IPv6

Matt Clemons

Topology

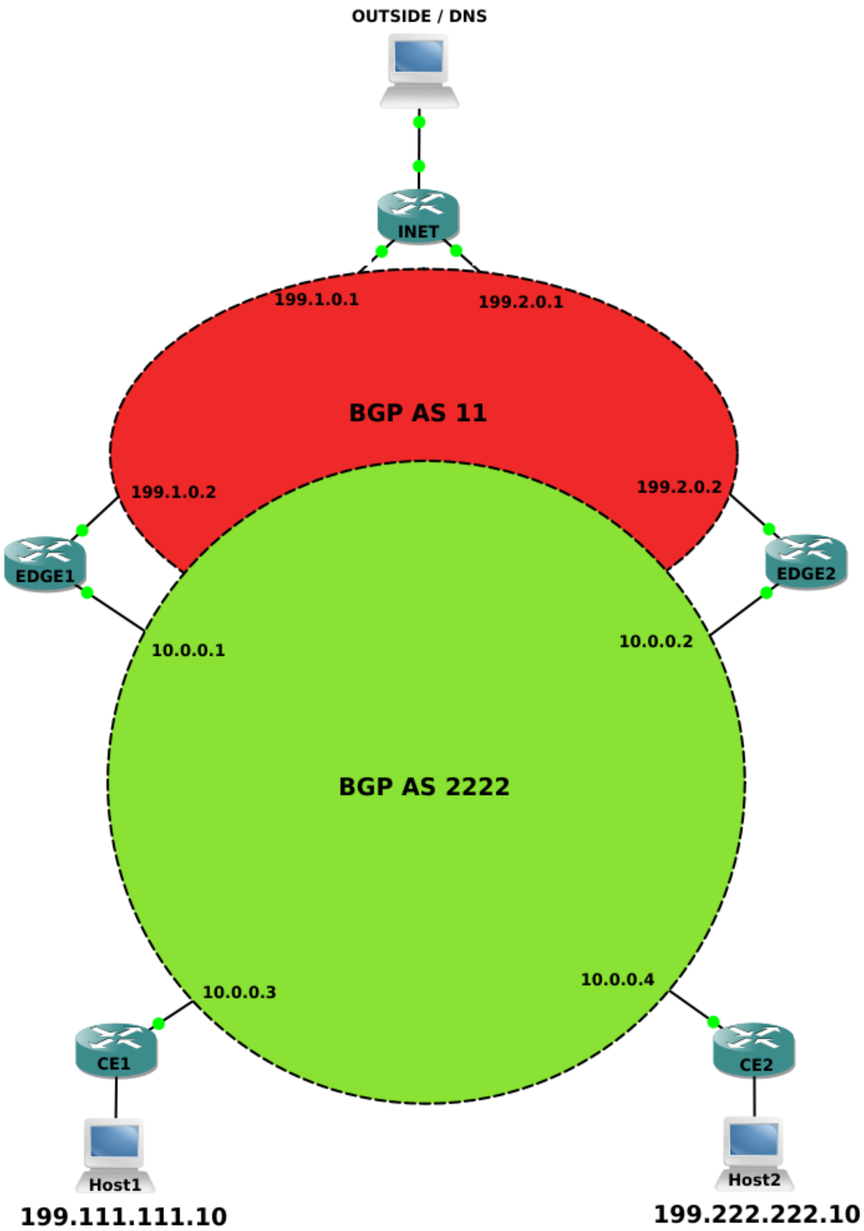

OUTSIDE / DNS
INET
199.1.0.1
199.2.0.1
BGP AS 11
199.1.0.2
199.2.0.2
EDGE1
EDGE2
10.0.0.1
10.0.0.2
BGP AS 2222
10.0.0.3
10.0.0.4
CE1
CE2
Host1
Host2
199.111.111.10
199.222.222.10

Setup

Grab a vncviewer like:

http://uvnc.com/download/1082/1082viewer.html

Or

https://www.realvnc.com/download/viewer/

Connect where I tell you and enter the password to see my screen.

