

IPv6 Addressing

ISP Workshops



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)

Last updated 15th January 2019

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
 - Please email *workshop (at) bgp4all.com*

Philip Smith

Agenda

- Recap: how it worked with IPv4
- Getting IPv6 address space
- Constructing a scalable IPv6 address plan
- IPv6 address plan example

How it used to be



Looking back at IPv4

How did it work for IPv4?

- Up until 1994:
 - Operators applied to InterNIC for address space
 - 1993 onwards: included RIPE NCC and APNIC, the first two RIRs
 - Class A: Big organisations
 - Class B: Medium organisations
 - From 1992 onwards, multiple class Cs often handed out instead of single class Bs
 - Class C: Small organisations
- From 1994 onwards (classless Internet)
 - Address space distributed by InterNIC (replaced by ARIN in 1998) and the other RIRs
 - Distribution according to demonstrated need (not want)

IPv4 address plans (pre 1994)?

- Prior to 1994, doing an address plan in IPv4 was very simple
- Class C was used for one LAN
 - If entity had more than one LAN, they'd normally get a class B
- An organisation with a class B had 256 possible LANs
 - And that was more than most networks had in those days
- Organisations with more than 256 LANs tended to be Universities, big IT companies, etc
 - They either had multiple class Bs, or even a class A

Typical early 90s address plan

- Organisation was not connected to the Internet as such
 - But used TCP/IP internally
- Would generally use 10.0.0.0/8
 - Or any other class A that InterNIC had not handed out
- 10.X.Y.Z was a typical layout, where:
 - X = building number
 - Y = floor number
 - Z = host address
 - Where each subnet was a /24 (like a class C)
- When these organisations connected to the Internet, they had to renumber
 - Often into a /19 (the minimum allocation then)

IPv4 address plans (post 1994)?

- In the classful Internet days, IP address planning didn't really exist
 - The address space was big enough for most needs, as the number of devices and LANs were small
- With the arrival of classless Internet, and IPv4 run out in the early 90s
 - IP address planning was needed
 - Organisations got address space according to demonstrated need
 - A previous class B might now only get a /19
 - LANs no longer were automatically /24s
 - etc

IPv4 address plans (post 1994)?

- Advent of NAT assisted with delaying IPv4 runout
 - End-user got single public address, and NATed on to that address
 - (End-users could get lazy again)
- Operators became more careful:
 - RIR policy required “demonstrated need”
 - Further allocations made only when existing allocations were proven to be mostly used up
 - Started assigning address space across backbone according to the needs of the infrastructure
 - No gaps, but still no real plan
 - /30s for point-to-point links etc
 - Although the “plans” often separated infrastructure address space from what went to customers

IPv4 address plans (today)

- Chaotic?
- Unstructured?
- Undocumented?
- With IPv4 address space almost all depleted
 - Operators becoming ever more creative
 - Operators extracting the last “drops” from their address space holdings
 - It is a scramble just to keep network infrastructure addressed with public IPv4
 - Some operators even use NAT within their backbones
 - Some operators are reclaiming IPv4 address space loaned to their customers
 - This hotch potch cannot and does not lead to good planning

IPv4 address plans (today)

- More serious issues – because of the lack of structure, lack of planning:
 - **Infrastructure security filters become very hard to manage**
 - Adding yet another small block of IPv4 addresses to perimeter and control plane filters
 - **Traffic engineering is more challenging**
 - Lots of small blocks of address space to manage and manipulate
 - With impacts on size of the global routing table too!
 - **Infrastructure addressing is difficult to manage**
 - Loopbacks and backbone point-to-point links no longer out of one contiguous block
 - **Access address pool resizing**
 - Broadband access pools renumbering, reassigning, etc

IPv6

- IPv6 changes all this
- Address space delegations are generous
 - Reminders of the “old days” of classful IPv4
- No NAT
- Address planning is very possible
- Address planning is very necessary
- Documentation is very necessary
- Operators accustomed to handling IPv4 in the 1980s and early 1990s might be able to use those old skills for IPv6 !

