What You Make Possible

TOMORROW
starts here.

Enterprise IPv6 Deployment

BRKRST-2301

TOMORROW starts here.

Reference Materials

- Deploying IPv6 in the Internet Edge: http://www.cisco.com/en/US/docs/solutions/Enterprise/Borderless_Networks/Internet_Edge/InternetEdgeIPv6.html
- Deploying IPv6 in Campus Networks: http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/CampIPv6.html
- Deploying IPv6 in Branch Networks: http://www.cisco.com/en/US/docs/solutions/Enterprise/Branch/BrchIPv6.html
- New/Updated IPv6 Cisco Sites: http://www.cisco.com/go/ipv6 http://www.cisco.gom/go/entipv6
- Cisco Network Designs: http://www.cisco.com/go/designzone
- Smart Business Architecture – IPv6 Guides:http://www.cisco.com/en/US/netsol/ns982/networking_solutions_program_home.html
- IPv6 Knowledge Base Portal: http://www.cisco.com/web/solutions/netsys/ipv6/knowledgebase/index.html

Recommended Reading

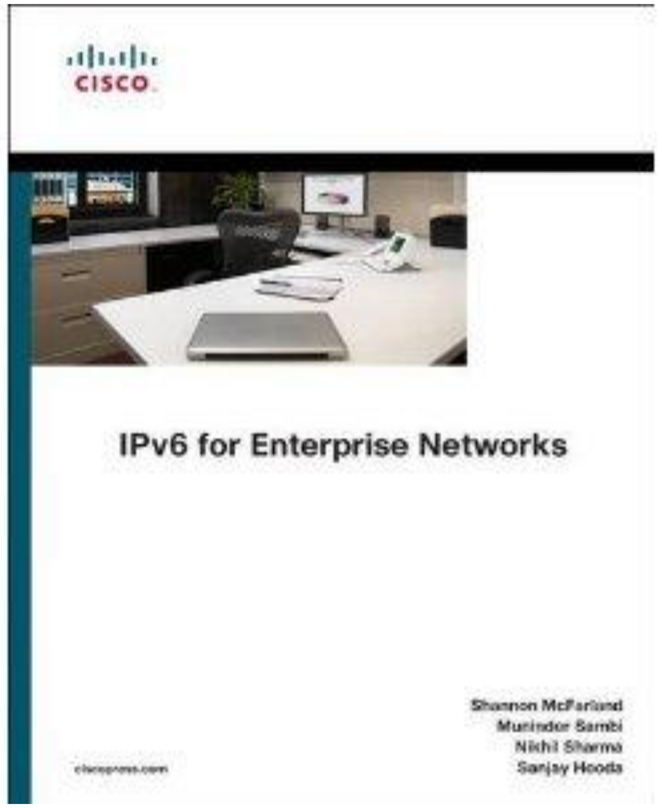

Deploying IPv6 in Broadband Networks - Adeel Ahmed, Salman Asadullah ISBN0470193387, John Wiley & Sons Publications®

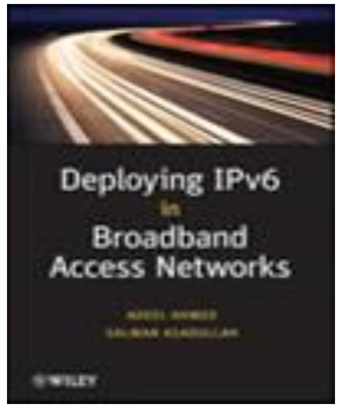

Search Session Builder: “ipv6”

IPv6 Related Sessions at Cisco Live

Session	Title
BRKRST-1069	Understand IPv6
BRKRST-2301	Enterprise IPv6 Deployment
BRKRST-2311	IPv6 Planning, Deployment and Operation Considerations
BRKSEC-2003	IPv6 Security Threats and Mitigations
BRKSPG-2604	Deploying Carrier Grade IPv6 using CGSE
COCRST-2464	Inside Cisco IT: Making the Leap to IPv6
TECRST-2661	Hands on Experience with IPv6

Agenda

- Why are we here?
- Planning and Deployment Summary
- IPv6 Address Considerations
- Infrastructure Deployment
  - Campus
  - Data Centre/Internet Edge
  - WAN/Branch

Dramatic Increase in Enterprise Activity

Why?

Business Continuity

External Pressure	Growth/Protection	• Enterprise that is or will be expanding into new markets • Address exhaustion – Protecting your revenue
	Partnership	• Enterprise that partners with other companies/organisations doing IPv6 • Governments, enterprise partners, contractors
Internal Pressure	OS/Apps	• Microsoft Windows 7, Server 2008 • Microsoft DirectAccess
	Fixing Old Problems	• Mergers & Acquisitions • NAT Overlap
	New Technologies	• High Density Virtual Machine environments (Server virtualisation, VDI) • SmartGrid

Innocent W2K3 -to- W2K8 Upgrade

Windows 2003

```
C:\>ping svr-01

Pinging svr-01.example.com [10.121.12.25] with 32 bytes of data:
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
```

px

Upgraded Host to Windows 2008

```
C:\>ping svr-01

Pinging svr-01 [fe80::c4e2:f21d:d2b3:8463%15] with 32 bytes of data:
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
```

No.	Time	Source	Destination	Protocol	Info
3969	244.938775	fe80::c4e2:f21d:d2b3:8463	ff02::1:3	UDP	Source port: 63828 Destination port: llmnr
3970	244.938958	10.121.12.25	224.0.0.252	UDP	Source port: 53753 Destination port: llmnr

...........**svr-01**........

Mergers & Acquisitions

- Unique blend of technical and business problems
- Colliding RFC1918 space
- Common options
  - If you don’t collide then leave as-is until renumbering is complete
  - NAT overlap pools (into non-colliding space) until renumbering is complete
  - IPv6 as an overlay network
  - IPv6 added as a native protocol (dual stack)
- This is a growing issue and IPv6 ends up being a perfect tool for resolving the technical issues

NAT Overlap + IPv6 Overlay Network

Build an overlay network to encapsulate IPv6 over IPv4

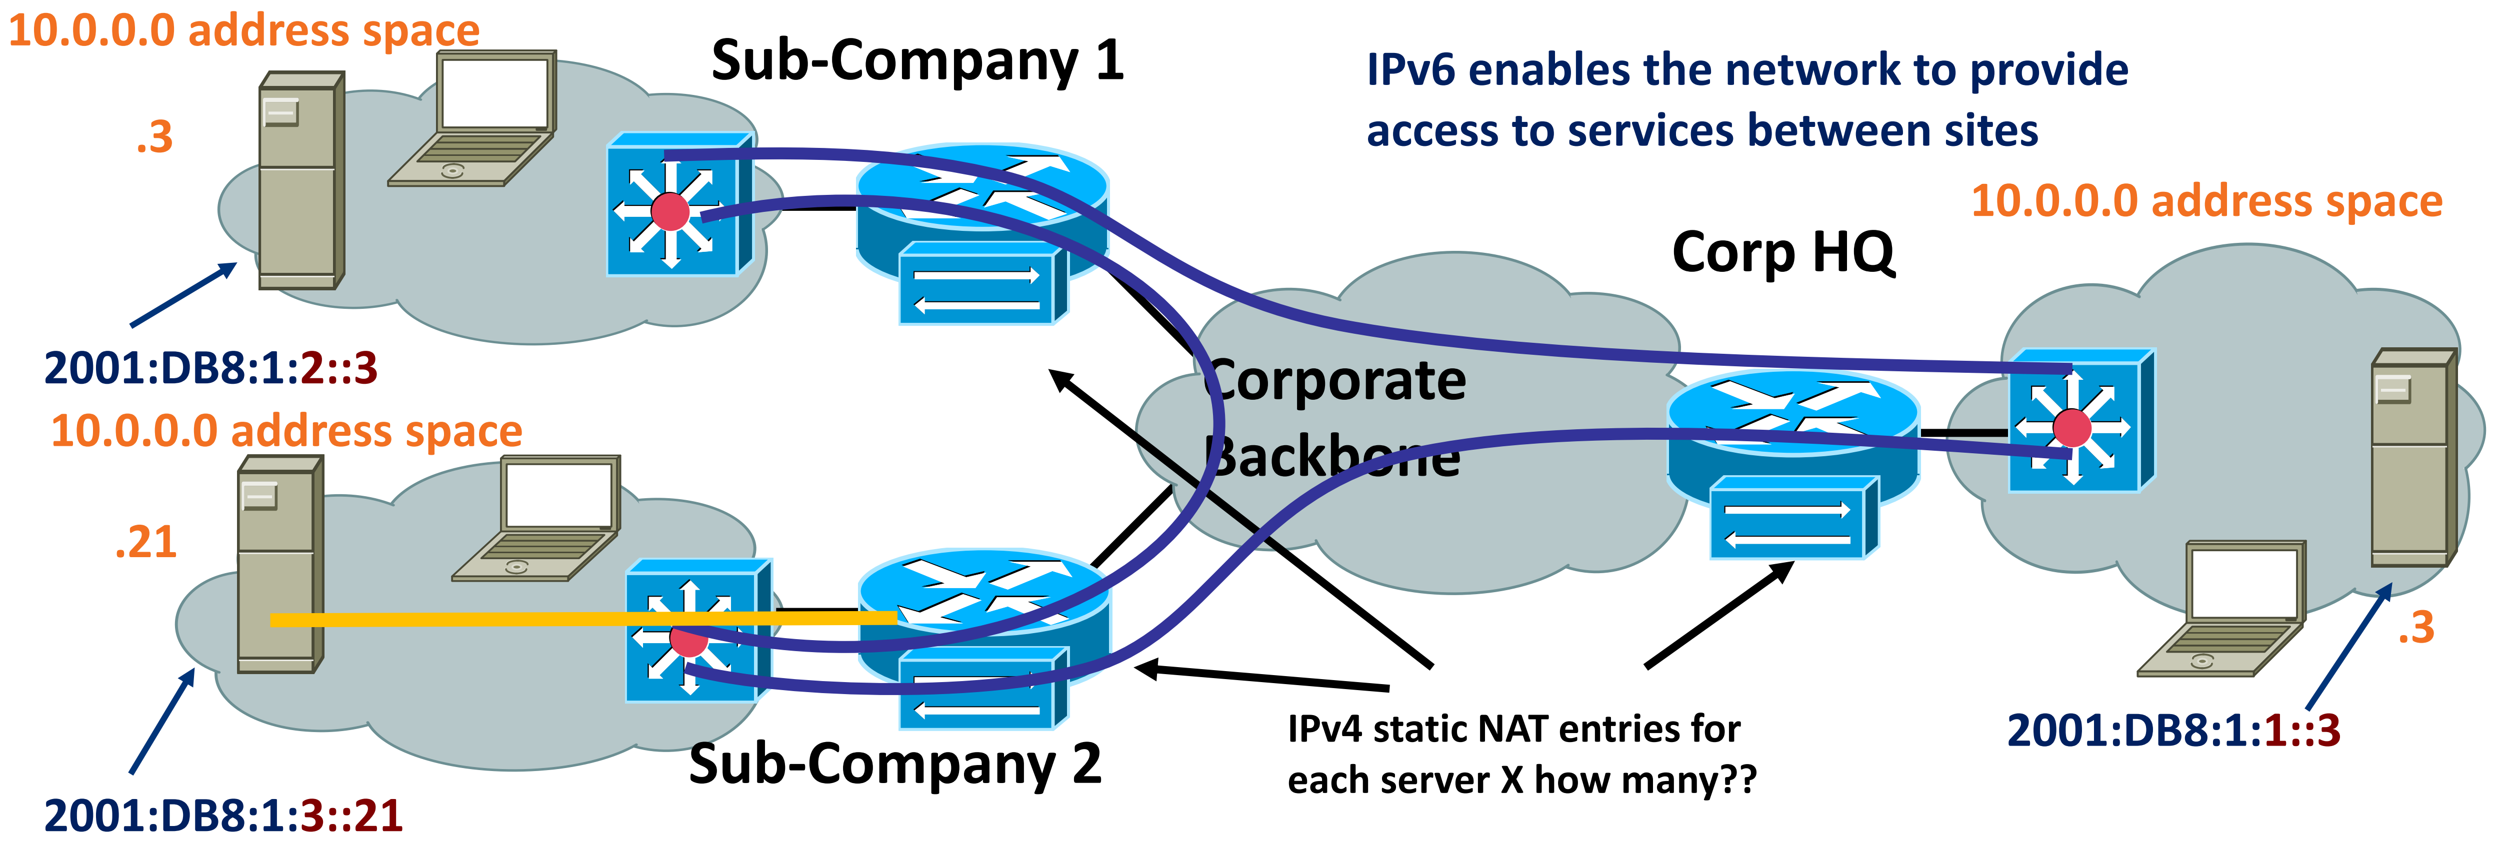

- IPv6 is deployed only at those sites and for specific hosts that need end-to-end routability between entities
- Can be very operationally difficult to maintain in large environments
- May be a show stopper if you have to get a lot of tunnels past a bunch of IPv4 NAT

Partial Overlay + Partial Dual Stack

Combine overlay network with dual stack

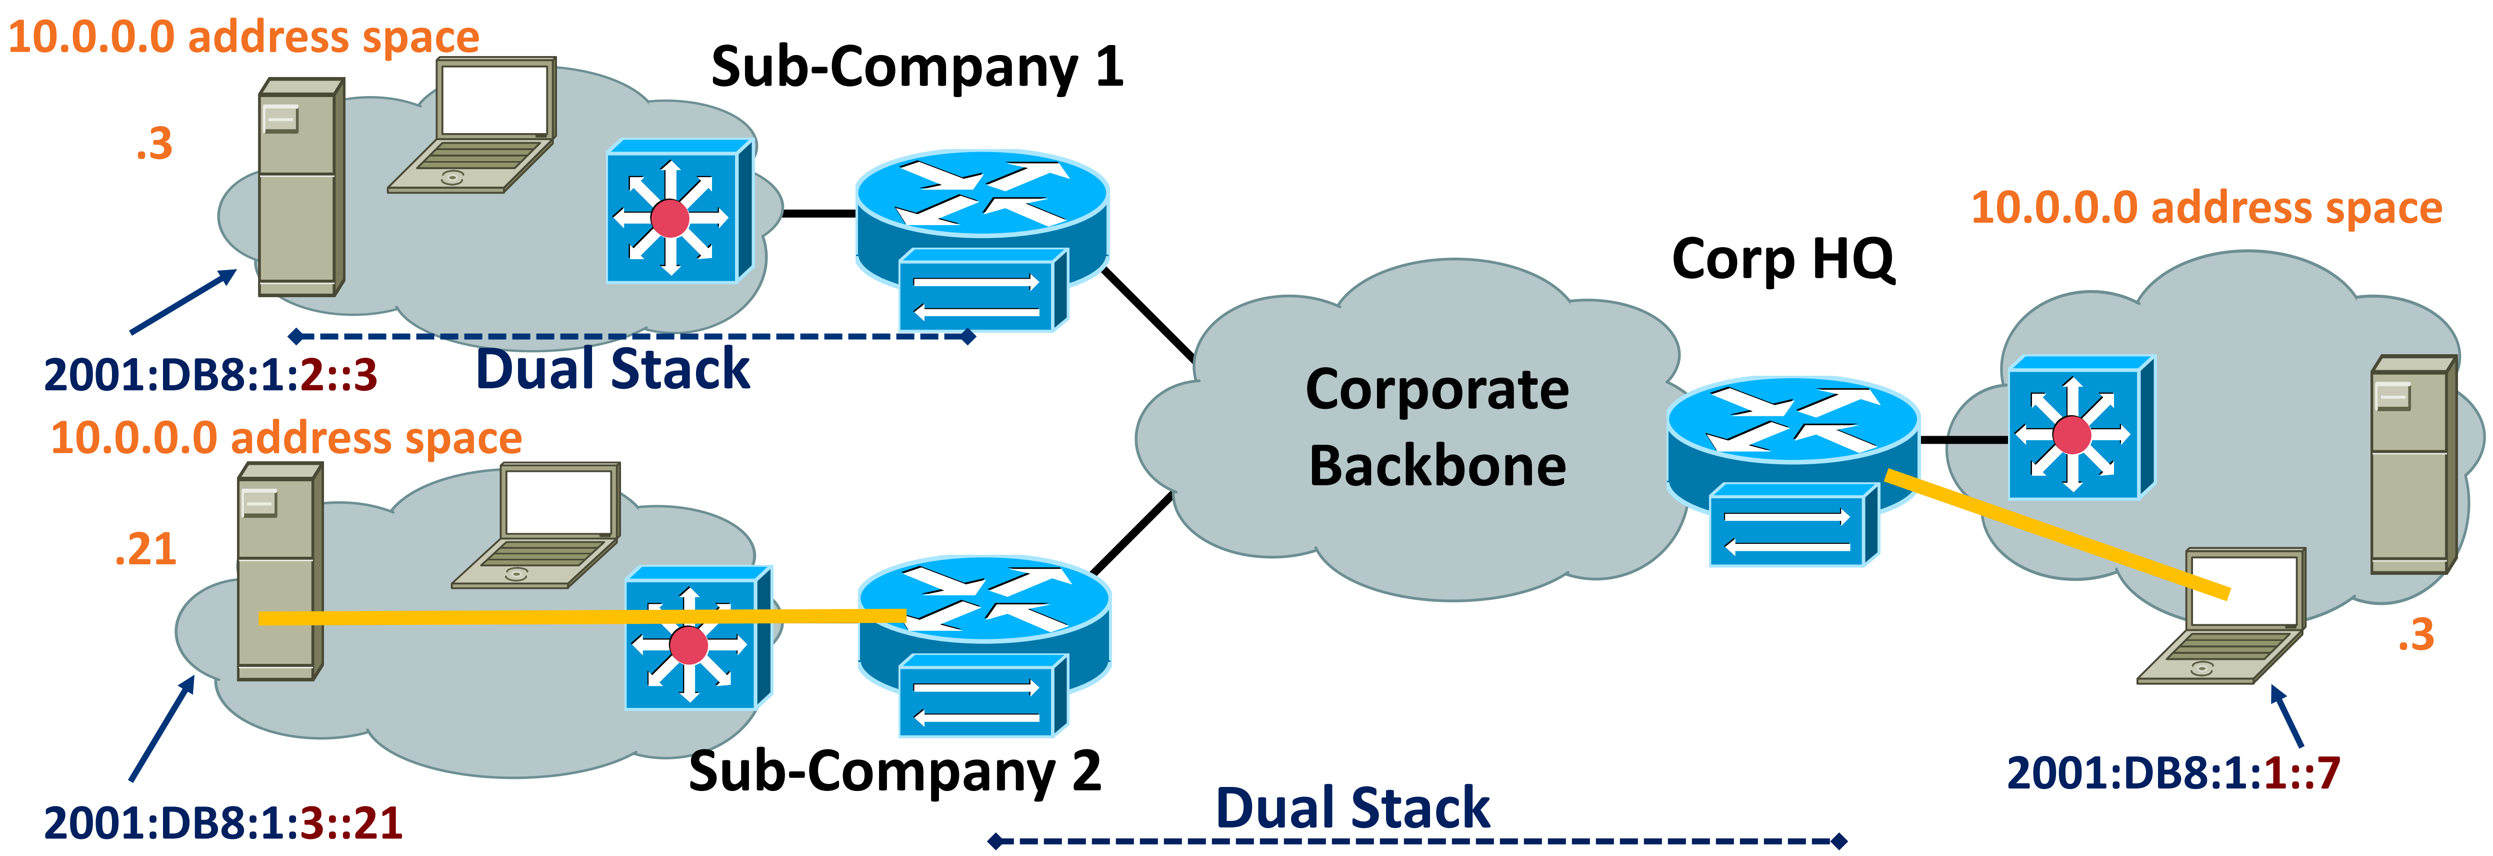

- Build as much dual stack as you can – tunnel only when you have to
- You won't want to keep this up forever – goal is dual stack to all places that need end-to-end connectivity between sites/orgs

Dual Stack Everywhere

Dual stack everywhere – there is nothing else to say ;-)

