

IPv6 address planning

Ijtsch van Beijnum

RIPE NCC::Educa IPv6-only, 8 June 2020

Our plan

- Hexadecimal: learn it [[WIKIPEDIA](#)], love it! 🥰
- Importance of IPv6 address planning
- IPv6 address types, sizes and subnetting
- IPv6 address structure
- Planning the subnet bits
- Configuring the local bits
- DNS server and router addresses
- Organizations with subdivisions
- Questions?

Importance of IPv6 address planning

- You will have to make an IPv6 address plan
 - the only question is how many...
- Ideal world:
 1. create the perfect IPv6 address plan
 2. request address space
 3. roll out IPv6
 4. profit

Importance of IPv6 address planning

- Real world: you will make mistakes, so try to build in flexibility and adjust quickly
- At least you have some address space to waste
 - so err on too big rather than too small
 - get rid of "IPv4 thinking"!
- Change is hard, so it's only worth it to make *big* changes
- Or try out IPv6 in a small way first to figure it out
 - but have the discipline to throw out your test setup and start from scratch!

IPv6 address types

- Link-local: not unique
 - created and used automatically
 - do not try to manage or use these yourself
- Global unicast: "regular" IPv6 addresses
 - you use these 99% (100%?) of the time
- Unique Site Local (ULA): unique, but not routable over the internet
 - a bit like RFC 1918 addresses but without NAT
 - very specific use cases

Assignment size

- (Assignment is RIPE-speak for the address block you get from your ISP)
- Home users often get /56, /60 or even /64
- For organizations, default size is **/48**
 - that means: 48 of the 128 address bits are given/fixed
 - you can fill in the remaining 80 bits yourself
 - even if you really don't need that much: smaller than /48 makes address planning harder
- Also: ISPs usually only accept /48 and larger blocks in BGP
 - so provider independent addresses must be at least /48

Subnetting

- IPv6 is classless: routers can deal with any size
- But: IPv6 addressing architecture [[RFC 4291](#)] says that the host part of the address must be 64 bits
- So 48 bits are given and 64 bits are used to number devices, this leaves $128 - 48 - 64 = 16$ bits to number subnets

IPv6 address structure

(Remember IPv6 address notation [\[RFC 5952\]](#))

| 0 | 16 | 32 | 48 | 64 | 80 | 96 | 112 | 127 |
|-------|------|------|------|------|----|----|-----|-----|
| 2001: | db8: | 188: | 301: | 145: | 0: | 2: | 10 | |


| 0 | 16 | 32 | 48 | 64 | 80 | 96 | 112 | 127 |
|-------|-------|-------|-------|-------|-------|-------|------|-----|
| 3ffe: | 4700: | 1f0b: | 1289: | cd06: | e4b7: | 247e: | 1cfe | |

IPv6 address structure

(Every digit in the IPv6 address is exactly 4 bits)



IPv6 address structure



Planning the subnet bits

- Why do we need to split our network into subnets?
 - to allow efficient / effective / robust **routing**, like:
 - each floor its own subnet
 - each rack in the datacenter its own subnet
 - each subnet must be confined to one location
 - for **security**, like:
 - guest network subnet(s)
 - work station subnet(s)
 - front end server subnet(s)
 - back end server subnet(s)

The easy way: VLAN IDs

- If it's nice outside and you want to leave work early instead of address planning the rest of the day...
- Put the VLAN ID in the subnet bits, like:
 - VLAN **1**: 2001:db8:edca:**1**::/64
 - VLAN **28**: 2001:db8:edca:**28**::/64
 - VLAN **3040**: 2001:db8:edca:**3040**::/64
- Still leaves all subnet values with a - f in them and above 4095
- Of course pay attention to your VLAN numbering!

Subnetting examples

| Location | Type | Instance |
|------------------------|----------------------|-----------------|
| Old city center office | Guest Wi-Fi | Floor 0 |
| | BYOD Wi-Fi | |
| | Managed workstations | |
| | Printers | Floor 0 and 1 |
| | Guest Wi-Fi | Floor 1 |
| | BYOD Wi-Fi | |
| | Managed workstations | |
| New tower office | Guest Wi-Fi | Floor 23 |
| | BYOD Wi-Fi | |
| | Managed workstations | |
| | Printers | Floor 23 and 24 |
| | BYOD Wi-Fi | Floor 24 |
| | Managed workstations | |

| Location | Type | Instance |
|----------------|-------------------|----------|
| Datacenter DC1 | Front end servers | Rack 11 |
| | Database servers | |
| | Storage servers | |
| | Front end servers | Rack 13 |
| | Database servers | |
| | Storage servers | |
| Datacenter DC2 | Front end servers | Rack A4 |
| | Database servers | |
| | Storage servers | |
| | Front end servers | Rack A5 |
| | Database servers | |
| | Storage servers | |

