IPv6 Courses

December 20, 2010

Table of Contents

G6

- Group of IPv6 actors in France (researchers, engineers...)
- Academic & industrial partners
  - CNRS, Institut TELECOM, INRIA, Universities...
  - AFNIC, 6Wind, Bull...
- Launched in 1995 by:
  - Alain Durand
  - Bernard Tuy
- Is today a legal association under French Law (1901)
  - Laurent Toutain, President
- For further information: http://www.g6.asso.fr/

G6Charter

- Share experience gained from IPv6 experimentations and deployment
- Spread IPv6 information
  - Tutorials and trainings (ISPs, Engineers, netadmins...)
  - Online book (in French), "IPv6, Théorie et pratique": http://livre.g6.asso.fr/
- Initiate research activities around IPv6
- Active in RIPE & IETF working groups
- Promotion of IPv6: French Task Force

Hypertext Symbols

- Several symbols are used in this document:
  - All RFCs and Internet Drafts are hypertext links.
    - Check that there is no more recent version of the document.
  -  is a link to a *Techniques de l'Ingénieur* article on the subject (in French, access may be restricted).
  -  is a link to the online edition of *IPv6, Théorie et Pratique* (in French)
  - W is a link to other information on the web.
- Material concerning IPv6 is taken from the G6 tutorial and copyrighted from G6.

Addresses

IPv6 Benefits

- Larger address space from 2^{32} to 2^{128}
  - Permanent address
- Stateless auto-configuration of hosts
  - Layer 3 "Plug & Play" Protocol
- Simple header $\Rightarrow$ Efficient routing
  - No checksum
  - No fragmentation by routers
  - Enhanced extension system
- Better support of mobility

Addresses

Notation

IPv6 addresses

F2C:544:9E::2:EF8D:6B7 F692:: A:1455::A:6E0 D:63:D::4:3A:55F B33:C::F2 7:5059:3D:C0::

9D::9BAC:B8CA:893F:80 1E:DE2:4C83::4E:39:F35:C875 2:: A:FDE3:76:B4F:D9D:: D6::

369F:9:F8:DBF::2 DD4:B45:1:C42F:BE6:75:: 9D7B:7184:EF::3FB:BF1A:D80 FE9::B:3

EC:DB4:B:F:F11::E9:090 83:B9:08:B5:F:3F:AF:B84 E::35B:8572:7A3:FB2 99:F:9:8B76::BC9

D64:07:F394::BDB:DF40:08EE:A79E AC:23:5D:78::233:84:8 F0D:F::F4EB:0F:5C7

E71:F577:ED:E:9DE8:: B::3 1D3F:A0AA:: 70:8EA1::8:D5:81:2:F302 26::8880:7 93:: F::9:0

E:2:0:266B:: 763E:C:2E:1EB:F6:F4:14:16 E6:6:F4:B6:A888:979E:D78:09

9:754:5:90:0A78:A1A3:1:7 2:8:: 97B:C4::C36 A40:7:5:7E8F:0:32EC:9A:D0 8A52::575

D::4CB4:E:2BF:5485:8CE 07:5::41 6B::A9:C 94FF:7B8::D9:51:26F 2::E:AE:ED:81 8241:: 5F97::

AD5B:259C:7DB8:24:58:552A:: 94:4:9FD:4:87E5:: 5A8:2FF:1::CC EA:8904:7C::

7C::D6B7:A7:B0:8B DC:6C::34:89 6C:1::5 7B3:6780:4:B1::E586 412:2:5E1:6DE5:5E3A:553:3::

7F0:: B39::1:B77:DB 9D3:1F1:4B:3:B4E6:7681:09:D4A8 61:520::E0 1:28E9:0:095:DF:F2::

1B61:4::1DE:50A 34BC:99::E9:9EFB E:EF:: BDC:672A:F4C8:A1::4:7:9CB7 C697:56AD:40:8:0::62

Don't Worry

Addresses are not random numbers. . . they are often easy to handle and even to memorize sometimes

Notation

- Base format (a 16-octet Global IPv6 Address):
  - `2001:0db8:beef:0001:0000:0000:cafe:deca`
- Compact Format:

2001:0db8:beef:0001:0000:0000:cafe:deca

1.
2.