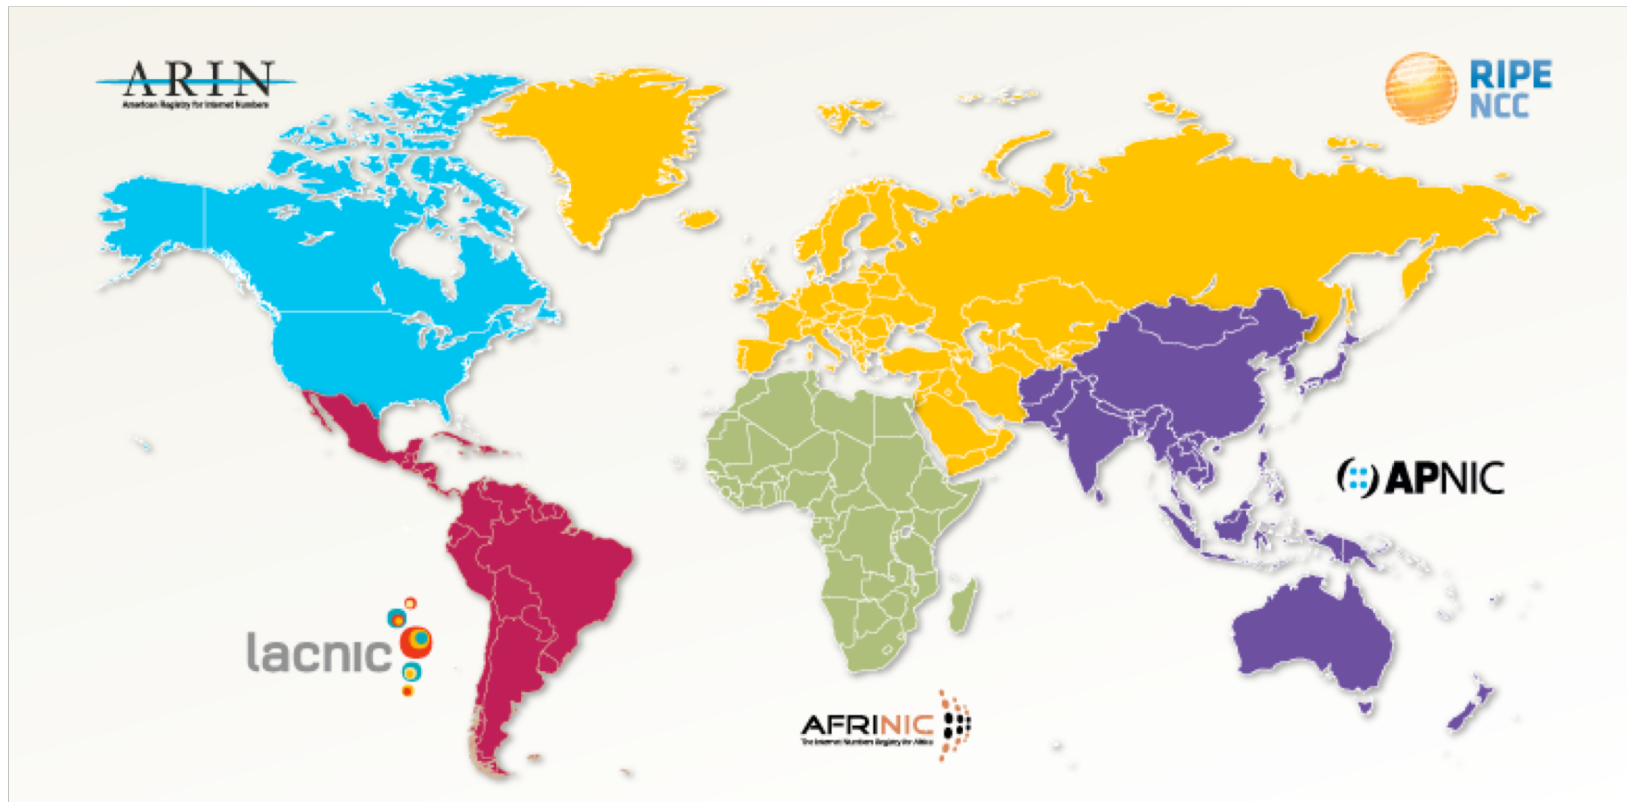
IPv6 Address Planning



Where to get IPv6 addresses

- Your upstream ISP
- Africa
 - AfriNIC – <http://www.afrinic.net>
- Asia and the Pacific
 - APNIC – <http://www.apnic.net>
- North America
 - ARIN – <http://www.arin.net>
- Latin America and the Caribbean
 - LACNIC – <http://www.lacnic.net>
- Europe and Middle East
 - RIPE NCC – <http://www.ripe.net/info/ncc>

Internet Registry Regions



Getting IPv6 address space (1)