Location or type first

- Location bits come first, then type bits:
 - smaller routing tables but larger firewall tables
- Type bits come first, then location bits:
 - smaller firewall tables but larger routing tables
- In general: routers can handle large tables better than firewalls

Location or type first

Location first

| Subnet | Location | Type | Instance |
|--------|----------|----------------|--------------|
| 1C1 | 1 = DC1 | C = Front ends | 1 = Rack 11 |
| 1C2 | | | 2 = Rack 13 |
| 1D1 | | D = Databases | 1 = Rack 11 |
| 1D2 | | | 2 = Rack 13 |
| 1E1 | | E = Storage | 1 = Rack 11 |
| 1E2 | | | 2 = Rack 13 |
| 2C1 | 2 = DC2 | C = Front ends | 1 = Rack A04 |
| 2C2 | | | 2 = Rack A05 |
| 2D1 | | D = Databases | 1 = Rack A04 |
| 2D2 | | | 2 = Rack A05 |
| 2E1 | | E = Storage | 1 = Rack A04 |
| 2E2 | | | 2 = Rack A05 |

Type first

| Subnet | Type | Location | Instance |
|--------|----------------|----------|--------------|
| C11 | C = Front ends | 1 = DC1 | 1 = Rack 11 |
| C12 | | | 2 = Rack 13 |
| C21 | | 2 = DC2 | 1 = Rack A04 |
| C22 | | | 2 = Rack A05 |
| D11 | D = Databases | 1 = DC1 | 1 = Rack 11 |
| D12 | | | 2 = Rack 13 |
| D21 | | 2 = DC2 | 1 = Rack A04 |
| D22 | | | 2 = Rack A05 |
| E11 | E = Storage | 1 = DC1 | 1 = Rack 11 |
| E12 | | | 2 = Rack 13 |
| E21 | | 2 = DC2 | 1 = Rack A04 |
| E22 | | | 2 = Rack A05 |

Let's assume 4 bits = 0 - 15 / 0 - F for everything right now. This works well in small networks, will need to use the right number of bits in larger networks.

Location or type first

Location first

| Subnet | Location | Type | Instance |
|--------|----------|----------------|--------------|
| 1C1 | 1 = DC1 | C = Front ends | 1 = Rack 11 |
| 1C2 | | | 2 = Rack 13 |
| 1D1 | | D = Databases | 1 = Rack 11 |
| 1D2 | | | 2 = Rack 13 |
| 1E1 | | E = Storage | 1 = Rack 11 |
| 1E2 | | | 2 = Rack 13 |
| 2C1 | 2 = DC2 | C = Front ends | 1 = Rack A04 |
| 2C2 | | | 2 = Rack A05 |
| 2D1 | | D = Databases | 1 = Rack A04 |
| 2D2 | | | 2 = Rack A05 |
| 2E1 | | E = Storage | 1 = Rack A04 |
| 2E2 | | | 2 = Rack A05 |

Type first

| Subnet | Type | Location | Instance |
|--------|----------------|----------|--------------|
| C11 | C = Front ends | 1 = DC1 | 1 = Rack 11 |
| C12 | | | 2 = Rack 13 |
| C21 | | 2 = DC2 | 1 = Rack A04 |
| C22 | | | 2 = Rack A05 |
| D11 | D = Databases | 1 = DC1 | 1 = Rack 11 |
| D12 | | | 2 = Rack 13 |
| D21 | | 2 = DC2 | 1 = Rack A04 |
| D22 | | | 2 = Rack A05 |
| E11 | E = Storage | 1 = DC1 | 1 = Rack 11 |
| E12 | | | 2 = Rack 13 |
| E21 | | 2 = DC2 | 1 = Rack A04 |
| E22 | | | 2 = Rack A05 |

Let's assume 4 bits = 0 - 15 / 0 - F for everything right now. This works well in small networks, will need to use the right number of bits in larger networks.