IPv4 Addressing

IPv4 is 4 "octets"

Example: 192.168.33.1/24

Which in binary is

11000000.10101000.00100001.00000001

The bit boundary is 24 so the network address is 192.168.33. and gives us 255 hosts -2 for network and broadcast

11000000.10101000.00100001. as the network address and

8 bits for hosts

IPv6 Addressing

IPv6 is 8 Sets of 16 totaling 128 bits

Example: 2600:1200:0000:0000:0000:0000:0000:0000/24

Addresses are in Hex and each digit is worth 4 bits

In binary it looks like:

0010011000000000:0001001000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000

Shorthand

Example: 2600:1200:0010:0000:0000:0000:0000:0000/48

Sets of all 0’s can be shortened with “::”

Leading zeros can also be removed

2600:1200:10::/48

Types of IPv6 Addresses

Global (Public IP):

2002:4088:daf6:0001:0000:0000:0000:0102

Shorthand address for this would be

2002:4088:daf6:1::102

Link Local (Local Segment):

fe80:0000:0000:0000:d6a7:256e:abdd:9ac1

Shorthand address for this would be

fe80::d6a7:256e:abdd:9ac1

Created through EUI-64 – NEXT SLIDES

EUI-64 RFC 2373

Conversion process is two steps

1. Convert the 48-bit MAC address to a 64-bit value

Break the MAC address into its two 24-bit halves: the Organizationally Unique Identifier (OUI) and the NIC specific part

The 16-bit hex value 0xFFFE is then inserted between these two halves to form a 64-bit address

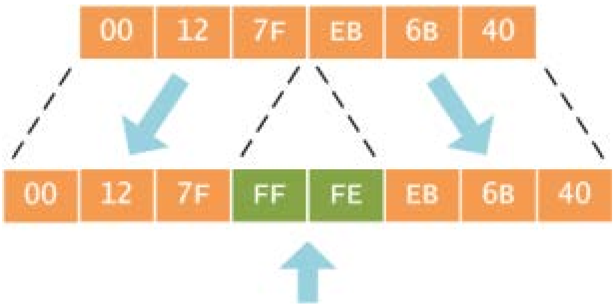

EUI-64

2. The second step is to flip the universal/local (U/L) flag (bit 7) in the OUI portion of the address

The U/L bit is inverted when using an EUI-64 address as an IPv6 interface ID

Then simply append FE80 to the EUI-64 address

EUI-64 Example

Example mac address is CA05.0EE6.0008

break it in half and add FFFE in the middle

CA.05.0E.FF.FE.E6.00.08

Convert it to binary:

11001010.00000101.00001110.11111111.11111110.11100110.00000000.00001000

Then flip bit 7:

11001000

When converted back to hex, it would be C805.0EFF.FEE6.0008

Append FE80 to it, use shorthand (remove 0's), and Your IPv6 Link Local address will be:

FE80::C805:EFF:FEE6:8

Broadcasts?

IPv6 no longer uses broadcasts and has moved to multicast

Anything that starts with FF is a multicast address.

Every interface in an IPv6 segment joins the FF02::1 multicast group.

Common multicast groups

Address	Description
ff02::1	All nodes on the local network segment
ff02::2	All routers on the local network segment
ff02::5	OSPFv3 All SPF routers
ff02::6	OSPFv3 All DR routers
ff02::8	IS-IS for IPv6 routers
ff02::9	RIP routers
ff02::a	EIGRP routers
ff02::d	PIM routers
ff02::16	MLDv2 reports (defined in RFC 3810)
ff02::1:2	All DHCP servers and relay agents on the local network segment (defined in RFC 3315)
ff02::1:3	All LLMNR hosts on the local network segment (defined in RFC 4795)
ff05::1:3	All DHCP servers on the local network site (defined in RFC 3315)
ff0x::c	Simple Service Discovery Protocol
ff0x::fb	Multicast DNS
ff0x::101	Network Time Protocol
ff0x::108	Network Information Service
ff0x::181	Precision Time Protocol (PTP) version 2 messages (Sync, Announce, etc.) except peer delay measurement
ff02::6b	Precision Time Protocol (PTP) version 2 peer delay measurement messages
ff0x::114	Used for experiments

There’s no ARP – RFC 1122

NDP operates in the Link Layer

Responsible for address auto configuration of nodes, discovery of other nodes on the link, determining the addresses of other nodes, duplicate address detection, finding available routers and Domain Name System (DNS) servers, address prefix discovery, and maintaining reachability information of other active neighbor nodes

