HURRICANE ELECTRIC
INTERNET SERVICES

IPv6 -- A light at the end of the tunnel

Owen DeLong
owend@he.net

HURRICANE ELECTRIC
INTERNET SERVICES

IPv6 -- A light at the end of the tunnel

Owen DeLong
owend@he.net

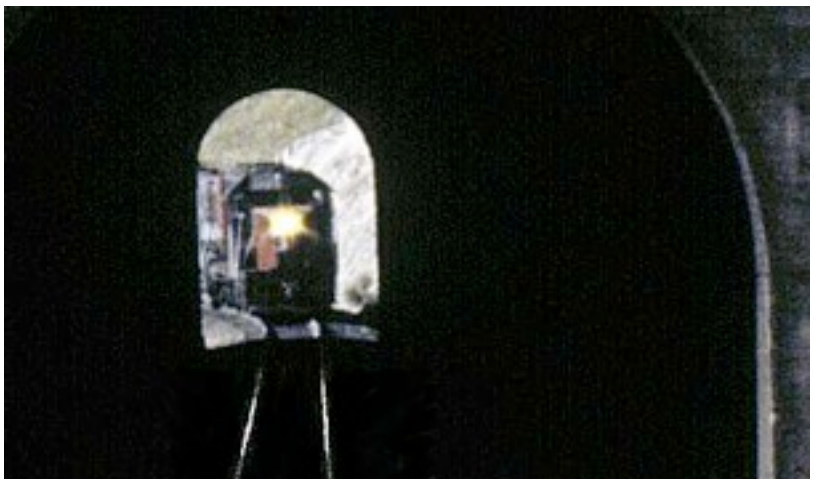

HURRICANE ELECTRIC
INTERNET SERVICES

IPv6 -- A light at the end of the tunnel

Owen DeLong
owend@he.net

HURRICANE ELECTRIC
INTERNET SERVICES

IPv6 -- A light at the end of the tunnel

Owen DeLong
owend@he.net

How soon do we need IPv6

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?
- Who thinks that all of the RIRs still have unused IPv4 space?

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?
- Who thinks that all of the RIRs still have unused IPv4 space?
- Who thinks they still will in:

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?
- Who thinks that all of the RIRs still have unused IPv4 space?
- Who thinks they still will in:
  - June?

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?
- Who thinks that all of the RIRs still have unused IPv4 space?
- Who thinks they still will in:
  - June?
  - October?

How soon do we need IPv6

- Who thinks there's still a global IPv4 free pool?
- Who thinks that all of the RIRs still have unused IPv4 space?
- Who thinks they still will in:
  - June?
  - October?
  - 2012?

How close are you to the end of the tunnel?

- How many of you are thinking Eye Pea Vee What?
- How many have IPv6 in a lab?
- How many have IPv6 deployed in production at all?
- How many have full IPv6 deployment to your backbone?
- How many offer production IPv6 services to your customers?