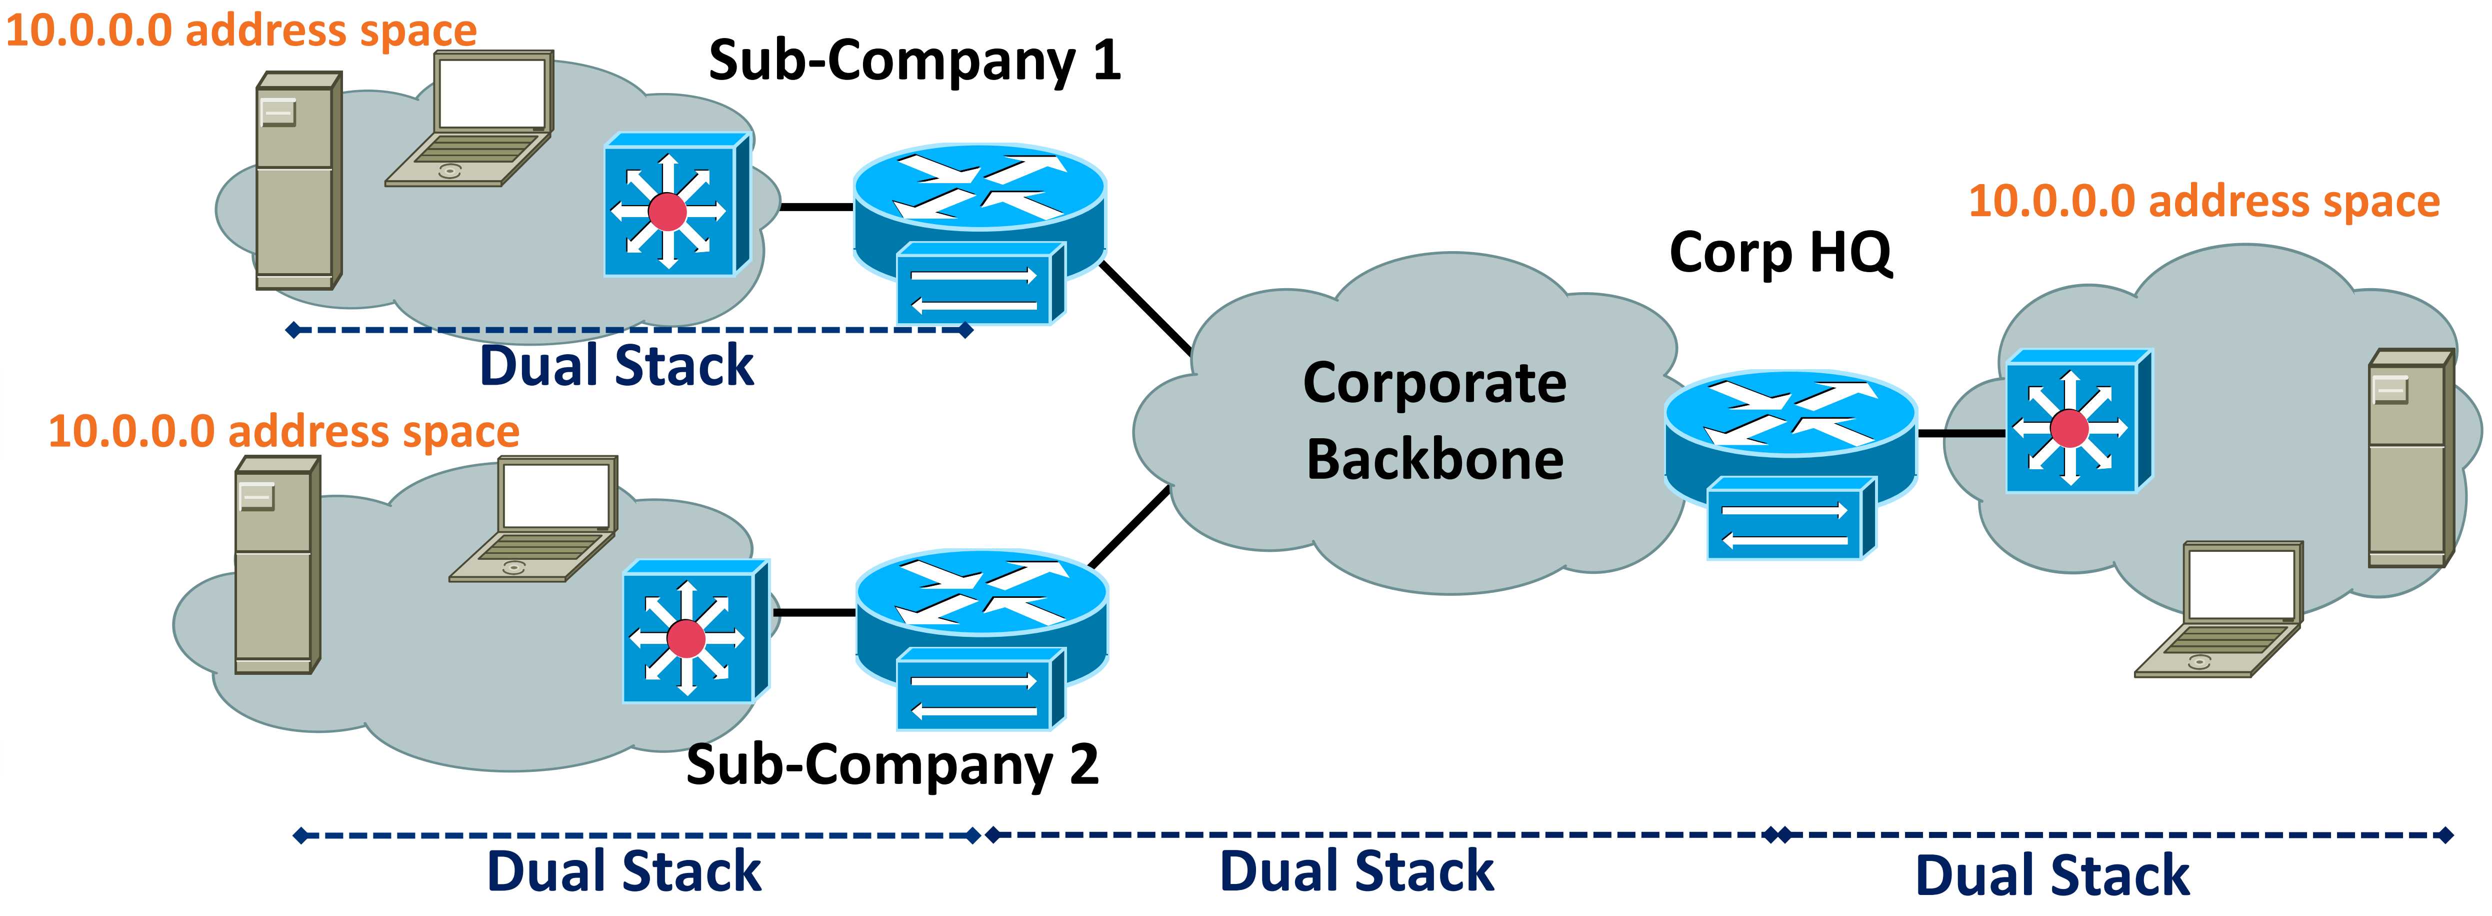

- We will discuss the deployment of dual stack and other end-to-end considerations for the rest of this talk

Planning and Deployment Summary

Architectural Scope of IPv6 Deployment

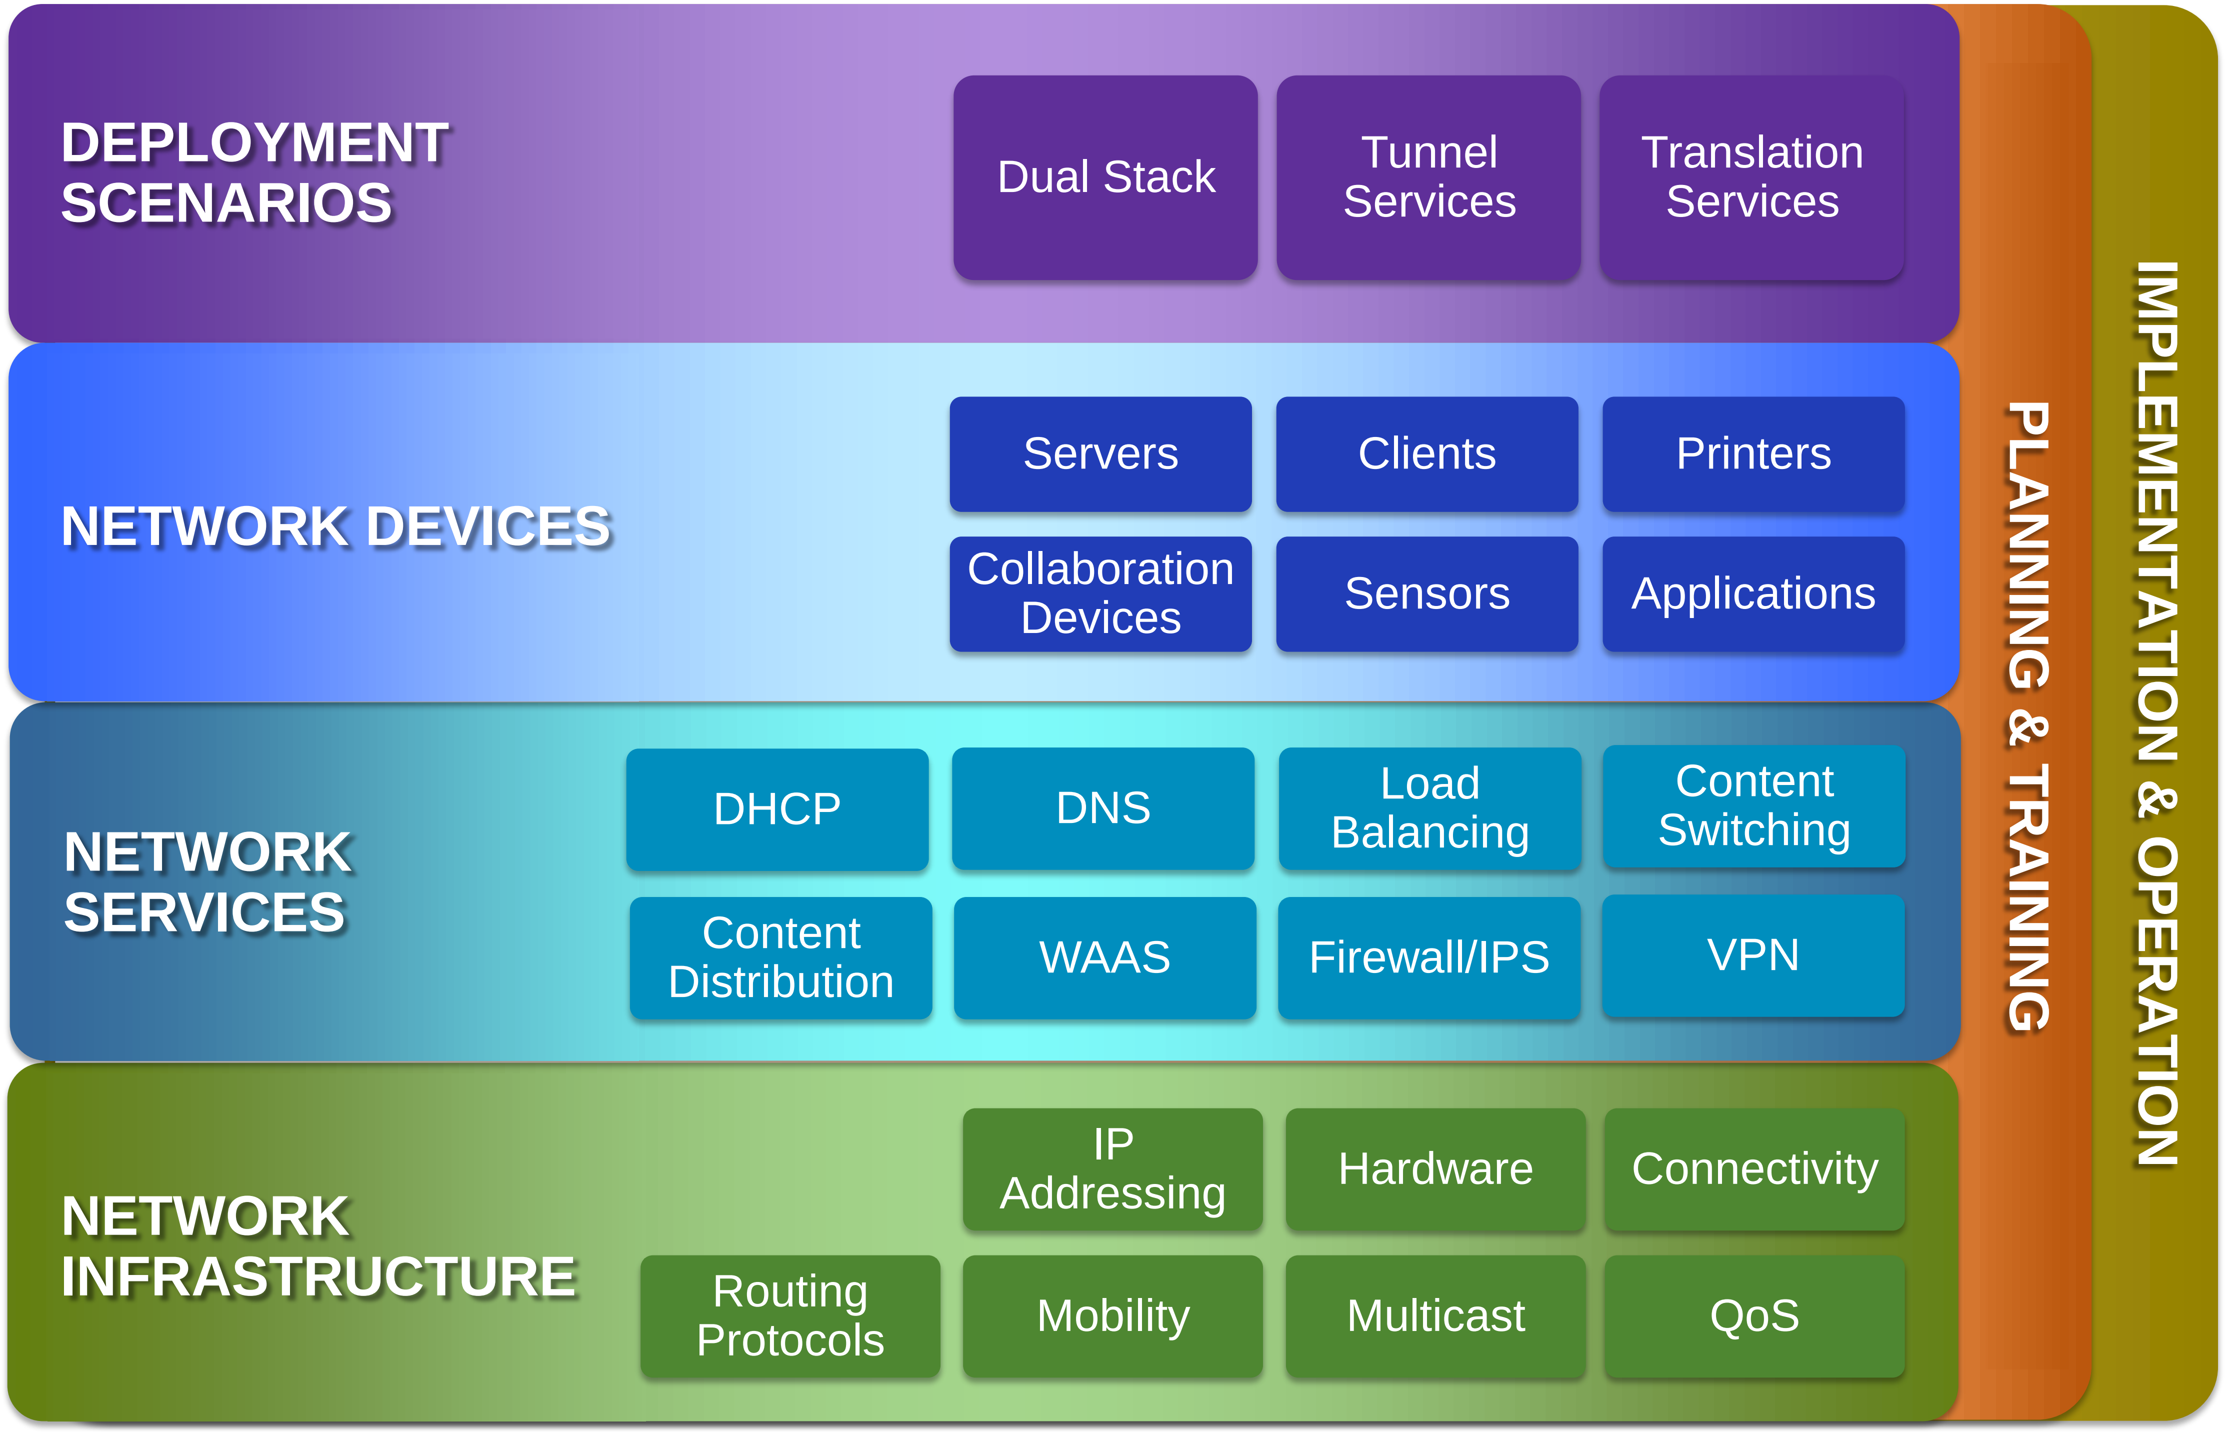

IPv6 Integration Outline

Pre-Deployment Phases	Deployment Phases
• Establish the network starting point • Importance of a network assessment and available tools • Obtain addressing • Build initial addressing architecture • What content are you serving?	• Peering capabilities • Internet Edge (ISP, Apps) • Campus IPv6 integration options • Data Centre integration options • WAN IPv6 integration options • Execute on gaps found in assessment

Where do I Start?

- Based on Timeframe/Use case
- Core-to-Edge – You are in no rush
- Edge-to-Core – Connect labs/developers now
- Internet Edge – Business continuity

