for subnetting

Agenda

- Quick Review of IPv6 Addressing
- Quick Review of Some IPv4 Addressing Techniques
- Choosing Prefix(es)
- Addressing Plan
  - IPv4 Mapping to IPv6
  - New Plan
- Summary, Q&A, References

Advanced IPv4 Addressing: Bit Assignments

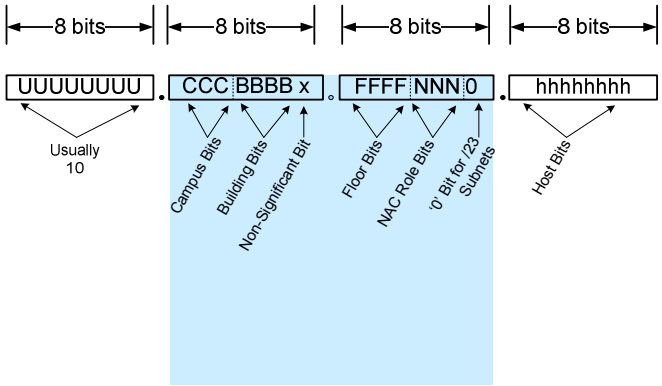

Advanced IPv4 Addressing: Bit Assignments

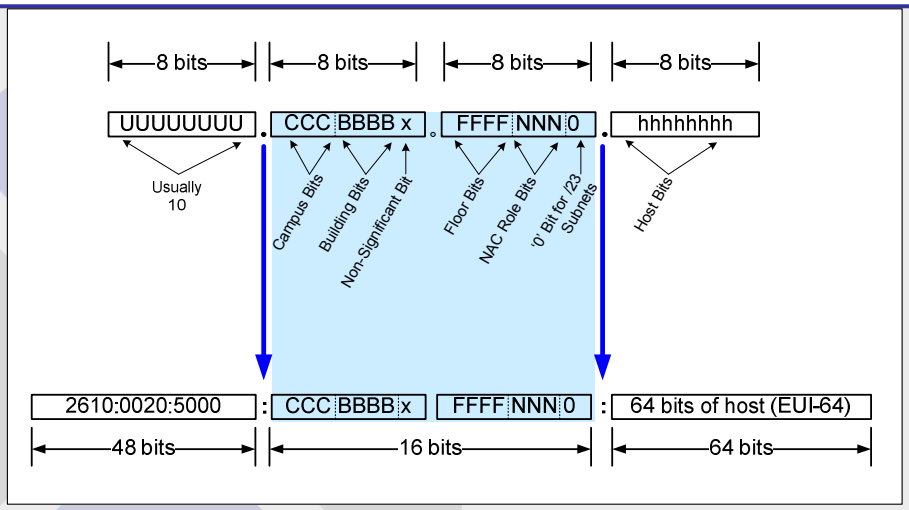

You can do the same sort of thing with IPv6 subnet bits

Advanced IPv4 Addressing Approach

- **Summarizable NAC-ready addressing**
  - **Old technique from OSPF area assignment: assign subnets in a coherent way using bit assignments**
  - **Site / bldg bits, floor bits, closet bits, NAC / service bits …**
    - **Yes, not enough bits for many NAC roles**
      - **But that's arguably a good thing!**
    - **Cisco TrustSec may obviate the need for NAC roles in addresses, a little simpler – but not be vendor-neutral**
  - **Advanced version: two or three variants, tuned to fit small / medium / large site (lead prefix bits identify the site and implicitly the bit assignments)**
- **Maximizes geographic summarization**
- **Other alternatives: service- or role-based (voice, guest, etc.)**
  - **Larger routing tables (geographically summarizes per-service)**

Agenda

- Quick Review of IPv6 Addressing
- Quick Review of Some IPv4 Addressing Techniques
- Choosing Prefix(es)
- Addressing Plan
  - IPv4 Mapping to IPv6
  - New Plan
- Summary, Q&A, References

Summarizable Global Prefixes; Multi-Homing

- The IPv6 protocol has had several interesting real-world oversights as far as I'm concerned (one or two mentioned later)…
  - They did well overall, complex situation, small gaps
- One was assuming multi-homing via a single provider
  - (Almost nobody in the U.S. wants to do that)
- The issue with multiple providers is that when one fails, the path through the other uses different addresses
  - Losing session state
- shim6 was proposed as a small tweak to deal with that
  - Initial effort had security, source routing, and other issues
  - Protocol refined and now specified
  - Now under testing
- Current ARIN policy is to allocate blocks of addresses for multi-homing under the same policies as for IPv4
  - Some fear "the routing swamp"
  - Chance to get non-provider address block before policy changes

Prefix Choices

- Choice:
  - Get one or several provider-supplied blocks
    - If several, form multi-homing strategy based on them
      - Where's IPv6 NAT when you need it?
      - Heresy!
      - NAT does create problems for app developers, makes firewalls more complex
  - Get an allocation from ARIN
    - Is a /48 enough?
    - How much to ask for?
    - Ask for bigger block to provide more subnets (see below)?

