IPV6 OVERVIEW

Why, When, Where, What, and How of IPv6

James Small, Sr. Consultant and Chief IPv6 Evangelist at CDW

Michigan!/usr/group

mug.org – A Free and Open Source Michigan Community

OVERVIEW OBJECTIVES

- Why IPv6
- IPv6 Current Landscape
- IPv6 Technical Overview
- IPv6 Pilot Plan

Q&A throughout, I may postpone questions until the end depending on time

WHY IPV6

- Address space
  - » Should be a virtually unlimited supply – think street addresses
  - » Facilitates communication/collaboration

- Innovation
  - » NAT Gateways make innovation harder (mainly driven by insufficient address space)
  - » Productivity (easy communication/collaboration) is a key business objective which NAT impedes

ROADMAP

- Why IPv6
- ***IPv6 Current Landscape***
- IPv6 Technical Overview
- IPv6 Pilot Plan

DRIVERS

- IPv4 Address Depletion
  - » IANA Free Pool Depleted – February, 2011
  - » APNIC Depletion – April, 2011
  - » RIPE Depletion – September, 2012
  - » ARIN Depletion – Predicted for June, 2014
  - » Price for public IPv4 addresses going up

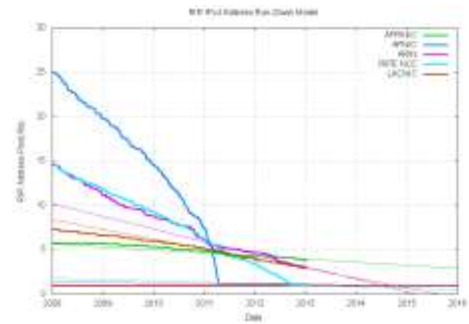

- Depletion Facts
  - » As of January 2013, there are < 113 million IPv4 addresses remaining before all global registries enter depletion mode
  - » Last year, 114 million addresses were allocated (down from a peak of 249 million in 2010)
  - » For this year (2013), based on 2012 allocations another 76 million addresses will be used

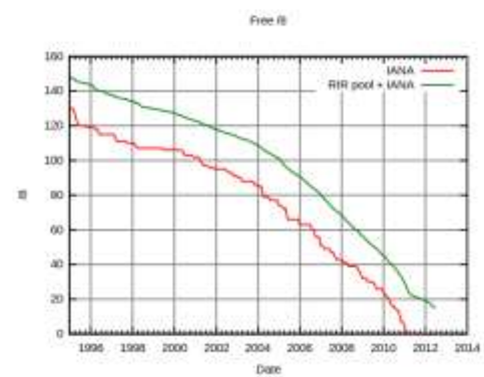

DRIVERS

- Geometric Growth of Internet Connected Devices
  - 2015 – 15 billion unique nodes
  - 2016 – 19 billion unique nodes
  - 2020 – 50 billion unique nodes
- World Internet Users – 2.5 billion out of 7 billion (36%)
  - This will double in 5 years
- The number of connected mobile devices is now greater than the world's population
  - By 2016 there will be over 10 billion
- Over 10 billion new micro-controllers are shipped each year with more and more networked - The Internet of Things

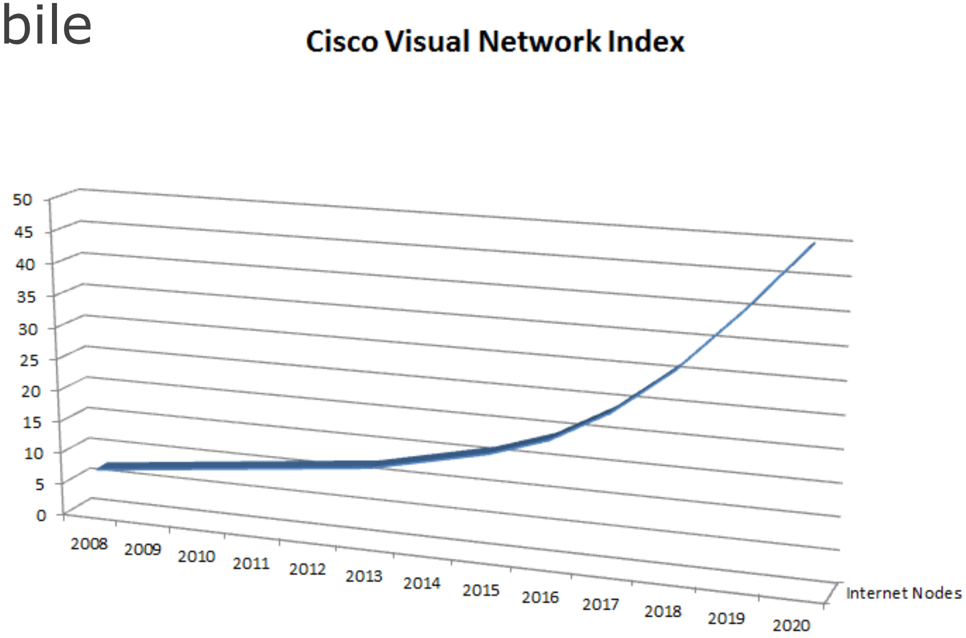

DRIVERS

- Explosion of mobile devices connected to the Internet
  - » 2015 – Average American will have 10 networked devices
    - Laptop
    - Smartphone
    - Smart Pill Dispensers
    - Exercise Monitors
    - Clothing sensors
    - Tablet
    - Umbrella with weather forecasting
    - Cars connected to ITS
    - Glasses with Internet Video
    - Wearable computing

- US Federal Government Deployment
  - » 39% deployed (Internet facing)
  - » Internal deployment by September, 2014
  - » VA & DOD – IPv6 only by 2015

DRIVERS

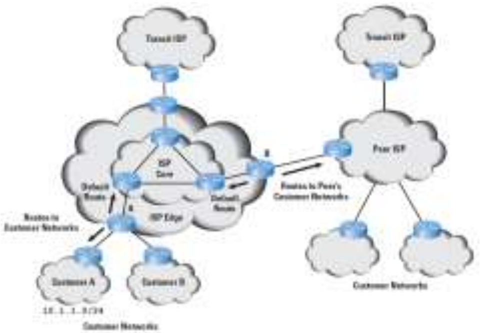

- ISP Deployment
  - » Comcast
    - 2.5% of traffic now IPv6
    - US Residential Deployment to be complete in June, 2013
  - » Time Warner
    - CGN costs $40/user per year
    - CGN will drive up costs for IPv4 Internet access (13-21%)
  - » Time Warner, AT&T, Verizon Wireless all exceeded 1% of traffic using IPv6 in June (2012)

- LTE/4G Deployment
  - » IPv6 preferred protocol
  - » Verizon LTE using IPv6 now
  - » T-Mobile new Windows 8 4G phones using IPv6 now
  - » AT&T, Sprint 4G/LTE IPv6 coming soon

DRIVERS

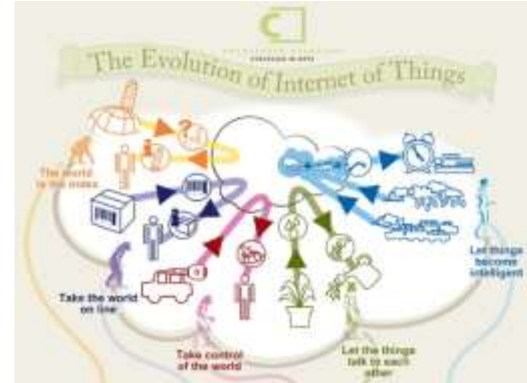

- Internet of Things
  - » Smart Grid (Meters) largely IPv6 – Consumers Energy in Michigan)
  - » 6LoWPAN (802.15.4) – New Wireless Light Bulbs, Sensors, Building Automation
  - » Intelligent Transportation System – US DOT
  - » Health Monitoring
  - » Machine to Machine communication for any electronic device
  - » Gartner – Top 10 Strategic Technologies for 2012 and 2013

Gartner®

MAJOR INTERNET SITES HAVE IPV6 ADDRESSES

- 86% of Top Level Domain Names support IPv6
- 41% of Globally Registered Domains have IPv6 Addresses
- Since World IPv6 Launch in June, many major Internet sites now advertise IPv6 Addresses in DNS:

```
C:\>nslookup www.google.com
(…)
Name:    www.l.google.com
Addresses:
  2607:f8b0:4002:802::1010
  173.194.73.106
(…)
Aliases:  www.google.com
```

```
C:\>nslookup www.facebook.com
(…)
Name:    www.facebook.com
Addresses:
  2a03:2880:10:1f03:face:b00c:0:25
  69.171.237.32
```