Warning:

`2001:db8:3::/40` is in fact `2001:db8:0003::/40` and not `2001:db8:0300::/40`

Notation

Addresses
Notation
Addressing scheme
Address Format
Kind of addresses

- Base format (a 16-octet Global IPv6 Address):
  - `2001:0db8:beef:0001:0000:0000:cafe:deca`
- Compact Format:

2001:db8:beef:1:0:0:cafe:deca

1. Remove 0 on the left of each word
2.

Warning:

`2001:db8:3::/40` is in fact `2001:db8:0003::/40` and not `2001:db8:0300::/40`

©G6 Association December 20, 2010 11 / 37

Notation

- Base format (a 16-octet Global IPv6 Address):
  - `2001:0db8:beef:0001:0000:0000:cafe:deca`
- Compact Format:

2001:db8:beef:1::cafe:deca

1. Remove 0 on the left of each word
2. To avoid ambiguity, substitute ONLY one sequence of zeros by `::`

Warning:

`2001:db8:3::/40` is in fact `2001:db8:0003::/40` and not `2001:db8:0300::/40`

Notation

Addresses
Notation
Addressing scheme
Address Format
Kind of addresses

- Base format (a 16-octet Global IPv6 Address):
  - `2001:0db8:beef:0001:0000:0000:cafe:deca`
- Compact Format:

2001:db8:beef:1::cafe:deca

1. Remove 0 on the left of each word
2. To avoid ambiguity, substitute ONLY one sequence of zeros by `::`

- an IPv4 address may also appear : `::ffff:192.0.2.1`

Warning:

`2001:db8:3::/40` is in fact `2001:db8:0003::/40` and not `2001:db8:0300::/40`

©G6 Association December 20, 2010 11 / 37

Is it enough for the future ?

- Address length
  - About $3.4x10^{38}$ addresses
  - 60 000 trillion trillion addresses per inhabitant on earth
  - Addresses for every grain of sands in the world
  - IPv4: 6 addresses per US inhabitant, 1 in Europe, 0.01 in China and 0.001 in India
- Justification of a fixed-length address

Warning:

- An address for everything on the network and not an address for everything
- No addresses for the whole life:
  - Depends on your position on the network
  - ISP Renumbering may be possible

Is it enough for the future ?

- Hop Limit:
  - Should not be a problem
  - Count the number of routers used to reach a destination
  - Growth will be in-width more than in-depth
- Payload Length
  - 64 Ko is not a current hard limit
  - Ethernet is limited to 1.5 Ko, evolution can use until 9Ko.
  - Use Jumbogram for specific cases

Addresses

Addressing scheme

Addressing scheme 

- RFC 4291 defines current IPv6 addresses
  - loopback (`::1`)
  - `link local (fe80::/10)`
  - `global unicast (2000::/3)`
  - `multicast (ff00::/8)`
- Use CIDR principles:
  - Prefix / prefix length notation
  - `2001:db8:face::/48`
  - `2001:db8:face:bed:cafe:deca:dead:beef/64`
- Interfaces have several IPv6 addresses
  - at least a link-local and a global unicast addresses

Addresses

Address Format

Addressing Space Utilization


```
0000::/8 Reserved by IETF [RFC4291]
0100::/8 Reserved by IETF [RFC4291]
0200::/7 Reserved by IETF [RFC4048]
0400::/6 Reserved by IETF [RFC4291]
0800::/5 Reserved by IETF [RFC4291]
1000::/4 Reserved by IETF [RFC4291]
2000::/3 Global Unicast [RFC4291]
4000::/3 Reserved by IETF [RFC4291]
6000::/3 Reserved by IETF [RFC4291]
8000::/3 Reserved by IETF [RFC4291]
a000::/3 Reserved by IETF [RFC4291]
c000::/3 Reserved by IETF [RFC4291]
e000::/4 Reserved by IETF [RFC4291]
f000::/5 Reserved by IETF [RFC4291]
F800::/6 Reserved by IETF [RFC4291]
fc00::/7 Unique Local Unicast [RFC4193]
fe00::/9 Reserved by IETF [RFC4291]
fe80::/10 Link Local Unicast [RFC4291]
fec0::/10 Reserved by IETF [RFC3879]
ff00::/8 Multicast [RFC4291]
```

http://www.iana.org/assignments/ipv6-address-space

Address Format

Global Unicast Address:

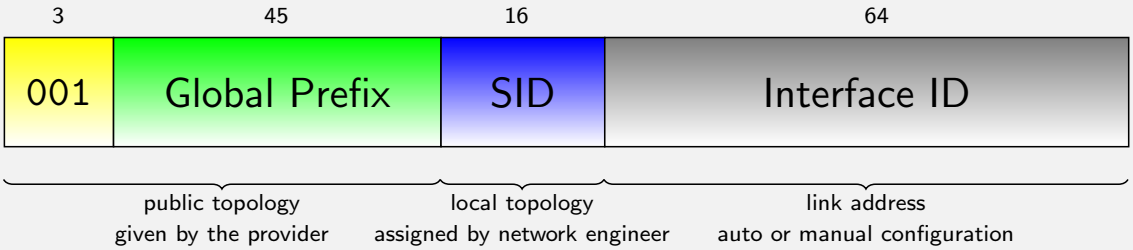

Link-Local Address:

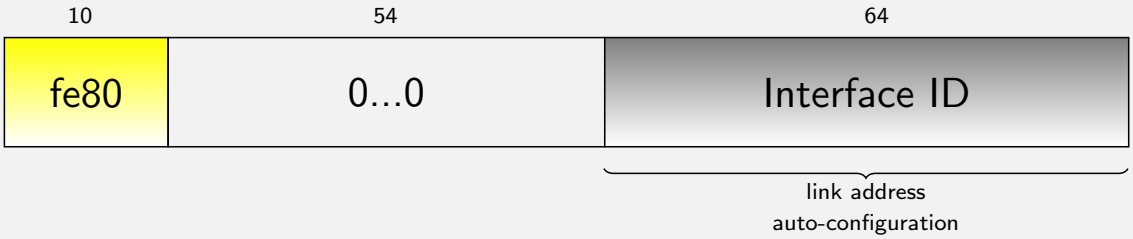

SID Values

- 16-bit length up to 65 535 subnets
  - Large enough for most companies
  - Too large for home network ?
  - May be a /56 or /60 GP will be allocated depending on the ISP
- There is no strict rules to structure SID:
  - sequencial : 1, 2, ...
  - use VLAN number
  - include usage to allow filtering, for instance, for a University:

4bits : Community	**8bits**	**4bits**
0 : Infrastructure	*Specific addresses*	
1 : Tests	*Specific addresses*	
6 : Point6	*Managed by Point6*	
8 : Wifi guests	*Specific addresses*	
A : Employees	Entity	Sub-Network
E : Students	Entity	Sub-Network
F : Other (Start-up, etc.)	*Specific addresses*	

Interface Identifier

Interface ID can be selected differently

- Derived from a Layer 2 ID (I.e. MAC address) :
  - for Link Local address
  - for Global Address : plug-and-play hosts
- Assigned manually :
  - to keep same address when Ethernet card or host is changed
  - to remember easily the address
    - 1, 2, 3, ...
    - last digit of the v4 address
    - the IPv4 address (for nostalgic system administrators)
    - ...

Interface Identifier

Interface ID can be selected differently

- Random value :
  - Changed frequently (e.g, every day, per session, at each reboot...) to guarantee anonymity
- Hash of other values (experimental) :
  - To link address to other properties
  - Public key
  - List of assigned prefixes
  - ...