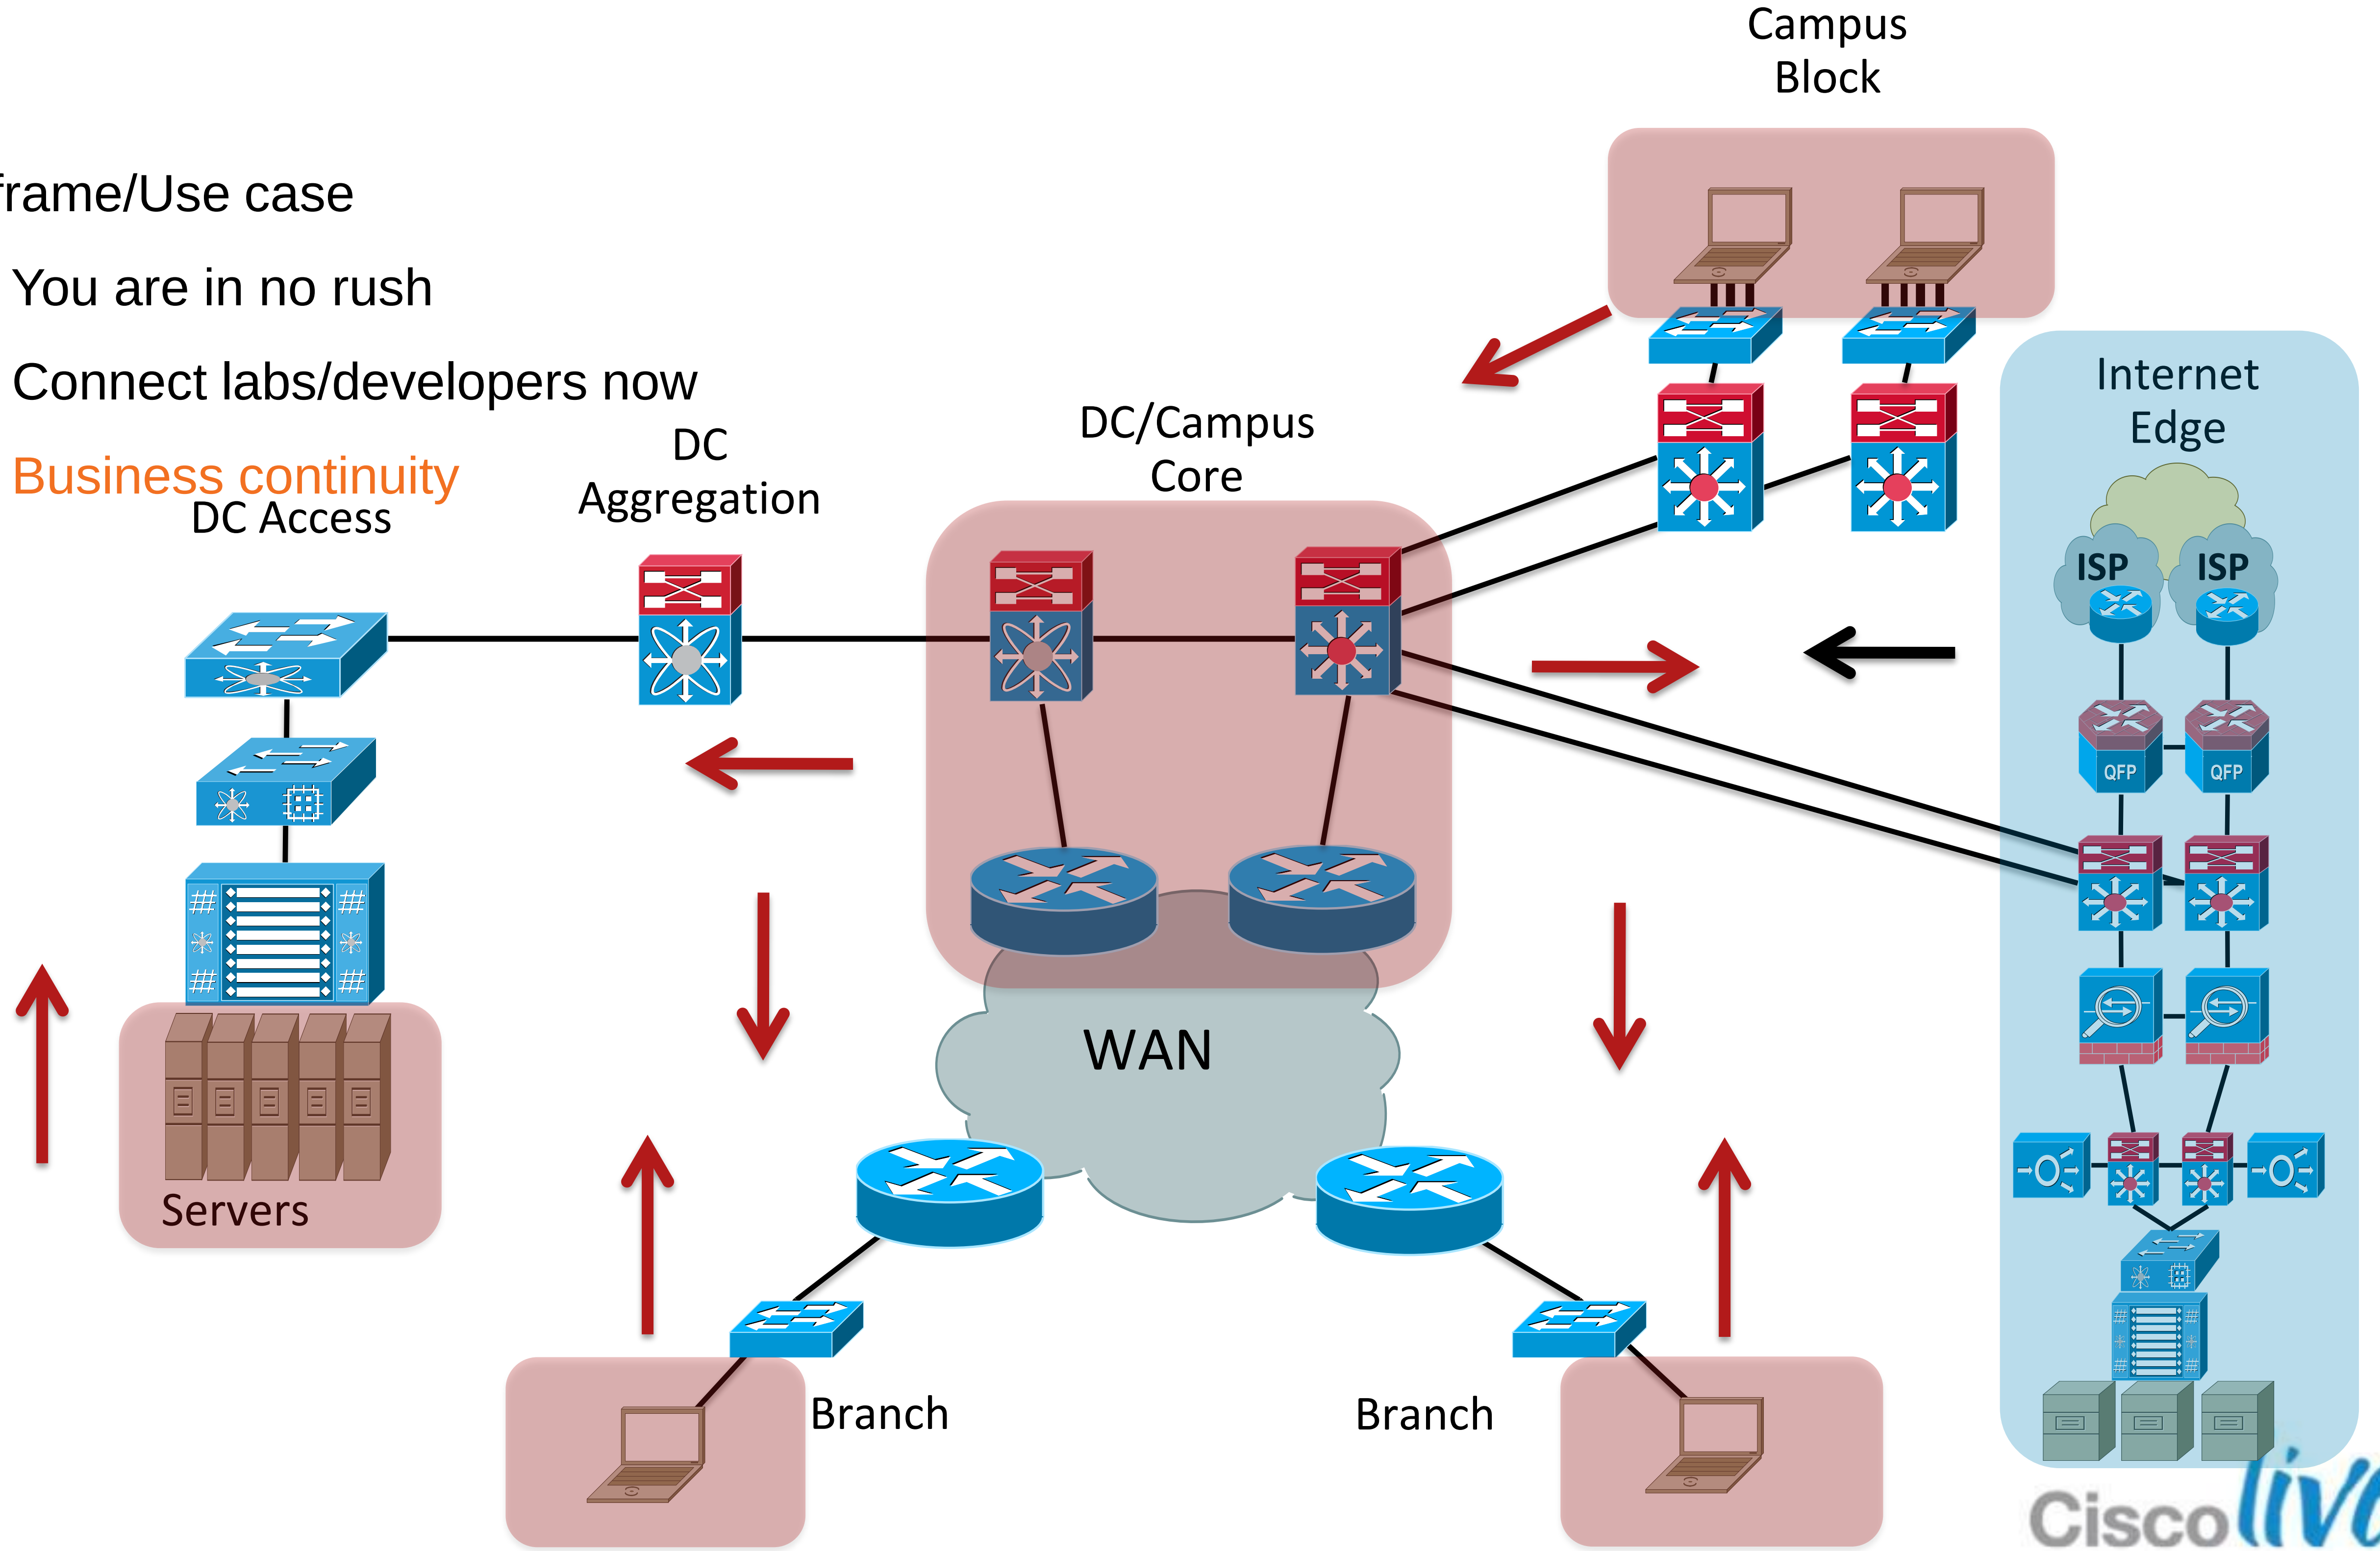

IPv6 Address Considerations

- http://www.cisco.com/web/strategy/docs/gov/IPv6_WP.pdf

IPv6 Addresses

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters
  - Separated by a colon (:)

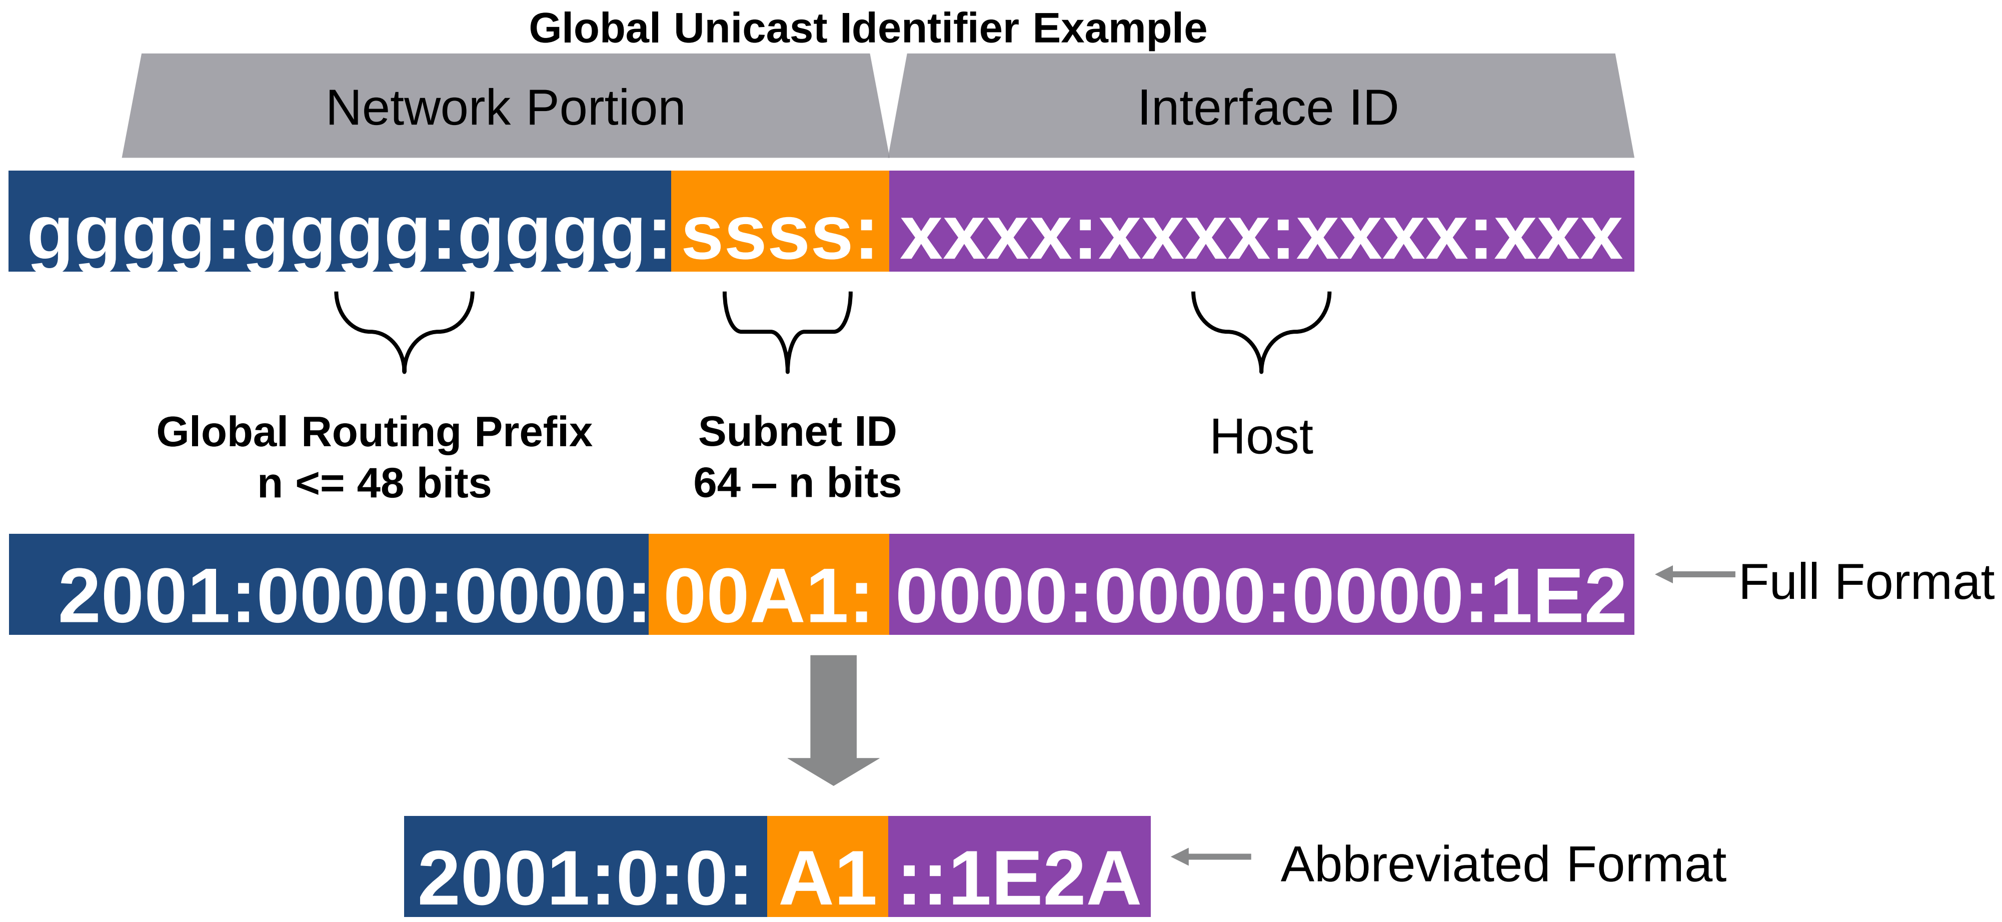

PI and PA Allocation Process

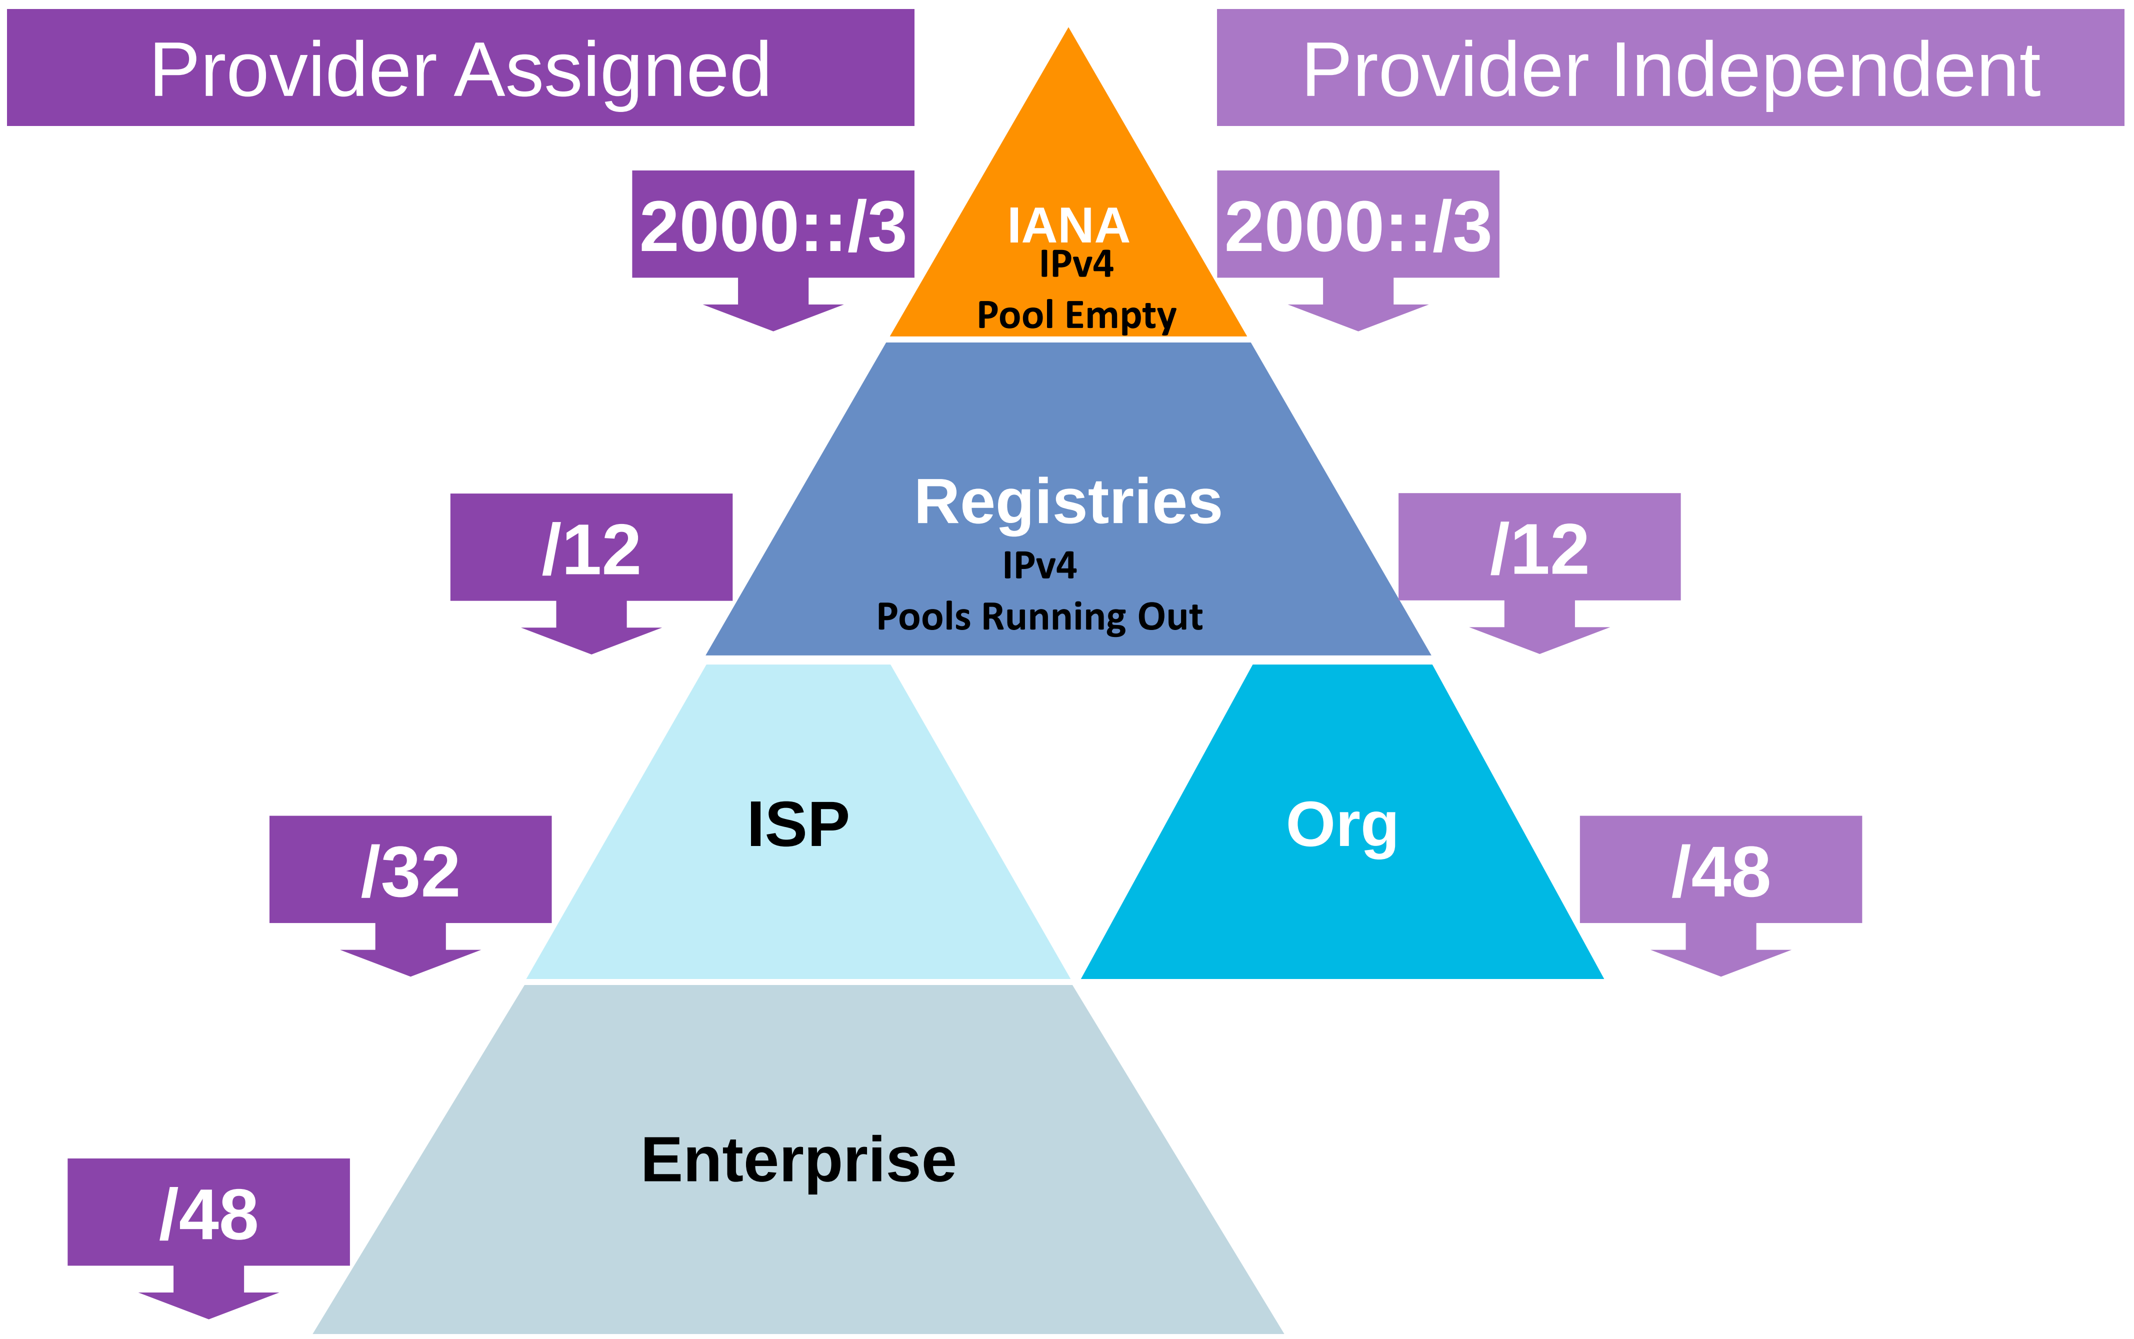

Obtaining IPv6 Address – The Provider Independent Method

- APNIC Kickstart IPv6: http://www.apnic.net/services/apply-for-resources/kickstart-your-ipv6
- IPv6 Address Allocation and Assignment Policy: http://www.apnic.net/policy/ipv6-address-policy
- Section 5.9 – IPv6 Portable Assignments with multihoming
- Proposal 101 (consensus reached) – Remove multihoming requirement: http://www.apnic.net/policy/proposals/prop-101

ULA, ULA + Global or Global-only

- What type of addressing should I deploy internal to my network? It depends:
  - ULA-only—Today, no IPv6 NAT is useable in production so using ULA-only will not work externally to your network
  - ULA + Global allows for the best of both worlds **but** at a price— much more address management with DHCP, DNS, routing and security—SAS does not always work as it should
  - Global-only—Recommended approach but the old-school security folks that believe topology hiding is essential in security will bark at this option
- Let's explore these options…

Unique-Local Addressing
(RFC4193)

- Used for internal communications, inter-site VPNs
  - Not routable on the internet—basically RFC1918 for IPv6 only better—less likelihood of collisions
- Default prefix is /48
  - /48 limits use in large organisations that will need more space
  - Semi-random generator prohibits generating sequentially 'useable' prefixes—no easy way to have aggregation when using multiple /48s
  - Why not hack the generator to produce something larger than a /48 or even sequential /48s?
  - Is it 'legal' to use something other than a /48? Perhaps the entire space? Forget legal, is it practical? Probably, but with dangers—remember the idea for ULA; internal addressing with a slim likelihood of address collisions with M&A. By consuming a larger space or the entire ULA space you will significantly increase the chances of pain in the future with M&A
- Routing/security control
  - You must always implement filters/ACLs to block any packets going in or out of your network (at the Internet perimeter) that contain a SA/DA that is in the ULA range— today this is the **only** way the ULA scope can be enforced
- Generate your own ULA: http://www.sixxs.net/tools/grh/ula/

```
Generated ULA= fd9c:58ed:7d73::/48

    * MAC address=00:0D:9D:93:A0:C3 (Hewlett Packard)
    * EUI64 address=020D9Dfffe93A0C3
    * NTP date=cc5ff71943807789 cc5ff71976b28d86
```