IPV6 SUPPORT FOR ISPS

Residential ISPs, Largest First

Cable (1 Million+ Subscribers)	xDSL (1 Million+ Subscribers)
Comcast – Deploying since 2011	AT&T – Deploying since 2011
Time Warner – Deploying	Verizon – Deploying
Cox – Residential trials this year	CenturyLink (Qwest/Savvis) – Deploying
Charter – Deploying	Frontier - ? (Yes for Business)
Cablevision - Testing	Windstream - ? (Yes for Business)

Tier 1 ISPs – Fully Deployed

AT&T	Verizon Business (UUNET)
Level 3	CenturyLink (Qwest/Savvis)
Sprint	XO Communications
Top 20 ISPs	Inteliquent (Tenet)

CONTENT PROVIDERS WITH IPV6

- Akamai – Delivers 15-30% of all web traffic
- Of Top 10 US Sites, 5 have IPv6 enabled:
  1) Google 2) Facebook
  3) YouTube 4) Yahoo
  6) Wikipedia

- Netflix – Up to 32.7% of Internet bandwidth
- 22% of top Alexa 500 sites including
  Bing, AOL, XBOX, WebEx,
  US News, USDA, NYU, …

CONSUMER IPV6 STATE

- All current Operating Systems have IPv6 on by default
- ISP Customer Premise Equipment now supports IPv6 and is on by default

- Where AT&T and Comcast enable IPv6 up to 40% of user traffic switches to IPv6

- All LTE/4G devices will use IPv6 with mobile devices outnumbering PCs this year

WATCHING IPV6 INTERNET GROWTH

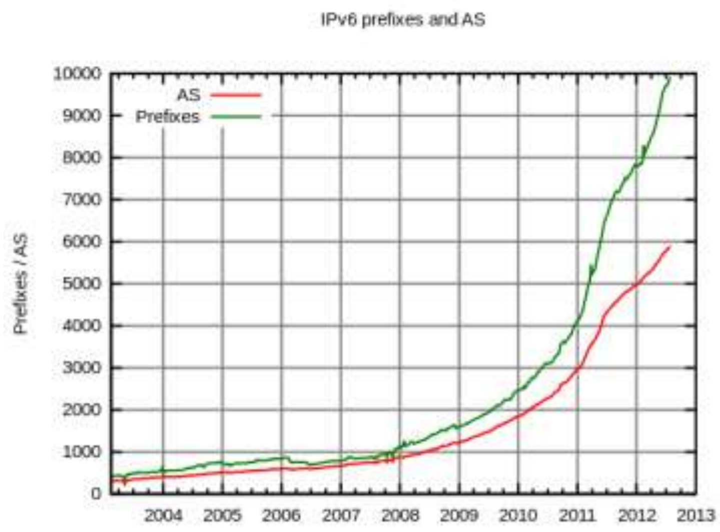

Growth of IPv6 Networks

Ramp up of IPv6 traffic globally at Google (150% increase/year) Exponential growth since 2010