Why is this important? – Last Year

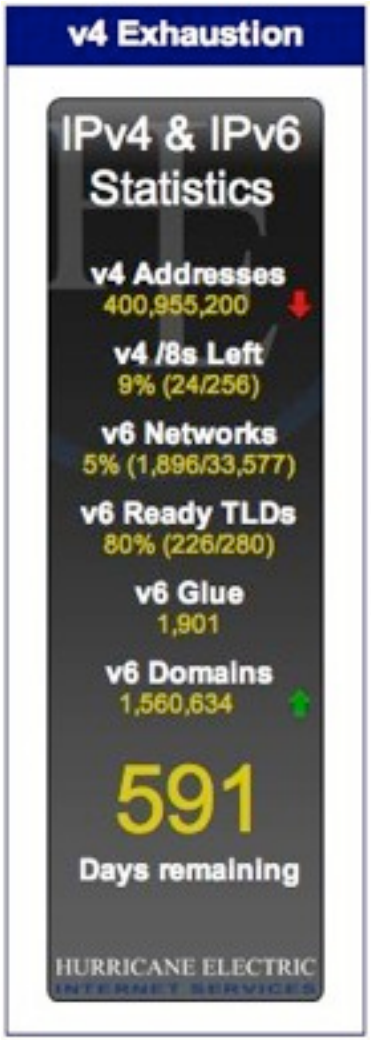

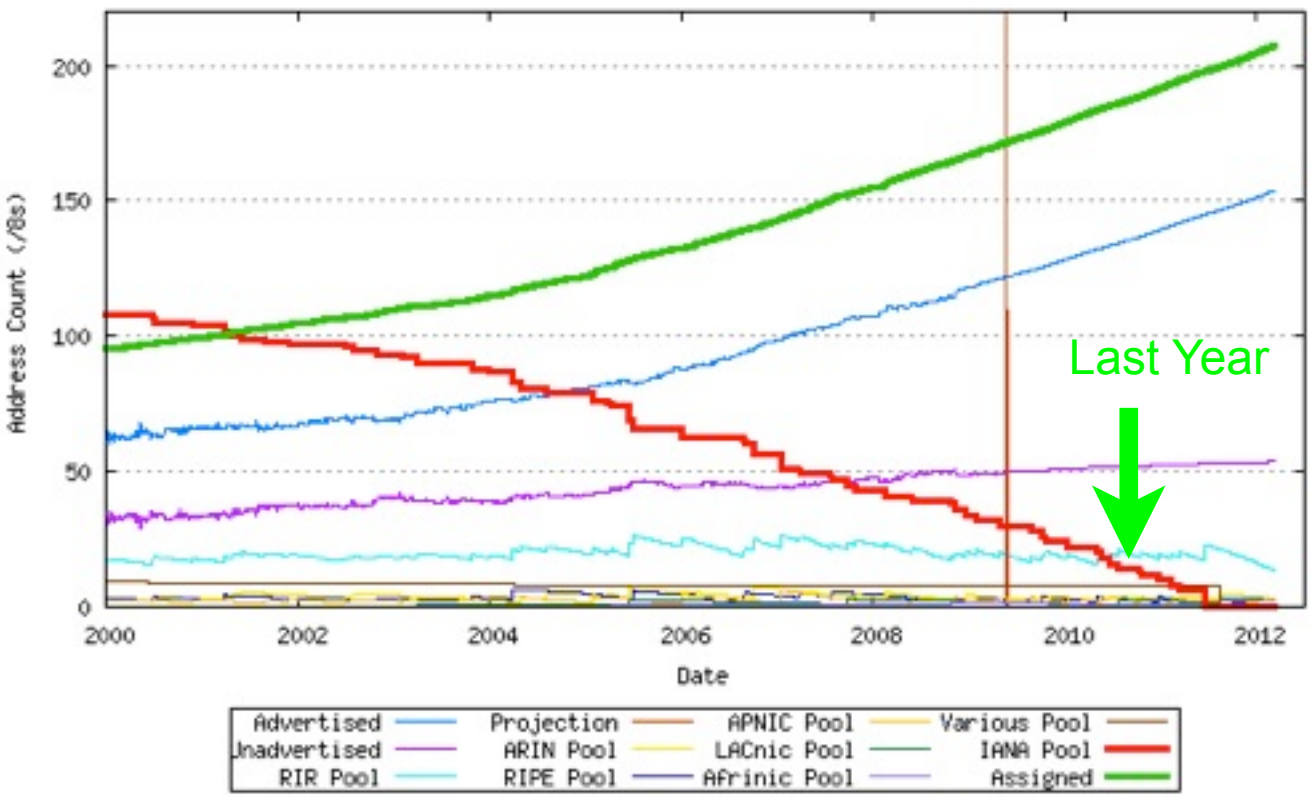

Why is this important? – Today

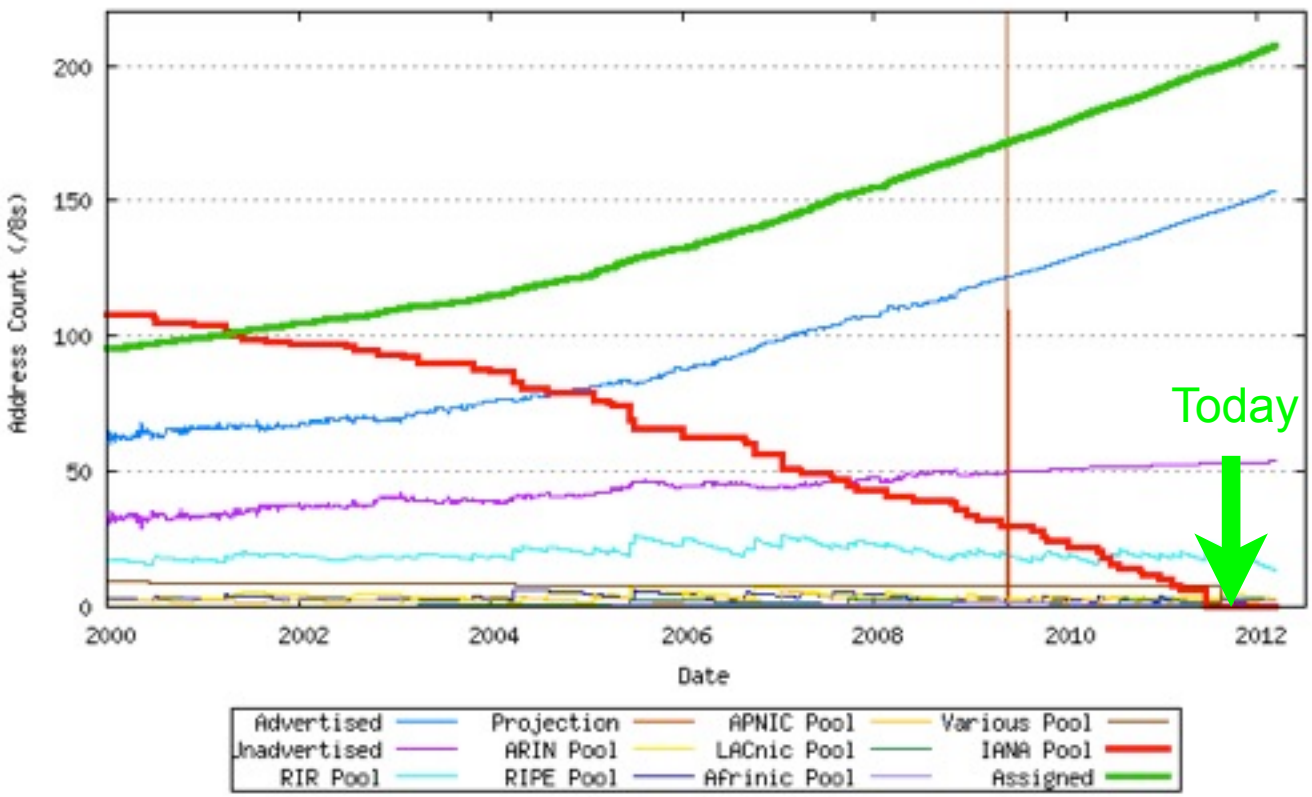

A review of the last 14 months in IPv4

IPv4 Address Space Consumption											January 2010				
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Allocated to RIR	IANA Free Pool	Other Uses