How to Construct an IID from MAC Address

- 64 bits is compatible with EUI-64 (i.e. IEEE 1394 FireWire, ...)
- IEEE propose a way to transform a MAC-48 to an EUI-64
- U/L changed for numbering purpose

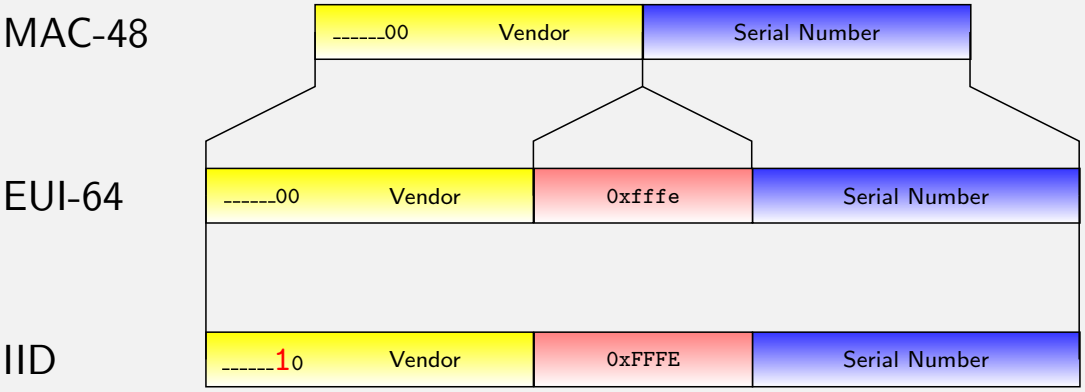

- There is no conflicts if IID are manually numbered: 1, 2, 3, ...

Example : Mac / Unix


```
%ifconfig
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
        inet6 ::1 prefixlen 128
        inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
        inet 127.0.0.1 netmask 0xff000000
en1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500

        inet6 fe80::216:cbff:febe:16b3%en1 prefixlen 64 scopeid 0x5

        inet 192.168.2.5 netmask 0xffffff00 broadcast 192.168.2.255
        inet6 2001:660:7307:6031:216:cbff:febe:16b3 prefixlen 64
        autoconf

        ether 00:16:cb:be:16:b3
        media: autoselect status: active
        supported media: autoselect
```