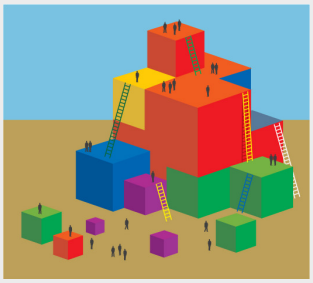

Is a /48 enough?

- If you get a /48, you have 64-48 =16 bits of subnetting to work with
  - Is 65,000 subnets enough?
  - Comparable to 10.a.b.0 with /24 subnets
  - If large, may not be enough for site / building / floor / NAC bit usage
  - Such schemes are inherently wasteful of address space, BUT conserve human time
- It may be a short-term-only issue
  - Other new security schemes?
- Your alternative: ask ARIN for more?

Agenda

- Quick Review of IPv6 Addressing
- Quick Review of Some IPv4 Addressing Techniques
- Choosing Prefix(es)
- Addressing Plan
  - IPv4 Mapping to IPv6
  - New Plan
- Summary, Q&A, References

Addressing Plans

- Choices:
  - Use IPv4 mapped directly to IPv6
  - Use tunneling in some form
  - Use NAT-PT (scalability???)
  - Dual-stack and “real” addresses seem better to me
  - Use an ad hoc IPv4 mapping to IPv6
    - Is your current addressing summarizable?
  - Come up with a new addressing plan
    - Preferably with better summarizability and / or NAC-readiness
    - Won’t help with NAC much unless you retrofit IPv4 … in which case might ad hoc map v4 to 6?
  - Summarization really helps
    - Rapid routing convergence
    - Human manageability of the network

IPv6 Subnets

- Some (including RFC's) say only do /64's
  - Later thinking seems mixed
  - I hate disagreeing with Jeff Doyle…
  - /64 is rather desirable for end station subnets
  - And /128 for loopbacks
- Do you mind /126's? (I like them for infrastructure links)
  - The interface ID u/g bit fine point:

```
0         0 0         1 1         2
0         7 8         5 6         3
+----+----+----+----+----+----+----+----+
|cccc|ccug|cccc|cccc|cccc|cccc|cccc|cccc|
+----+----+----+----+----+----+----+----+
```


How IPv4 "Ad Hoc" Mapping Works

- 10.x.y.0 /24: Take x.y and convert to hex
- Organization prefix as /48, tack on hex version of xy (16 bits = 4 hex digits). Result is a /64
- Map /30 to /126 as hex in last byte
  - Watch the u/g bits… if map last byte, no problem, lots of lead 0's
- Embedded XLS below, just uses string functions (could do hex math if necessary)

Sample Prefix		**2610:0040:1000**				**/38**			Treated as a /48 (just using part)
10.162.20.0	10	162	20	0	0A	A2	14	00	2610:0040:1000:A214::/55
Last octet carryover for a /30:									
10.162.0.5	10	162	0	5	0A	A2	00	05	2610:0040:1000:A200::05/126

Forming a new IPv6 Addressing Plan

- Take your prefix and divide up the bits as in our IPv4 approach
- Routing summarizability matters
  - (not overwhelmingly)
- Have to choose whether to do all /64's, or take /126's out of the plentiful space from

your-prefix:16 more bits:xx00: and-48-more-bits

There are only 2^46 ~= 7 x 10^13 or 70,000,000,000,000 /126 blocks this way

Dealing with /64 Subnets

- EUI-64 calls for 2^64
  - Ridiculously large number of host addresses
  - You have to decide what to do with them
  - I like automatically assigned EUI-64 "extended MAC" auto-addressing
  - Even though EUI-64 "wastes" 16 bits, for no useful purpose, as far as I can see
- IEEE exhibited human tendency to get lax with resources perceived as plentiful (e.g. bandwidth?)
  - IEEE tendency to tweak things? (Think Ethernet headers, VRRP vs. HSRP, etc.)