A review of the last 14 months in IPv4

Allocated to RIR in 2010

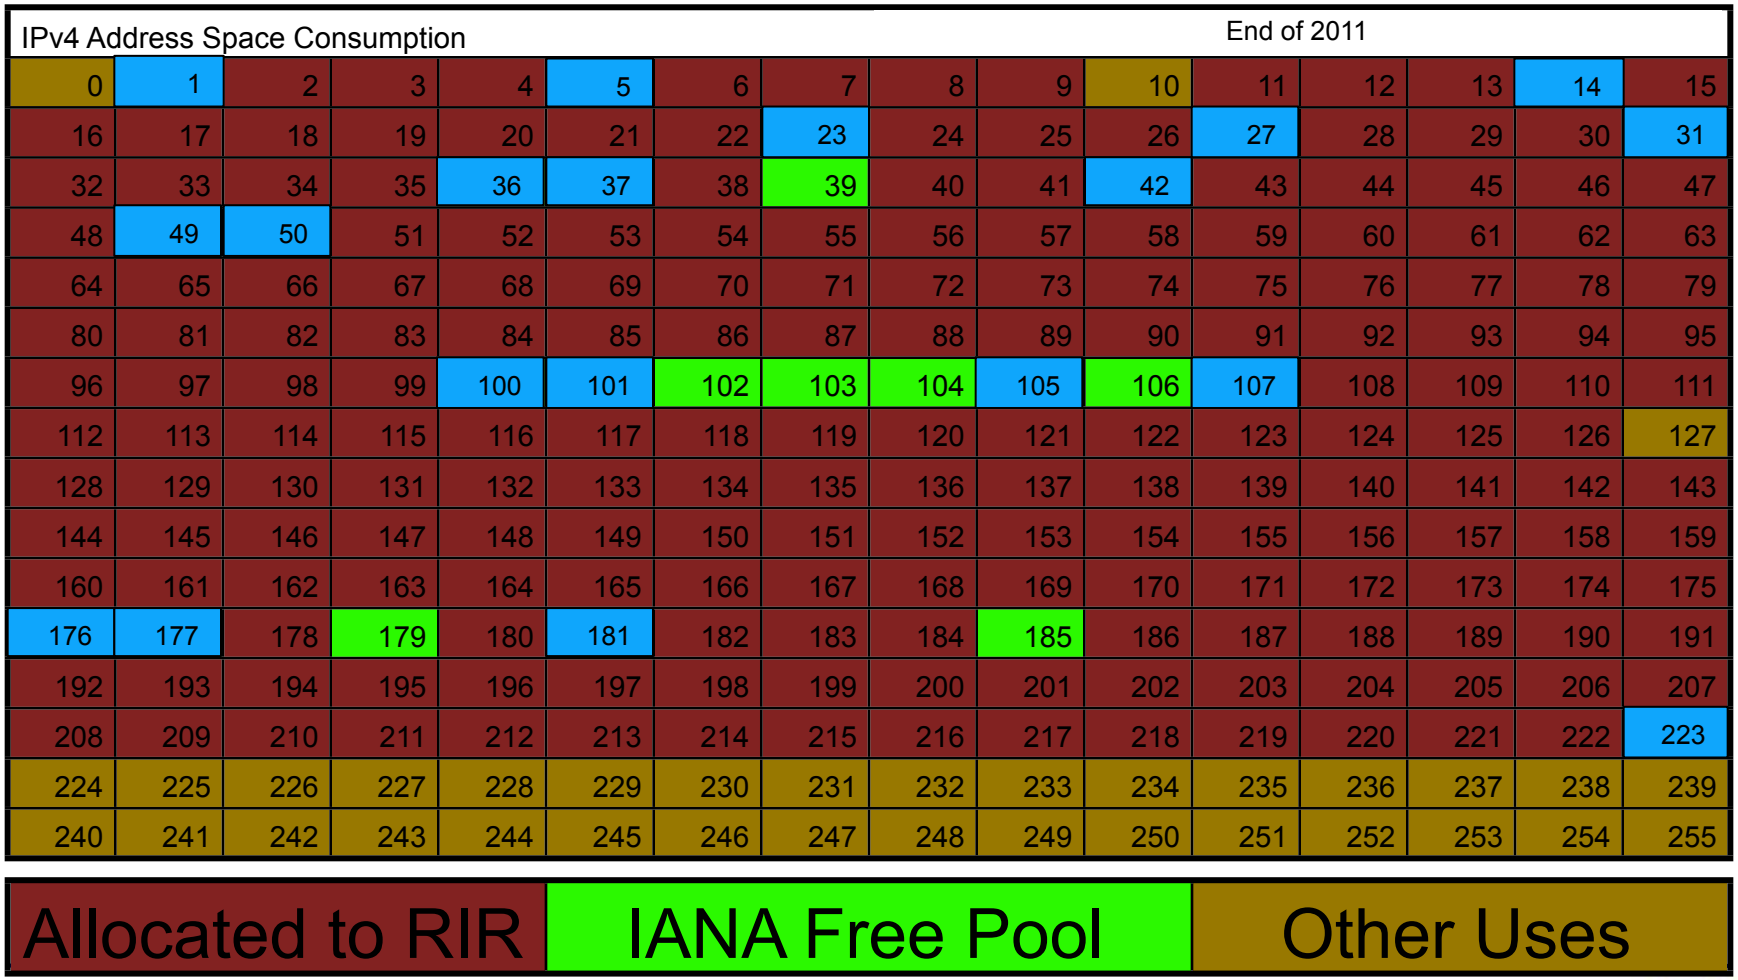

A revi… last 14 months in IPv4

APNIC receives final normal IANA allocation 31 January, 2011

Allocated to RIR in 2010

IPv4 Address Space Consumption — End of 2011

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

Allocated to RIR | IANA Free Pool | Other Uses

A revi… last 14… in IPv4…

APNIC receives final normal IANA allocation 31 January, 2011

IANA Free Pool Ends 3 February, 2011

Allocated to RIR in 2010