Windows 7

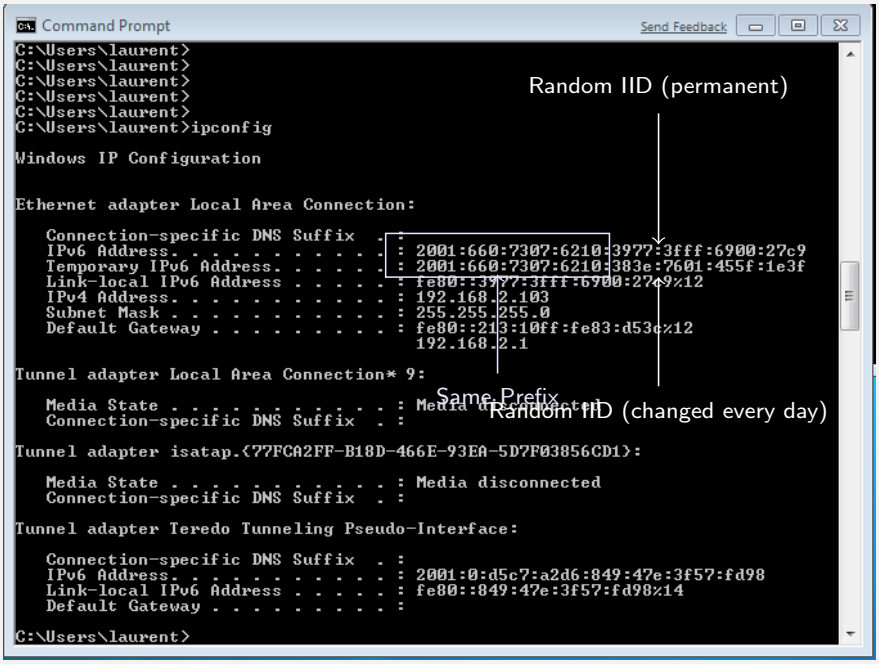

Address Lifetime

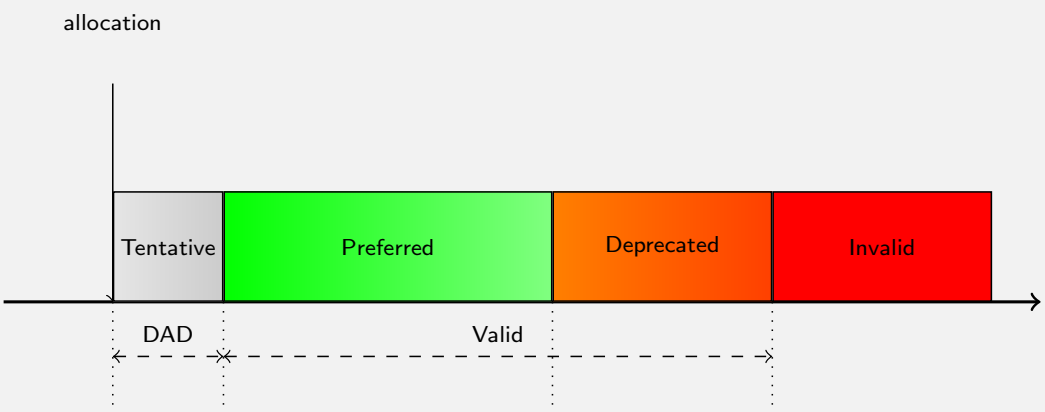

Addresses

Kind of addresses

Link-Local Scoped Addresses 

- Global Address, the prefix designates the exit interface
- Link-Local address, the prefix is always fe80::/10
  - The exit interface is not defined
  - A %iface, can be added at the end of the address to avoid ambiguity
- Example:

```
Routing tables

Internet6:

Destination                    Gateway                         Flags      Netif Expire
default                        fe80::213:c4ff:fe69:5f49%en0    UGSc         en0
```

Other kind of addresses : ULA (RFC 4193)

- Equivalent to the private addresses in IPv4
- But try to avoid same prefixes on two different sites:
  - avoid renumbering if two company merge
  - avoid ambiguities when VPN are used
- These prefixes are not routable on the Internet

Unique Local IPv6 Unicast Addresses:

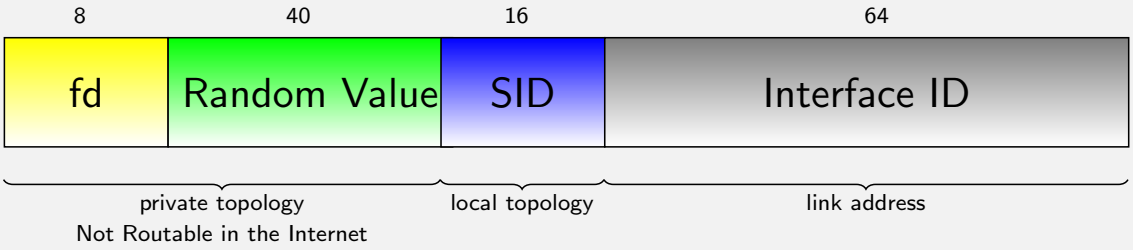

http://www.sixxs.net/tools/grh/ula/ to create your own ULA prefix.

Multicast

Generic Format:

- T (Transient) 0: well known address - 1: temporary address
- P (Prefix) 1 : assigned from a network prefix (T must be set to 1)
- R (Rendez Vous Point) 1: contains the RP address (P & T set to 1)
- Scope :
  - 1 - interface-local
  - 2 - link-local
  - 3 - reserved
  - 4 - admin-local
  - 5 - site-local
  - 8 - organisation-local
  - e - global
  - f - reserved

Some Well Known Multicast Addresses


```
ff02:0:0:0:0:0:0:1 All Nodes Address (link-local scope)
ff02:0:0:0:0:0:0:2 All Routers Address
ff02:0:0:0:0:0:0:5 OSPFIGP
ff02:0:0:0:0:0:0:6 OSPFIGP Designated Routers
ff02:0:0:0:0:0:0:9 RIP Routers
ff02:0:0:0:0:0:0:fb mDNSv6
ff02:0:0:0:0:0:1:2 All-dhcp-agents
ff02:0:0:0:0:1:ffxx:xxxx Solicited-Node Address
ff05:0:0:0:0:0:1:3 All-dhcp-servers (site-local scope)
```

http://www.iana.org/assignments/ipv6-multicast-addresses

Solicited Multicast Addresses

- Derive a Multicast Address from a Unicast Address
  - Widely used for stateless auto-configuration
  - Avoid the use of broadcast