Five different ICMPv6 packet types to perform functions for IPv6 similar to the ARP

Neighbor discovery

1. Router Solicitation (Type 133) – Any routers out there?

2. Router Advertisement (Type 134) – I'm a router!

3. Neighbor Solicitation (Type 135) – What is neighbor A's link address? Sourced from link local address

4. Neighbor Advertisement (Type 136) – Here's my link address!

5. Redirect (Type 137) – Use router B instead!

Solicited-node multicast

For each Unicast and Anycast address configured there is a corresponding solicited-node multicast

Used in neighbor solicitation messages

Solicited-node multicast consists of

FF02::1:FF/104 (Lower 24 bits from IPv6 Unicast interface ID)

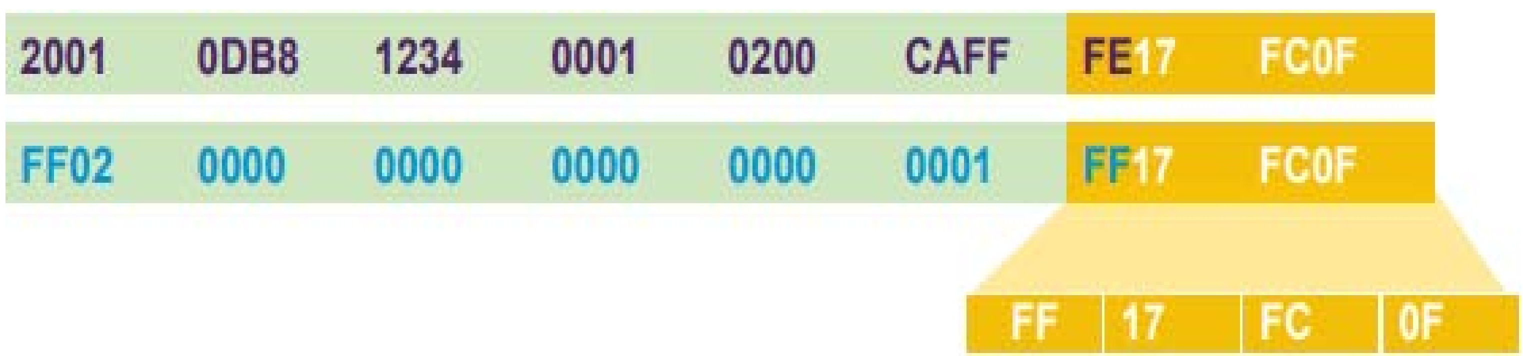

Solicited-node multicast group

Each interface joins its own group as well as the all hosts FF02::1

If your IPv6 Link local IP is: fe80::3285:a9ff:fe41:b8a3/64

Your solicited-node multicast group is:

FF02:0000:0000:0000:0000:0001:FF41:b8a3

Or

FF02::1:FF41:b8a3

Solicited-node multicast

How IPv6 NDS works:

Source host:

clear ipv6 neighbors

ping ff02::1

! select output interface

! Sent to all hosts on the network

1. Echo request from source Link Local address to ff02::1

2. NS from each Link Local interface on the network to source solicited-node multicast group. Derived from the MAC address in the Ethernet Header

3. NA from Source Link local address to each host on the network

4. Echo reply from dests to source

Duplicate Address Detection

Neighbor Solicitation ICMP Type 135

IPv6 Source = Unspecified

IPv6 Dest = Solicited-Node Multicast group A
FF02::1:FF00:A

Data = FE80::address of A

Query = Anyone using A?