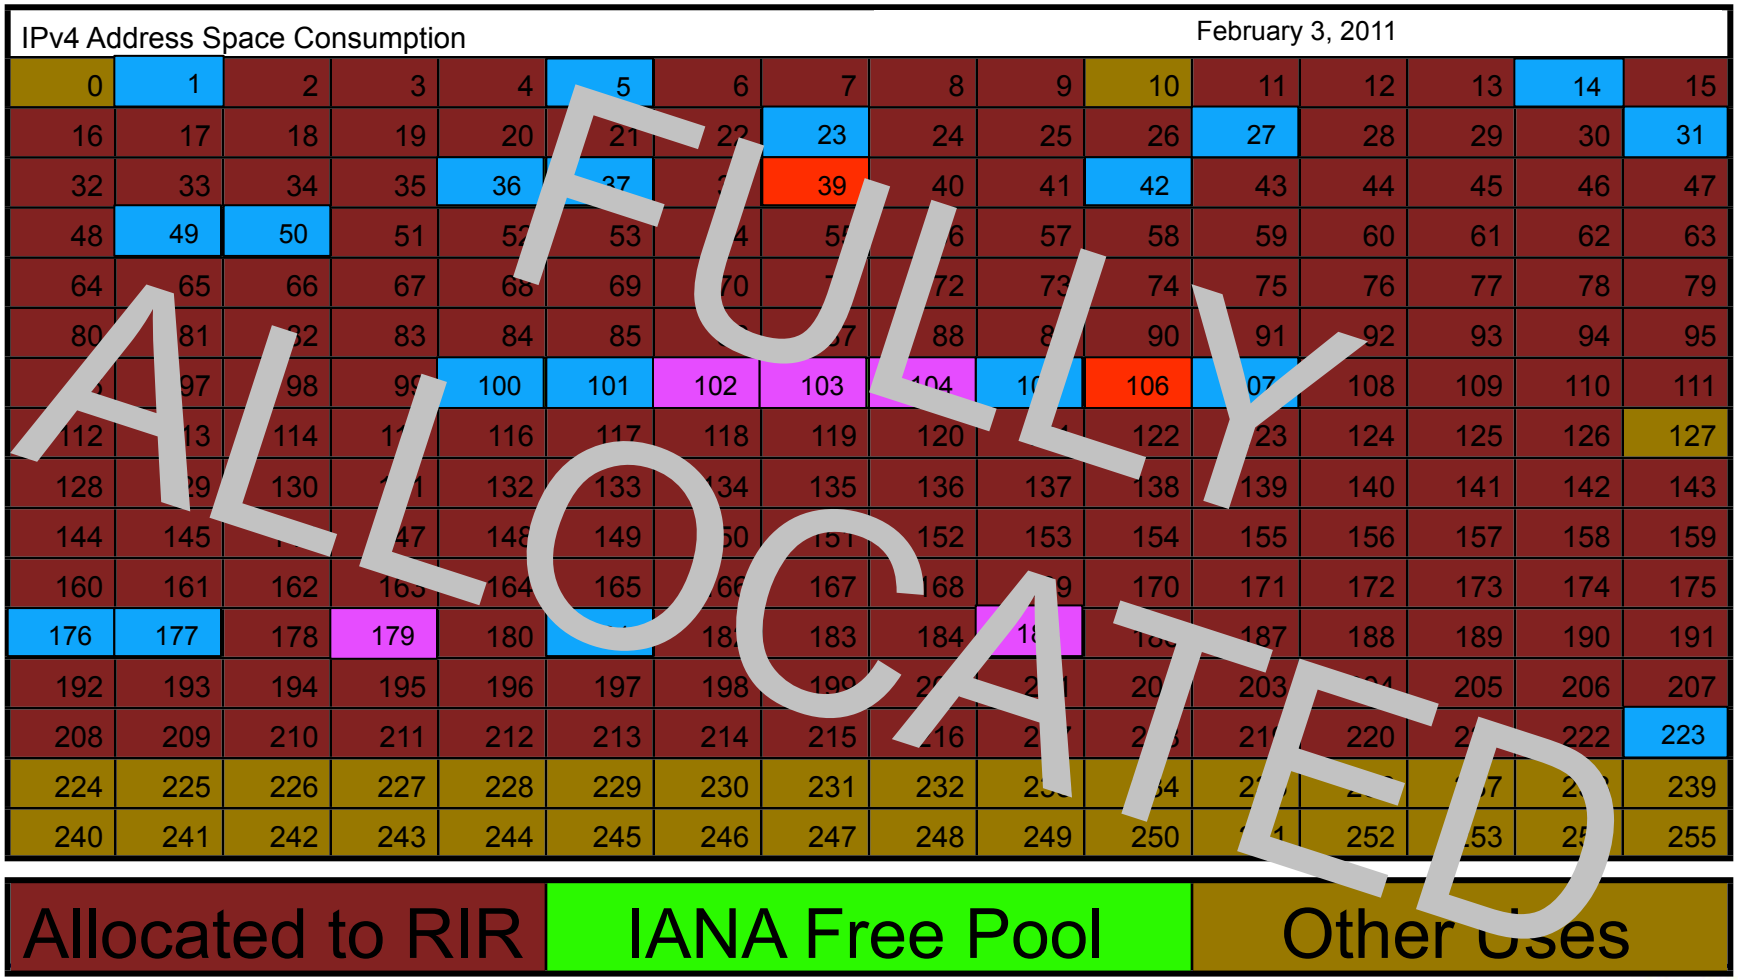

Tuesday, March 29, 2011

Light at the End of the Tunnel

- An apropos metaphor:
  - Could be good (the end of a long dark tube)
  - Could be bad (an oncoming high speed train)
- As applied to IPv6, some of each:
  - Good
    - Much larger address space
    - New Autoconfiguration Features
    - Improved IPSEC support
    - Simplified Header
    - Better Mobility support
  - Bad
    - Requires effort and investment
    - Software updates
    - Hardware upgrades (in some cases)
    - Staff Training
    - Procedure Updates