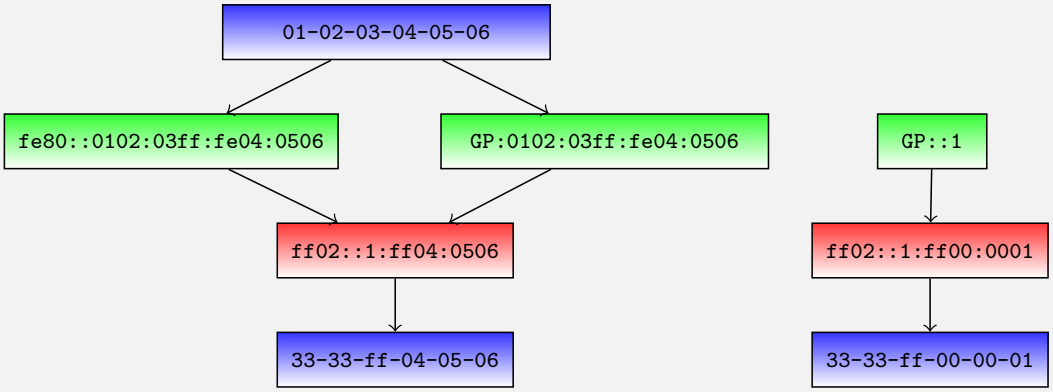

Example


```
Vlan5 is up, line protocol is up
IPv6 is enabled, link-local address is fe80::203:fdff:fed6:d400
Description: reseau C5
Global unicast address(es):
  2001:660:7301:1:203:fdff:fed6:d400, subnet is 2001:660:7301:1::/64

Joined group address(es):
  ff02::1  <- All nodes
  ff02::2  <- All routers
  ff02::9  <- RIP
  ff02::1:ffd6:d400  <- Solicited Multicast
```


Anycast 

- In the same addressing space as unicast
- No way to distiguish them
- Two anycast families:
  - Same prefix on Internet
    - same a IPv4 anycast for DNS or 6to4
  - Same address on the link
    - Must avoid DAD
    - Some IID values are reserved
    - All IID bits to 1 except last byte
    - Only 0x7E Mobile Home Agent
  - May more addresses with Wireless Sensor Network ?
    - Temperature, presence. . .

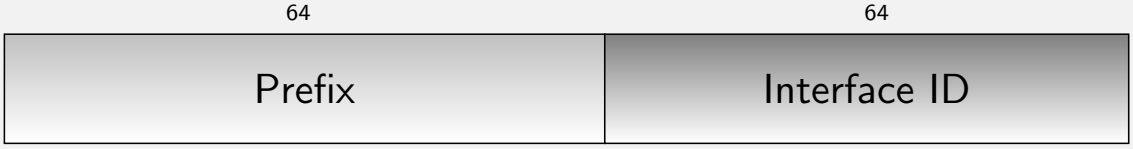

Anycast on prefix : Example from Renater


```
#traceroute6 2001:500:2f::f
traceroute6 to 2001:500:2f::f (2001:500:2f::f) from 2001:660:7301:3103:223:6cf
30 hops max, 12 byte packets
1 2001:660:7301:3103::1 4.774 ms 1.198 ms 2.764 ms
2 2001:660:7301:3036::1 3.364 ms 2.215 ms 1.417 ms
3 vl856-gi9-9-rennes-rtr-021.noc.renater.fr 2.892 ms 6.794 ms 2.195 ms

4 te4-1-caen-rtr-021.noc.renater.fr 7.706 ms 5.1 ms 4.193 ms
5 te4-1-rouen-rtr-021.noc.renater.fr 6.527 ms 6.296 ms 6.661 ms
6 te0-0-0-1-paris1-rtr-001.noc.renater.fr 8.702 ms 10.26 ms 8.696 ms
7 F-root-server.sfinx.tm.fr 8.495 ms 8.607 ms 8.664 ms
8 f.root-servers.net 8.738 ms 9.171 ms 8.702 ms
```