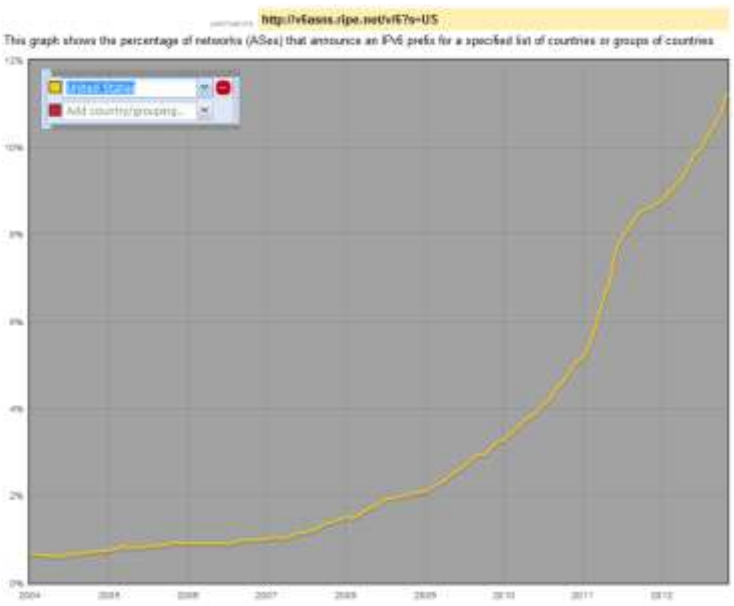

% US Administrative Network Domains with IPv6

ENTERPRISE IPV6 ADOPTION

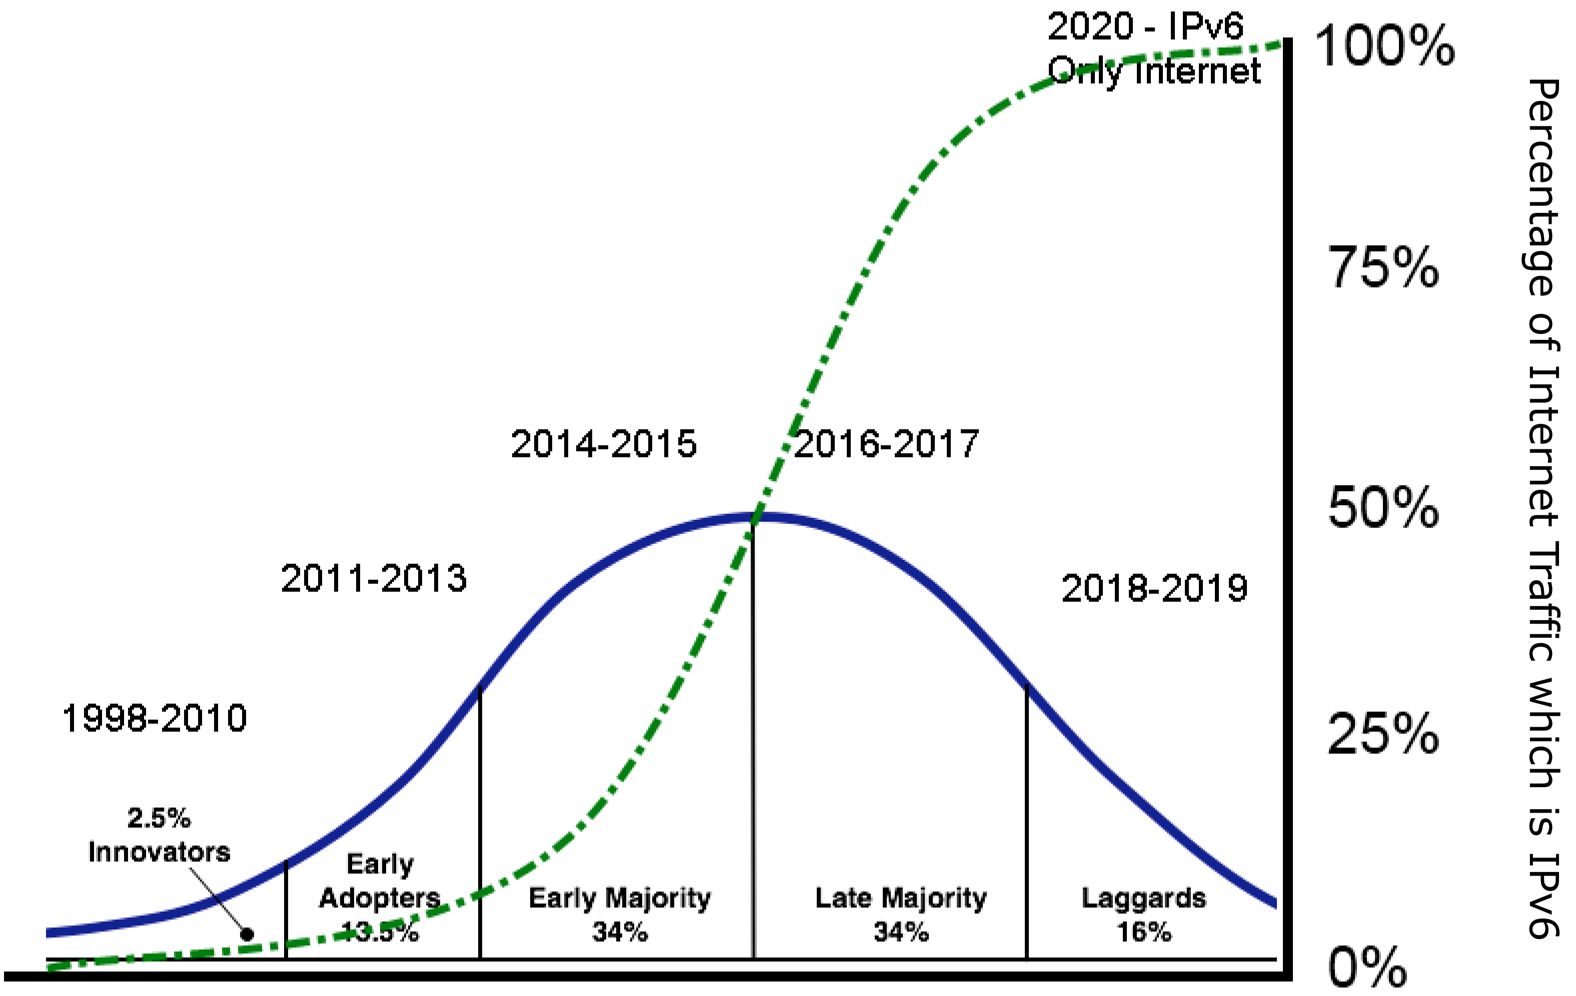

US IPv6 Adoption Lifecycle

AMERICAN IPV6 GROWTH

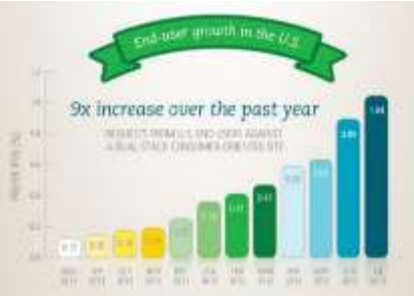

- US IPv6 capable users: 2.18%
- Over 57% of US transit AS support IPv6
- Over 29% of US viewed Internet content is available via IPv6
- US will hit critical mass for IPv6 (10%) in 2014
- At current growth rate, 50% of global traffic will be IPv6 by 2017
- Gartner
  - » Between now (2012) and 2015 organizations should be connected to IPv6 Internet
  - » By 2015 17% of Internet will use IPv6
  - » By 2015 28% of new Internet users will be running IPv6
  - » (Above stats are global, will be much higher in US)
- Typical timeline for Enterprise IPv6 Deployment: 3 – 5 years
- IPv6 is real – don't get left behind

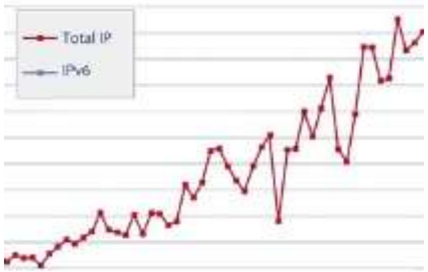

ROADMAP

- Why IPv6
- IPv6 Current Landscape
- ***IPv6 Technical Overview***
- IPv6 Pilot Plan

