

IPv6 Address Planning Shifting Paradigms for a new Internet

Owen DeLong
owend@he.net

Acknowledgements

- Special thanks for:
 - Content and graphics:
 - Mukom Akong Tamon (AfriNIC)
 - Nishal Goburdhan (AfriNIC)
 - Research, Data, and graphics
 - Geoff Huston (APNIC)
 - Inviting me to present this
 - Srinivas Chendi (APNIC)
 - Attending
 - All of you



IPv6 -- The basics

Anatomy of a Global Unicast address

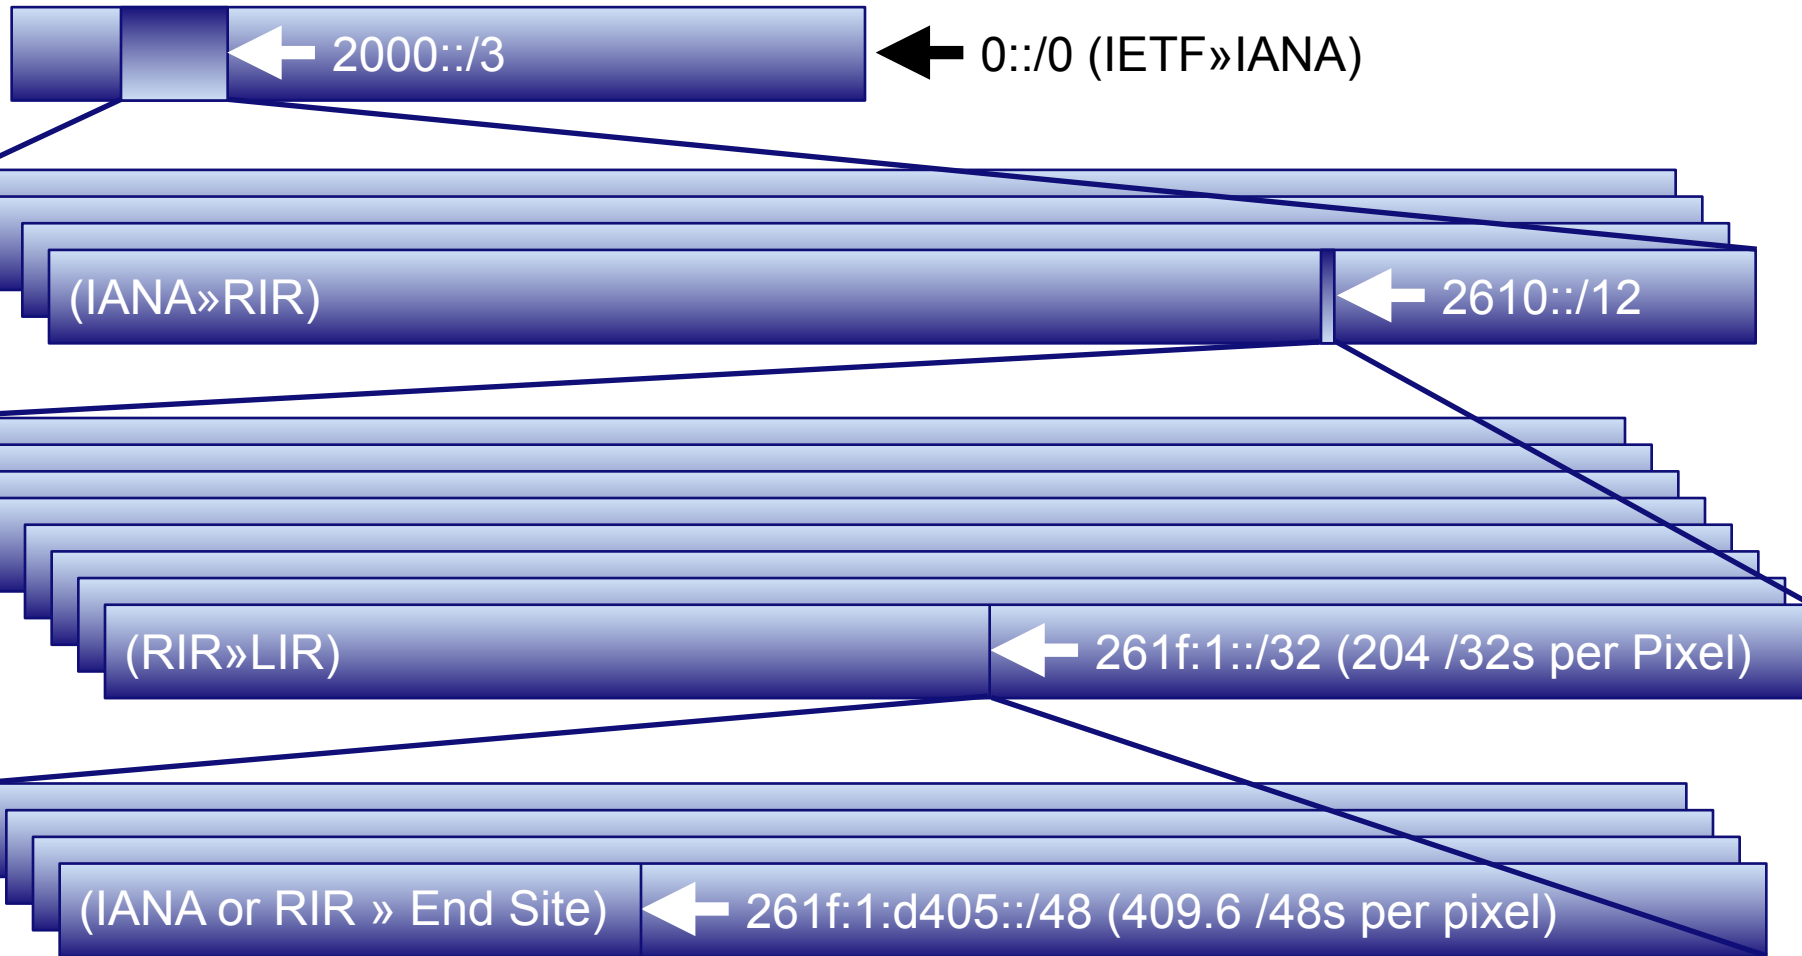
3 bits	9 bits	20 bits	16 bits	16 bits	64 bits
001	IANA to RIR	RIR to ISP	ISP to End Site	Net	Interface ID
001	IANA to RIR	RIR to End Site		Net	Interface ID
3 bits	9 bits	36 bits		16 bits	64 bits