- **From your Regional Internet Registry**
 - Become a member of your Regional Internet Registry and get your own allocation
 - Membership open to all organisations who are operating a network
 - Address allocation policies listed on the individual RIR website
 - Minimum allocation is a /32 (or larger if you will have more than 65k /48 assignments)

Getting IPv6 address space (2)

- **From your upstream ISP**
 - Receive a /48 from upstream ISP's IPv6 address block
 - Receive more than one /48 if you have more than 65k subnets
- **If you need to multihome:**
 - Apply for a /48 assignment from your RIR
 - Multihoming with the provider's /48 will be operationally challenging
 - Provider policies, filters, etc

Do NOT use 6to4

- Some entities still use 6to4
 - **Deprecated in May 2015 due to serious operational and security issues**
 - Read <https://tools.ietf.org/rfc/rfc7526.txt> (BCP196) for the reasoning why

- FYI: 6to4 operation:
 - Take a single public IPv4 /32 address
 - 2002:<ipv4 /32 address>::/48 becomes your IPv6 address block, giving 65k subnets
 - Requires a 6to4 gateway
 - 6to4 is a means of connecting IPv6 islands across the IPv4 Internet

Nibble Boundaries

- IPv6 offers network operators more flexibility with addressing plans
 - Network addressing can now be done on nibble boundaries
 - For ease of operation
 - Rather than making maximum use of a very scarce resource
 - With the resulting operational complexity
- A nibble boundary means subnetting address space based on the address numbering
 - Each number in IPv6 represents 4 bits = 1 nibble
 - Which means that IPv6 addressing can be done on 4-bit boundaries

Nibble Boundaries – example

- Consider the address block 2001:DB8:0:10::/61
 - The range of addresses in this block are:

2001:0DB8:0000:0010:0000:0000:0000:0000
to
2001:0DB8:0000:0017:FFFF:FFFF:FFFF:FFFF



- Note that this subnet only runs from 0010 to 0017.
- The adjacent block is 2001:DB8:0:18::/61

2001:0DB8:0000:0018:0000:0000:0000:0000
to
2001:0DB8:0000:001F:FFFF:FFFF:FFFF:FFFF

- The address blocks don't use the entire nibble range

Nibble Boundaries – example

- Now consider the address block 2001:DB8:0:10::/60
 - The range of addresses in this block are:

2001:0DB8:0000:0010:0000:0000:0000:0000
to
2001:0DB8:0000:001F:FFFF:FFFF:FFFF:FFFF



- Note that this subnet uses the entire nibble range, 0 to F
- Which makes the numbering plan for IPv6 simpler
 - This range can have a particular meaning within the ISP block (for example, infrastructure addressing for a particular PoP)

Addressing Plans – Infrastructure

- All Network Operators should obtain a /32 from their RIR
- Address block for router loop-back interfaces
 - Number all loopbacks out of **one** /64
 - /128 per loopback
- Address block for infrastructure (backbone)
 - /48 allows 65k subnets
 - /48 per region (for the largest multi-national networks)
 - /48 for whole backbone (for the majority of networks)
 - Infrastructure/backbone usually does NOT require regional/geographical addressing
 - Summarise between sites if it makes sense

Addressing Plans – Infrastructure

- What about LANs?
 - /64 per LAN
- What about Point-to-Point links?
 - Protocol design expectation is that /64 is used
 - /127 now recommended/standardised
 - <http://www.rfc-editor.org/rfc/rfc6164.txt>
 - (reserve /64 for the link, but address it as a /127)
 - Other options:
 - /126s are being used (mimics IPv4 /30)
 - /112s are being used
 - Leaves final 16 bits free for node IDs
 - Some discussion about /80s, /96s and /120s too
 - Some equipment doesn't support /127s ☹

Addressing Plans – Infrastructure

□ NOC:

- ISP NOC is “trusted” network and usually considered part of infrastructure /48
 - Contains management and monitoring systems
 - Hosts the network operations staff
 - take the last /60 (allows enough subnets)

□ Critical Services:

- Network Operator’s critical services are part of the “trusted” network and should be considered part of the infrastructure /48
- For example, Anycast DNS, SMTP, POP3/IMAP, etc
 - Take the second /64
 - (some operators use the first /64 instead)

Addressing Plans – Link from ISP to End-Site

□ Option One:

- Use ipv6 unnumbered
- Which means no global unicast ipv6 address on the point-to-point link
- Router adopts the specified interface's IPv6 address
 - Router doesn't actually need a global unicast IPv6 address to forward packets

```
interface loopback 0
  ipv6 address 2001:DB8::1/128
interface serial 1/0
  ipv6 address unnumbered loopback 0
```