Host Anonymity in IPv6

- There are schemes for host anonymity
  - Random 64 bits used
  - Changes over time
  - Creates problems if security, RIAA, etc. trigger finding out which device did something in the past (forensics)
    - Would require at least historical Neighbor MAC table snapshots, timestamps, lookup technique
  - If everyone wears a mask and hoodie, theft is easy?
- Scanning for hosts could be hard with a /64
  - Virii have to work hard to find 64 bit addresses that are randomly distributed
  - That may confound common security approaches
- Servers can't be anonymous: DNS, fixed addresses
- Routers can be more anonymous, but that can make management harder
  - Your choice

DHCP

- Not needed for address per se ("stateless")
- Can be useful for providing other info, e.g. DNS server IP's, preferred default gateway choice

IPv6 Router Addressing

- We've covered /126 and /128
- Router interfaces don't need to be addressed out of global space
- Link-local addressing suffices for routing
  - Can be a bit painful to learn to troubleshoot
  - Can break network management tools that ping interfaces directly (link-local addresses can't be routed, not necessarily unique either)
  - Provides external security: not globally reachable
  - Could do link-locals ending in 1 or 2…
  - Perhaps assign global /128's to loopbacks, so can SSH from remotely
  - I think it better to never underestimate the flaws and limitations of network management software (the programmers usually lack real networking experience)

FHRP's: HSRP, VRRP, GLBP

- A bit late to Cisco IOS
  - I.e. check your code level, don't assume…
- Can use Router Discovery priority
- Implementation allows control of active router by means of
  - Interface tracking
  - Enhanced Object Tracking
- Not needed if use VSS on Cisco 6500 (one virtual router, two chassis)
- Still needed with Cisco Nexus 7000 series

Agenda

- Quick Review of IPv6 Addressing
- Quick Review of Some IPv4 Addressing Techniques
- Choosing Prefix(es)
- Addressing Plan
  - IPv4 Mapping to IPv6
  - New Plan
- Summary, Q&A, References

Summary

- IPv6 addressing is similar to but a little different from IPv4 (as all of IPv6 is)
  - We can use good ideas from IPv4 with IPv6
- Need to think about obtaining a prefix
  - Go to ARIN?
  - Obtain from one or more providers
- Need to form a policy re how you will address loopbacks, infrastructure links
- Need to think about host, server, router addressing
  - DHCP
  - First Hop Routing Protocol

Any Questions?

Reference URLs

- Basic IPv6 addressing
  - http://tools.ietf.org/html/rfc4291
- Unicast address assignment considerations
  - http://tools.ietf.org/html/rfc5375
- Enterprise network analysis (transition)
  - http://tools.ietf.org/html/rfc4852
- Use of /127 Prefix Length Between Routers Considered Harmful
  - http://tools.ietf.org/html/rfc3627
- Somewhat useful link, provider-addressing (mostly)
  - http://www.getipv6.info/index.php/IPv6_Addressing_Plans
- IPv6 block allocation tool
  - http://www.ipv6book.ca/allocation.html
- Shim6 links:
  - http://tools.ietf.org/wg/shim6/charters
  - http://www.shim6.org/
  - http://www.rfc-editor.org/rfc/rfc5533.txt (and others)

My IPv6 Web Content

- http://www.netcraftsmen.net/resources/archived-articles/280-introduction-to-ipv6-part-1.html
- http://www.netcraftsmen.net/resources/archived-articles/281-introduction-to-ipv6-part-2.html
- http://www.netcraftsmen.net/resources/archived-articles/279-introduction-to-ipv6-part-3.html
- http://www.netcraftsmen.net/resources/archived-articles/387-securing-ipv6-networks.html
- http://www.netcraftsmen.net/resources/archived-articles/657-preparing-for-ipv6.html

About Chesapeake Netcraftsmen

- For a copy of the presentation, email me at pjw@netcraftsmen.net
- About Chesapeake Netcraftsmen:
  - Cisco Premier Partner (have the certifications for Gold status)
  - Cisco Customer Satisfaction Excellence rating
  - We wrote the original version of the Express Foundations courses required for VAR Premier Partner status (and took and passed the tests)
  - We rewrote the DESGN / ARCH (CCDA / CCDP courses)
  - Cisco Advanced Specializations:
    - Advanced Route & Switch (10+ CCIEs on staff)
    - Advanced Unified Communications (and IP Telephony)
    - Advanced Wireless
    - Advanced Security (4 double R&S/Sec CCIE's now)
    - Advanced Data Center
  - We do Network / Security / Data Center / Wireless / Unified Communications design, assessment, deployment
  - Expertise and experience in many other areas as well

CISCO PARTNER Premier Certified