A brief overview of Address Policy

- Where do IP Addresses come from?
- IANA
- RIRs
- NIRs
- LIRs
- RIR Policy Process
- Global Policy

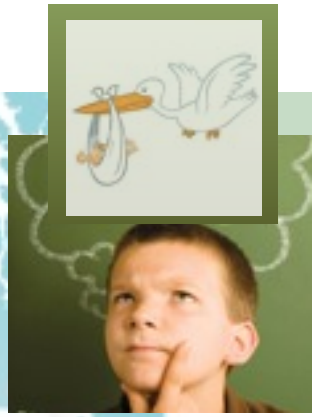

IPv6 -- The basics
How Global Unicast is Allocated

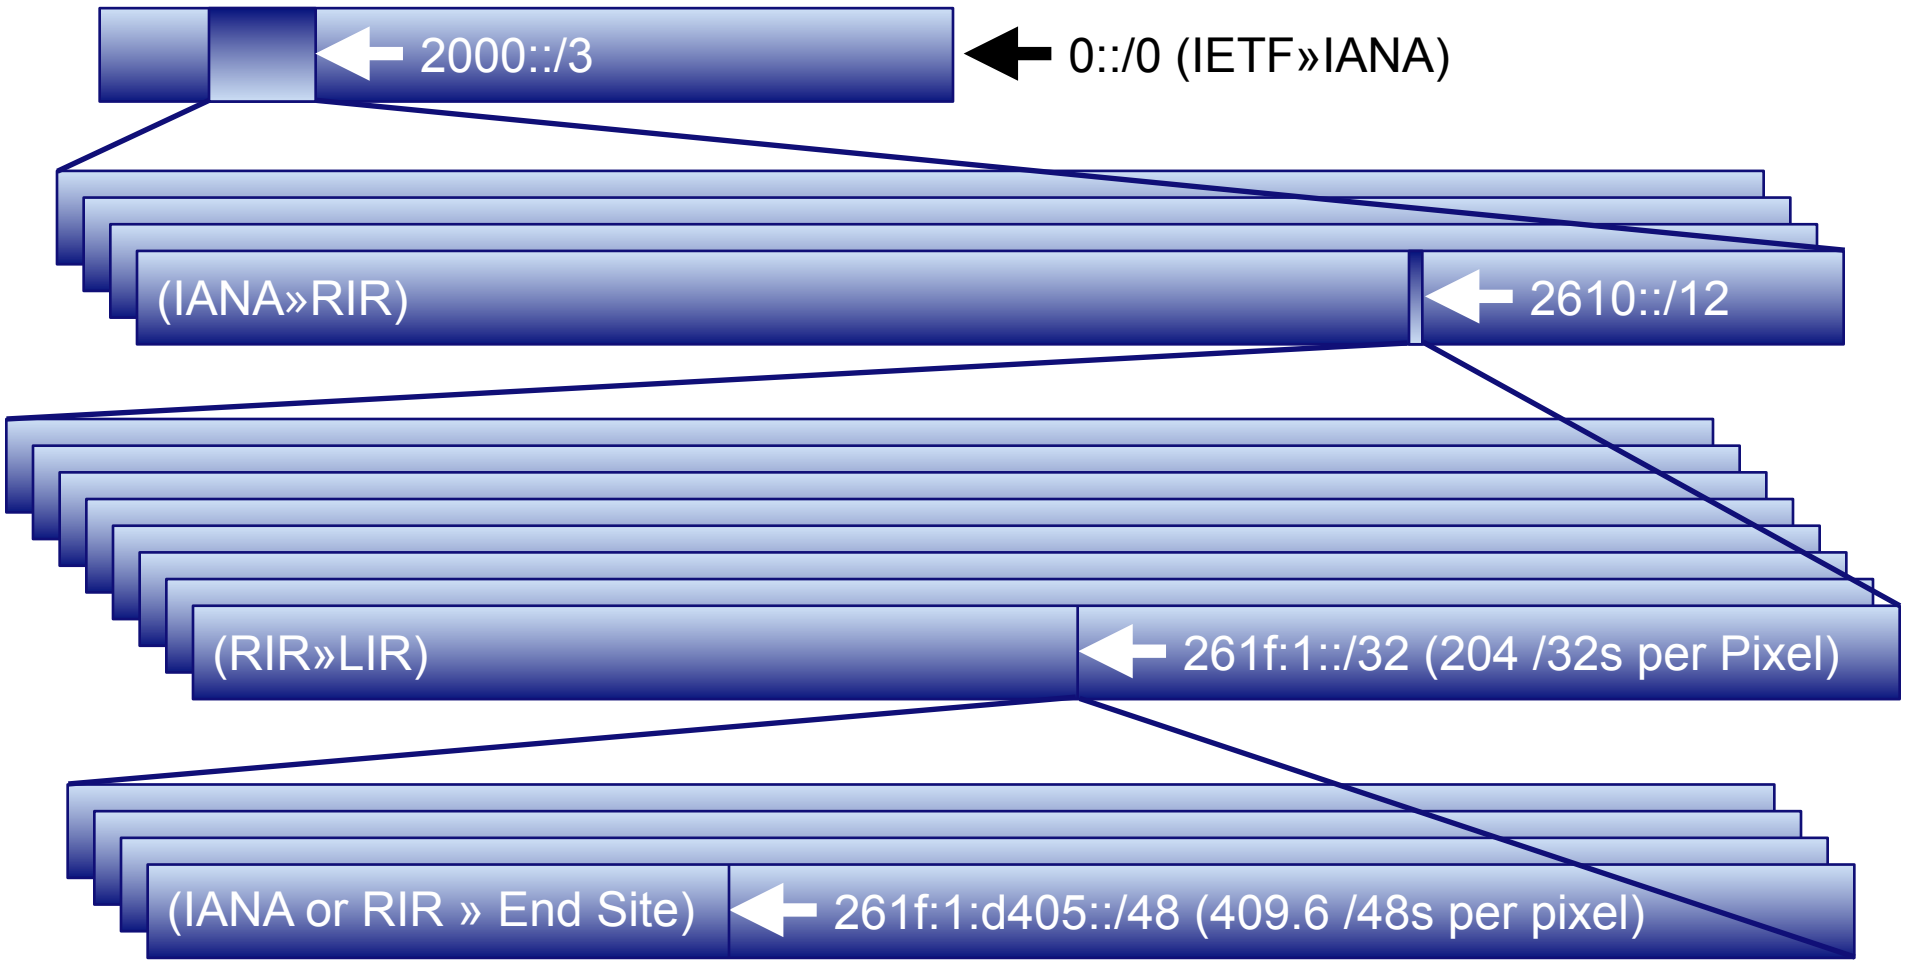

IPv6 -- The basics
How Global Unicast is Allocated

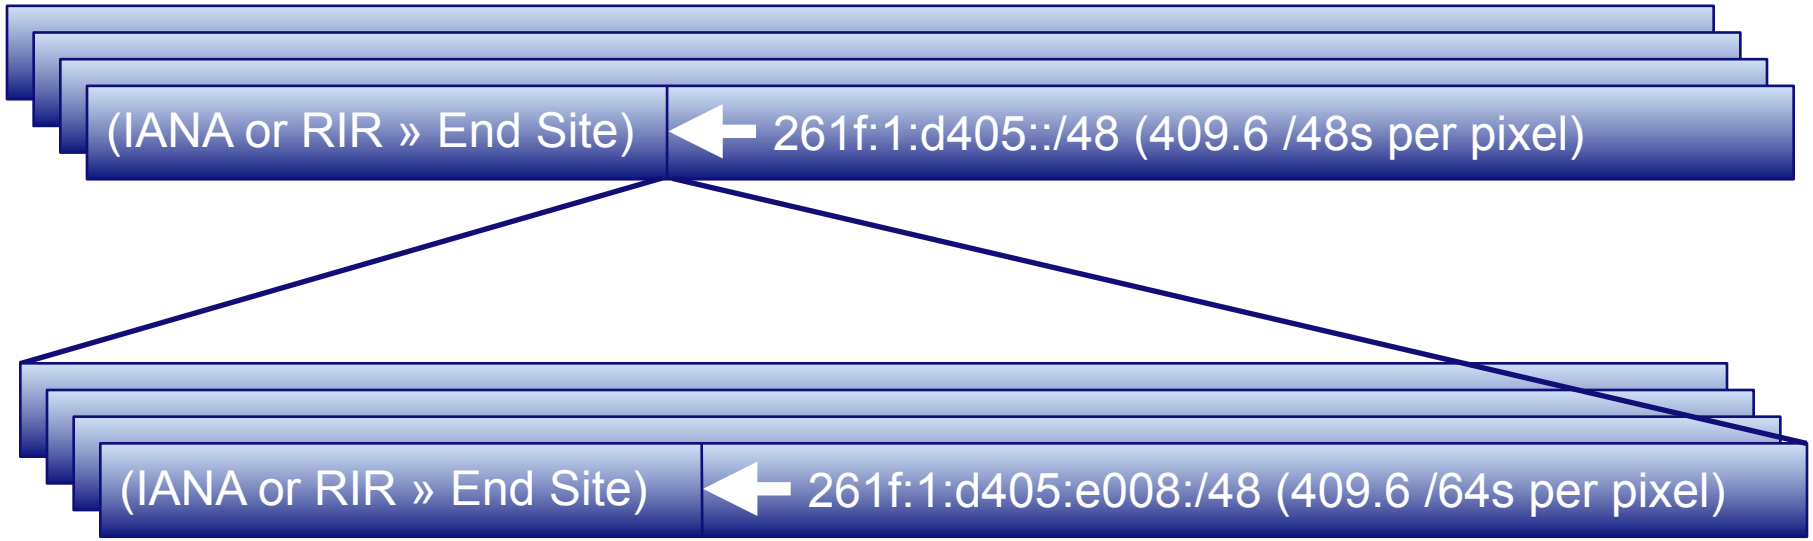

- The Numbers:
  - 8 /3s, one of which is in use
  - 512 /12 allocations to RIRs in first /3 (6 used so far)
  - 1,048,576 LIR /32s in each RIR /12
  - 65,536 /48 Assignments in each /32

IPv6 -- The Basics
Global Unicast in perspective

- The Numbers (cont.)
  - The first /12 assigned to each RIR can support 68,719,476,736 /48 End Sites
  - There are 506 /12s remaining if that's not enough for any particular region.
  - Many ISPs will require more than a /32, but, even if we figure a /28 for every ISP on average, that's still enough addresses for 65,536 ISPs in each RIR region without exhausting their first /12. (There are currently fewer than 30,000 BGP speaking ISPs worldwide)
  - In short... There is more than enough address space for liberal assignments under current and any likely policy.

What is IPv6

- Vastly larger address space
  - 128 bits (IPv4 is 32 bits)
- What does that mean?
- IPv4 4,294,967,296 addresses
- IPv6 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses
- Graphically (if each IP address were a unit of mass):

7 liters of water (IPv4)
Earth 5.9742×10^{24} kg (IPv6)

(1kg = 2.2 Pounds,
15 Pounds of water)

Address Size Comparison (continued)

- Number of networks
  - In IPv4, the generic network chunk is a /24 (254 possible hosts)
  - In IPv6, the generic chunk is a /64 (18,446,744,073,709,551,616 or 18+ quintillion possible hosts)
  - In IPv4, there is room for roughly 14,614,528 /24s.
  - In IPv6, there is room for 18,446,744,073,709,551,616 /64s.