If no host responds

Node can start using A

If there is a duplicate address

Your OS will let you know

IPv6 Addressing

Compressed Address: 2600:1200::/24

Expanded Address:

2600:1200:0000:0000:0000:0000:0000:0000/24

Prefix: ffff:ff00:0000:0000:0000:0000:0000:0000

Range: 2600:1200:0:0:0:0:0:0

2600:12ff:ffff:ffff:ffff:ffff:ffff:ffff

IPv6 Addressing

The bit boundary is 24 bits which sets the network to

0010011000000000:0001001000000000:00000000

And gives us 104 host bits to work with

00000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000

Let’s try to understand the magnitude of the amount of addresses that is

Number of networks and hosts

20,282,409,603,651,670,423,947,251,286,016 total addresses

Typically, /64 networks are given out to entities that don’t require VLANs

We have 1,099,511,627,776 /64 networks

Each /64 has 18,446,744,073,709,551,616 hosts

Every employee could have over 9 million /64 subnets

IPv6 Subnetting

Splitting up that /24

We can make multiple networks by moving up 4 bits to a /28

2600:1200:0000:0000:0000:0000:0000:0000/28

0010011000000000:000100100000 |
0000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000:0000000000000000

This gives us 100 host bits to work with.

Each /28 has 68,719,476,736 /64 networks

Each /64 has 18,446,744,073,709,551,616 hosts

/28

Compressed Address: 2600:1200::/28

Expanded Address: 2600:1200:0000:0000:0000:0000:0000:0000/28

Prefix: ffff:fff0:0000:0000:0000:0000:0000:0000

Range: 2600:1200:0:0:0:0:0:0

2600:120f:ffff:ffff:ffff:ffff:ffff:ffff

Compressed Address: 2600:12f0::/28

Expanded Address: 2600:12f0:0000:0000:0000:0000:0000:0000/28

Prefix: ffff:fff0:0000:0000:0000:0000:0000:0000

Range: 2600:12f0:0:0:0:0:0:0

2600:12ff:ffff:ffff:ffff:ffff:ffff:ffff

IPv6 Subnetting

Typically, since there are so many addresses, users will subnet with prefix lengths in multiples of 4 which makes it easy. You don't even need a calculator.

For example:

2600:1200:1234:1234::/64

2600:1200:1234:123X::/60

2600:1200:1234:12XX::/56

2600:1200:1234:1XXX::/52

2600:1200:1234:XXXX::/48

2600:1200:123X:XXXX::/44

2600:1200:12XX:XXXX::/40

IPv6 Subnetting

From the network end, it would look like this

2600:1200::/24

2600:12X0::/28

2600:12XX::/32

2600:12XX:X::/36

2600:12XX:XX::/40

2600:12XX:XXX::/44

2600:12XX:XXXX::/48

Simulated network

Just an example

Dual stack

Simulated Internet with one external host

Simulated Dual Edge Routers

Two Simulated CE Routers with two internal hosts

IPv4 Connectivity end to end

EBGP peering with the Internet

IBGP peering with all CE and Edge routers

Edge1 OUTSIDE IP is 2600:1200::2/28

Edge2 OUTSIDE IP is 2600:12f0::2/28

Outside Host is 4444:4:4::4 (4.4.4.4) which is also serving DNS

Simulated network

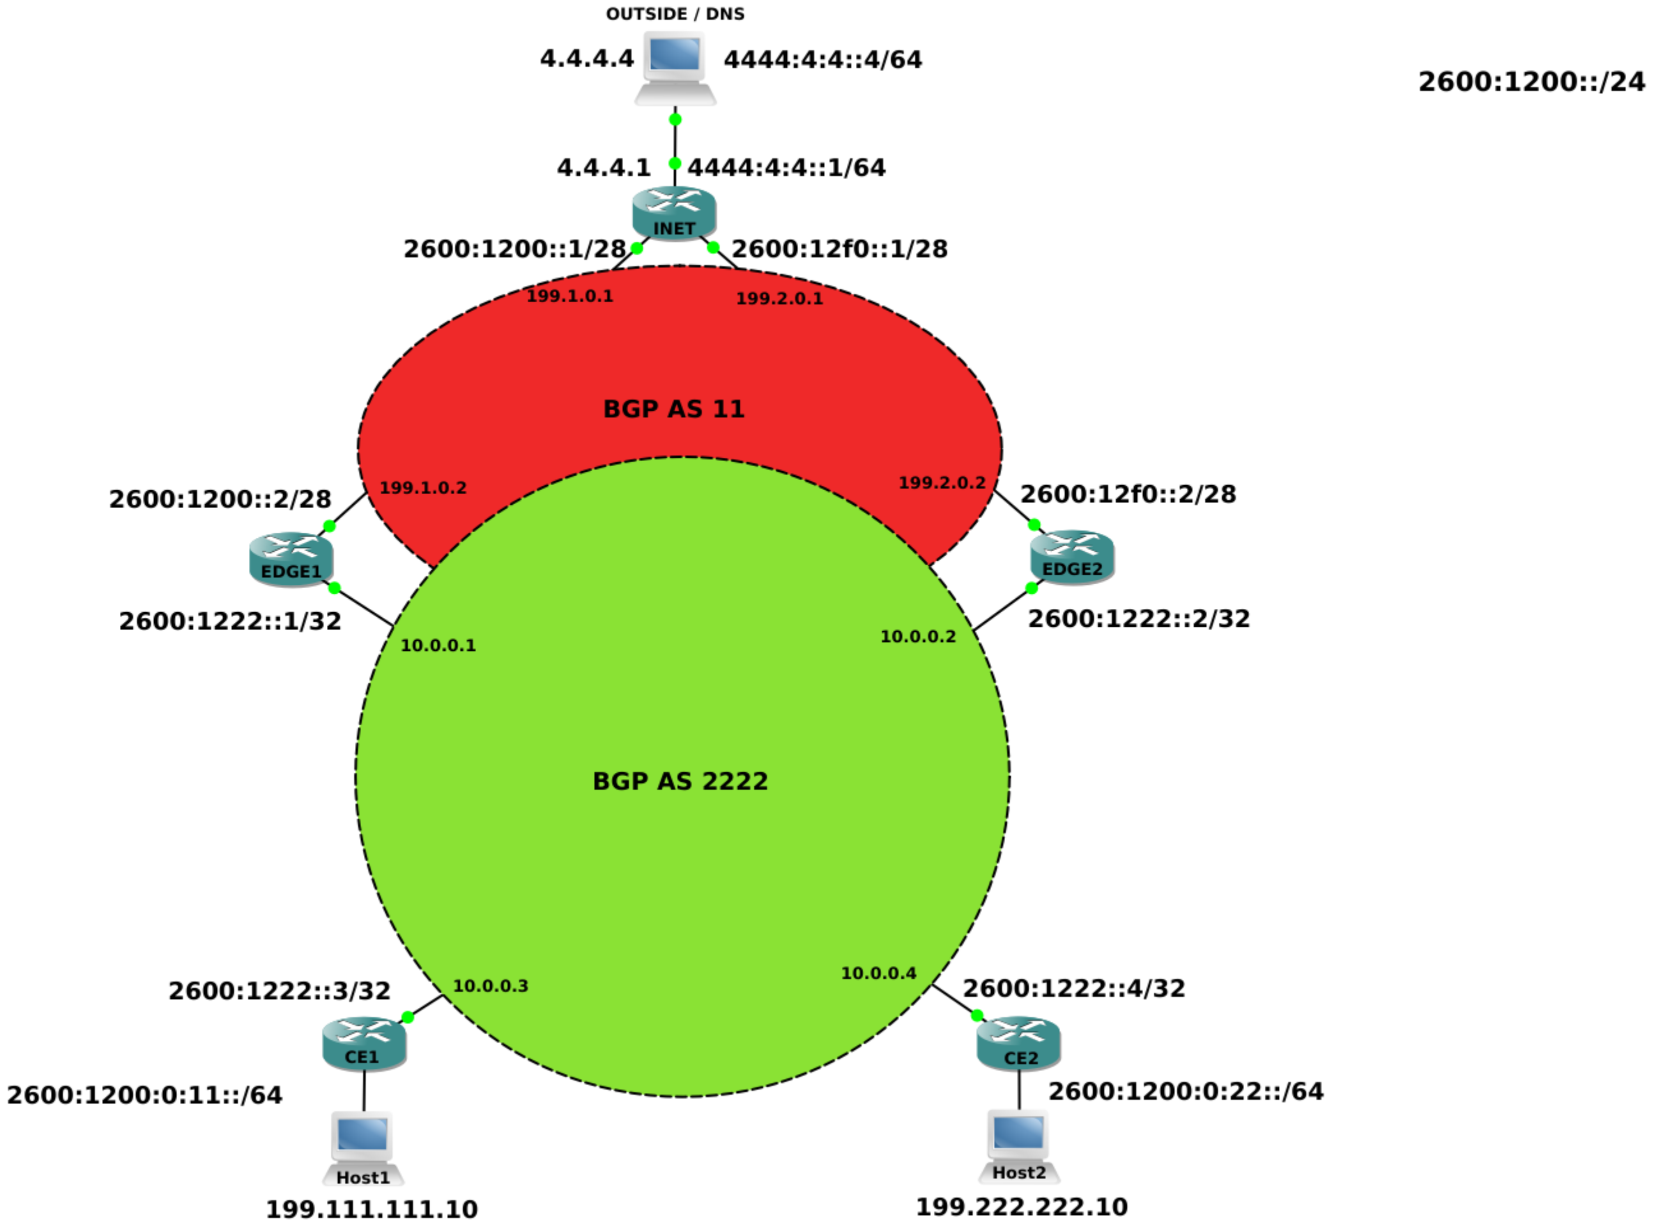

Let’s do some subnetting

One subnet for all internal routers by moving up 4 bits

2600:12XX::/32

Compressed Address: 2600:1222::/32

Expanded Address: 2600:1222:0000:0000:0000:0000:0000:0000/32

Prefix: ffff:ffff:0000:0000:0000:0000:0000:0000

Range: 2600:1222:0:0:0:0:0:0

2600:1222:ffff:ffff:ffff:ffff:ffff:ffff

Let’s do some subnetting

We’ll set them as follows

2600:1222::1/32 Edge1 INSIDE interface

2600:1222::2/32 Edge2 INSIDE interface

2600:1222::3/32 CE1 OUTSIDE interface

2600:1222::4/32 CE2 OUTSIDE interface

Let's do some subnetting

Next we'll make a subnet for each CE internal network, moving up multiples of 4 to a /64

2600:1200:0000:0011::/64

Compressed Address: 2600:1200:0:11::/64

Range: 2600:1200:0:11:0:0:0:0

2600:1200:0:11:ffff:ffff:ffff:fff

2600:1200:0000:0022::/64

Compressed Address: 2600:1200:0:22::/64

Range: 2600:1200:0:22:0:0:0:0