Addressing Plans – Link from ISP to End-Site

□ Option Two:

- Use the second /48 for point-to-point links
- Divide this /48 up between PoPs
- Example:
 - For 10 PoPs, dividing into 16, gives /52 per PoP
 - Each /52 gives 4096 point-to-point links
 - Adjust to suit!
- Useful if ISP monitors point-to-point link state for customers
 - Link addresses are **untrusted**, so do not want them in the first /48 used for the backbone &c
- Aggregate per router or per PoP and carry in iBGP (not ISIS/OSPF)

Addressing Plans – End-Site

- RFC6177/BCP157 describes assignment sizes to end-sites
 - Original (obsolete) IPv6 design specification said that end-sites get one /48
 - Operators now must recognise that end-sites need to get enough IPv6 address space (multiples of /64) to address all subnets for the foreseeable future
- **In typical deployments today:**
 - /64 if end-site will only ever be a LAN (not recommended!!)
 - /56 for small end-sites (e.g. home/office/small business)
 - /48 for large end-sites
- **Observations:**
 - RFC7084 specifies Basic Requirements for IPv6 Customer Edge Routers
 - Including ability to be able to request at least a /60 by DHCPv6-PD
 - Don't assume that a mobile end-site needs only a /64 – 3GPP Release 10 introduces DHCPv6-PD for tethering
 - Some operators are distributing /60s to their smallest customers!!

Addressing Plans – End-Site

- Broadband Example:
 - DHCPv6 pool is a /48
 - DHCPv6 hands out /56 per customer
 - Which allows for 256 customers per pool
 - If BRAS has more than 256 customers, increase pool to a /47
 - This allows for 512 customers at /56 per customer
 - The whole nibble (/44) allows for 4096 delegations
 - In all cases, BRAS announces entire pool as one block by iBGP

Addressing Plans – End-Site

- Business “leased line”:
 - /48 per customer
 - One stop shop, no need for customer to revisit ISP for more addresses until all 65k subnets are used up
- Hosted services:
 - One physical server per vLAN
 - One /64 per vLAN
 - How many vLANs per PoP?
 - /48 reserved for entire hosted servers across backbone
 - Internal sites will be subnets and carried by iBGP

Addressing Plans – End-Site

- Geographical delegations to Customers:
 - Network Operator subdivides /32 address block into geographical chunks
 - E.g. into /36s
 - Region 1: 2001:DB8:1xxx::/36
 - Region 2: 2001:DB8:2xxx::/36
 - Region 3: 2001:DB8:3xxx::/36
 - etc
 - Which gives 4096 /48s per region
 - For Operational and Administrative ease
 - Benefits for traffic engineering if Network Operator multihomes in each region

Addressing Plans – End-Site

- Sequential delegations to Customers:
 - After carving off address space for network infrastructure, Network Operator simply assigns address space sequentially
 - Eg:
 - Infrastructure: 2001:DB8:0::/48
 - Customer P2P: 2001:DB8:1::/48
 - Customer 1: 2001:DB8:2::/48
 - Customer 2: 2001:DB8:3::/48
 - etc
 - Useful when there is no regional subdivision of network, no regional multihoming, or traffic engineering needs

Addressing Plans – Routing Considerations

- Carry Broadband pools in iBGP across the backbone
 - Not in OSPF/IS-IS
- Multiple Broadband pools on one BRAS should be aggregated if possible
 - Reduce load on iBGP
- Aggregating leased line customer address blocks per router or per PoP is undesirable:
 - Interferes with ISP's traffic engineering needs
 - Interferes with ISP's service quality and service guarantees

Addressing Plans – Traffic Engineering


- Smaller providers will be single homed
 - The customer portion of the ISP's IPv6 address block will usually be assigned sequentially
- Larger providers will be multihomed
 - Two, three or more external links from different providers
 - Traffic engineering becomes important
 - Sequential assignments of customer addresses will negatively impact load balancing

Addressing Plans – Traffic Engineering

- ISP Router loopbacks and backbone point-to-point links make up a small part of total address space
 - And they don't attract traffic, unlike customer address space
- Links from ISP Aggregation edge to customer router needs one /64
 - Small requirements compared with total address space
 - Some ISPs use IPv6 unnumbered
- Planning customer assignments is a very important part of multihoming
 - Traffic engineering involves subdividing aggregate into pieces until load balancing works