Network Size and Number of networks (The tasty version)

One IPv4 /24 -- 254 M&Ms

One IPv6 /64 -- Enough M&Ms to fill all 5 of the great lakes.

Full Address Space, One M&M per /24 covers 70% of a football field

Full Address Space, One M&M per /64 fills all 5 great lakes.

Comparison based on Almond M&Ms, not plain. Caution! Do not attempt to eat a /64 worth of any style of M&Ms.

If it ain't broke, why fix it?

- It has been broken for years, we've just gotten used to working around it.
  - Various workarounds for NAT
  - NAT itself is a workaround for not enough addresses
  - Huge routing table (300,000+ routes) due to disaggregation from slow-start and other address conservation tradeoffs
  - Poor implementations of address mobility and IPSEC

That doesn't seem like enough for such a major change

- Going from IPv4 only to IPv6 only would be a major change.
- Going from IPv4 only to IPv4/IPv6 dual stack isn't such a major change (but it's not completely minor, either).
- When we run out of IPv4 addresses, the internet will not stop growing. There will be hosts added which do not have directly workable IPv4 addresses.
- Which major change(s) do you want?

The choice of change(s)

- IPv6/Dual Stack -- Continued connectivity to everything.
  - Maybe DS-Lite
  - Maybe 6rd
  - Maybe NAT64/DNS64

- Choices without IPv6
  - LSN/CGN/NAT444(4444...)
  - IPv4 business as usual (while it works)
  - The Mayan Calendar Solution

Alternatives to IPv6

- The only alternative to IPv6 with any traction at all at this point is what is known as “Carrier Grade NAT”.
- Very few test implementations
- None of the test implementations work with instant messenger services (Yahoo, AIM, Jabber, Skype, IRC ALL break)
- VOIP severely impaired or non-functional in all implementations.
- The internet is more than the web and email. CGN does not support much outside of these services.

Cost Benefit Analysis

- Two sets of alternatives to consider:
  - IPv6 vs. CGN
  - IPv6 now vs. IPv6 later
- IPv6 vs. CGN
  - What is the opportunity cost of incredibly poor user experience (virtually guaranteed by CGN)?
  - CGN is complex to set up, more complex to maintain, and, even harder to troubleshoot. What does that cost?
  - Will it even scale?

Cost Benefit continued

- IPv6
  - Unless hardware is extremely old, likely no required upgrades for IPv6 support.
  - Can be relatively simple to deploy by overlaying existing IPv4 technology.
  - Temporarily requires duplicate maintenance efforts for peering sessions, access control lists, prefix filters, etc.
  - Compared to the likely costs of CGN, IPv6 looks cheap in almost every case.

Cost Benefit (Continued)

- IPv6 Now vs. IPv6 later
  - IPv6 offers real savings in the long run
  - Beginning implementation now allows a slow, steady progression to full integration in a controlled manner (planned spending, research, time to seek best pricing).
  - Implementing IPv6 later may require significantly accelerated deployment (emergency spending, increased shipping costs, no time to negotiate)
  - Getting staff exposure to IPv6 while it's not mission critical pays off by reducing training costs and service-affecting outages.

The Ultimate Business Case for IPv6

- There is no “Killer Application”
- There is no “ROI case”
- So, why do it?
- For the same reasons you buy insurance, invested in Y2K compliance and have a disaster recovery plan (you do have one, right?):
  - If you don’t have IPv6 when IPv4 runs out, you will be at an ever increasing disadvantage compared to your competitors that do!

This is the Internet

This is the Internet on NAT444/LSN

Any quesitons?

Customer Premise Equipment and the equipment it connects to.

- The most urgent thing from a carrier perspective is to get vendors to make IPv6-ready CO and CPE solutions available.
- The details depend largely on your topology
  - BPON/GPON -- Better off than DSL, still behind.
  - DSL -- Vendors, start your engines!!! NOW!!
  - DOCSIS -- DOCSIS 3 ready, but CPE issues
  - other?

How to move forward

- Start with a test lab for each phase.
- Deploy IPv6 at one of your peering edges.
- Add IPv6 to the rest of your backbone and your other edges as you connect them.
- Public facing content and interfaces (web, email, etc.)
- In-house “customer” trials
- Consider adding IPv6 capabilities to your enterprise where it makes sense (low hanging fruit)

The test lab

- You don’t need a lot.
- Simulating the full internet is not necessary.
- A small number of routers and end systems is probably sufficient.
- Test configuration elements and become familiar with the configurations and gotchas of whatever vendors apply to your network
- Try out microcosms of various deployment scenarios and break-fix.

Planning Your IPv6 Address Space

- IPv6 is NOT IPv4
- IPv4 -- Driving force in planning was address scarcity with aggregation as a somewhat secondary concern.
- IPv6 -- No scarcity. Get what you need to be able to maximize aggregation without regard for utilization density.
- IPv4 -- Scale was based on hosts.
- IPv6 -- Scale based on networks.

IPv4-think -- Avoid these common mistakes

- Over-conservatism
  - Don't assign various size subnets to stuff. Just accept that a network is a /64, even if it is a point-to-point. There are many advantages to this.
- Disaggregation for density optimization
  - Assign the same size chunk to each site. (Usually a /48 internal, perhaps a /36 or /40 for customers).
  - A few sites may require multiple chunks, that's OK.
  - ROUND UP!!