IPV4 AND IPV6 HEADER COMPARISON

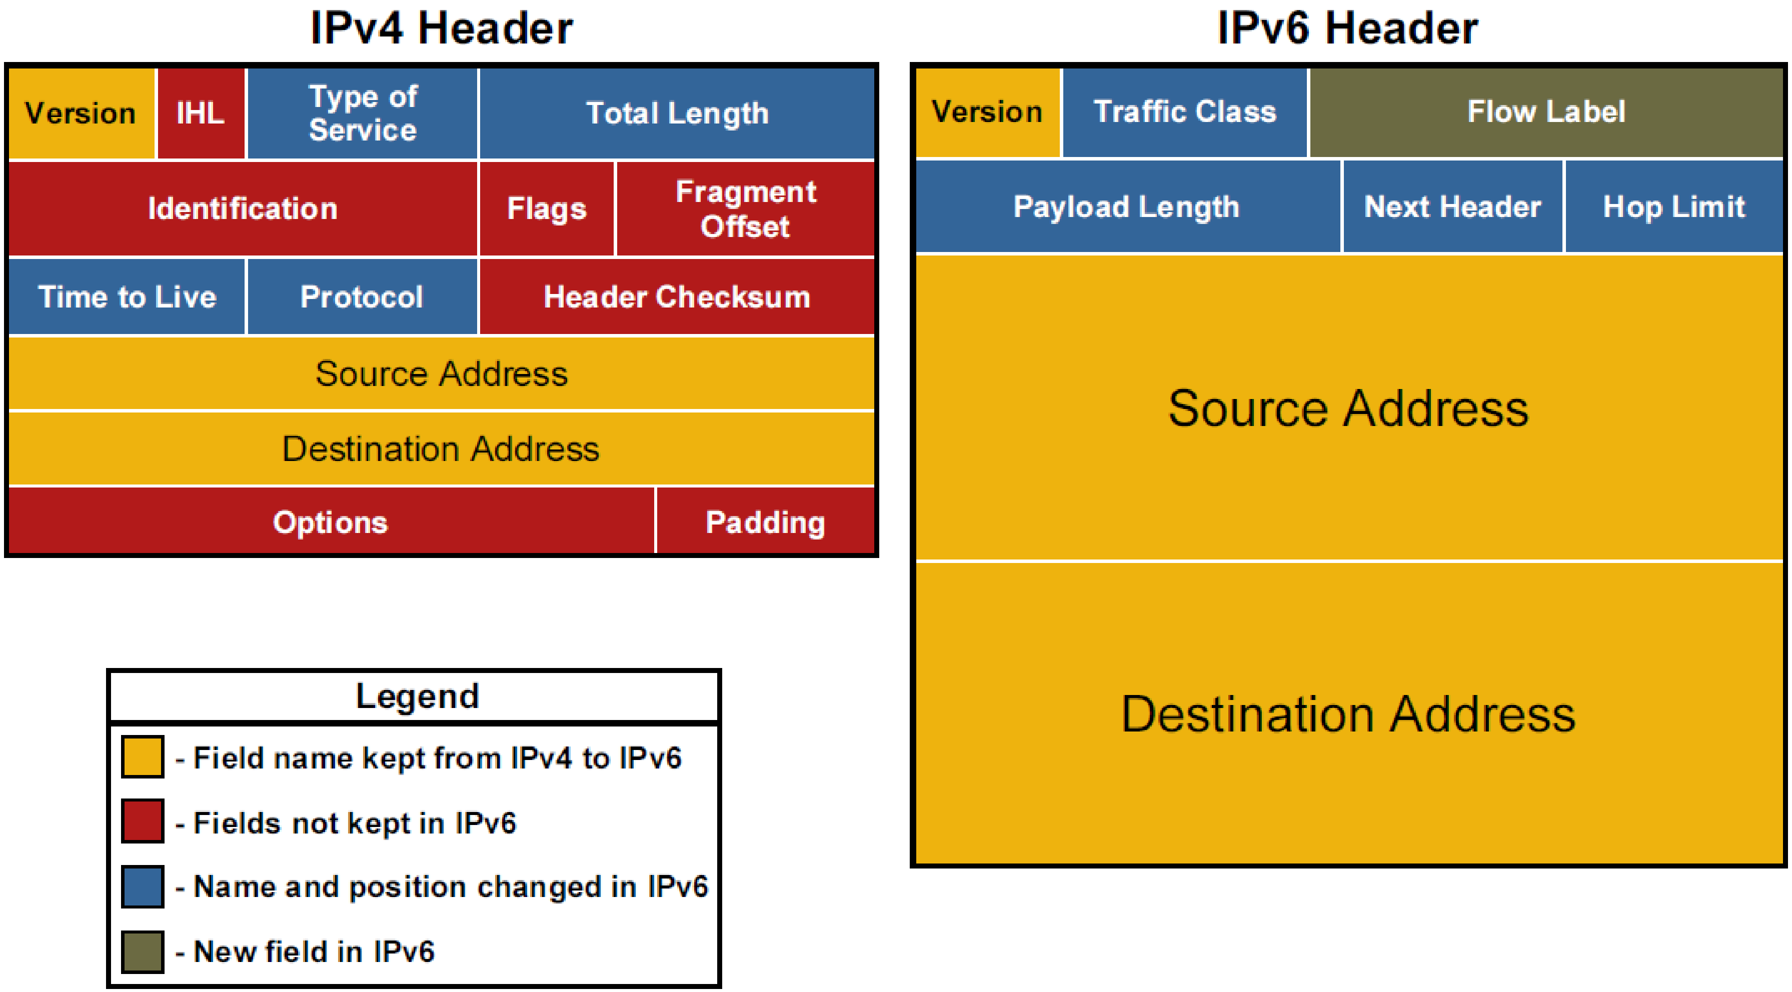

IPV6 EXTENSION HEADERS

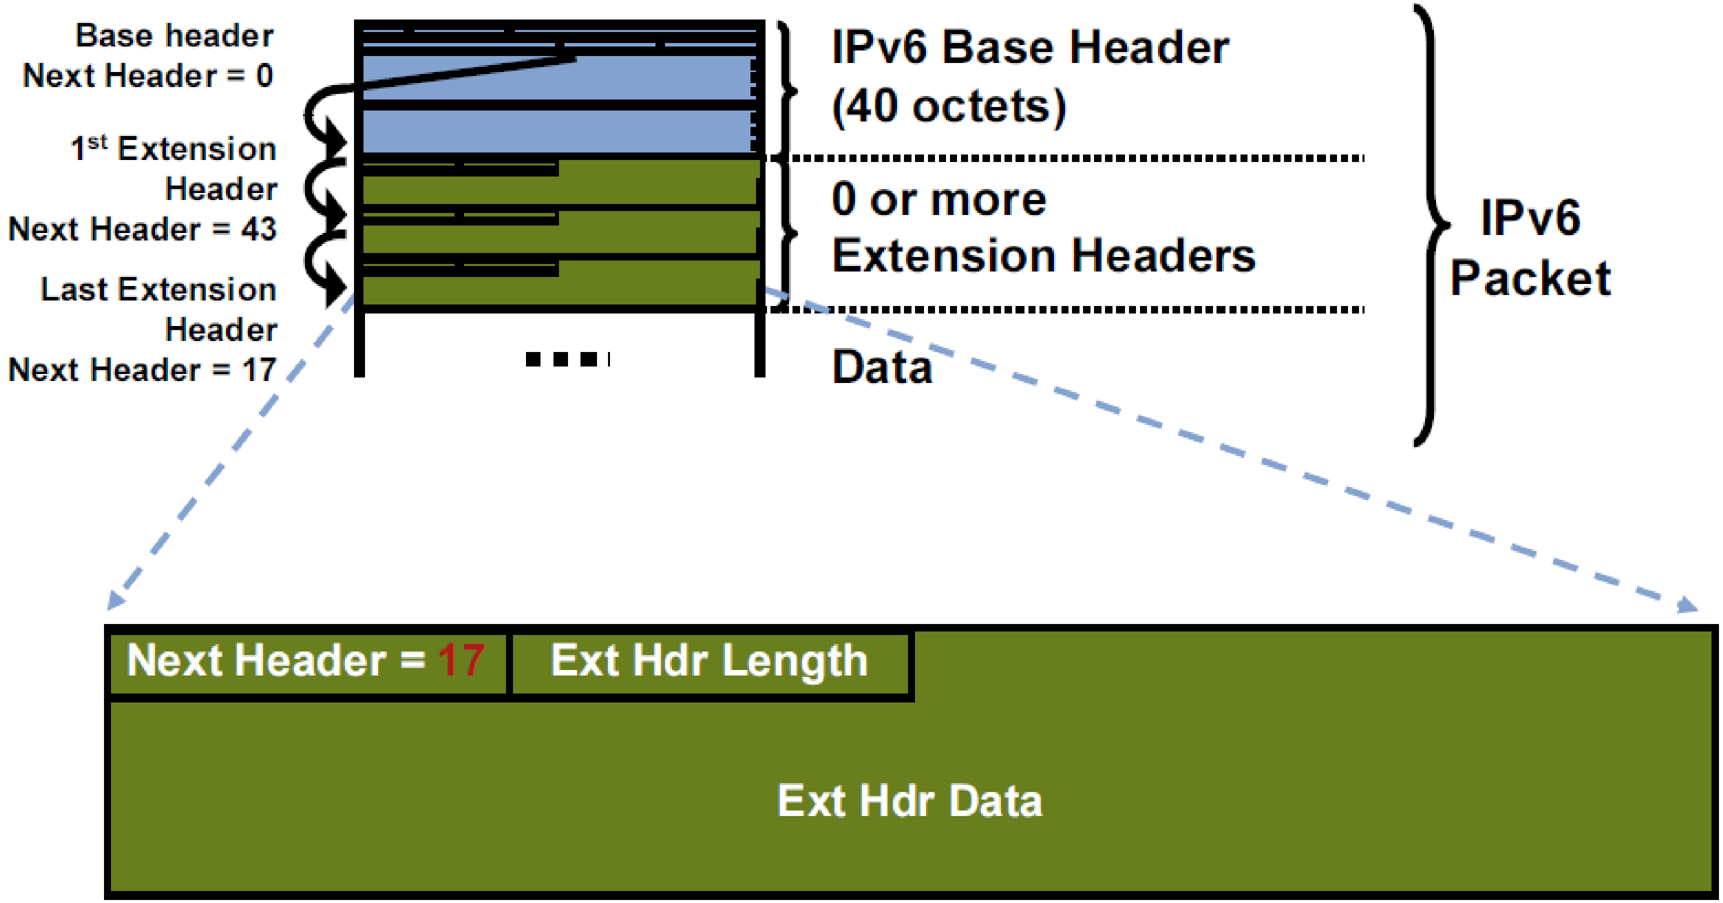

EXTENSION HEADERS AND THEIR ORDER

Type	Number
Hop-by-Hop	0
Destination Options (w/ routing header)	60
Routing	43
Fragment	44
Authentication (IPsec)	51
ESP (IPsec)	50
Mobility	135
Destination Options	60
ICMPv6	58
No Next Header	59
(Upper Layer)	e.g. TCP=6, UDP=17

IPV6 CHANGES FROM IPV4

- ~~Broadcasts~~
- Fragmentation only done by end nodes, not by routers
- ARP replaced by NDP, a subset of ICMPv6
- IGMP replaced by MLD, a subset of ICMPv6
- DHCP replaced by RAs (subset of NDP) + DHCPv6
- Nodes can auto-configure address with SLAAC (use RAs)
- NetBIOS/WINS do not function with IPv6
- UNC paths based on address must use ipv6-literal.net space
  - » \\2001-db8-28-3-f8a-5b31-67b7-6ef.ipv6-literal.net\docs
- Blocking ICMPv6 will completely break IPv6!!!
  - » Careful with firewall/route/switch/operating system ACLs
- Minimum MTU changes from 68 to 1280

WHAT'S NEW WITH IPV6 ADDRESSING?

- Size
  - » Address is 4x as big – essentially limitless address space
- Scalable Multicast
  - » Assigned IPv6 prefix provides globally unique range
  - » Internet multicast services now possible
  - » Rendezvous Point can be embedded in address
- Scoping
  - » Link-local – only meaningful per link
  - » ULA – non-Internet-routable, scalable, non-overlapping space
  - » Multicast scopes from host local to Internet global

WHAT’S NEW WITH IPV6 ADDRESSING?