ULA-Only

Not Recommended Today

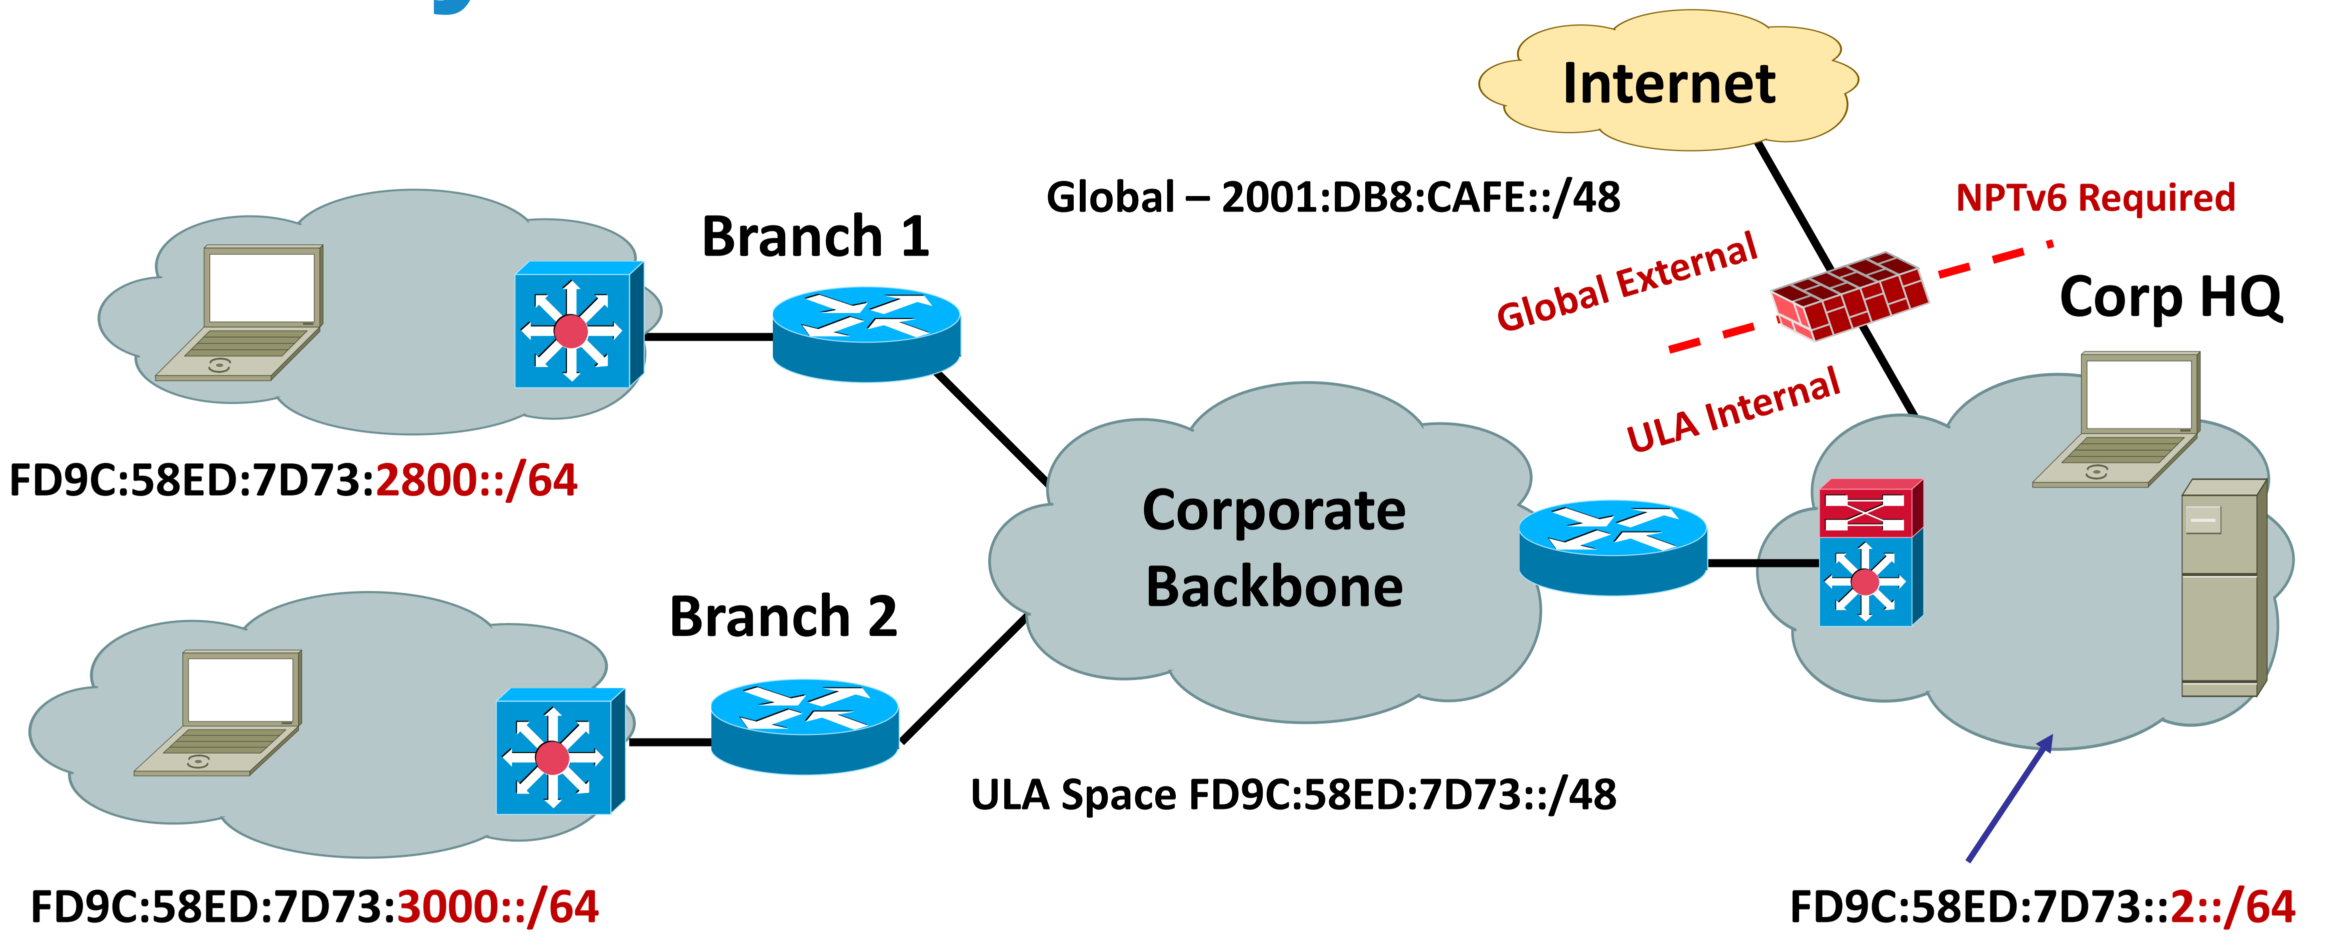

- Everything internal runs the ULA space
- A NAT supporting IPv6 or a proxy is required to access IPv6 hosts on the internet
- Is there a NAT66? RFC6296 (Network Prefix Translation (NPTv6)
- draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-xx
- Removes the advantages of not having a NAT (i.e. application interoperability, global multicast, end-to-end connectivity)

ULA + Global

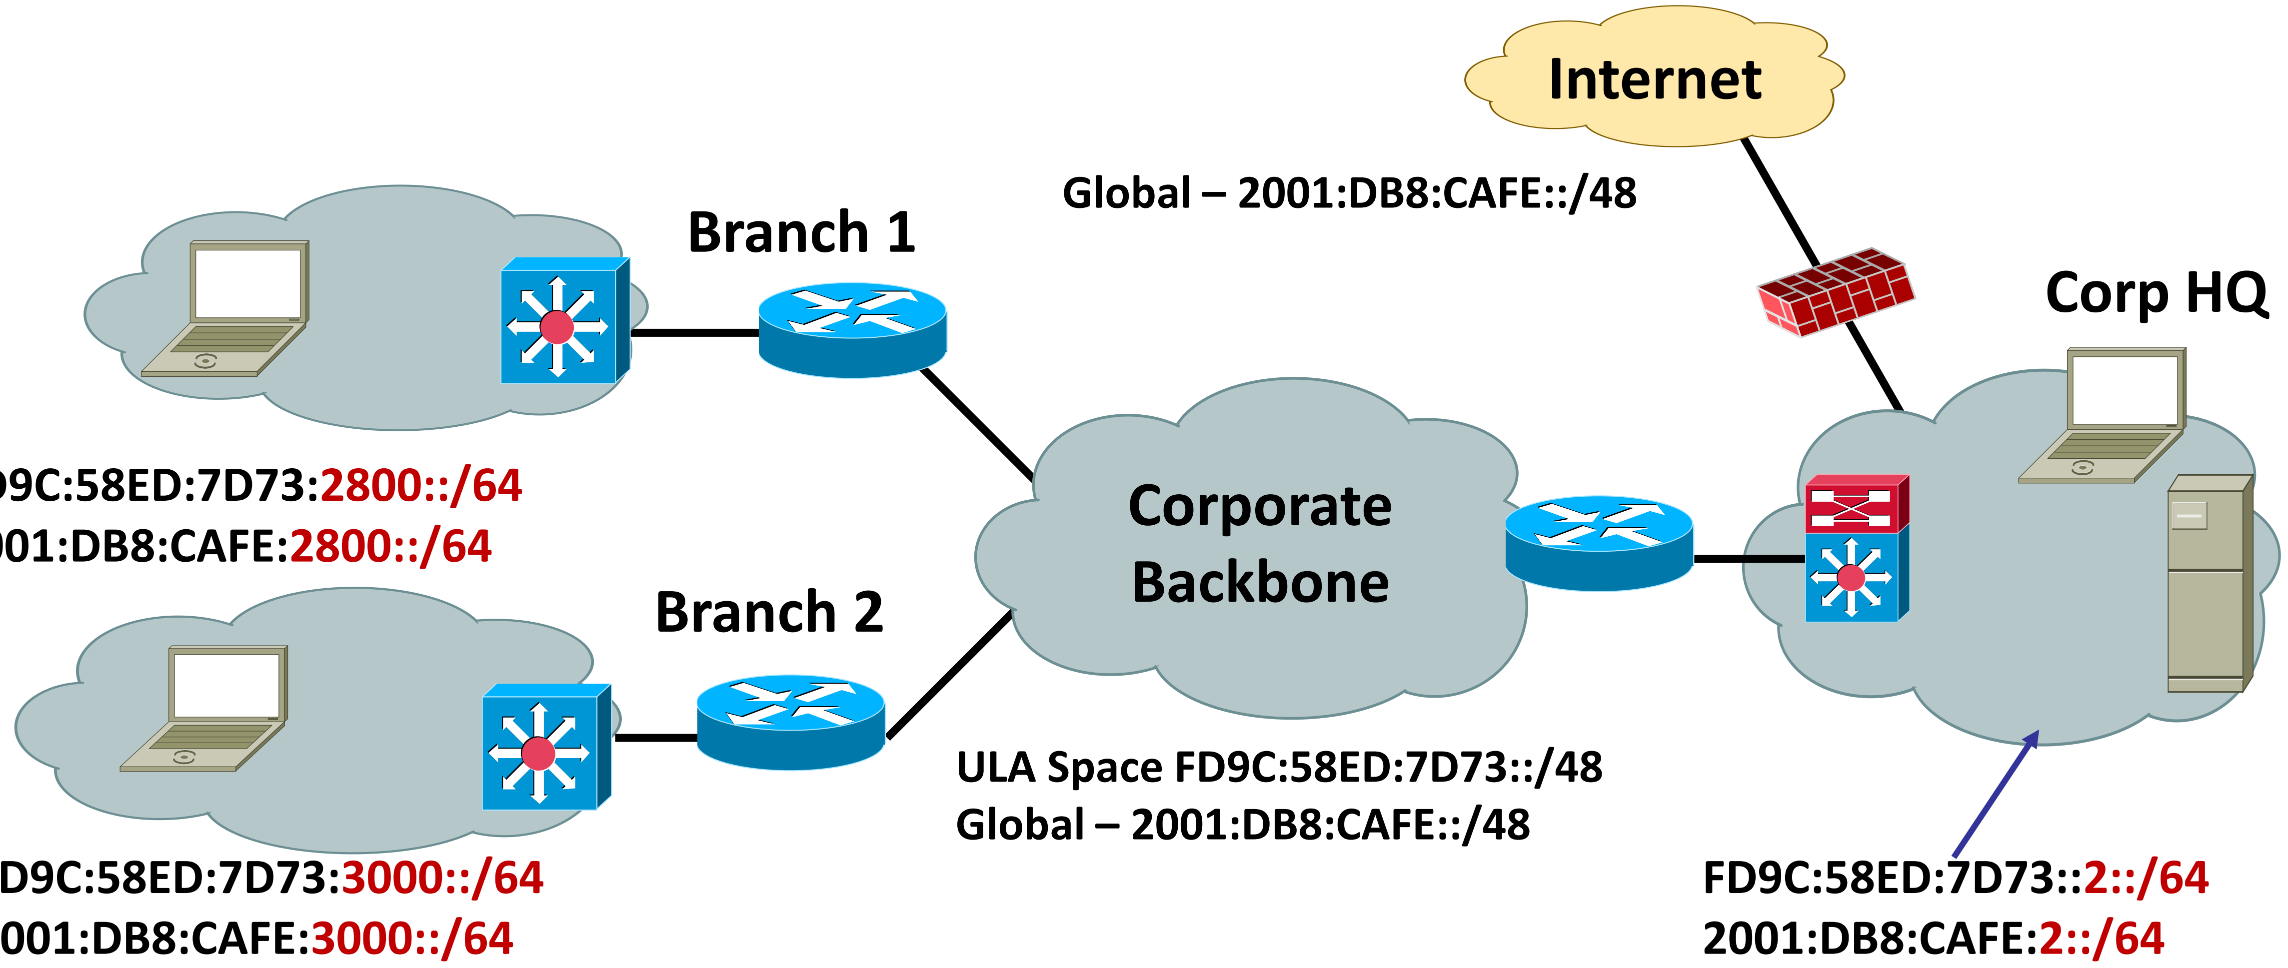

- Both ULA and Global are used internally except for internal-only hosts
- Source Address Selection (SAS) is used to determine which address to use when communicating with other nodes internally or externally
- In theory, ULA talks to ULA and Global talks to Global—SAS ‘should’ work this out
- ULA-only and Global-only hosts can talk to one another internal to the network
- Define a filter/policy that ensures your ULA prefix does not ‘leak’ out onto the Internet and ensure that no traffic can come in or out that has a ULA prefix in the SA/DA fields
- Can be a management NIGHTMARE for DHCP, DNS, routing, security, etc…

Considerations—ULA + Global

- Use DHCPv6 for ULA and Global—apply different policies for both (lifetimes, options, etc..)
- Check routability for both—can you reach an AD/DNS server regardless of which address you have?
- Any policy using IPv6 addresses must be configured for the appropriate range (QoS, ACL, load-balancers, PBR, etc.)
- If using SLAAC for both—Microsoft Windows allows you to enable/disable privacy extensions globally—this means you are either using them for both or not at all!!!
- One option is to use SLAAC for the Global range and enable privacy extensions and then use DHCPv6 for ULA with another IID value (EUI-64, reserved/admin defined, etc.)

```
Temporary  Preferred  6d23h59m55s  23h59m55s 2001:db8:cafe:2:cd22:7629:f726:6a6b
Dhcp       Preferred  13d1h33m55s 6d1h33m55s fd9c:58ed:7d73:1002:8828:723c:275e:846d
Other      Preferred     infinite   infinite fe80::8828:723c:275e:846d%8
```

- Unlike Global and link-local scopes ULA is not automatically controlled at the appropriate boundary—you must prevent ULA prefix from going out or in at your perimeter
- SAS behaviour is OS dependent and there have been issues with it working reliably

ULA + Global Example

```
Addr Type  DAD State   Valid Life Pref. Life Address
---------  ----------- ---------- ---------- ------------------------
Dhcp       Preferred  13d23h48m24s 6d23h48m24s 2001:db8:cafe:2:c1b5:cc19:f87e:3c41
Dhcp       Preferred  13d23h48m24s 6d23h48m24s fd9c:58ed:7d73:1002:8828:723c:275e:846d
Other      Preferred     infinite   infinite   fe80::8828:723c:275e:846d%8
```

```
interface Vlan2
 description ACCESS-DATA-2
 ipv6 address 2001:DB8:CAFE:2::D63/64
 ipv6 address FD9C:58ED:7D73:1002::D63/64
 ipv6 nd prefix 2001:DB8:CAFE:2::/64 0 0 no-autoconfig
 ipv6 nd prefix FD9C:58ED:7D73:1002::/64 0 0 no-autoconfig
 ipv6 nd managed-config-flag
 ipv6 dhcp relay destination 2001:DB8:CAFE:11::9
```

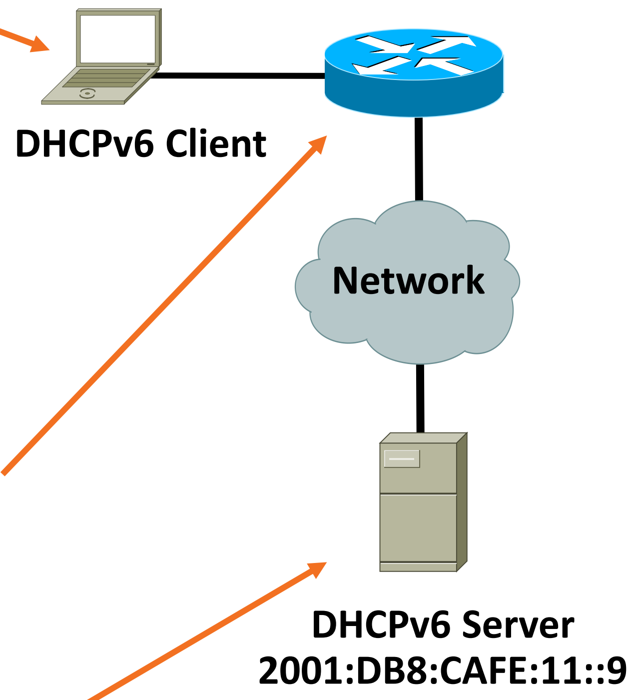

List DHCP Leases for Prefix *VLAN2*

	Address	State	Lookup Key	Flags	State Expi
	2001:db8:cafe:2:c1b5:cc19:f87e:3c41	leased	00:01:00:01:0d:7f:9c:f8:00:0d:60:84:2c:7a		Tue Sep 16

List DHCP Leases for Prefix *VLAN2-ULA*

	Address	State	Lookup Key	Flags	State Expira
	fd9c:58ed:7d73:1002:8828:723c:275e:846d	leased	00:01:00:01:0d:7f:9c:f8:00:0d:60:84:2c:7a		Tue Sep 16 1