Unplanned IP addressing


- ISP fills up customer IP addressing from one end of the range:



- Customers generate traffic
 - Dividing the range into two pieces will result in one /33 with all the customers and the ISP infrastructure the addresses, and one /33 with nothing
 - No loadbalancing as all traffic will come in the first /33
 - Means further subdivision of the first /33 = harder work

Planned IP addressing

- If ISP fills up customer addressing from both ends of the range:



- Scheme then is:
 - First customer from first /33, second customer from second /33, third from first /33, etc
- This works also for residential versus commercial customers:
 - Residential from first /33
 - Commercial from second /33

Planned IP Addressing


- This works fine for multihoming between two upstream links (same or different providers)
- Can also subdivide address space to suit more than two upstreams
 - Follow a similar scheme for populating each portion of the address space
- Consider regional (geographical) distribution of customer delegated address space
- Don't forget to always announce an aggregate out of each link

Addressing Plans – Advice

- Customer address assignments should not be reserved or assigned on a per PoP basis
 - Follow same principle as for IPv4
 - Subnet aggregate to cater for multihoming needs
 - Consider regional delegation
 - ISP iBGP carries customer nets
 - Aggregation within the iBGP not required and usually not desirable
 - Aggregation in eBGP is very necessary
- Backbone infrastructure assignments:
 - Number out of a **single** /48
 - Operational simplicity and security
 - Aggregate to minimise size of the IGP


Addressing Plans – Scheme

□ Looking at Infrastructure:




Addressing Plans – Scheme

□ Looking at Infrastructure:



□ Alternative:



Addressing Plans Planning

- Registries will usually allocate the next block to be contiguous with the first allocation
 - (RIRs use a sparse allocation strategy – industry goal is aggregation)
 - Minimum allocation is /32
 - Very likely that subsequent allocation will make this up to a /31 or larger (/28)
 - So plan accordingly

Addressing Plans (contd)

- Document infrastructure allocation
 - Eases operation, debugging and management
- Document customer allocation
 - Customers get /48 each
 - Prefix contained in iBGP
 - Eases operation, debugging and management
 - Submit network object to RIR Database

Addressing Tools

□ Examples of IP address planning tools:

- NetDot netdot.uoregon.edu
- OpenNetAdmin opennetadmin.com
- HaCi sourceforge.net/projects/haci
- Racktables racktables.org
- IPAT nethead.de/index.php/ipat
- freeipdb home.globalcrossing.net/~freeipdb/

□ Examples of IPv6 subnet calculators:

- ipv6gen code.google.com/p/ipv6gen/
- sipcalc www.routemeister.net/projects/sipcalc/

Example Address Plan



Example Address Plan

- Generic Network Operator
 - Has 2001:DB8::/32 address block
 - Takes first /48 for network infrastructure
 - First /64 for loopbacks
 - Last /60 for NOC
 - Takes second /48 for point-to-point links to customer sites
 - Remainder of address space for delegation to customers, content hosting and broadband pools
- Network Operator has 20 locations (Points of Presence) around the country

Example: Loopback addresses

- 2001:DB8:0::/48 is used for infrastructure
- Out of this, 2001:DB8:0:0::/64 is used for loopbacks
 - Each loopback is numbered as a /128
- Scheme adopted is:
 - 2001:DB8::XXYY/128
 - Where XX is the PoP number (01 through FF)
 - Where YY is the router number (01 through FF)
 - Scheme is good for:
 - 255 PoPs
 - 255 routers per PoP
 - keeping addresses small/short

Loopbacks Example

PoP 1 Routers	Loopbacks	PoP 10 Routers	Loopbacks
cr1	2001:DB8::101/128	cr1	2001:DB8::A01/128
cr2	2001:DB8::102/128	cr2	2001:DB8::A02/128
br1	2001:DB8::103/128	sr1	2001:DB8::A05/128
br2	2001:DB8::104/128	sr2	2001:DB8::A06/128
sr1	2001:DB8::105/128	ar1	2001:DB8::A10/128
sr2	2001:DB8::106/128	ar2	2001:DB8::A11/128
ar1	2001:DB8::110/128	gw1	2001:DB8::A20/128
ar2	2001:DB8::111/128	gw2	2001:DB8::A21/128
gw1	2001:DB8::120/128	etc...	
gw2	2001:DB8::121/128		
etc...			