- Every end site gets a /48
- Global Unicast currently being allocated from 2000::/3
 - Top: Provider assigned
 - Bottom: Provider Independent



IPv6 -- The basics

How Global Unicast is Allocated



IPv6 -- The basics

How Global Unicast is Allocated

(IANA or RIR » End Site) ← 261f:1:d405::/48 (409.6 /48s per pixel)

(IANA or RIR » End Site) ← 261f:1:d405:e008:/48 (409.6 /64s per pixel)

- 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 -- Address Planning

Don't oversimplify too much!

- There are lots of people saying “ISPs get /32s, end sites get /48s.”
- That's an unfortunate oversimplification.
- ISPs get AT LEAST a /32 and can get whatever larger allocation they can justify.
- End sites should get at least a /48 and should be given whatever larger assignment they can justify.



Plan a Trial vs. Plan a Deployment

- There's really nothing to be saved by planning your trial address structure separately.
- Take your best stab at planning your real deployment and use that.
 - If you're right, you don't have to renumber a bunch of customers to go from trial to production.
 - If you're wrong, you probably got better data about how you were wrong and why.



IPv6 -- Address Planning Methodology

- Don't start with a predetermined size and figure out how to make your needs fit within it.
- Start by analyzing your needs and apply for a prefix that will meet those needs.
- In your analysis, it's worth while to try and align allocation units to nibble boundaries. A nibble boundary is a single hex digit, or, a number 2^n such that n is a multiple of 4. (e.g. 16, 256, 4096, 65536...)



IPv6 Address Planning Analysis (ISP version)

- Start with the number of end sites served by your largest POP. Figure a /48 for each. Round up to the a nibble boundary. (if it's 3,000 end sites, round up to 4096, for example... a /36 per POP.
- Next, calculate the number of POPs you will have. Include existing POPs and likely expansion for several years. Round that up to a nibble boundary, too. (140 POPs, round up to 256).



IPv6 Address Planning Analysis (ISPs)

- Now that you have an address size for each POP ($4096 = 12$ bits in our example) and a number of POPs ($256 = 8$ bits in our example), you know that you need a total of $\text{POP} \times \text{nPOPs} / 48\text{s}$ for your network ($4096 \times 256 = 1,048,576$ or $12 + 8 = 20$ bits).
- 48 bits - 20 bits is 28 bits, so, you actually need a $/28$ to properly number your network.
- You probably could squeeze this into a $/32$, but, why complicate your life unnecessarily?



IPv6 Address Planning Analysis (End-User Version)

- What's an end-site?
 - A single building, structure, or tenant in a multi-structure building.
- How much do I need for my end-site?
 - This is actually pretty simple in most cases.
 - Up to ~48,000 subnets needed, just give each end-site a /48.
 - If you have an end-site that needs more than 48,000 unique subnets, then assign the necessary number of /48s.



IPv6 Address Planning Analysis (End-user)

- Take the total number of /48s you need for all of your end-sites and round-up to a nibble boundary (if your local RIR policy permits. Currently ARIN is the only RIR that explicitly permits this).
- Once you receive your /48s it is worth considering distributing them to end sites using sparse allocation to the extent practicable.
- Though the RIR will provide a single aggregable prefix, each end site can be an independent /48 and should be administered accordingly. However, when possible, routing should be aggregated.