Recommended

Global-Only

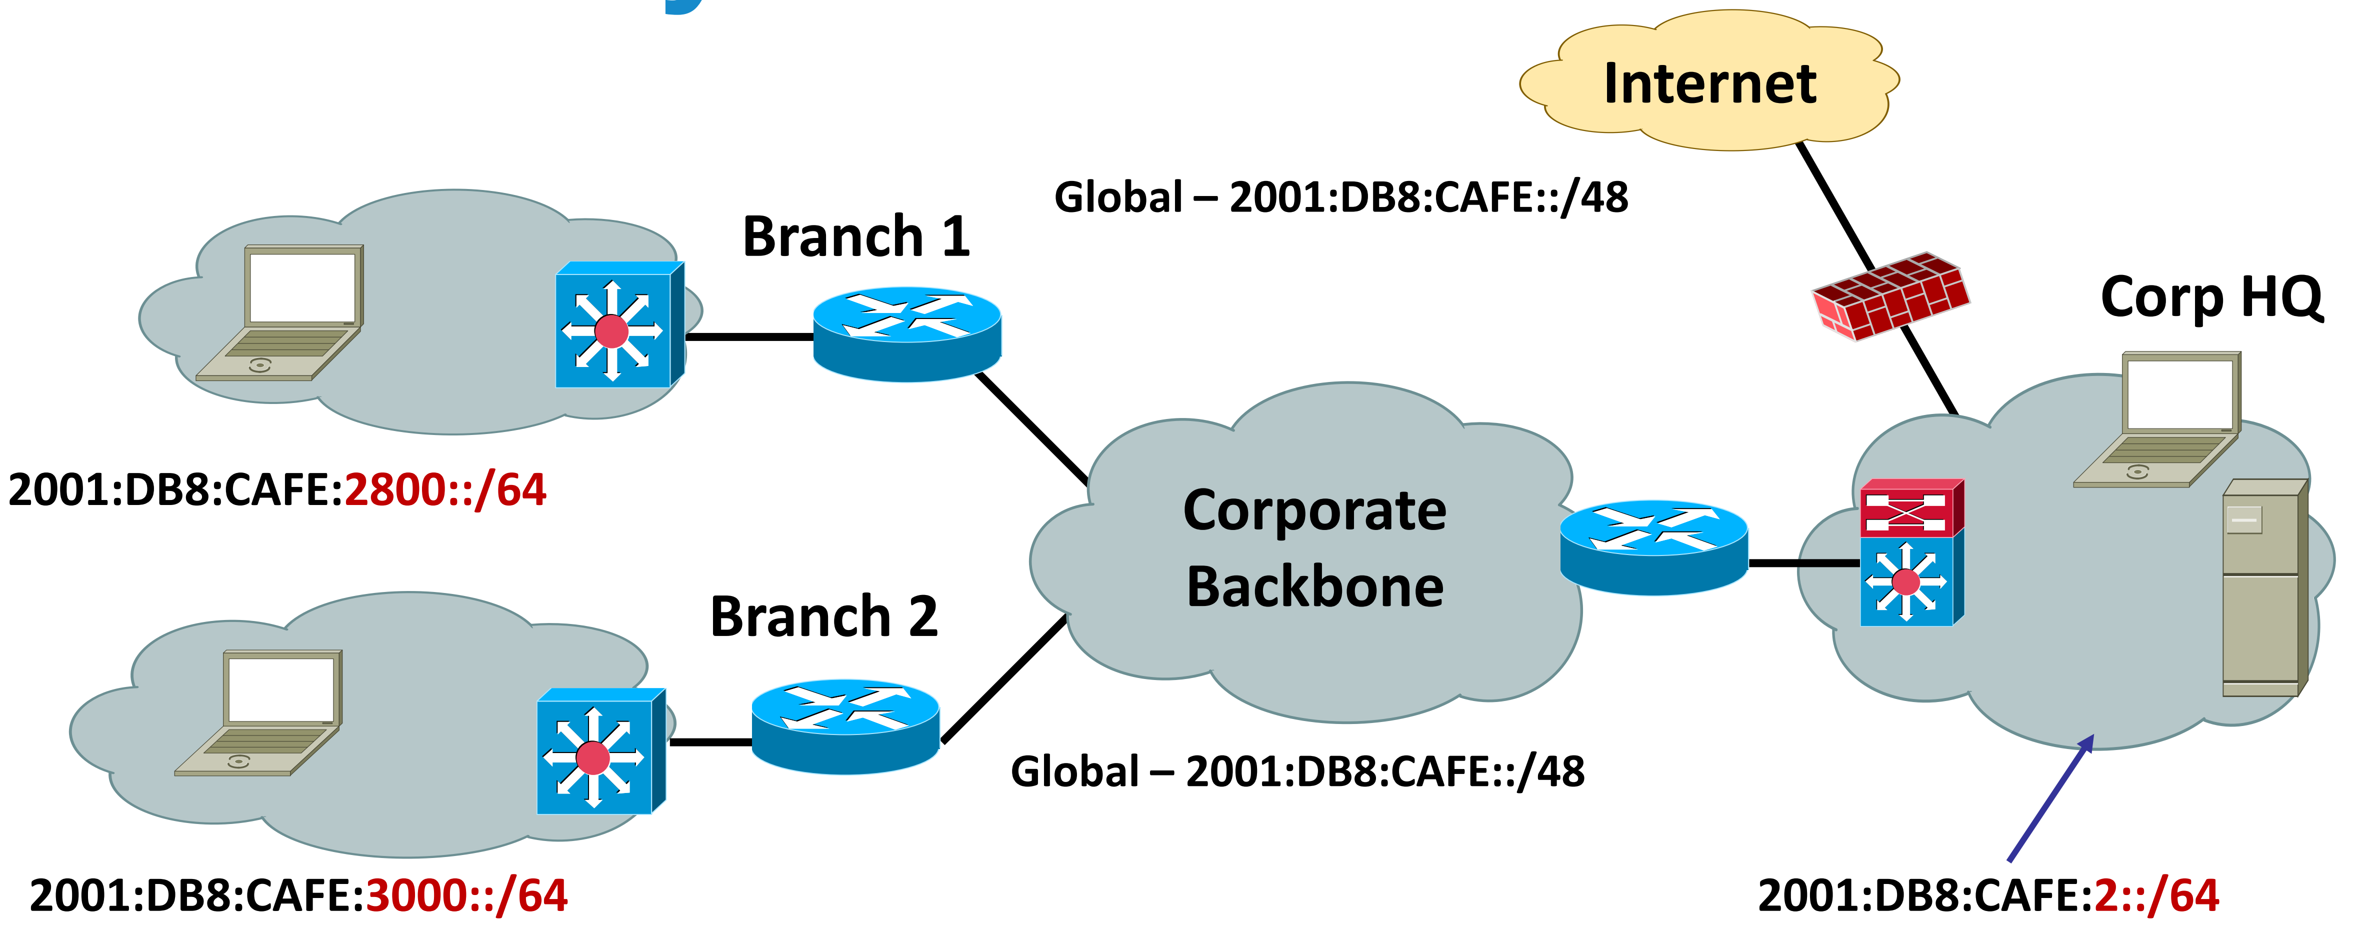

- Global is used everywhere
- No issues with SAS
- No requirements to have NAT for ULA-to-Global translation—but, NAT may be used for other purposes
- Easier management of DHCP, DNS, security, etc.
- Only downside is breaking the habit of believing that topology hiding is a good security method ☺

Randomised IID and Privacy Extensions

- Enabled by default on Microsoft Windows
- Enable/disable via GPO or CLI

```
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
netsh interface ipv6 set privacy state=disabled store=persistent
```

- Alternatively, use DHCP (see later) to a specific pool
- Randomised address are generated for non-temporary autoconfigured addresses including public and link-local—used instead of EUI-64 addresses
- Randomised addresses engage Optimistic DAD—likelihood of duplicate LL address is rare so RS can be sent before full DAD completion
- Windows W7/2008 send RS while DAD is being performed to save time for interface initialisation (read RFC4862 on why this is ok)

Link Level—Prefix Length Considerations

64 bits	> 64 bits
▪ Recommended by RFC3177 and IAB/IESG ▪ Consistency makes management easy ▪ MUST for SLAAC (MSFT DHCPv6 also) ▪ Significant address space loss (18.466 Quintillion)	▪ Address space conservation ▪ Special cases: /126—valid for p2p /127—valid for p2p if you are careful – RFC6164 (RFC3627) /128—loopback ▪ Must avoid overlap with specific addresses: Router Anycast (RFC3513) Embedded RP (RFC3956) ISATAP addresses

- /64 everywhere
- /64 + /126
  - 64 on host networks
  - 126 on P2P
- /64 + /127
  - 64 on host networks
  - 127 on P2P
- /128 on loopback

Using Link-Local for Non-Access Connections

- What if you did not have to worry about addressing the network infrastructure for the purpose of routing?
  - IPv6 IGPs use LL addressing
  - Only use Global or ULA addresses at the edges for host assignment
  - For IPv6 access to the network device itself use a loopback
- What happens to route filters? ACLs?—Nothing, unless you are blocking to/from the router itself
- Stuff to think about:
  - Always use a RID
  - Some Cisco devices require “ipv6 enable” on the interface in order to generate and use a link-local address
  - Enable the IGP on each interface used for routing or that requires its prefix to be advertised

Using LL + Loopback Only

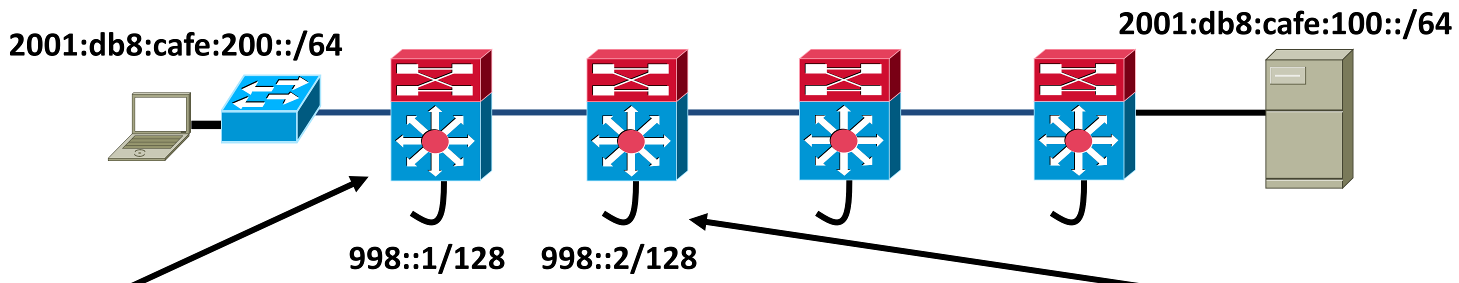


```
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8:CAFE:998::1/128
 ipv6 eigrp 10
!
interface Vlan200
 ipv6 address 2001:DB8:CAFE:200::1/64
 ipv6 eigrp 10
!
interface GigabitEthernet1/1
 ipv6 enable
 ipv6 eigrp 10
!
ipv6 router eigrp 10
 router-id 10.99.8.1
 no shutdown
```

```
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8:CAFE:998::2/128
 ipv6 eigrp 10
!
interface GigabitEthernet3/4
 ipv6 eigrp 10
!
interface GigabitEthernet1/2
 ipv6 eigrp 10
!
ipv6 router eigrp 10
 router-id 10.99.8.2
 no shutdown
```

```
IPv6-EIGRP neighbors for process 10
0   Link-local address:     Gi1/2
    FE80::212:D9FF:FE92:DE77
```

SLAAC & Stateful/Stateless DHCPv6

- StateLess Address AutoConfiguration (SLAAC) – RA-based assignment (a MUST for Mac prior to Lion)
- Stateful and stateless DHCPv6 server
  - Cisco Network Registrar: http://www.cisco.com/en/US/products/sw/netmgtsw/ps1982/
  - Microsoft Windows Server 2008/2012
  - DNSMASQ
- DHCPv6 Relay—supported on routers and switches

```
interface FastEthernet0/1
 description CLIENT LINK
 ipv6 address 2001:DB8:CAFE:11::1/64
 ipv6 nd prefix 2001:DB8:CAFE:11::/64 0 0 no-autoconfig
 ipv6 nd managed-config-flag
 ipv6 dhcp relay destination 2001:DB8:CAFE:10::2
```

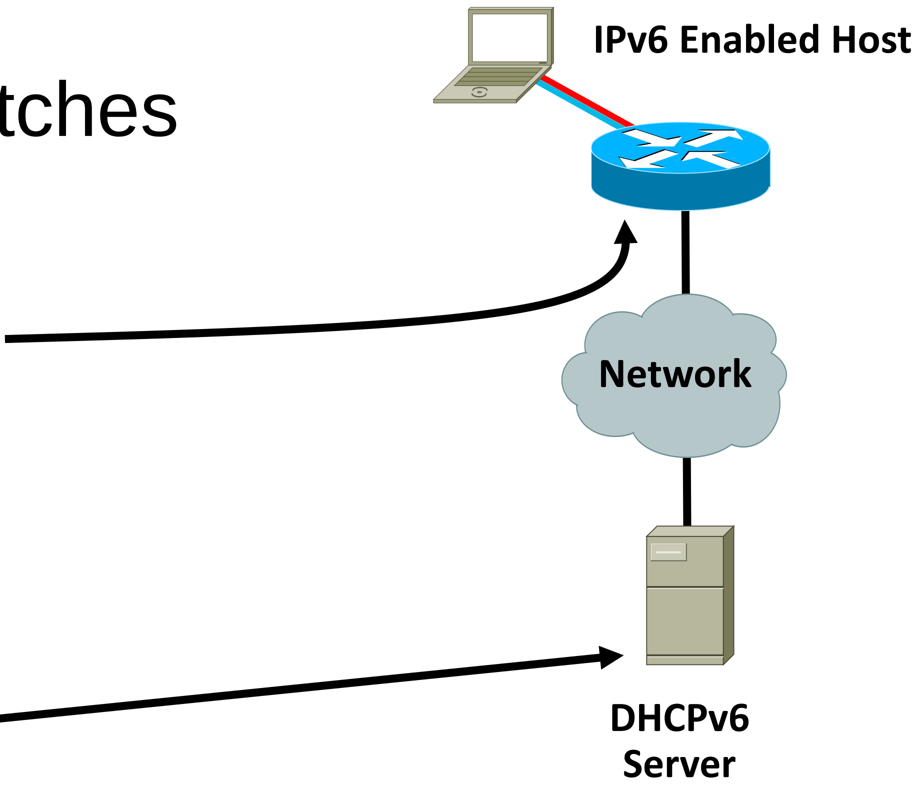

DNS Basic Steps - 1

- Add AAAA records in your DNS server for the hostnames of the devices that can be reached through the IPv6 protocol.
- Add pointer (PTR) records in your DNS server for the IP addresses of the devices that can be reached through the IPv6 protocol.
- Enable IPv6 access to the authoritative DNS servers. Be sure that TCP/53 and UDP/53 can be accessed through IPv6.
- Enable IPv6 connectivity to the external full-service resolvers that send DNS queries to authoritative servers in the world.

DNS Basic Steps - 2

- Make sure that the full-service resolver is configured with both IPv4 and IPv6 glue for the root servers in the world.
- Enable IPv6 on the recursive resolver so that it responds to DNS requests over IPv6 as well as IPv4.
- Enable IPv6 on the node that sends queries so that it can send DNS requests to the recursive resolver.
- Configure the stub resolver on the node that sends queries so that it uses IPv6 to send DNS queries, either statically or using Dynamic Host Configuration Protocol Version 6 (DHCPv6).
- Review policies for flows and make sure that both TCP/53 and UDP/53 can be accessed over IPv4 and IPv6

General Network Considerations

HSRP for IPv6

- Many similarities with HSRP for IPv4
- Changes occur in Neighbour Advertisement, Router Advertisement, and ICMPv6 redirects
- No need to configure GW on hosts (RAs are sent from HSRP active router)
- Virtual MAC derived from HSRP group number and virtual IPv6 link-local address
- IPv6 Virtual MAC range:
  - 0005.73A0.0000 - 0005.73A0.0FFF (4096 addresses)
- HSRP IPv6 UDP Port Number 2029 (IANA Assigned)

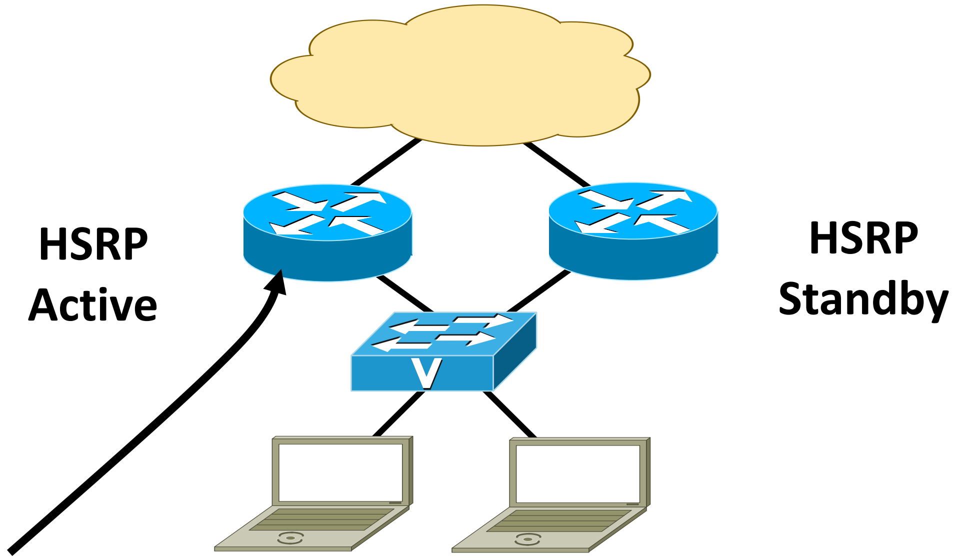


```
track 2 interface FastEthernet0/0 line-protocol
interface FastEthernet0/1
 ipv6 address 2001:DB8:66:67::2/64
 standby version 2
 standby 2 ipv6 autoconfig
 standby 2 timers msec 250 msec 800
 standby 2 preempt
 standby 2 preempt delay minimum 180
 standby 2 authentication cisco
 standby 2 track 2 decrement 10
```

Host with GW of Virtual IP

```
#route -A inet6 | grep ::/0 | grep eth2
::/0      fe80::5:73ff:fea0:1                 UGDA  1024   0        0 eth2
```

IPv6 Multicast

- Multicast Listener Discovery (MLD)
  - – Equivalent to IGMP
- PIM Group Modes: Sparse Mode, Bidirectional and Source Specific Multicast
- RP Deployment: Static, Embedded, Anycast-RP

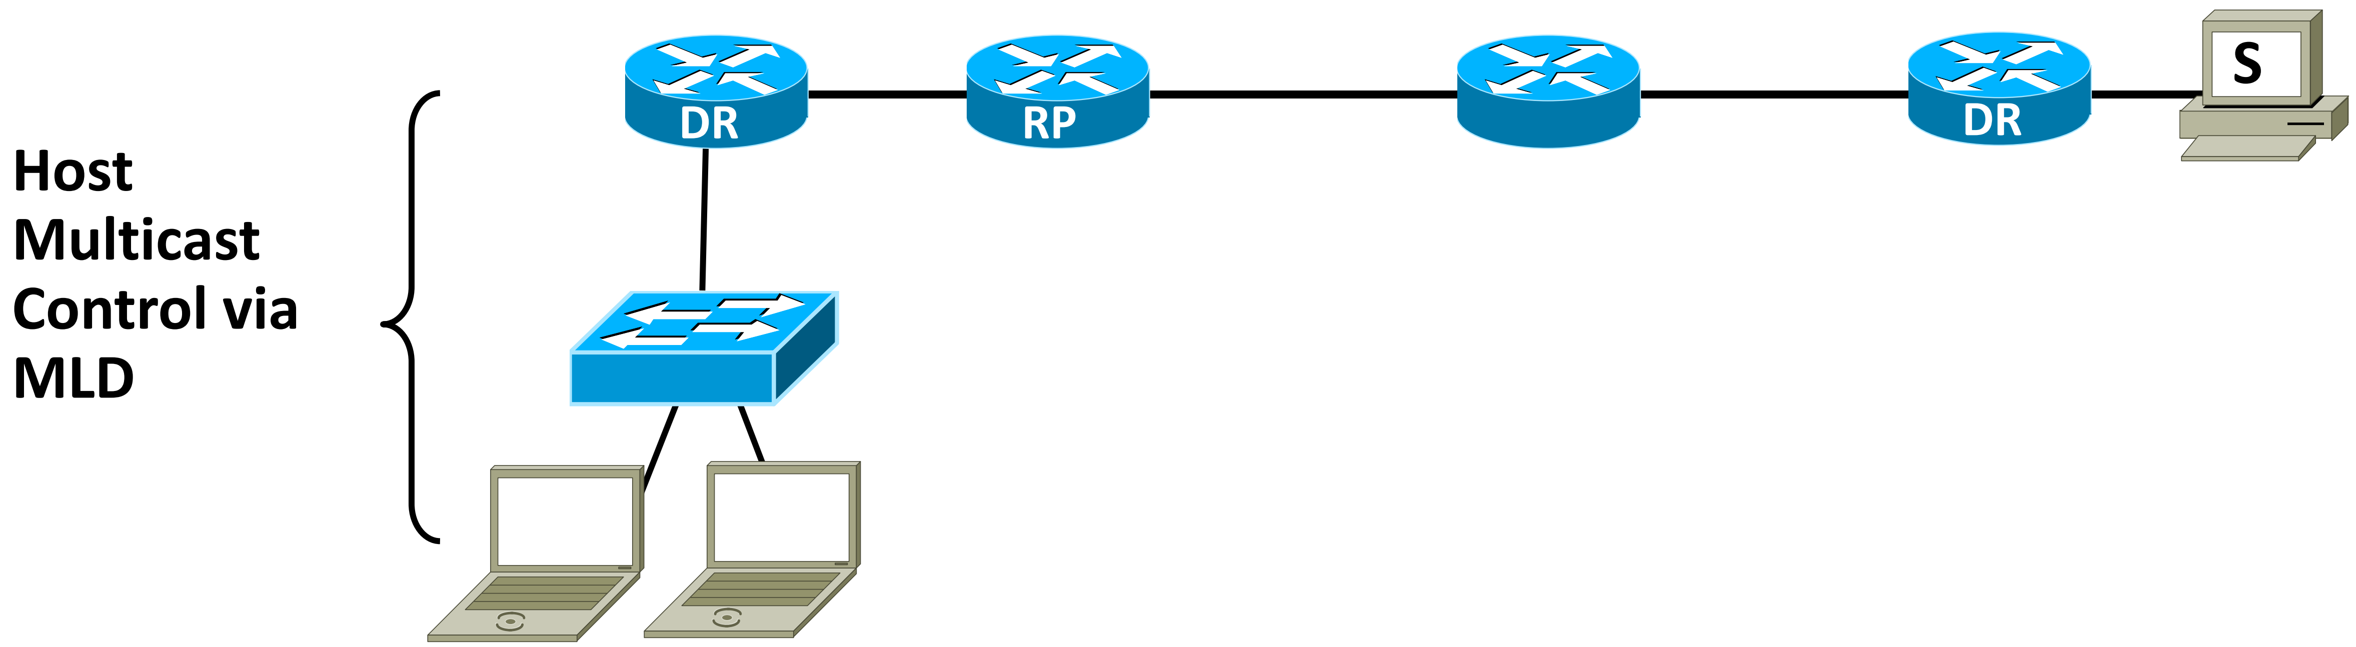

Multicast Listener Discovery: MLD

Multicast Host Membership Control

- MLD is equivalent to IGMP in IPv4
- MLD messages are transported over ICMPv6
- MLD uses link local source addresses
- MLD packets use “Router Alert” in extension header (RFC2711)
- Version number confusion:
  - MLDv1 (RFC2710) like IGMPv2 (RFC2236)
  - MLDv2 (RFC3810) like IGMPv3 (RFC3376)
- MLD snooping

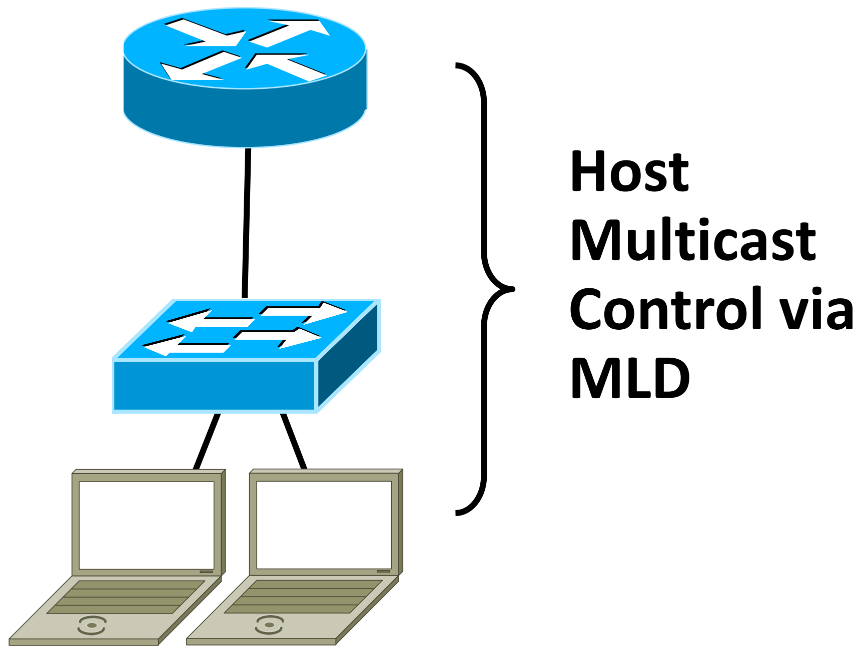

PIMv6 Anycast-RP (RFC4610)