- Multiple Addresses per Interface
  - All IPv6 enabled interfaces will now have at least 2 addresses, more not unusual
  - Addresses will include link-local and either ULA or GUA
  - Address lifetimes make hot migrations possible
- Plug-in-play
  - Like IPX and DECNet, IPv6 provides automatic addressing
  - When you configure a router on a link/subnet, all IPv6 capable hosts will automatically configure themselves and start using (and preferring) IPv6
- Possibilities
  - No real "usable" security or other benefits yet, but…
  - Many great ideas for security and flow management – expect to see innovation here

IPV4 VERSUS IPV6 ADDRESSING

Type	IPv4	IPv6
Unspecified	0.0.0.0	::
Loopback	127.0.0.1	::1
“Link-Local”	169.254.251.1	fe80::584c:7cf5:4a0e:3ce9%17
“Private”	10.152.16.87	fd8d:8b76:0494:659b::4ae:3ce9
Public	12.203.95.104	2001:46e:cae8::a7:1afe:e91d
Multicast	224.0.0.18	ff02::1:ff94:e774
“Broadcast”	255.255.255.255	ff02::1

IPV6 ADDRESS MAGNITUDE

IPv4 32-bits

IPv6 128-bits

$2^{32} = 4,294,967,296$

$2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456$

$2^{128} = 2^{32} * 2^{96}$

$2^{96} = 79,228,162,514,264,337,593,543,950,336$

times the number of possible IPv4 addresses

(79 trillion trillion)

IPV6 ADDRESS STRUCTURE

- Hexadecimal digits (0-9, a-f)
- 8 groups of 16 bit (4 digit) numbers
- Groups separated by colons (:)
- Digits not case sensitive, but lower case preferred
- Abbreviations are possible
  - » Can omit leading zeroes
  - » Can substitute longest string of zeroes with double colon (::)

- Examples:
  - » 2001:db8:cafe:f0a2:a8b0:2:ffe1:a90b
  - » fd92:e075:819c:7a2::fc9a:105
  - » fe80::f413:9b1e:9f3:feb6
  - » ff05::fb

IPV6 ADDRESSING BASICS

- IPv6 prefixes/addresses always use CIDR notation:
  - IPv4 – 192.0.2.0/24
  - IPv6 – 2001:db8:101::/48

- IPv6 addresses can omit leading zeroes, but not trailing ones:
  - 2001:0db8:0101:00a0:0000:0000:0d20:9ce5
    - Becomes:
      - 2001:db8:101:a0:0:0:d20:9ce5
    - But Not (Just like 010 = 10, but 010 ≠ 1, 010 ≠ 01)
      - ~~2001:db8:101:a:0:0:d2:9ce5~~
      - ~~2001:db8:101:00a:0:0:0d2:9ce5~~

- IPv6 addresses can substitute longest string of zeroes with a double colon:
  - 2001:db8:101:a0::d20:9ce5

IPV6 ADDRESS SCOPE, LIFETIME, AND TYPES

- Link Local – fe80::/10 (link only significance)
- Unique Local – f700::/7 (not routable on Internet)
- Global – 2000::/3 (currently IANA allocated address space)

- Address Lifetimes
  - Valid – How long address can be used for connections
  - Preferred – How long address can be used to initiate connections

- Address Types
  - Unicast – As in IPv4
  - ~~Broadcast~~ Multicast – From IPC to Global Communication
  - Anycast – Same service on multiple devices with closest one selected

IPV6 GLOBAL ADDRESS BIT HIERARCHY

3 bits	9 bits	20 bits	16 bits	16 bits	64 bits
001	IANA to RIR	RIR to ISP	ISP to End Site	Net	Interface ID
001	IANA to RIR	RIR to End Site		Net	Interface ID
3 bits	9 bits	36 bits		16 bits	64 bits

- Global unicast addresses being allocated from 2000::/3
  - Top – Provider Assigned (PA) address space
  - Bottom – Provider Independent (PI) address space
- End sites get a /48 (65k networks)

From Owen DeLong@he.net's Intro to IPv6 presentation pg. 21

IPV6 LINK LOCAL ADDRESS BIT HIERARCHY

16 bits	48 bits	64 bits
fe80	0	Interface ID

- Always fe80::/64 on every link
- Never forwarded/routed to another network (link scope)
- Must be present for interface to participate in IPv6, auto-configured
- ZoneID used to uniquely identity links:
  - » UNIX: ping6 fe80::101:1%eth0 (ifconfig -a)
  - » Windows: ping fe80::101:1%12 (ipconfig/all)

From Owen DeLong@he.net's Intro to IPv6 presentation pg. 22

IPV6 ADDRESS SPACE

Address Range	Description
0000-00FF::/8	Reserved and Special Purpose
0100-1FFF::/4	Reserved
2000-3FFF::/3	Global Unicast Addresses
2000-2CFF	Allocated
2D00-3FFF	Unallocated
4000-FBFF	Reserved (*5 more /3s*, a /4, /5, & /6)
FC00-FDFF::/7	Unique Local Unicast Addresses
FC00-FCFF::/8	Reserved for centralized allocations
FD00-FDFF::/8	Unrestricted – no registry/registration
FE00-FE7F::/9	Reserved
FE80-FEBF::/10	Link-local Addresses
FEC0-FEFF::/10	Reserved (Formerly Site-Local)
FF00-FFFF::/8	Multicast

IPV6 INTERFACE ADDRESSING

- Always 64 bits – Why?
  - IEEE hands out 48 bit (EUI-48) and 64 bit (EUI-64) MACs
  - Original idea was to use MAC to generate IID like in IPX
  - For 48 bit MACs, you insert 0xFFFE in the middle to generate a 64 bit address (IEEE Rule)
  - The "U" bit is also flipped (Modified EUI-64)
    This is so if you create a local address you can use 2001:db8::1 instead of 2001:db8::0200:0:0:1
    Wasn't that nice of them?

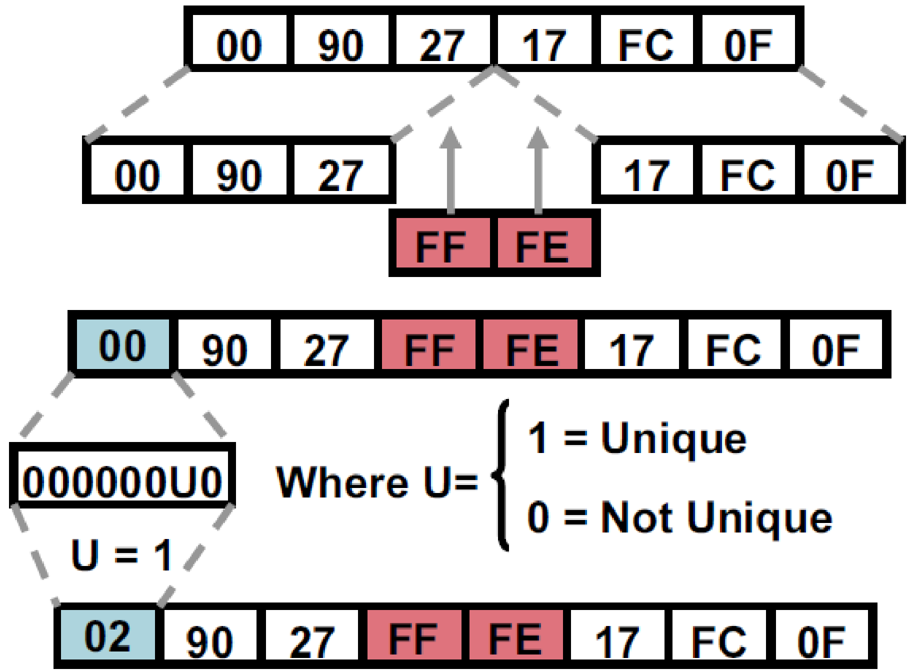

IPV6 INTERFACE ADDRESSING

- Modified EUI-64 was originally used for Stateless Auto Address Configuration (SLAAC)
  - When a node received an RA it used the prefix(es) and M-EUI-64 to create addresses
- Unfortunately M-EUI-64 based addresses can be used as a super cookie
- To address this, privacy extensions were created (RFC 4941)
  - With privacy extensions, the IID is essentially a random number that can periodically change
- Privacy addresses are great for consumers but problematic for the enterprise
  - DHCPv6 relays don't include client MAC address
  - Privacy addresses make accountability/security difficult because addresses periodically rotate/change

IPV6 MULTICAST ADDRESSING

- FF00::/12 – Well known (e.g. OSPFv3)
- FF10::/12 – Locally defined
- FF30::/12 – Locally defined using global unicast prefix
- FF70::/12 – Locally defined using prefix with embedded RP
- Bits:

1111-1111-0RPT-SSSS-0000-IIII-LLLL-LLLL-<64 bit prefix>-<32 bit GID>

R=Embedded RP
P=Global Prefix
T=Locally Defined
S=Scope Bits
I=RP Address
L=Prefix Length

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	If Permanent
1	If Temporary

Scope	
1	Node
2	Link
5	Site
8	Organization
E	Global

IPV6 SPECIAL ADDRESSES

IPv4 Compatible - ::<IPv4> (deprecated)

IPv4 Mapped - ::FFFF:<IPv4>

Discard Prefix - 0100::/64 (Implement RTBH)

Well Known Prefix - 64:ff9b::/96 (NAT64)