IPv6 Address Planning

Apply for your addresses

- Now that you know what size block you need, the next step is to contact your friendly neighborhood RIR (Regional Internet Registry) and apply.
- Most RIRs provide either an email-based template or a web-based template for you to fill out to get addresses.
- If you are a single-homed end-user, you usually should get your addresses from your upstream rather than an RIR.



IPv6 Address Planning

The bad news

- The addressing methodology I described above may not be consistent with RIR policy in all regions (yet)*.
- This means you might have to negotiate to a smaller block.
- All RIRs have an open policy process, so, you can submit a proposal to enable this kind of allocation, but, that may not help you immediately.

* Prop-096 in APNIC this week, Adopted 2011-3 in ARIN, mostly permitted in RIPE, not yet discussed in AfriNIC or LACNIC.



IPv6 Address Planning

The good news

- Having things on nibble boundaries is convenient, but, not necessary.
 - ip6.arpa DNS delegations
 - Human Factors
 - Routing Table management
 - Prefix lists
- The techniques that follow should work either way.



IP Address Planning

Carving it up

- These examples are for ISPs.
- For the most part, you've already done this.
- Take the number you came up with for the nPOPs round-up and convert that to a number of bits ($256 = 8$ bits in our example).
- Now, take what the RIR gave you ($/28$ in our example) and add that number to the above number ($28+8 = 36$) and that's what you need for each POP (a $/36$ in our example).



This is the Internet



This is the Internet on IPv4 (2012)



Any quesitons?



IPv6 Address Planning

Carving it up

- Now let's give address segments to our POPs.
- First, let's reserve the first /48 for our infrastructure. Let's use 2000:db80 - 2000:db8f as our example /28.
- Since each POP gets a /36, that means we have 2 hex digits that designate a particular POP.
- Unfortunately, in our example, that will be the last digit of the second group and the first digit of the third group.



IPv6 Addressing

Carving it up


- Strategy
 - Sequential Allocation
 - Advantage: Simple, easy to follow
 - Advantage: POP Numbers correspond to addresses
 - DisAdvantage: Complicates unexpected growth
 - Allocation by Bisection
 - Advantage: Simplifies growth
 - Advantage: Greatest probability of Aggregation
 - Disadvantage: “Math is hard. Let’s go shopping!”



IPv6 Addressing

Allocation by Bisection

- Bisection? What does THAT mean?
- Simple... It means to cut up the pieces by taking the largest remaining piece and cutting in half until you have the number of pieces you need.
- Imagine cutting up a pie into 8 pieces...



First, we cut it in half.
Then we cut it in half again.
And finally a fourth cut

IPv6 Addressing

Allocation by Bisection

- It's a similar process for IPv6 addresses.
 - Let's start with our 2001:db80::/28 prefix.
 - We've already allocated 2001:db80:0000::/48
 - Our available space is now 2001:db80:0001:: to 2001:db8f:ffff:ffff:ffff:ffff:ffff:ffff. Cutting that in half we get 2001:db88:0000::/36 as our first POP address.
 - That leaves the largest chunk at 2001:db88:1000:: to 2001:db8f:ffff:ffff:ffff:ffff:ffff:ffff. Cutting that in half, we get 2001:db8c:0000::/36 as our next POP



IPv6 Address Planning

Allocation by Bisection

- After repeating this for 19 POP allocations, we have a table that looks like this:

Infrsastructure	2001:db80:0000:/48	POP1	2001:db88:0000:./36
POP12	2001:db80:8000:./36	POP13	2001:db88:8000:./36
POP8	2001:db81:0000:./36	POP9	2001:db89:0000:./36
POP4	2001:db82:0000:./36	POP5	2001:db8a:0000:./36
POP14	2001:db83:0000:./36	POP15	2001:db8b:0000:./36
POP2	2001:db84:0000:./36	POP3	2001:db8c:0000:./36
POP16	2001:db84:8000:./36	POP17	2001:db8c:8000:./36
POP10	2001:db85:0000:./36	POP11	2001:db8d:0000:./36
POP6	2001:db86:0000:./36	POP7	2001:db8e:0000:./36
POP18	2001:db87:0000:./36	POP19	2001:db8f:0000:./36



IPv6 Address Planning

Allocation by Bisection

- Notice how by doing that, most of the /36s we created have 15 more /36s before they run into allocated space and all have at least 7.
- Notice also that if any POPs get larger than we expect, we can expand them to /35s, /34s, /33s, and most all the way to a /32 without having to renumber.
- By default, at /36, each pop has room for 4096 /48 customers. End sites that need more than a /48 should be extremely rare*.



IPv6 Address Planning Allocation by Bisection

- * End Site means a single customer location, not a single customer. Many customers may need more than a /48, but, with 65,536 /64 subnets available, even the largest building should be addressable within a /48.





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



The end

Thank you

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