- Support began in 15.1(3)S and XE 3.4S:http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-1s/ip6-pimv6-anycast-rp.html http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/xe-3s/ip6-pimv6-anycast-rp.html
- Similar to original IPv4 PIM Anycast-RP only without MSDP requirement
- Two interfaces defined (usually loopbacks): 1) Used for Anycast-RP address 2) Used for Anycast-RP peer interface
- DRs are configured just like they were in IPv4 Anycast-RP – rp-address entry to Anycast-RP address
- RPs define Anycast-RP “set” with Anycast-RP address and peer address: **ipv6 pim anycast-rp {rp-addresss peer-address}**

PIMv6 Anycast-RP Register – Normal Operation

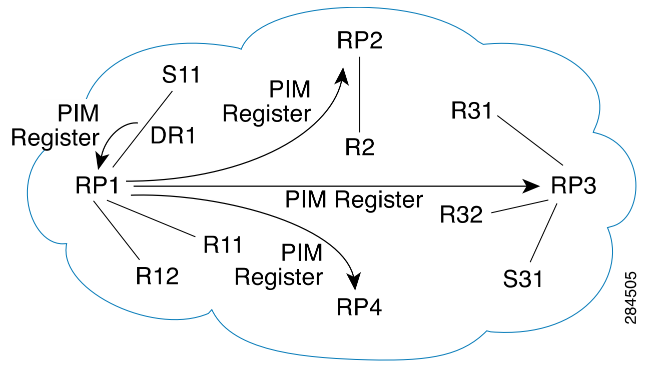

*See notes for walkthrough

PIMv6 Anycast-RP Failover

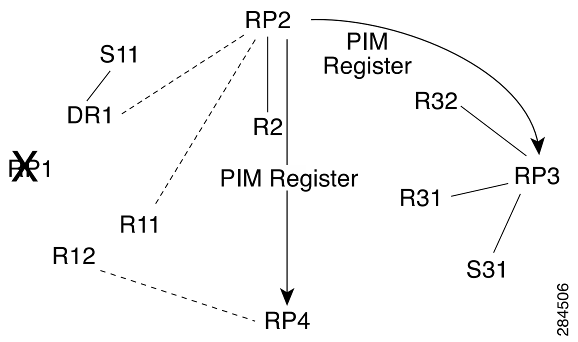

*See notes for walkthrough

RP1

```
ipv6 multicast-routing
!
interface Loopback0
 description ANYCAST-RP
 no ip address
 ipv6 address 2001:DB8:CAFE:999::3/128
 ipv6 eigrp 1
!
interface Loopback1
 description ANYCAST-RP LOOP
 no ip address
 ipv6 address 2001:DB8:CAFE:998::1/128
 ipv6 eigrp 1
!
ipv6 pim rp-address 2001:DB8:CAFE:998::1
ipv6 pim anycast-rp 2001:DB8:CAFE:998::1 2001:DB8:CAFE:999::4
ipv6 pim anycast-rp 2001:DB8:CAFE:998::1 2001:DB8:CAFE:999::3
```

```
RP-1#sh ipv6 pim anycast-rp
Anycast RP Peers For 2001:DB8:CAFE:998::1   Last Register/Register-Stop received
  2001:DB8:CAFE:999::4                      00:01:13/00:00:17
  2001:DB8:CAFE:999::3                      00:39:17/00:39:17
```

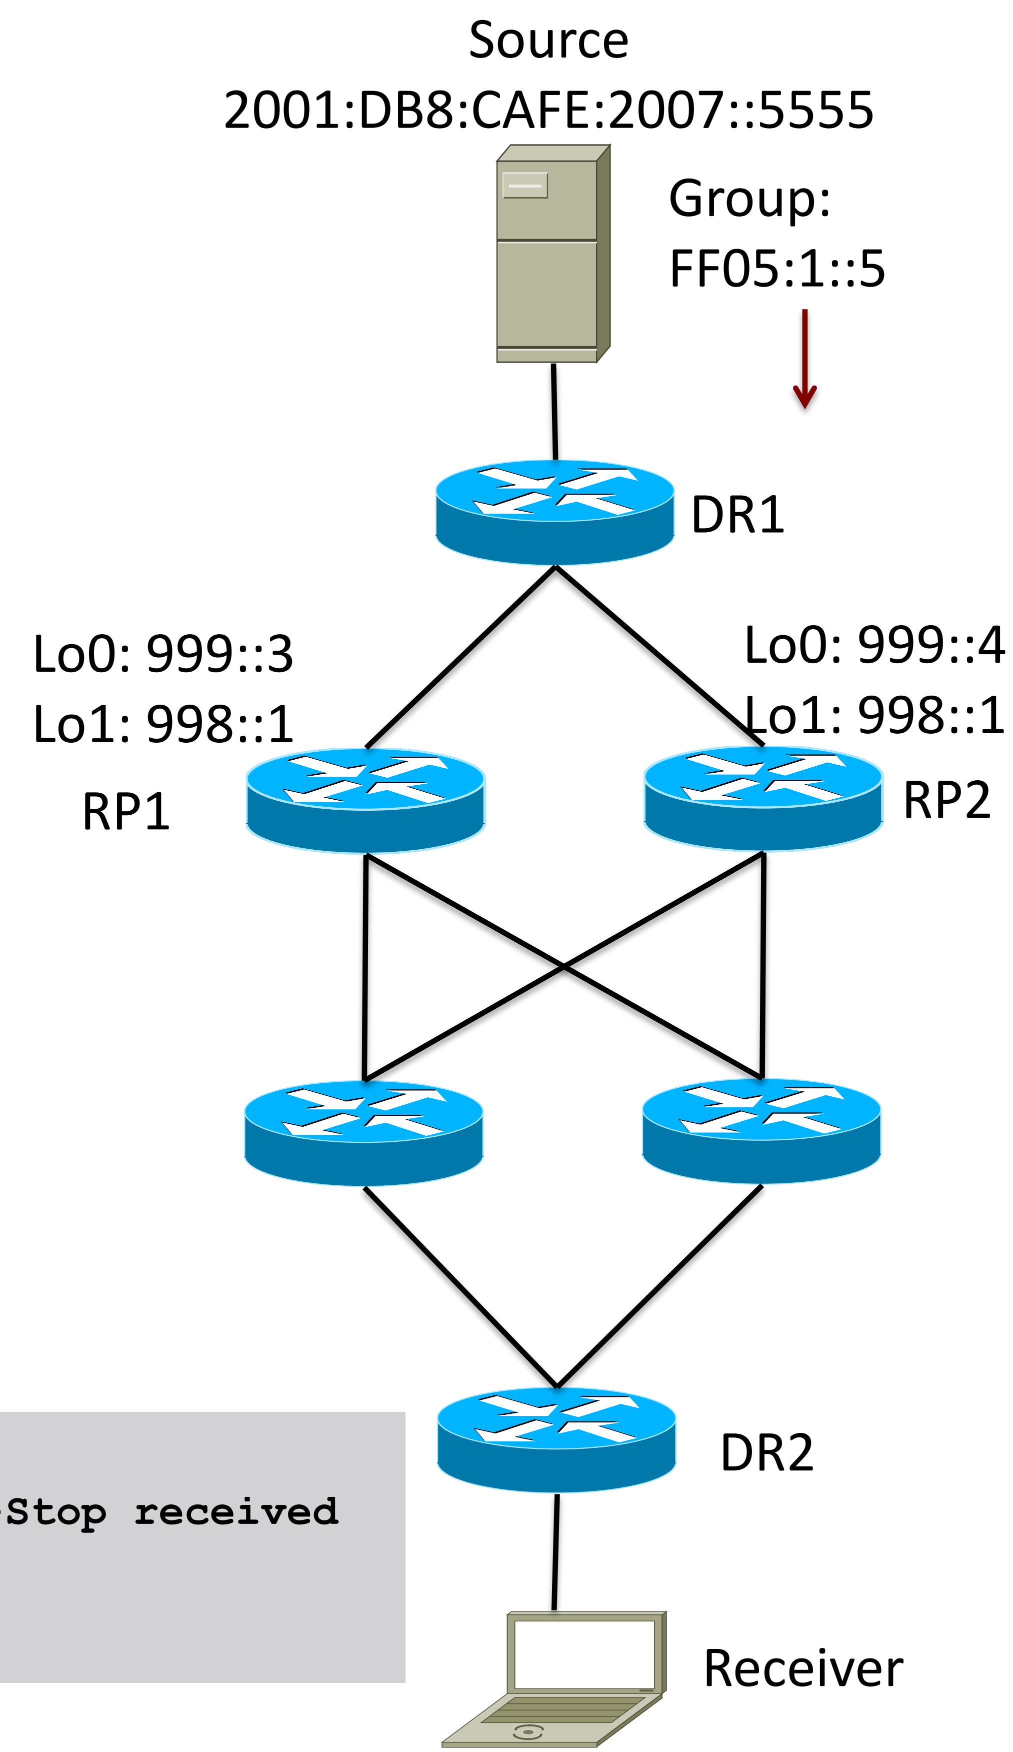

RP2

```
ipv6 multicast-routing
!
interface Loopback0
 description ANYCAST-RP
 no ip address
 ipv6 address 2001:DB8:CAFE:999::4/128
 ipv6 eigrp 1
!
interface Loopback1
 description ANYCAST-RP LOOP
 no ip address
 ipv6 address 2001:DB8:CAFE:998::1/128
 ipv6 eigrp 1
!
ipv6 pim rp-address 2001:DB8:CAFE:998::1
ipv6 pim anycast-rp 2001:DB8:CAFE:998::1 2001:DB8:CAFE:999::4
ipv6 pim anycast-rp 2001:DB8:CAFE:998::1 2001:DB8:CAFE:999::3
```

Only change is Lo0

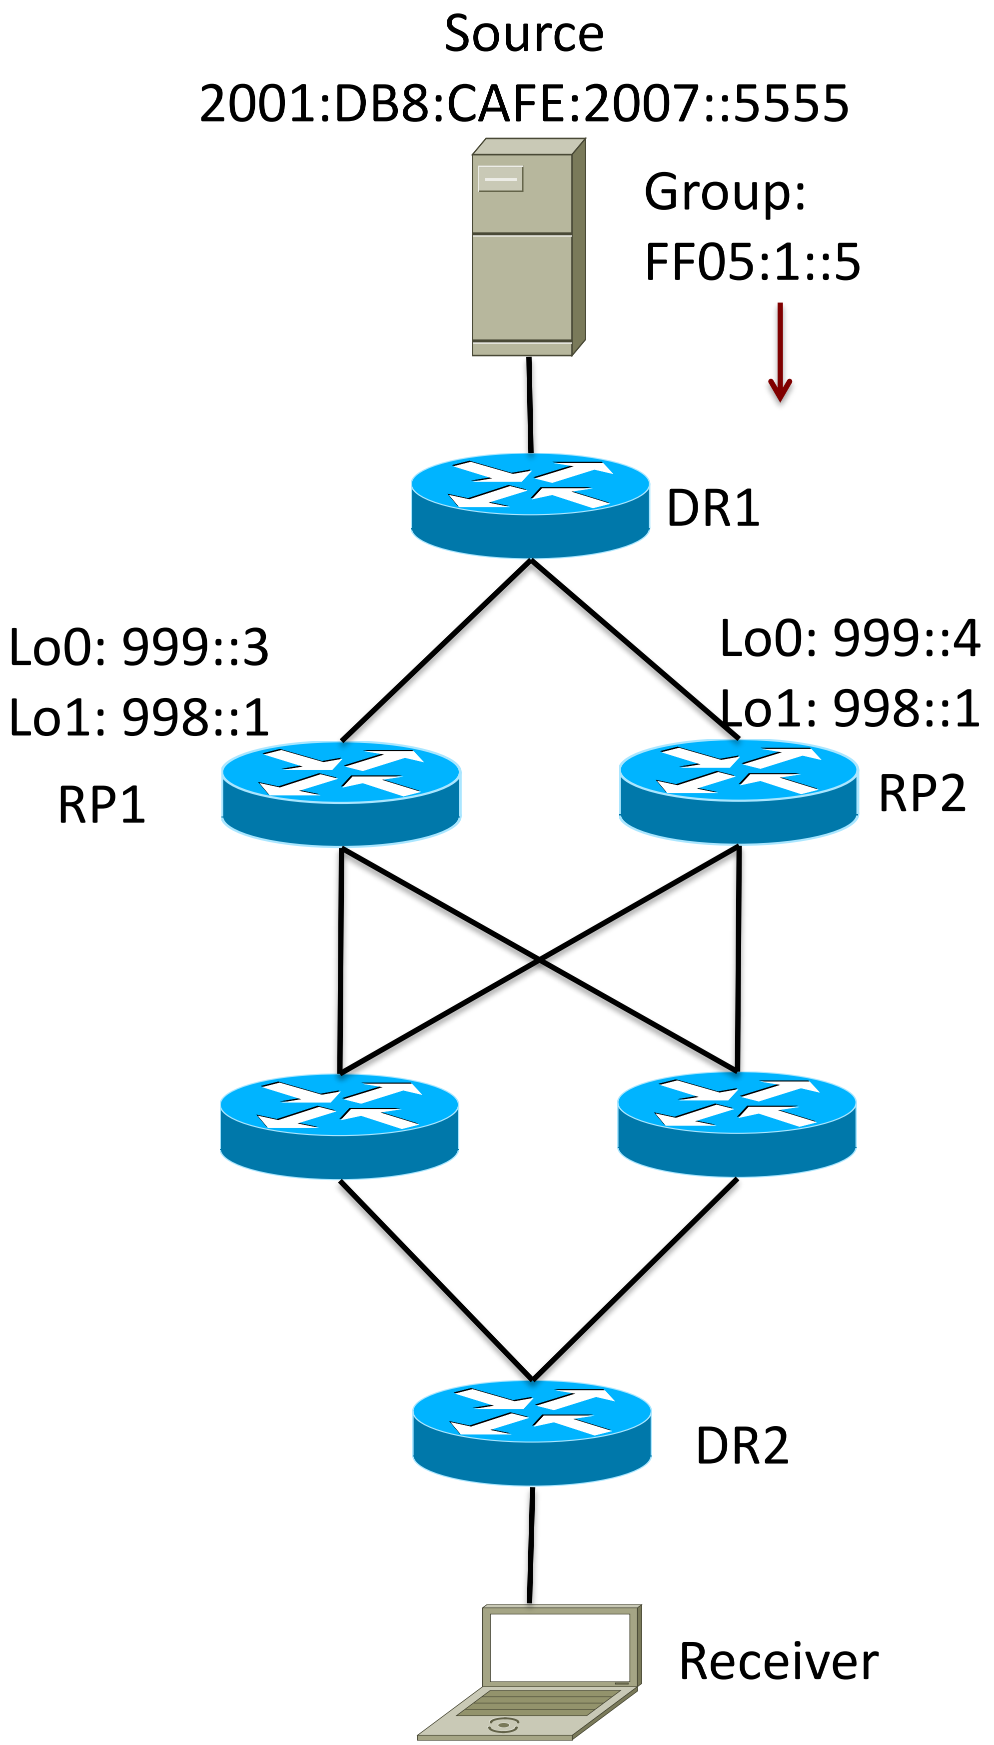

DR1

DR1

```
ipv6 multicast-routing
!
ipv6 pim rp-address 2001:DB8:CAFE:998::1
```

DR2

```
ipv6 multicast-routing
!
ipv6 pim rp-address 2001:DB8:CAFE:998::1
```

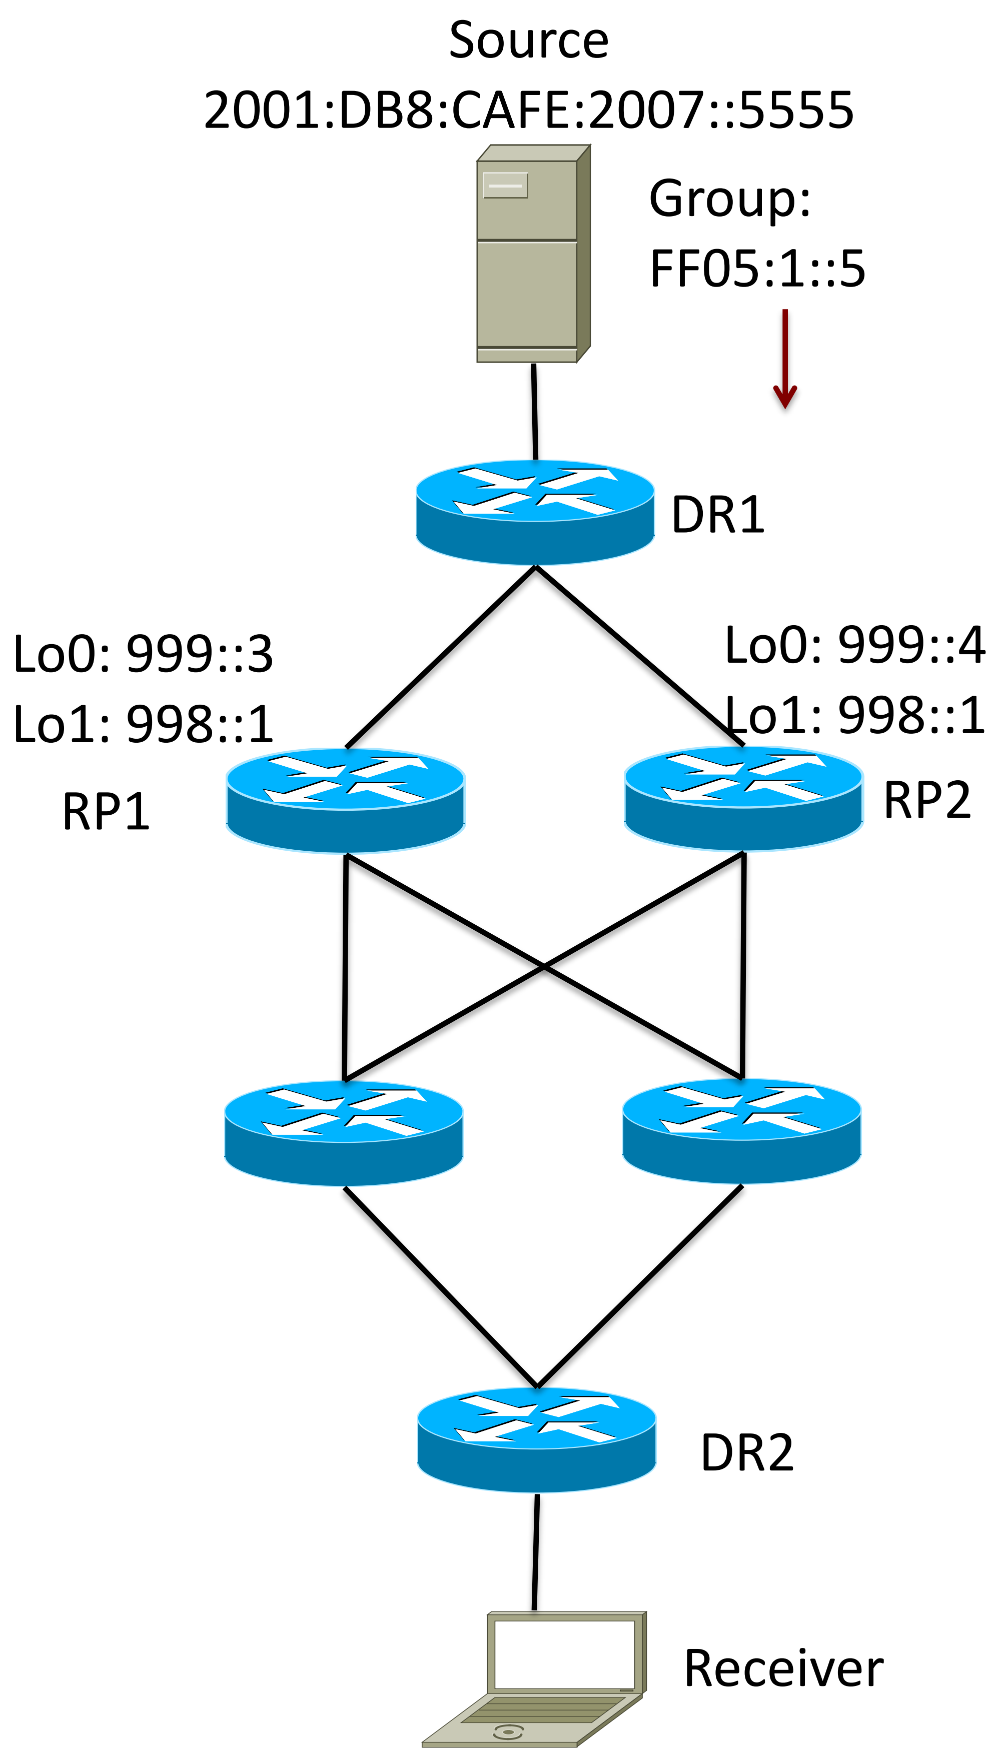

Summarised Debug – Source/Receiver

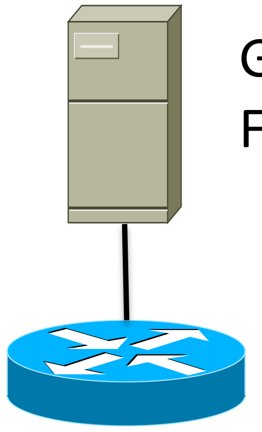

DR1 – Debug summarised (Source active)

```
(2001:DB8:CAFE:2007::5555,FF05:1::5) Create entry
. . . Output Summarized
(2001:DB8:CAFE:2007::5555,FF05:1::5) Start registering to 2001:DB8:CAFE:998::1
```

DR2 – Debug summarised (Receiver active)

```
(2001:DB8:CAFE:2007::5555,FF05:1::5/128) MRIB update (t=1)
. . . Output Summarized
(2001:DB8:CAFE:2007::5555,FF05:1::5) Ethernet0/0 Local state changed from Null to Join
(2001:DB8:CAFE:2007::5555,FF05:1::5) Ethernet0/0 FWD state change from Prune to Forward
```