2600:1200:0:22:ffff:ffff:ffff:fffff

Stateful autoconfiguration

We'll enable DHCPv6 on CE1

The Dynamic Host Configuration Protocol version 6 is a network protocol for configuring IPv6 hosts with IP addresses, IP prefixes and other configuration data required to operate in an IPv6 network. It is the IPv6 equivalent of the Dynamic Host Configuration Protocol for IPv4.

Stateful autoconfiguration

Server config:

```
ipv6 dhcp pool ipv6dhcp
 address prefix 2600:1200:0:11::/64 lifetime infinite infinite
 link-address 2600:1200:0:11::1/64
 dns-server 4444:4:4::4
 domain-name test.com
!
interface GigabitEthernet1/0
 ipv6 address 2600:1200:0:11::1/64
 ipv6 enable
 ipv6 nd prefix default no-advertise
 ipv6 nd managed-config-flag
 ipv6 nd other-config-flag
 ipv6 nd ra suppress
 ipv6 dhcp server ipv6dhcp rapid-commit
!
```

Stateful autoconfiguration

Client config:

```
interface g1/0
    ipv6 address dhcp
```

Stateless autoconfiguration

We’ll configure SLAAC on CE2

Stateless Autoconfiguration is a method in which the host or router interface is assigned a 64-bit prefix, and then the last 64 bits of its address are derived by the host or router with help of EUI-64 process. SLAAC uses NDP protocol to work.

Stateless autoconfiguration

Server config:

```
ipv6 dhcp pool ipv6slaac
 dns-server 4444:4:4::4
 domain-name test.com
!
interface GigabitEthernet1/0
 ipv6 address FE80::1 link-local
 ipv6 address 2600:1200:0:22::1/64
 ipv6 enable
 ipv6 nd other-config-flag
 ipv6 dhcp server ipv6slaac
```

Stateless autoconfiguration

Client config:

```
interface g0/0
    ipv6 address autoconfig
    ipv6 enable
```

Anycast

An Anycast address is an IPv6 address that is used more than once in different locations