Teredo - 2001:0000::/32 (Tunnel through NAT)

Documentation - 2001:db8::/32

6to4 - 2002::/16 - 6to4 (Tunnel through IPv4)

- 2002:<Public IPv4 Address>::/48 (Private IPv4 Address undefined)

ISATAP (IntraSite Tunnel through IPv4)

- 64 bit Unicast Prefix:0:5efe:<Private IPv4 Address>
- 64 bit Unicast Prefix:200:5efe:<Public IPv4 Address>

Solicited-Node Multicast

- ff02::1::ff00:0/104 + last 24 bits of IPv6 Address

MULTI-PROTOCOL REALITIES

IPv4 L2 Cache:

```
# arp -a
athena.int.level14.net (192.168.234.142) at 00:0c:29:e9:bc:50 [ether] on eth0
honeydrop.local (192.168.234.180) at 00:26:2d:fc:05:9b [ether] on eth0
hsrp-vl101.int.level14.net (192.168.234.129) at 00:00:0c:9f:f0:65 [ether] on
eth0
```

IPv6 L2 Cache:

```
root@ubuntu12:~# ip -6 neigh show
fe80::101:2 dev eth0 lladdr a8:b1:d4:60:11:41 router STALE
2001:470:c4e9:fb1:3fd:6182:d9b6:b027 dev eth0 lladdr 00:26:2d:fc:05:9b
REACHABLE
fe80::101:1 dev eth0 lladdr 00:05:73:a0:00:66 router STALE
fe80::19:7ff:fe24:4fcb dev eth0 lladdr 02:19:07:24:4f:cb router STALE
2001:470:c4e9:fb1::101:a861 dev eth0 lladdr 00:0c:29:80:a8:61 STALE
```

MULTI-PROTOCOL REALITIES

IPv4 and IPv6 are ships in the night!

IPv4 Firewall

```
# iptables -nL
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     icmp --  0.0.0.0/0            0.0.0.0/0
ACCEPT     all  --  0.0.0.0/0            0.0.0.0/0            state RELATED,ESTABLISHED
ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0            tcp dpt:22
```

IPv6 Firewall

```
root@ubuntu12:~# ip6tables -nL
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     icmpv6    ::/0                 ::/0
ACCEPT     all      ::/0                 ::/0                 state RELATED,ESTABLISHED
ACCEPT     tcp      ::/0                 ::/0                 tcp dpt:80
```

IPV6 AND DNS

DNS has supported IPv6 for a long time – only a new resource record for the address and alternate pointer name space:

- IPv4 – A Record:
  - arin.net. IN A 192.149.252.75
- IPv6 – AAAA Record:
  - arin.net. IN AAAA 2001:500:4:13::80

- IPv4 – PTR Record:
  - 75.252.149.192.***in-addr.arpa.*** IN PTR www.arin.net.
- IPv6 – PTR Record:
  - 0.8.0.0.0.0.0.0.0.0.0.0.0.0.0.0.3.1.0.0.4.0.0.0.0.0.5.0.1.0.0.2.***ip6.arpa.*** IN PTR www.arin.net.

Transport agnostic - works equally well over IPv4 or IPv6, careful!

LINUX IPV6 TOOLS

General:

- ping → ping6
- traceroute → traceroute6
- tracepath → tracepath6

More:

- host/nslookup/dig – same tool, may have to specify IPv6 records (e.g. AAAA)
- telnet (still useful for raw connection to service) - same
- ssh - same

Network Analysis:

- tcpdump – only filtering options change
- wireshark – only filtering options change

LINUX IPV6 TOOL QUIRKS

IPv6 literals (RFC 3986)

- Generally means enclose the IPv6 address in brackets:
  - 2001:db8:fb::1a → [2001:db8:fb::1a]
- Necessary or many programs will interpret colons as port number delimeter
- Much more "interesting" if you use link local or multicast as you must specify the interface with a zone identifier

Examples:

- wget http://[2001:500:4:13::80] – works fine
- curl http://[2001:500:4:13::80] – doesn't work:
  - curl: (3) [globbing] error: bad range specification after pos 9
  - Known issue, must use "-g":
  - curl -g http://[2001:500:4:13::80]

LINUX IPV6 GUIDANCE

Start with the Linux IPv6 How To:

http://www.tldp.org/HOWTO/Linux+IPv6-HOWTO/index.html

Many things depend on the distro – check out wikis/documentation:

Ubuntu: https://wiki.ubuntu.com/IPv6

Debian: http://madduck.net/docs/ipv6/

MONITORING AND CONTROLLING IPV6

Service	Number	Description
IPv6 Encapsulation	IPv4/41	Tunnel IPv6 over IPv4
Generic Tunnel	IPv4/47	Tunnel anything over GRE
Teredo/Miredo	UDP/3544	Tunnel IPv6 over UDP (NAT Traversal)
Teredo/Miredo	Non-Standard	IPv6 destination starting with 2001:0000::/32 over UDP over IPv4
TSP	TCP|UDP/3653	IPv6 Tunnel Broker using the Tunnel Setup Protocol (RFC 5572)
AYIYA	TCP|UDP/5072	IPv6 Tunnel Broker using Anything in Anything (www.sixxs.net/tools/ayiya/)
Public 6to4 Anycast Relay	IPv4:192.88.99.1	Starting with IPv6 source address of 2002::/16 (6to4 is IPv6 over IPv4/41) Destined to 192.88.99.0/24 for IPv4
IPv6 Encapsulation	TCP/443	IPv6 over IPv4 SSL Tunnel, many variants
IPv6 Ethertype	0x86DD	Distinct from IPv4 Ethertype (0x0800)
DNS IPv6 Records	Several	AAAA, updated PTR records - can be transported over IPv4 or IPv6

Image source: gfi.com

IPV6 SECURITY

Common IPv6 L2 Security Issues and Options:

Issue	Solution
Spoofed/Illegitimate RAs	RA Guard (or PACL)
Spoofed NDP NA	MLD Snooping, DHCPv6 Snooping, NDP Inspection, ~~SeND~~
(Spoofed) Local NDP NS Flood	NDP Inspection, NDP Cache Limits, CoPP
(Spoofed) Remote NDP NS Flood	Ingress ACL, CoPP, NDP Cache Limits
(Spoofed) DAD Attack	MLD Snooping, NDP Inspection
(Spoofed) DHCPv6 Attack	DHCPv6 Guard
Spoofed/Illegitimate DHCPv6 Replies	DHCPv6 Guard

IPV6 ACCESS CONTROL

- Firewall Policy
  - » Don’t block all ICMPv6!!!
  - » Simple Examples for transit traffic, can get more granular:

Source Criteria:		Destination Criteria:		Service	Action
Source	…	… Destination	…		
ming rules)					
any6		any6		IPv6-Ops packet-too-big parameter-problem time-exceeded unreachable	Permit

Source Criteria:		Destination Criteria:		Service	Action
Source	…	… Destination	…		
ming rules)					
any4		any4		IPv4-Ops parameter-problem time-exceeded unreachable	Permit

  - » Reference NIST SP 800-119 (Section 3.5, Table 3-7)
  - » Reference RFC 4890 (Recommendations for Filtering ICMPv6 Messages in Firewalls)

ROADMAP

- Why IPv6
- IPv6 Current Landscape
- IPv6 Technical Overview
- ***IPv6 Pilot Plan***

PILOT PLAN – INITIAL SOFTWARE

Key Services

- Dual Stack DNS Server with DNS64 support
- Dual Stack DHCP/DHCPv6 Server
- Dual Stack File Server
- Dual Stack Web Server
- Key Applications (e.g. E-mail, Directory Services, User/Web Apps)

Bonus Items:

- IPAM Solution

DESIGN YOUR INITIAL PILOT TOPOLOGY

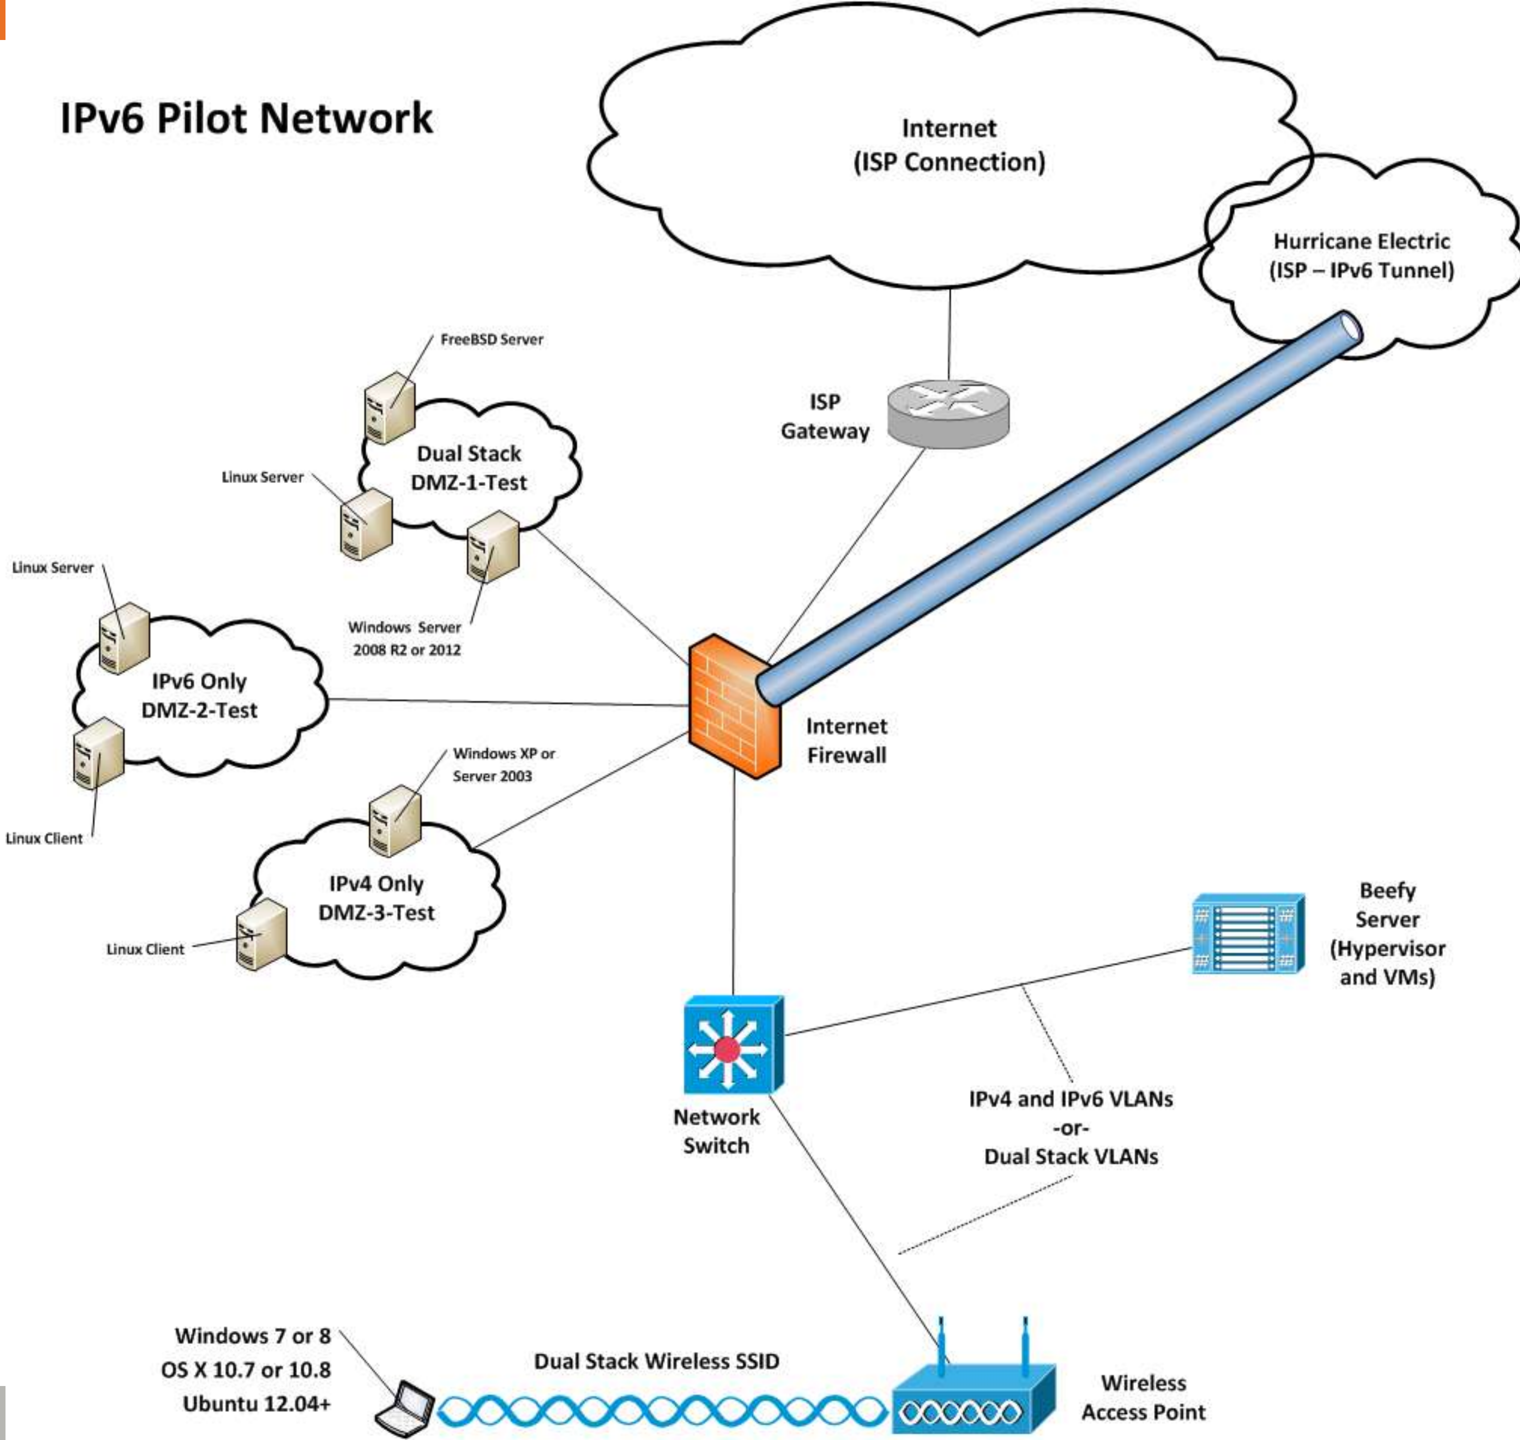

INITIAL LAB ROADMAP

- Obtain IPv6 /48 Prefix
- Pilot Addressing Plan
- Design and Build Out
- Address Provisioning
- DMZ Setup
- Internal Network Setup

OBTAIN AN IPV6 NETWORK ADDRESS

- Sign up for free IPv6 Internet access from Hurricane Electric (http://tunnelbroker.net)
- With your account, request a /48 prefix

- Q: Why start with Hurricane Electric?
- A: It works great, service is available from anywhere on the Internet, and you get a /48 all for free.

- Most important aspect of starting with HE:
  - » You need practice creating an addressing plan and deploying IPv6. It will take you at least 3 times to get your addressing plan right so let's get started…

PILOT ADDRESS PLAN GUIDELINES

Developing a great address plan takes practice

- Site - /48
- Loopback Network - /64
- Loopback Interface - /128
- Translation Services - /56
- Point-to-Point* - /126
- Everything else - /64

*Still good to set aside /64

1	NAME	STREET	CITY	STATE
2	LOMAR TAXIDERMY	1401 AUSTIN ST	ABILENE	TX
3	HAMMER'S TAXIDERMY	1711 ELMDALE RD S	ABILENE	TX
4	VERRIPS TAXIDERMY STUDIO	10655 US HIGHWAY 87 S	ADKINS	TX
5	NOURI'S TAXIDERMY	11465 FORD RD	ADKINS	TX
6	LAKE FORK TAXIDERMY	HIGHWAY 17 & COUNTY 1558	ALBA	TX
7	ROCKER B TAXIDERMY	401 S MAIN ST	ALBANY	TX
8	STACY HARGROVE	N US HIGHWAY 283	ALBANY	TX
9	L ADAMS WOODWORK PLUS	1198 N JOHNSON ST	ALICE	TX
10	S G BRISENO TAXIDERMIST	3717 W HIGHWAY 44	ALICE	TX
11	HIP-O TAXIDERMY	2801 E HIGHWAY 90	ALPINE	TX
12	C D TAXIDERMY	103 WOFFORD LN	ALVIN	TX

EXAMPLE HIGH LEVEL PILOT ADDRESS PLAN

Create your addressing plan on nibble boundaries:

- Split up your address allocation by Place In Network (e.g. 2001:db8:babe:***X000***::/52)
  - » 2001:db8:babe:0000::/52 – Management
    - 2001:db8:babe:0000::/64 – Loopbacks
  - » 2001:db8:babe:1000::/52 – Labs
  - » 2001:db8:babe:2000::/52 – DMZs
  - » 2001:db8:babe:3000::/52 – Servers
  - » 2001:db8:babe:4000::/52 – User/Desktop
  - » (...)
  - » 2001:db8:babe:F000::/52 – Special Purpose
    - 2001:db8:babe:FF00::/56 – Reserved for translation services

PILOT ADDRESS PLAN THOUGHTS

Prefixes

- Basic subnet plan - spreadsheet
- 65k prefixes per /48 - not scalable!