RP1 – Debug summarised (Source active)

```
Received Register from 2001:DB8:CAFE:2006::FACE
Send Register to Anycast-RP peer 2001:DB8:CAFE:999::4 from 2001:DB8:CAFE:999::3 length 48
. . . Output Summarized
Send Register-Stop to 2001:DB8:CAFE:2006::FACE
. . .
Send Register-Stop to 2001:DB8:CAFE:999::4
Received J/P on Ethernet0/0 from FE80::A8BB:CCFF:FE00:2110 target: FE80::A8BB:CCFF:FE00:2300 (to us)
J/P entry: Join root: 2001:DB8:CAFE:2007::5555 group: FF05:1::5 flags:  S
```

DR2 – Active Receiver

```
DR-2#show ipv6 mroute active
Active Multicast Sources - sending >= 4 kbps
Group: FF05:1::5
  Source: 2001:DB8:CAFE:2007::5555,
   SW Rate: 232 pps/87 kbps(1sec), 90 kbps(last 21 sec)
```

IPv6 QoS Policy & Syntax

- Unified QoS Policy (v4/v6 in same policy) or separate?
- IPv4 syntax has used “ip” following match/set statements
  - **Example:** **`match ip dscp, set ip dscp`**
- Modification in QoS syntax to support IPv6 and IPv4
  - **`New match criteria`**

    `match dscp — Match DSCP in v4/v6`

    `match precedence — Match Precedence in v4/v6`
  - **`New set criteria`**

    `set dscp — Set DSCP in v4/v6`

    `set precedence — Set Precedence in v4/v6`
- Additional support for IPv6 does not always require new Command Line Interface (CLI)
  - **Example—WRED**

Infrastructure Deployment

- IPv6 Knowledge Base Portal
- http://www.cisco.com/web/solutions/netsys/ipv6/knowledgebase/index.html

IPv6 Co-existence Solutions

Dual Stack

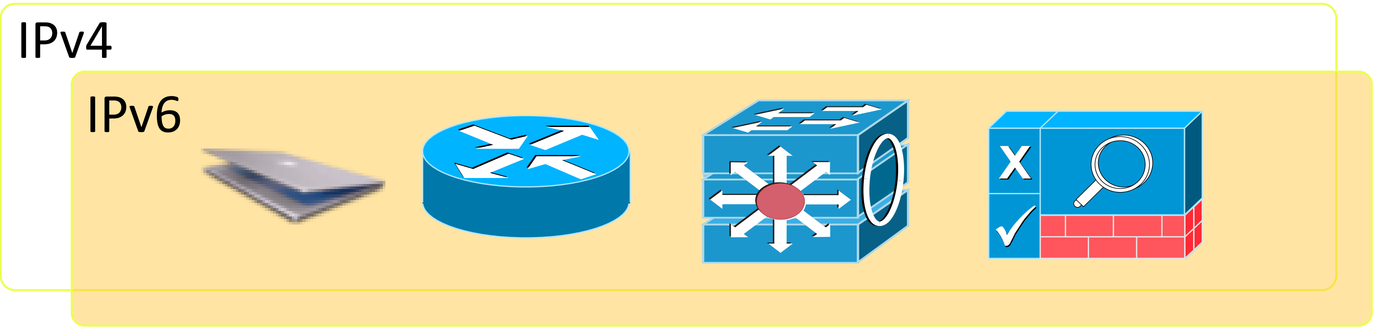

Recommended Enterprise Co-existence strategy

Tunnelling Services

Connect Islands of IPv6 or IPv4

Translation Services

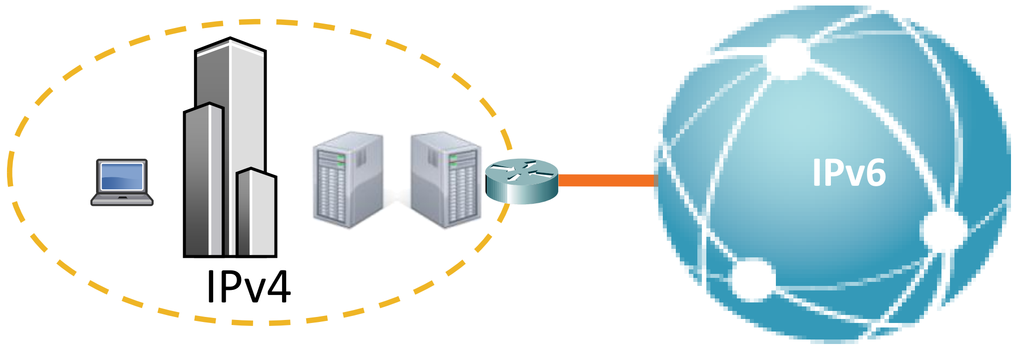

An interim approach to bridging the gap

Deploying IPv6 in Campus Networks

- Deploying IPv6 in Campus Networks: http://www.cisco.com/univercd/cc/td/doc/solution/campipv6.pdf

Campus IPv6 Deployment Options

Dual-Stack IPv4/IPv6

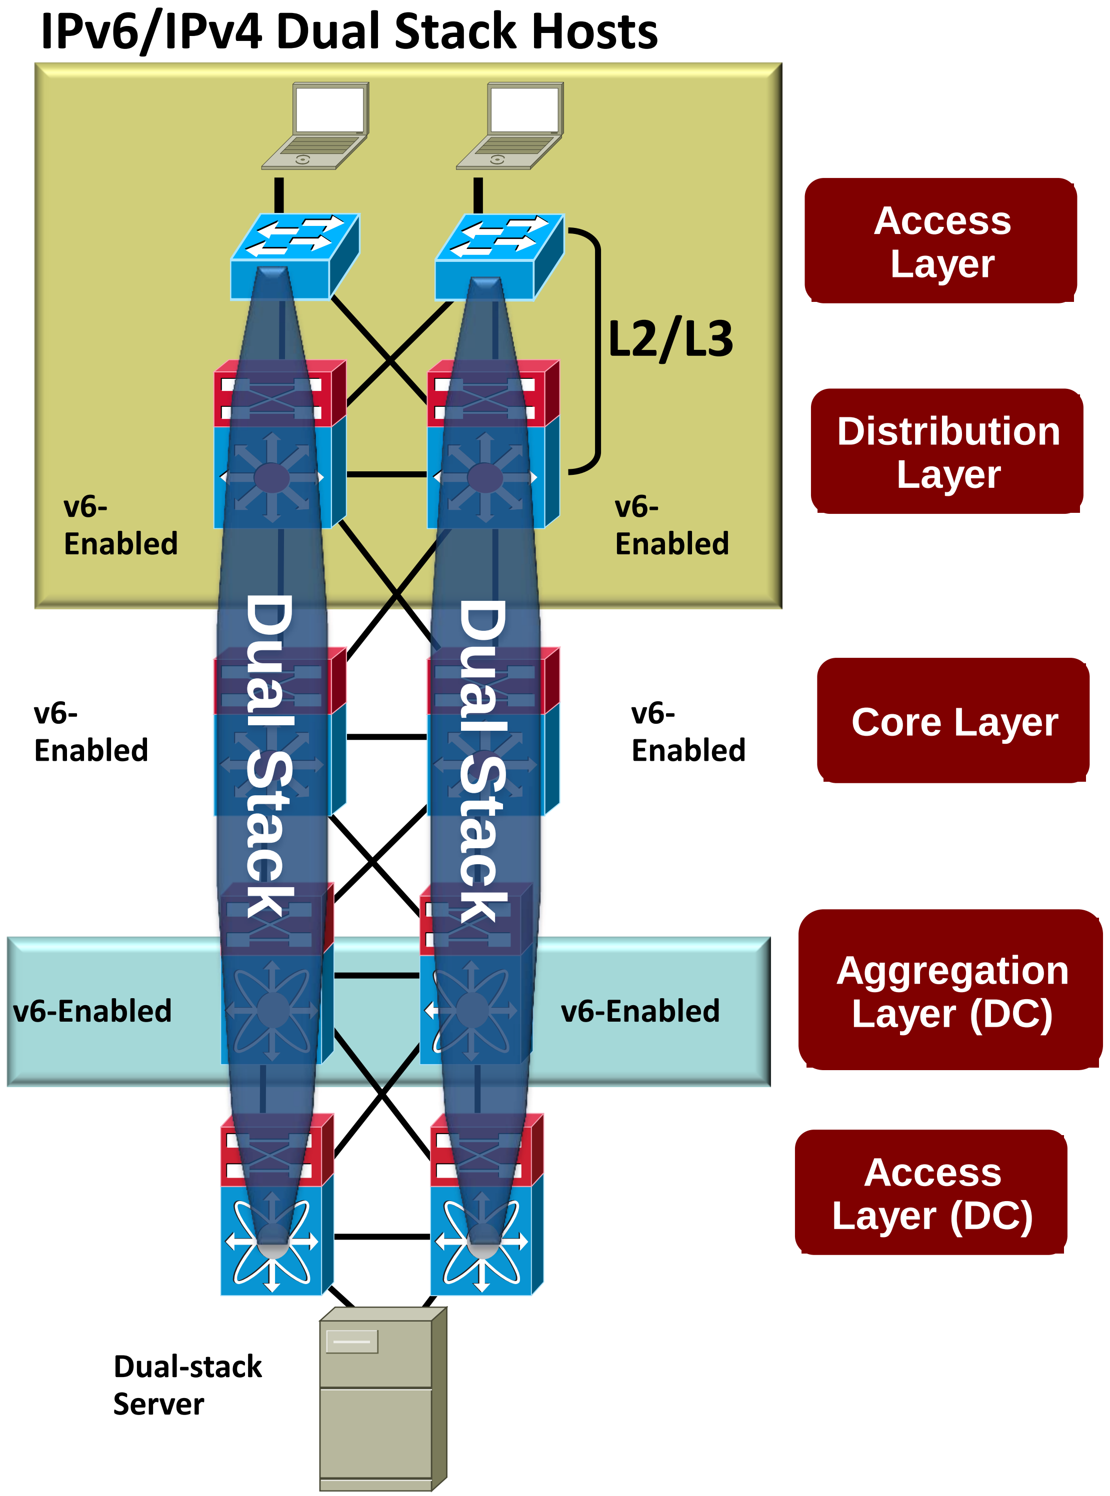

- Dual Stack = Two protocols running at the same time (IPv4/IPv6)
- #1 requirement—switching/ routing platforms must support hardware based forwarding for IPv6
  - 3560/3750* +
  - 4500 Sup6E +
  - 6500 Sup32/720 +
- IPv6 is transparent on L2 switches but consider:
  - L2 multicast—MLD snooping
  - IPv6 management—Telnet/SSH/HTTP/SNMP
  - Intelligent IP services on WLAN – CHECK OUT 7.2 CODE
- Layer 2 Security does change

*check HW limitations in non-E/X/C series

Distribution Layer: HSRP, EIGRP and DHCPv6-Relay (Layer 2 Access)

```
ipv6 unicast-routing
!
interface GigabitEthernet1/0/1
 description To 6k-core-right
 ipv6 address 2001:DB8:CAFE:1105::A001:1010/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
interface GigabitEthernet1/0/2
 description To 6k-core-left
 ipv6 address 2001:DB8:CAFE:1106::A001:1010/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
```

```
interface Vlan4
 description Data VLAN for Access
 ipv6 address 2001:DB8:CAFE:4::2/64
 ipv6 nd managed-config-flag
 ipv6 nd prefix 2001:DB8:CAFE:4::/64 0 0 no-autoconfig
 ipv6 dhcp relay destination 2001:DB8:CAFE:10::2
 ipv6 eigrp 10
 standby version 2
 standby 2 ipv6 autoconfig
 standby 2 timers msec 250 msec 750
 standby 2 priority 110
 standby 2 preempt delay minimum 180
 standby 2 authentication ese
!
ipv6 router eigrp 10
 no shutdown
 router-id 10.122.10.10
 passive-interface Vlan4
 passive-interface Loopback0
```

First Hop Security

```
ipv6 access-list HOST_PACL
 remark Deny Rogue DHCP
 deny udp any eq 547 any eq 546
 remark Deny RA From Client
 deny icmp any any router-advertisement
 permit ipv6 any any
!
interface GigabitEthernet1/0/6
 ipv6 traffic-filter HOST_PACL in
```

```
interface GigabitEthernet1/0/6
 ipv6 nd raguard
```

```
interface GigabitEthernet1/0/6
 ipv6 nd router-preference High
```

- L2/L3 Security
- Port ACL (PACL), RA Guard, SEND, etc…
- RA Preference “High”
- Much more to know

Attend: BRKSEC-2003

http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-first_hop_security.html

Wireless IPv6 Support - Pre-v7.2

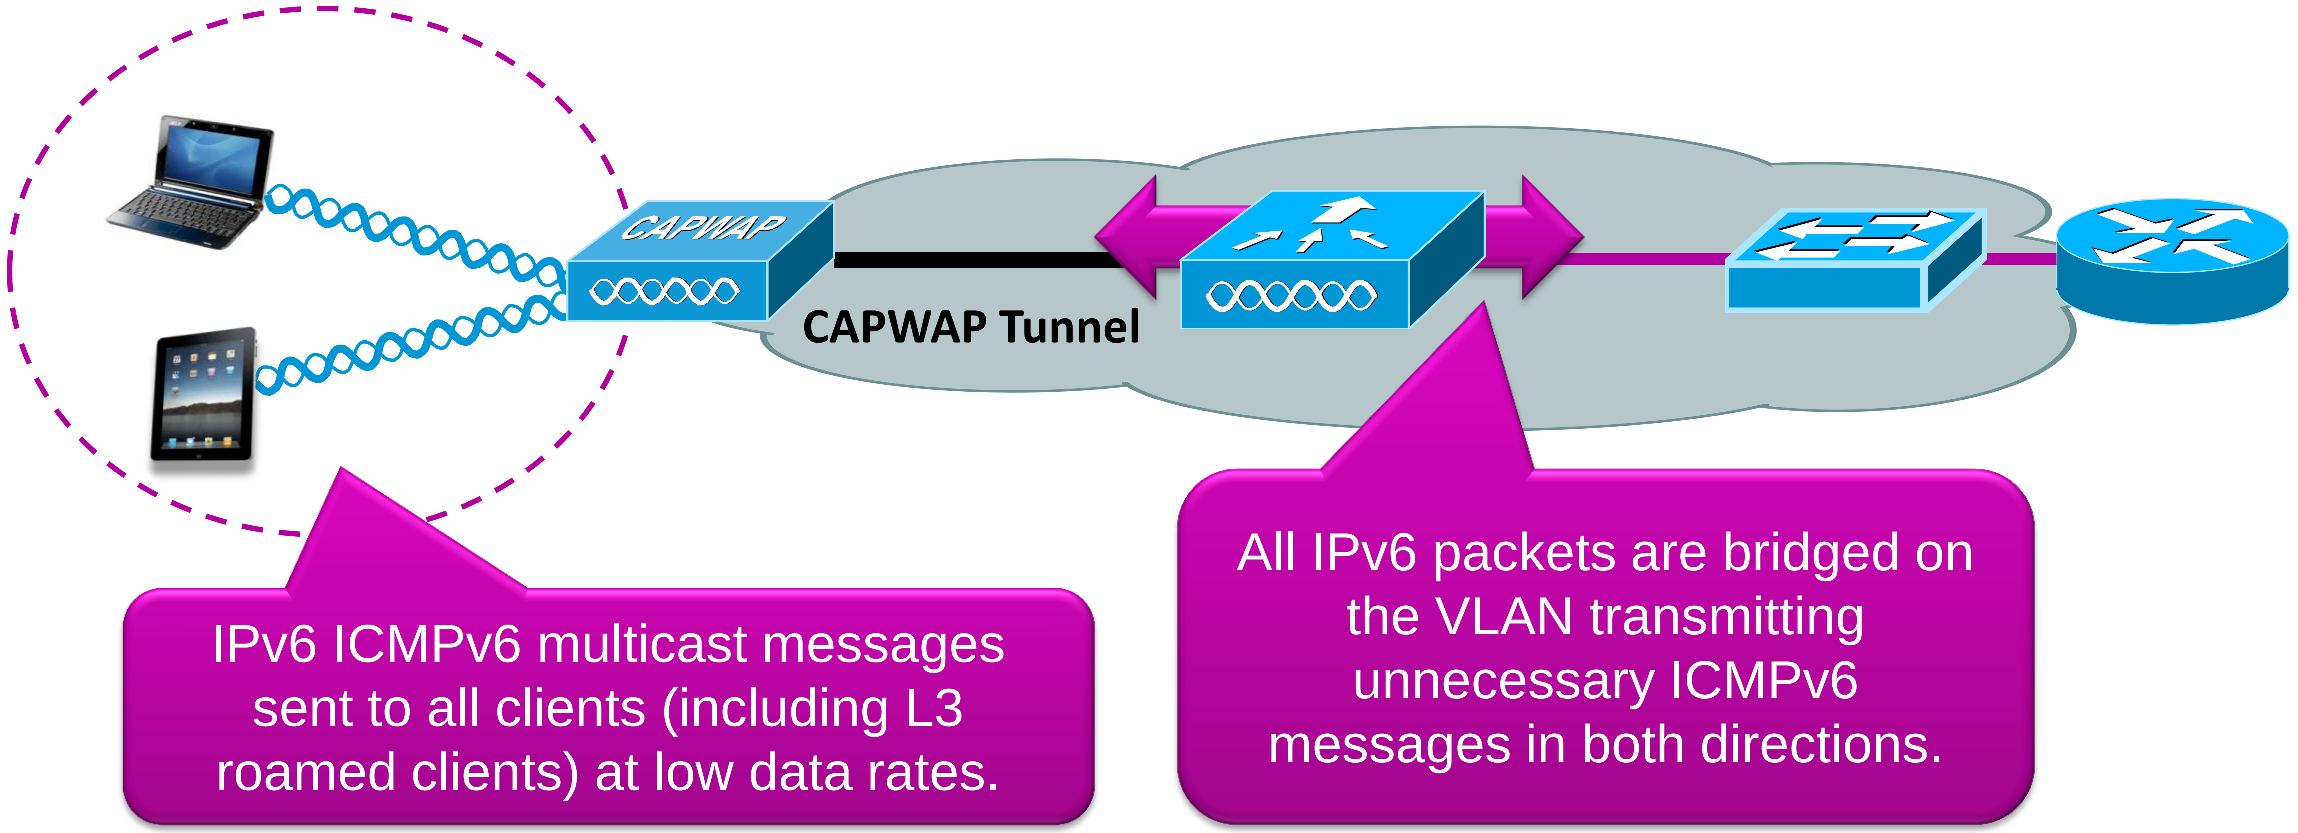

- In releases prior to 7.2, enabling IPv6 bridging provided a limited solution with no Layer 3 mobility and non-optimised delivery of essential ICMPv6 messages to clients.