Example: Backbone Point-to-Point links

- Backbone Point-to-Point links come out of Infrastructure block 2001:DB8:0::/48
 - Scheme adopted is:
 - 2001:DB8:0:XXYY::Z/64
 - Where
 - XX is the PoP number (01 through FF)
 - YY is the LAN number (00 through 0F)
 - YY is the P2P link number (10 through FF)
 - Z is the interface address (0 or 1)
 - Scheme is good for 16 LANs and 240 backbone PtP links per PoP, and for 255 PoPs

LANs and PtP Links Example

PoP 1	Subnet	PoP 14	Subnet
LAN1	2001:DB8:0:101::/64	LAN1	2001:DB8:0:E01::/64
LAN2	2001:DB8:0:102::/64	LAN2	2001:DB8:0:E02::/64
LAN3	2001:DB8:0:103::/64	LAN3	2001:DB8:0:E03::/64
PtP1	2001:DB8:0:111::/64	LAN4	2001:DB8:0:E04::/64
PtP2	2001:DB8:0:112::/64	LAN5	2001:DB8:0:E05::/64
PtP3	2001:DB8:0:113::/64	PtP1	2001:DB8:0:E11::/64
PtP4	2001:DB8:0:114::/64	PtP2	2001:DB8:0:E12::/64
PtP5	2001:DB8:0:115::/64	PtP3	2001:DB8:0:E13::/64
PtP6	2001:DB8:0:116::/64	etc...	
PtP7	2001:DB8:0:117::/64		
etc...			

Note: PtP links have /64 reserved but are addressed as /127s

Links to Customers

- Some ISPs use “ip unnumbered” for IPv4 interface links
 - So replicate this in IPv6 by using “ipv6 unnumbered” to address the links
 - This will not require one /48 to be taken from the ISP’s /32 allocation
- Other ISPs use real routable addresses
 - So set aside the second /48 for this purpose
 - Gives 65536 possible customer links, assuming a /64 for each link

Customer Links Example

Customer	Point to point link address
Customer 1	2001:DB8:1:0::/64
Customer 2	2001:DB8:1:1::/64
Customer 3	2001:DB8:1:2::/64
Customer 4 (link one)	2001:DB8:1:3::/64
Customer 4 (link two)	2001:DB8:1:4::/64
Customer 5	2001:DB8:1:5::/64
Customer 6	2001:DB8:1:6::/64
etc...	

Note1: PtP links are numbered out of 2001:DB8:1::/48

Note2: PtP links have /64 reserved but are addressed as /127s

Example: Allocations from the /32

- Master allocation documentation would look like this:

Category	Purpose
Single /64	Loopbacks
Single /60	NOC
Single /48	Backbone Point-to-Point links (/64 each)
Single /48	Customer Point-to-Point links (/64 each)
Single /40	65536 Broadband Customers in Region 1 (/56 each)
Single /40	256 Enterprise Customers in Region 1 (/48 each)
Single /40	65536 Broadband Customers in Region 2 (/56 each)
Single /40	256 Enterprise Customers in Region 2 (/48 each)
Etc...	

Example: Allocations from the /32

□ Detailed documentation:

- Region – /36
- Pool – /40
- Customer – /48
- NOC – /60
- LANs – /64
- PtPs – /127
- Loops – /128

Address Blocks	Purpose
2001:DB8:0::/48	Infrastructure (Loops, NOC, PtP)
2001:DB8:1::/48	Customer Point-to-Point Links
2001:DB8:0110::/48	Customer One in Region 1
2001:DB8:0111::/48	Customer Two in Region 1
2001:DB8:0112::/48	Customer Three in Region 1
2001:DB8:0200::/40	Broadband Pool 1 in Region 1
2001:DB8:0300::/40	Broadband Pool 2 in Region 1
2001:DB8:8110::/48	Customer One in Region 2
2001:DB8:8111::/48	Customer Two in Region 2
2001:DB8:8112::/48	Customer Three in Region 2
2001:DB8:8200::/40	Broadband Pool 1 in Region 2
2001:DB8:8300::/40	Broadband Pool 2 in Region 2

Summary

- First /48 for infrastructure
 - Out of that, first /64 for Loopbacks
- PoP structure within IPv6 addressing is very possible
 - Greater flexibility than with IPv4
 - Possible to come up with a simple memorable scheme
- Documentation vitally important!

Presentation Recap

- How it worked with IPv4
- Getting IPv6 address space
- Constructing a scalable IPv6 address plan
- IPv6 address plan example

IPv6 Addressing



ISP Workshops