Rules of Thumb for Address sizing

- Issuing to Customers:
  - Point-to-Point: /64
  - Small site: /~~56~~48 (residential, maybe small business)
  - Normal site: /48 (issue /48 on request without justification even to small site)
  - Multi-site customer: /48 per site
- Allocating to POPs and Facilities:
  - Point-to-Point: /64
  - POP: /36 or /40 (depending on whether you have large (/36) or small (/40) POPs)

Address Sizing (continued)

- POP Allocations
  - A /40 gives you 256 /48 customer assignments per POP. If you need more than that in more than a handful of POPs, go to /36 per POP.
  - A /36 gives you 4096 /48 customer assignments per POP, but, only 16 POPs fit in a /32 that way.
  - If you need to support more than 16 POPs, but, need /36s in most POPs, ask for a /28 instead of a /32. If you need more than a /28 to make it work, ask for a /24, a /20, or even a /16 if that's what you need. (However, expect to provide some serious justification).
  - Start at the bottom (customer assignments) and aggregate upward, rounding up to nibble boundaries at each level.
  - Preserve aggregation by reducing the likelihood for additional prefixes. Try to plan addressing on a 3-year horizon.

Routing Options

- Native IPv6
  - Best choice if available
  - May be uphill battle with upstream providers
  - Worth pushing your upstreams now
- Tunneled Solutions
  - Free tunnels such as http://tunnelbroker.net
  - Good for situations where you can't get native
  - Not ideal in terms of performance
  - Usual preference: 6in4, 6to4, Teredo in that order.

More about Tunnels -- 6in4

- Manual Configuration
- Defined Endpoints
- Essentially like GRE (in fact, can use GRE to tunnel dual-stack over either IPv4 or IPv6)
- Usually minimal “extra topology”
- Easier to troubleshoot (fewer moving pieces which are easier to find than auto-tunneled solutions).

More about Tunnels -- 6to4

- “Server Side” found by anycast
- Automatic, little or no manual configuration required.
- Anycast theoretically minimizes “extra topology”
- As 6to4 servers are deployed topologically closer, automatically migrates tunnel to closer server
- No provision for over/underloaded server balancing.

More about Tunnels -- Teredo

- Mechanism most likely to transit Firewall/NAT
  - Whether you want it to or not!
- Enabled by default on many Windows products
- HUGE security problem for IPv6-unaware enterprises
- Three-party NAT traversal tunneling solution
- Lots of moving parts, works automatically most of the time
- Hard to troubleshoot when it doesn't

What's ready

- Most Routers (Backbone, Core, Enterprise, Workgroup, etc.)
- Most hosts (Linux, BSD, MacOS, Windows*)
- Higher-end Switches (especially most L3 capable switches)
- Many ISPs (such as Hurricane Electric)
- Some Content Providers (NetFlix, Google, YouTube)

*Windows 2000+, but, no IPv6 DNS Resolver before Vista

What's not ready

- CPE
  - Very few consumer-grade residential gateways
  - DHCP-PD mostly unimplemented/untested
  - Consumer Electronics -- The biggest remaining gap!!
- Last-Mile
  - DSLAMs
  - BPON/GPON Concentrators
  - Other consumer aggregator technologies
- Infrastructure Management Systems
  - In-house software
  - Vendor-Provided software

Getting Ready -- Keeping Track

- Hurricane Electric: http://tunnelbroker.net
  - Training
  - Tunnels
  - Statistics
  - Forums
- ARIN IPv6 WIKI: http://www.getipv6.info
  - Status Information about most IPv6-ready products and services
  - User-updatable -- It's a wiki, contribute what you know!
  - Lots of IPv6 Advice and Help available

Getting Connected

- Start by demanding IPv6 from your upstreams. Renewal check-list item.
- If they tell you nobody else is asking for it, escalate. Some ISPs are saying that to everyone who asks.
- If they're not ready, push for a commit date. Consider alternatives if necessary.
- Implement via Tunnel at least to get your infrastructure up and tested.

Getting Connected

- If you are at an Exchange Point, leverage that
- Look for peers with open peering policy
- Hurricane Electric offers free IPv6 Transit as well as open peering for IPv4 and IPv6

Vendor Management

- If your vendor(s) aren't IPv6 ready, it's time to push them
- When possible, avoid new purchases of equipment that isn't IPv6 ready
- Make IPv6 a "checklist item" for product qualification
- TEST IPv6 capabilities, don't just trust the vendor "checklist" on the spec. sheet(s)
- Report Bugs as you encounter them

Vendor Management

- Use tools like Wiki to compare notes about vendors and to share information about vendor accomplishments and shortcomings
- Don’t hesitate to make “me too!” phone calls to vendors to raise the visibility of IPv6 as a priority
- Push on sales, marketing, and support
- Minimal operational experience means vendors are still figuring out IPv6 implementation priorities.

Managing your Management

- IPv6 explained for the CxO:
  - http://businessv6.he.net
- Start the dialogue now, if you haven't already. Let them know what IPv6 is and how it will affect your organization.
- Be honest. Explain why waiting until customers demand it is a recipe for failure.
- Be equally honest about the fact that this is like insurance or disaster recovery... One of those things with no immediate tangible ROI, but, you have to do it anyway.

Training Resources

- On-line
  - Free training such as at http://tunnelbroker.net
  - Bookshelf products such as http://safari.oreilly.com
  - Executive/Business Case: http://businessv6.he.net
- Books from
  - Juniper
  - Cisco Press
  - O'Reilly

Implementation Considerations

- Staff Training
- Prototyping and Development
- Staff Training -- So important I list it twice!
- Backbone Deployment
- Support Department Deployment
- Customer Trials
- Customer Deployment
- Start at an edge and expand, avoid islands where possible

More implementation considerations

- Software Updates
  - Provisioning Systems
  - IP Allocation Systems
  - SWIP/RWHOIS Management Systems
  - Logging/Reporting Systems
  - Monitoring/Alerting Systems
  - Other in-house software
  - Database Schemas
  - Parsers

Q&A

Copy of these and other slides available at:
http://owend.corp.he.net/ipv6/

Contact:

Owen DeLong
IPv6 Evangelist
Hurricane Electric
760 Mission Court
Fremont, CA 94539, USA
http://he.net/

owend at he dot net
+1 (408) 890 7992