Configuring the local bits

- IPv6 has all the address configuration mechanisms
 - stateless autoconfiguration
 - least stable address, but most automatic
 - hard to add to DNS and don't know which device has which address
 - DHCPv6
 - not in Android and dependency on DHCPv6 server
 - manual configuration
 - most stable address, but not automatic

Configuring the local bits

- Guest/BYOD etc. Wi-Fi:
 - stateless autoconfig
 - in order to be compatible with Android
- Managed work stations:
 - stateless autoconfig or DHCPv6
 - DHCPv6 for DNS registration or address logging
- Servers:
 - probably manual configuration

Configuring the local bits

- These are just suggestions to keep things simple
- Manual configuration:
 - ::1 for default route address (probably VRRP, with maybe :11 for router 1 and :12 for router 2)
 - use service port number: ::53 for DNS, ::80 for HTTP
 - matching IPv4:
 - 192.0.2.1 → 2001:db8:edca:8001:192:0:2:1
 - (but 2001:db8:edca:8001::192.0.2.1 is something different!)
- DHCPv6:
 - ::2000 - ::2fff keeps addresses short (allows for 4096 DHCPv6 addresses)

Local bits examples

- 2001:db8:edca:8001::1
- 2001:db8:edca:8001::11
- 2001:db8:edca:8001::12
- 2001:db8:edca:8001::80
- 2001:db8:edca:8001::2005
- 2001:db8:edca:8001:203::113:127
- 2001:db8:edca:8001:5054:18ff:fedb:d4a4
- 2001:db8:edca:8001:c139:b4c1:6850:12e5

Local bits examples

- 2001:db8:edca:8001::1 → default gateway (VRRP)
- 2001:db8:edca:8001::11 → router 1
- 2001:db8:edca:8001::12 → router 2
- 2001:db8:edca:8001::80 → HTTP server
- 2001:db8:edca:8001::2005 → DHCPv6
- 2001:db8:edca:8001:203::113:127 → manual from 203.0.113.127
- 2001:db8:edca:8001:5054:18ff:fedb:d4a4 → stateless autoconfig from MAC
- 2001:db8:edca:8001:c139:b4c1:6850:12e5 → stateless autoconfig + privacy

DNS server addresses

- DNS addresses are the only addresses you can't look up in the DNS!
 - (at least, those *should* be the only ones)
 - so need to be easy to type/remember and avoid renumbering them
- So give each their own /64 (so they can be moved independently) with low subnet #, such as:
 - DNS server 1: 2001:db8:edca:1::53/64
 - DNS server 2: 2001:db8:edca:2::53/64
 - DNS server 3: 2001:db8:edca:3::53/64


Router addresses

- Subnets with only internal routers:
 - OSPF etc. don't need global addresses, can use just link-local addresses
- Subnets with stateless autoconfig and/or DHCPv6:
 - EUI-64 addressing (or even link-local only)
- Subnets with manually configured hosts:
 - manually configured default gateway for hosts

Router EUI-64 addressing

- Each router has the same configuration, but they all create a unique IPv6 address from the subnet /64 and their MAC address

```
!  
interface Ethernet0  
  ipv6 address 2001:db8:edca:4::/64 eui-64  
!
```



Organizations with subdivisions

- Big organizations that are made up from different sub-organizations in different locations, such as:
 - a multinational enterprise
 - a national government
- Having one big prefix and a unified numbering plan can help with security
- But the sub-organizations probably need to connect to the internet on their own

Organizations with subdivisions

- Get an LIR prefix (such as /29) from RIPE NCC
- Give out /48s (or larger... talk to the NCC) to sub-organizations
 - smaller than /48 won't be routable
- Where do the security bits go?
 - "below" /48 = **location** before type = large **firewall** tables
 - "above" /48 = type before **location**, so each location must have multiple (/48?) prefixes
 - may not fit with RIPE IPv6 assignment policies
- There is an NCC contact for governments for these issues

Questions?

- Much of this based on the Surfnet IPv6 address planning guide:
 - https://www.ripe.net/support/training/material/IPv6-for-LIRs-Training-Course/Preparing-an-IPv6-Addressing-Plan.pdf/at_download/file
- Also available in Dutch:
 - https://wiki.surfnet.nl/download/attachments/11211103/rapport_201309_IPv6_numplan_NL.pdf
- Find me (and this presentation) at:
 - ipv6.iljitsch.com
 - www.inet6consult.com