Enabling Anycast disables DAD

Uses routing to find the closest host

Anycast

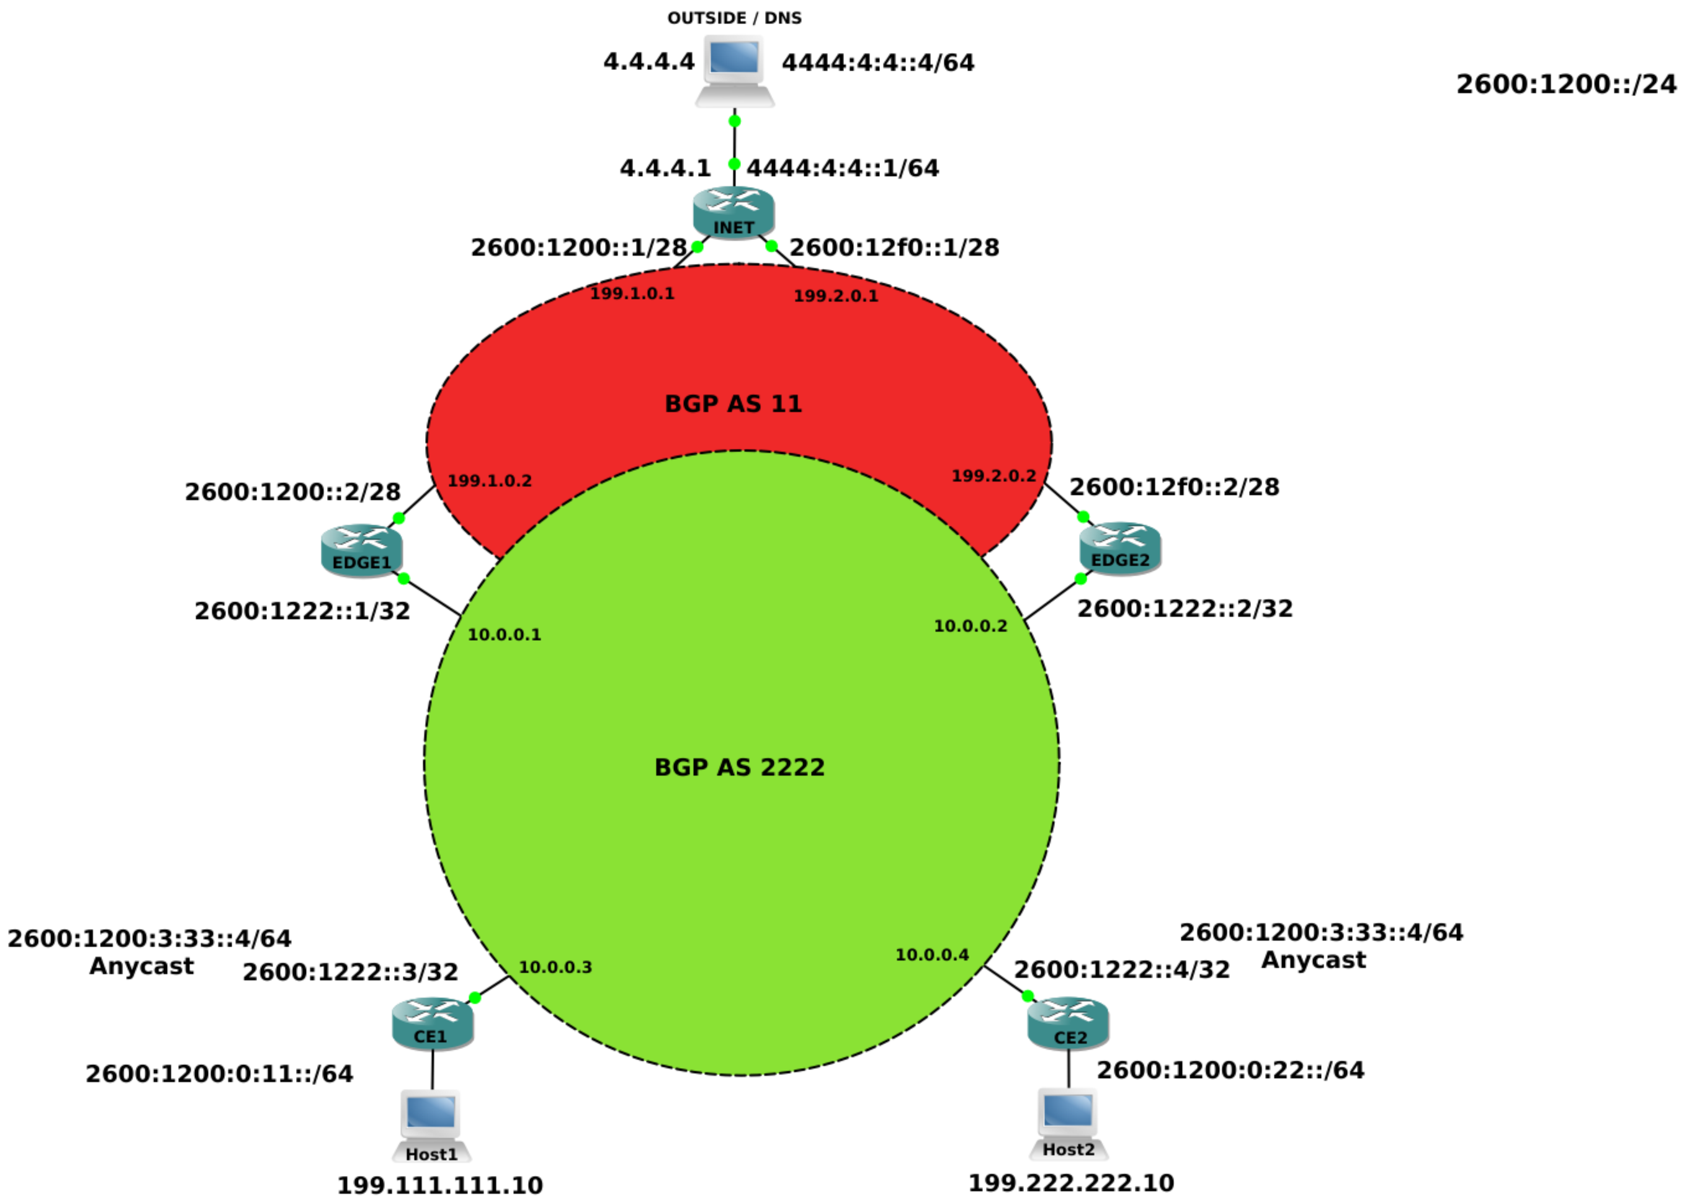
OUTSIDE / DNS
4.4.4.4
4444:4:4::4/64
2600:1200::/24
4.4.4.1
4444:4:4::1/64
INET
2600:1200::1/28
2600:12f0::1/28
199.1.0.1
199.2.0.1
BGP AS 11
2600:1200::2/28
199.1.0.2
199.2.0.2
2600:12f0::2/28
EDGE1
EDGE2
2600:1222::1/32
10.0.0.1
10.0.0.2
2600:1222::2/32
BGP AS 2222
2600:1200:3:33::4/64
Anycast
2600:1222::3/32
10.0.0.3
10.0.0.4
2600:1200:3:33::4/64
2600:1222::4/32
Anycast
CE1
CE2
2600:1200:0:11::/64
2600:1200:0:22::/64
Host1
Host2
199.111.111.10
199.222.222.10

Anycast config demo

CE1 & 2:

```
int g0/0
ipv6 address 2001:1200:4:44::4/64 anycast
!
```

Then, advertise the network in BGP

Privacy Extensions for IPv6

AKA SLAAC Temporary Addresses

RFC 4941

MAC addresses can be derived from Auto configured addresses

Generates a unique Link Local Address for use in generating global auto configured addresses

Interface no longer uses the MAC but instead picks a series of bits randomly, and fills in the last 48 bits

Use this for a short period (whatever the lifetime is) and then change it for another one.

Modern OS's have this enabled by default

What can you do to learn more?

Create a free tunnel at tunnelbroker.net

GNS3 https://www.gns3.com/

Free IPv6 certification: https://ipv6.he.net/certification/

You get a free tee shirt. I'm a Sage.

Books 24x7 IPv6: Theory, Protocol, and Practice

Wikipedia: https://en.wikipedia.org/wiki/IPv6

Questions?