Nodes

- > 18 quintillion possible per subnet
- Sizeable deployments - IPAM desirable

Reference:

IPv6 Subnetting Best Current Operational Practices

THOUGHTS ON INITIAL TOPOLOGY

- Network Types
  - » Dual Stack
  - » IPv4 Only
  - » IPv6 Only

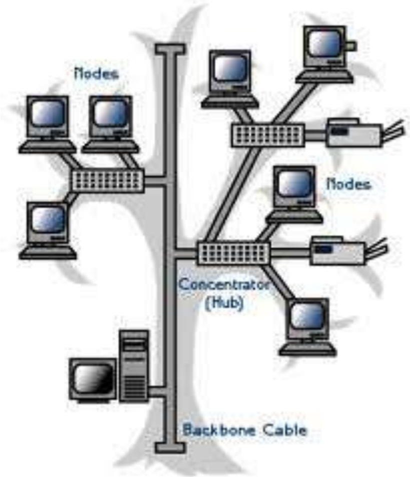

- Areas to Look at:
  - » Static/Dynamic Routing
  - » Load Balancing
  - » Proxying
  - » Tunneling
  - » NAT
  - » Dual data/control/management planes

A WORD OF CAUTION ON NAT

- NAT was invented for address conservation
- Address conservation not needed for IPv6
- Think carefully before using NAT
  - » What applications will this degrade or break?
  - » How much is operational complexity increasing?
  - » How difficult does support become?

BUILD OUT INITIAL LAB

- Infrastructure setup
- Hypervisor setup
- Physical and Virtual Nodes with representative Operating Systems
- Key Applications

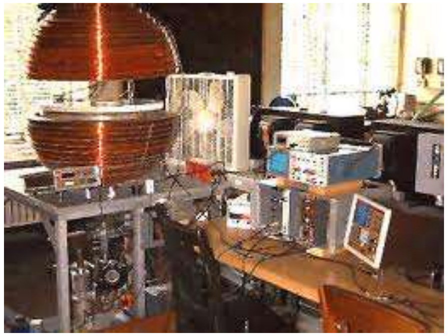

IPV6 SUPPORT INFRASTRUCTURE

- DNS
  - » Transport
  - » Accessibility
  - » Dynamic DNS
- DHCPv6
  - » Stateless
  - » Stateful
- WINS/NetBIOS
  - » Viability
  - » Recommendations

IPV6 ADDRESS PROVISIONING THOUGHTS

Address Options and Applicable Systems:

- Pure Static (Must disable SLAAC)
- Static with Options
- SLAAC, no DHCPv6
  - » Basic
  - » RDNSS
  - » Dynamic VLAN Assignment
- SLAAC with (Stateless) DHCPv6
- DHCPv6 (Stateful DHCPv6)
  - » Still requires SLAAC for default gateway

BUILD YOUR IPV6 DMZ

In order of preference:

- Option 1 – Dual Stack
- Option 2 – Load balanced (SLB64)
- Option 3 – Dual Stack Reverse Proxy
- Option 4 (Discouraged) – Use NAT64

BUILD YOUR IPV6 INTERNAL NETWORK

- Connect Internal IPv6 Network to IPv6 Internet
  - » Option 1 (Preferred) – Dual Stack
  - » Option 2 – Forward Proxy
  - » Option 3 – (Legacy) Tunneling
  - » Option 4 – Stateful NAT64 (IPv6 Only)

RECOMMENDED READING

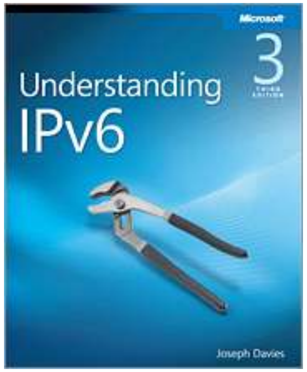

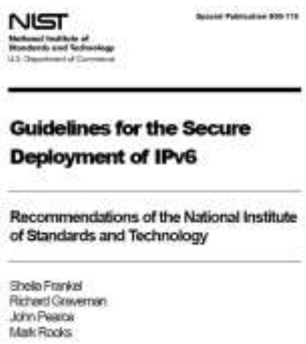

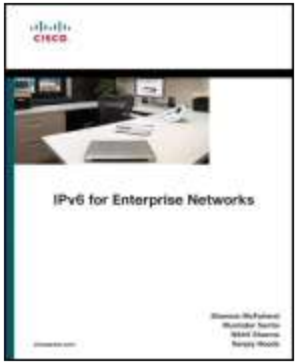

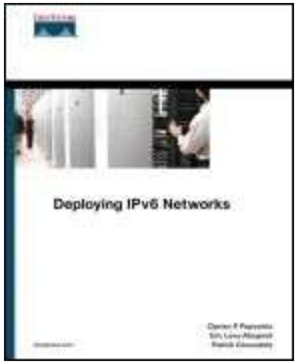

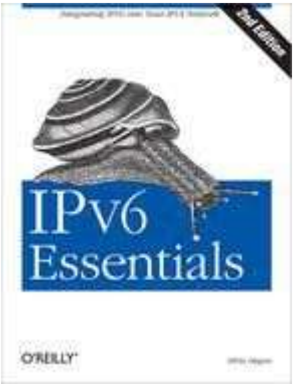

QUESTIONS

@netsec14

My IPv6 Blogs:
Packet Pushers

Appendix

IPV6 CONNECTIVITY OPTIONS

In order of preference:

- Native dual stack (e.g. Comcast XFinity)
- You have a direct public IPv4 address:
  - 6rd – Must be supported by your ISP
  - Tunnelbroker (6in4 tunnel) – Hurricane Electric
  - Unmanaged 6to4 tunnel – Works better if your ISP supports, but will work without too
- Behind a NAT gateway/CGN/LSN or can't terminate ISP connection:
  - AYIYA to Tunnelbroker (SiXXS)
  - TSP with Gogonet (Freenet6)
  - VPN or Tunnel Connection to someplace with IPv6 support
  - Use public Teredo/Miredo servers (but performance isn't great)

IPV6 PREFIX POLICIES

- When multiple transport protocols are used (IPv4 and IPv6), a method must exist to choose which one is used including:
  - » Use IPv4 or IPv6?
  - » Where multiple addresses exist:
    - Which destination address should be chosen?
    - Which source address should be chosen?

RFC 3484 - Default Address Selection handles this

Prefix policies (RFC 3484 implementation) may be viewed and changed:

Windows: netsh interface ipv6 show prefixpolicies

Linux: ip addrlabel show

WINDOWS IPV6 BASICS

New Windows Commands - netsh interface ipv6:

show addresses	Detailed information on IPv6 interface addresses
show destinationcache	Displays the contents of the destination cache, sorted by interface; the destination cache stores the next-hop addresses for destination addresses
show global	Shows global configuration parameters such as interface address randomization
show interfaces	Detailed interface list including index numbers/zone identifiers, also try level=verbose
show neighbors	Displays contents of the neighbor cache, sorted by interface; the neighbor cache stores the link-layer addresses of recently resolved next-hop addresses
show prefixpolicies	Shows prefix policy table (IPv6 versus IPv4 preference order)
show privacy	Shows interface address privacy configuration parameters

Note: netsh commands can be abbreviated:

- netsh interface ipv6 show interface

Abbreviate as:

- netsh int ipv6 sh int