Anycast on prefix : Example from Hawaï


```
#traceroute6 2001:500:2f::f
traceroute6 to 2001:500:2f::f (2001:500:2f::f) from 2001:1888:0:1:2d0:b7ff:fe7
64 hops max, 12 byte packets
1 apapane-fe0-0-1 1.169 ms 0.970 ms 0.947 ms
2 r1.mdtnj.ipv6.att.net 121.159 ms 121.737 ms 121.378 ms
3 bbr01-p1-0.nwrk01.occaid.net 130.468 ms 129.640 ms 130.845 ms
4 bbr01-g1-0.asbn01.occaid.net 131.372 ms 131.596 ms 131.421 ms
5 bbr01-g1-0.atln01.occaid.net 144.937 ms 144.550 ms 144.834 ms
6 bbr01-p1-0.dlls01.occaid.net 166.709 ms 196.177 ms 165.983 ms
7 dcr01-p1-5.lsan01.occaid.net 138.437 ms 138.690 ms 138.544 ms
8 bbr01-g0-2.irvn01.occaid.net 138.552 ms 137.956 ms 137.649 ms
9 dcr01-g1-2.psdn01.occaid.net 137.629 ms 138.030 ms 141.332 ms
10 bbr01-f1-5.snfc02.occaid.net 138.501 ms 138.511 ms 137.483 ms
11 exit.sf-guest.sfo2.isc.org 147.941 ms 144.929 ms 145.956 ms
12 f.root-servers.net 139.063 ms 139.715 ms 142.571 ms
```

OnLink Anycast: Example

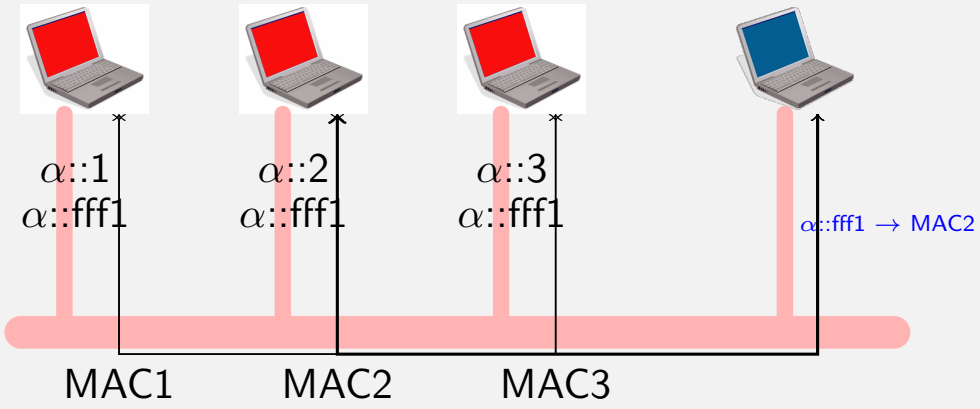

RFC 2526 Anycast values, all bit of IID set to 1 except last 8 bits:

- 0x7F: reserved
- 0x7E: Home Agent (Mobile IP)
- 0x00 to 0x7D: reserved

http://www.iana.org/assignments/ipv6-anycast-addresses

Anycast on prefix : Example from Hawaï


```
# ifconfig en3
en3: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500

inet6 fe80::223:6cff:fe97:679c
inet 192.168.103.177 netmask 0xffffff00 broadcast 192.168.103.255
inet6 2001:660:7301:3103:223:65ff:fe97:679c prefixlen 64 autoconf
ether 00:23:6c:97:67:9c
media: autoselect status: active
supported media: autoselect
# ifconfig en3 inet6 2001:660:7301:3103:FF::FF anycast
# ifconfig en3
en3: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500

inet6 fe80::223:6cff:fe97:679c
inet 192.168.103.177 netmask 0xffffff00 broadcast 192.168.103.255
inet6 2001:660:7301:3103:223:65ff:fe97:679c prefixlen 64 autoconf
inet6 2001:660:7301:3103:ff::ff prefixlen 64 anycast
ether 00:23:6c:97:67:9c
media: autoselect status: active
supported media: autoselect
```