Wireless IPv6 Support - Post-v7.2

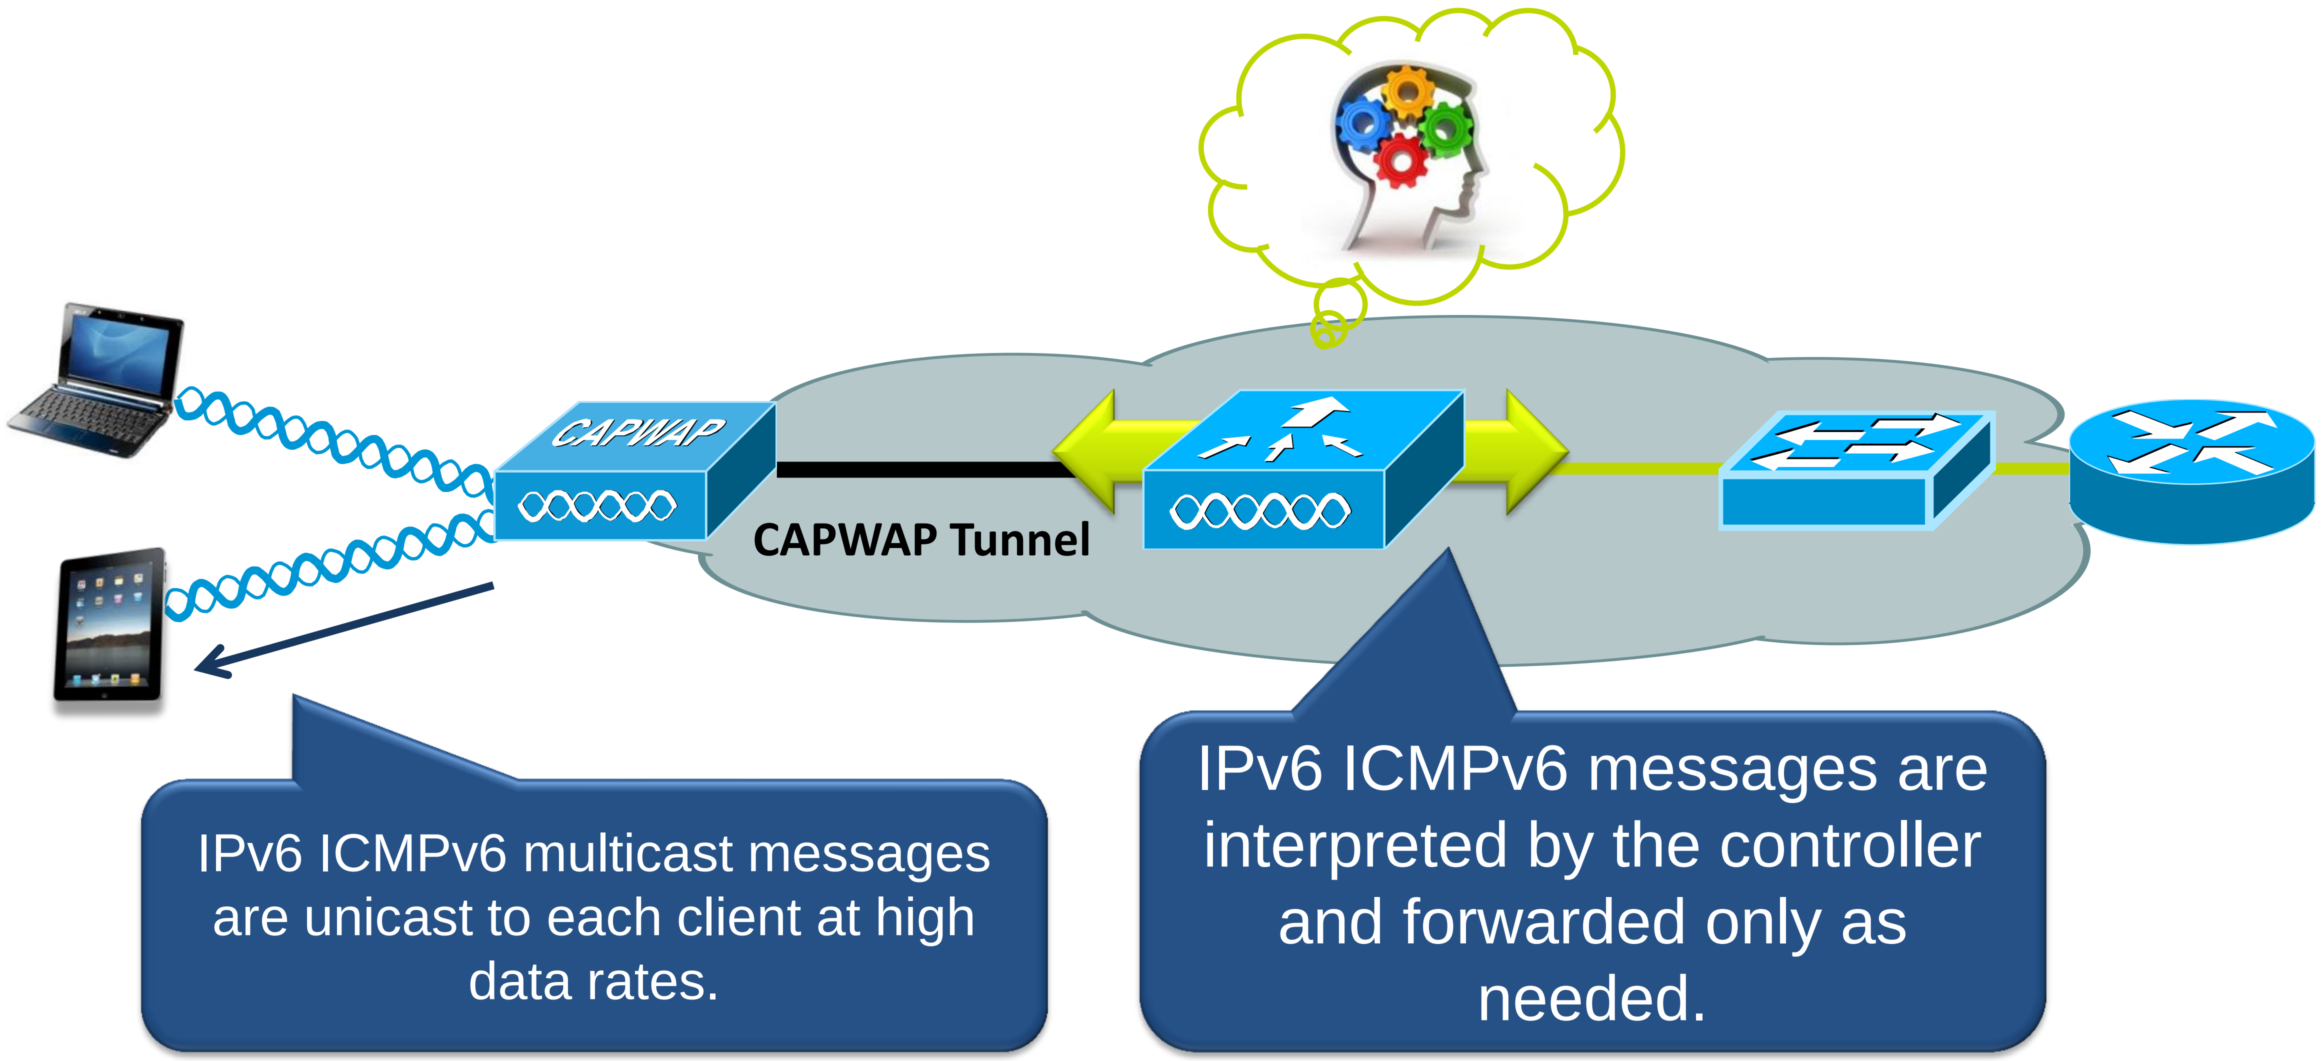

- In releases 7.2, the controller now processes ICMPv6 messages allowing for optimised delivery, Layer 3 mobility and first hop security.

Wireless IPv6 Client Support

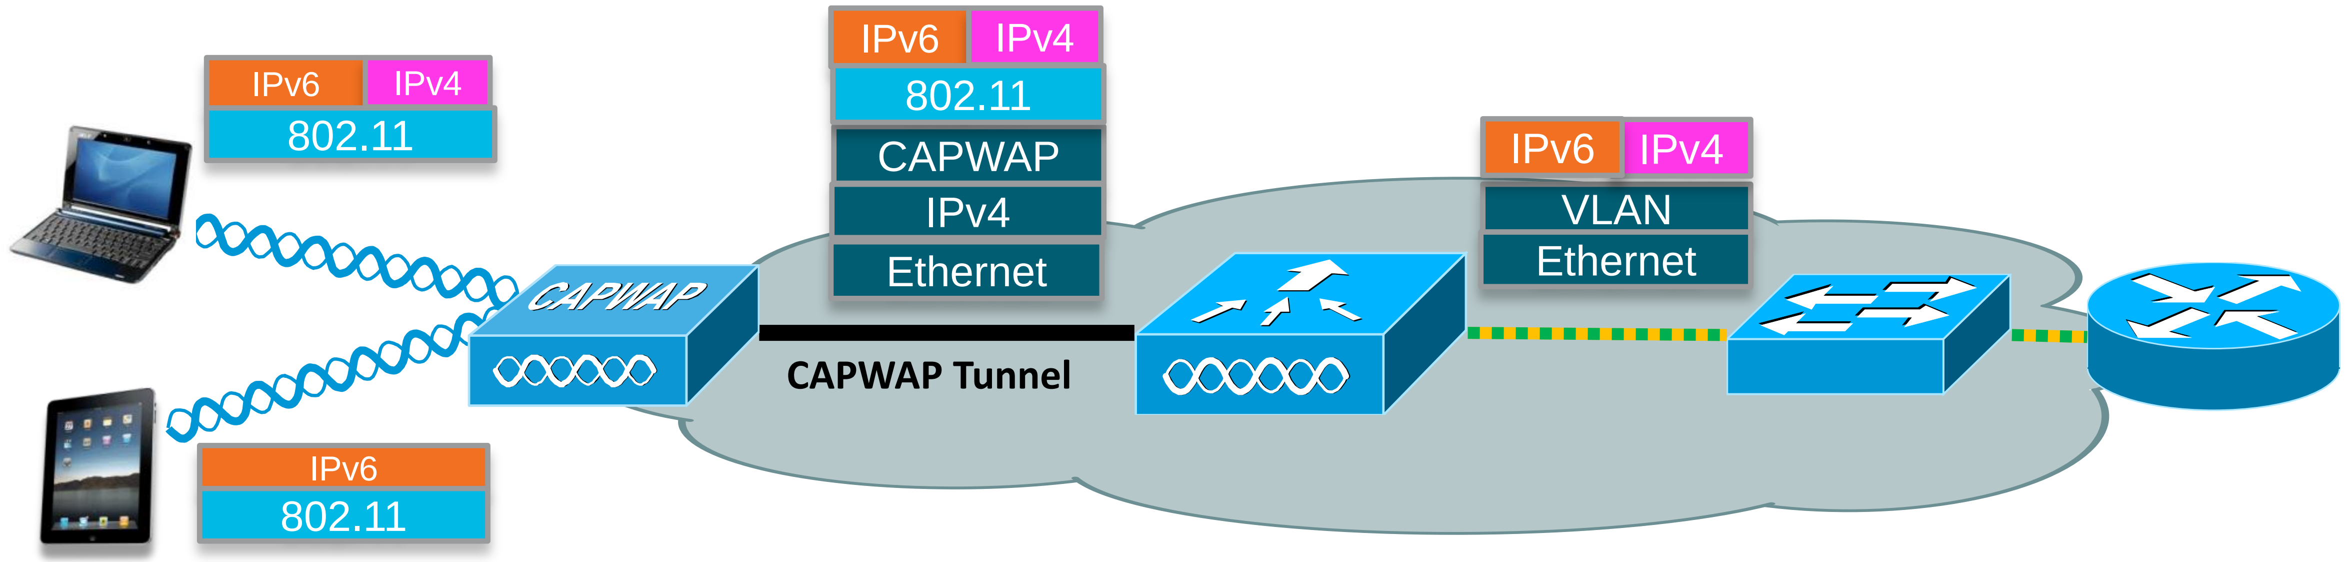

- Supports IPv4, Dual Stack and Native IPv6 clients on single WLAN simultaneously
- Supports the following IPv6 address assignment for wireless clients:
  – IPv6 Stateless Autoconfiguration [SLAAC]
  – Stateless, Stateful DHCPv6
  – Static IPv6 configuration
- Supports up to 8 IPv6 addresses per client
- Clients will be able to pass traffic once IPv4 or IPv6 address assignment is completed after successful authentication

IPv6 Client Connectivity on Multiple WLANs

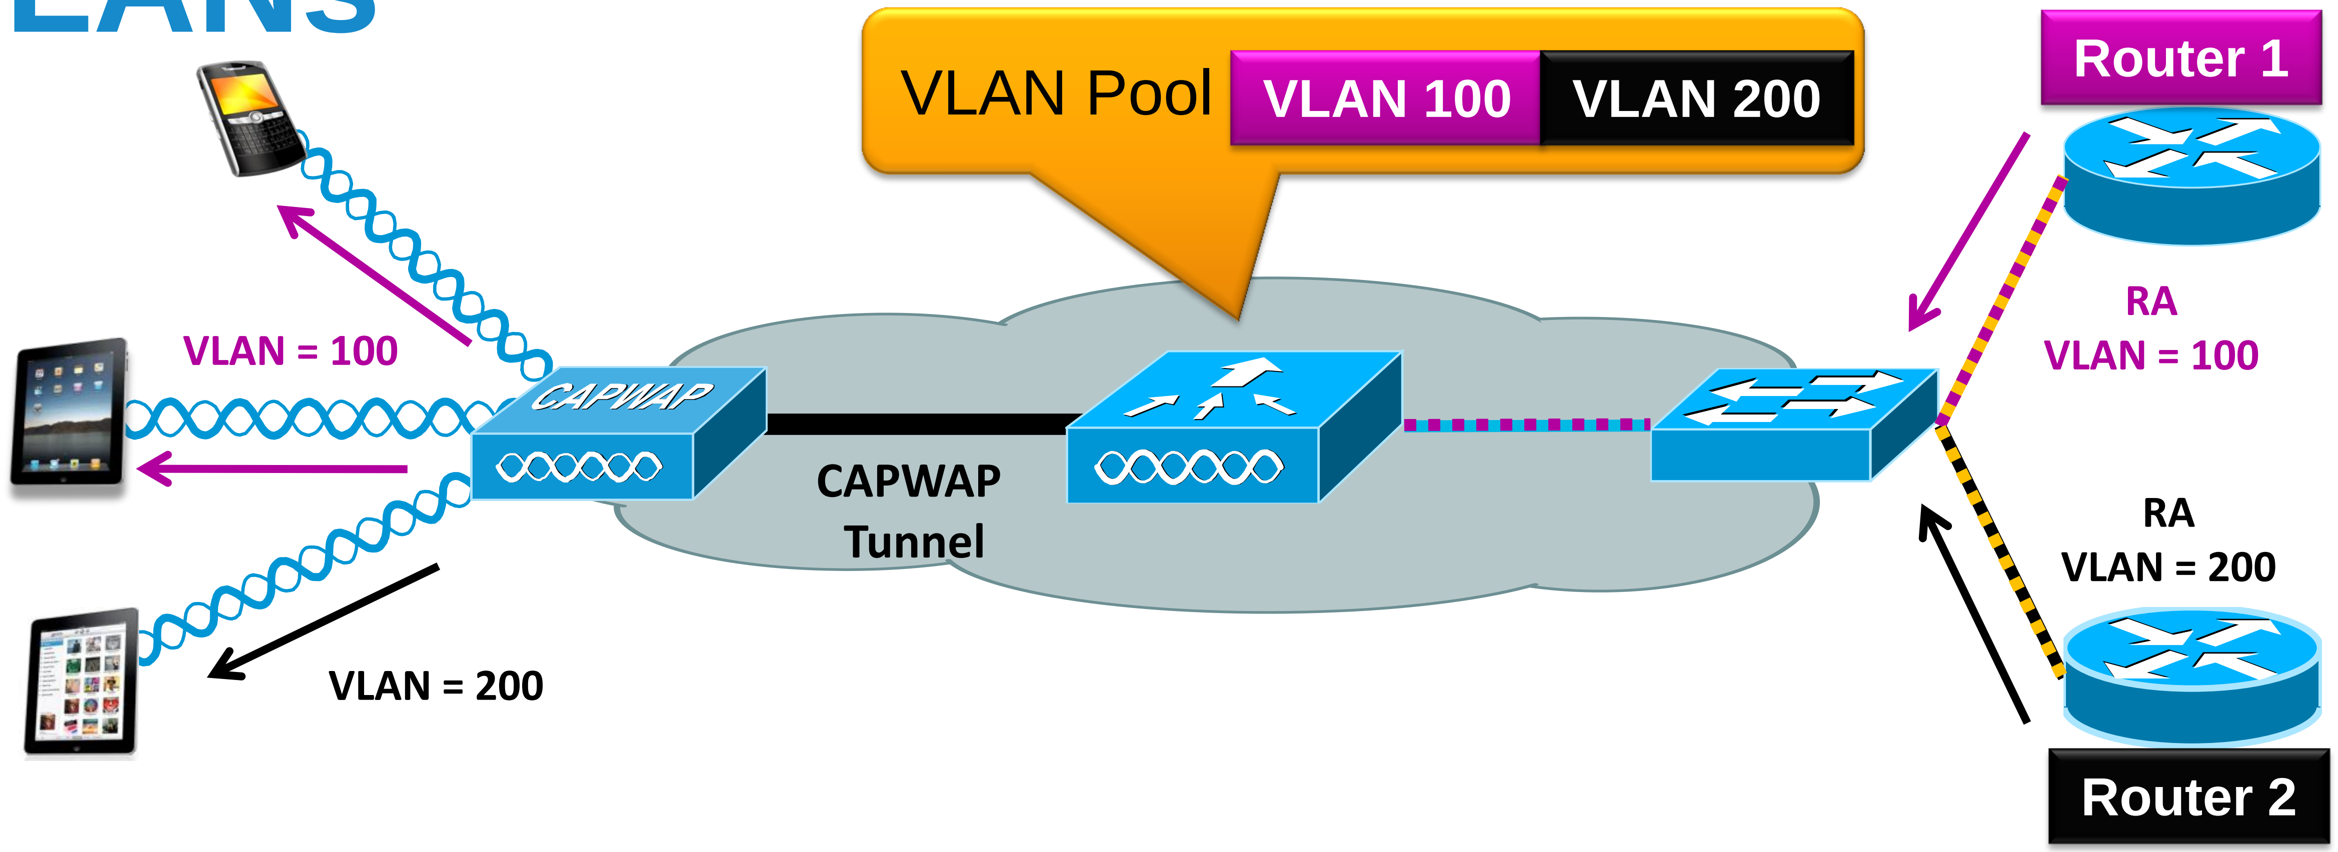

- Access Points keep track of individual clients and unicast the Router Advertisement to the clients depending on the WLAN they belong to
- Access Point support up to 16 WLANs/SSIDs for dual stack clients
- To maintain proper routing capability, mobile clients need to have proper global unique unicast prefix from router within their own network

Stateless Address Auto Configuration (SLAAC)

SLAAC – Stateless Address Auto Configuration

- Enabled on almost all IPv6 clients
  - Windows Vista/7, Windows 2008 Server, Apple iOS, Apple OSX, Linux
  - Installable on Windows XP, Windows 2003 Server
- Uses an advertised prefix of 64 bits and has two modes of generating the remaining 64-bits
  - Privacy mode – Generated using random host address (64-bits) appended to the prefix
  - EUI-64 mode – Generated using the MAC address of the network adapter

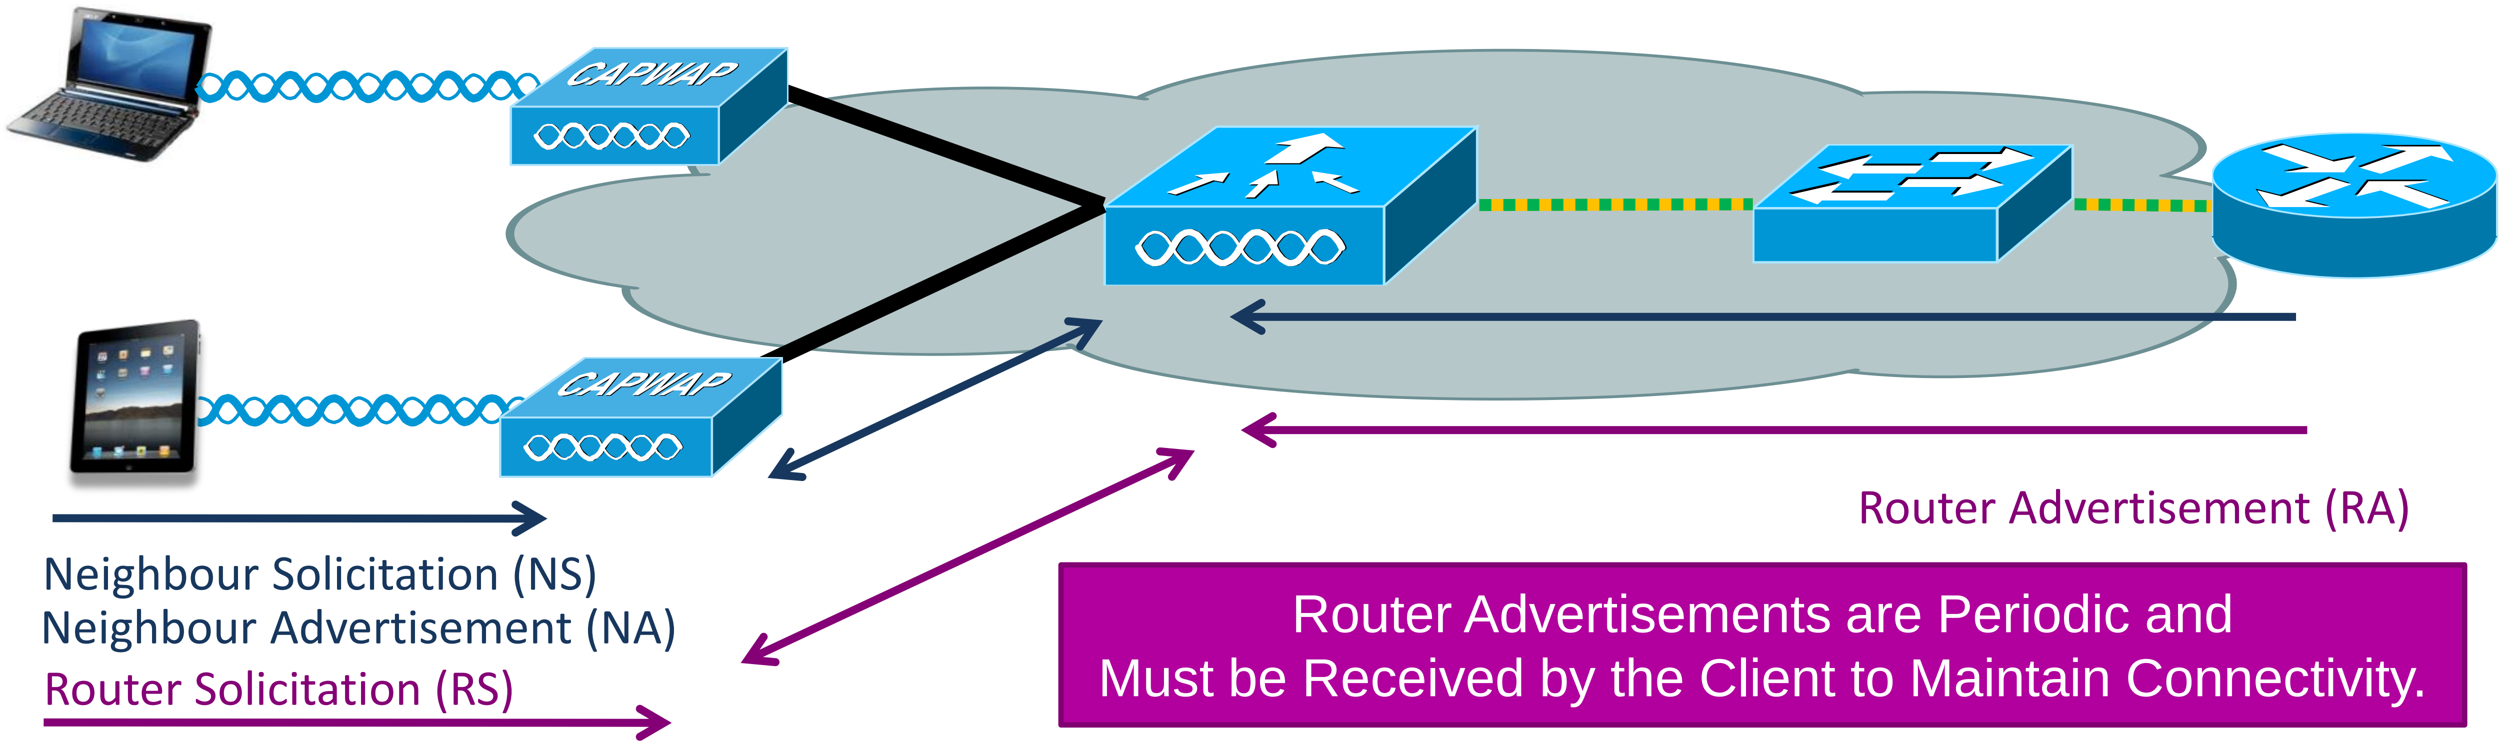

DHCPv6 Address Assignment

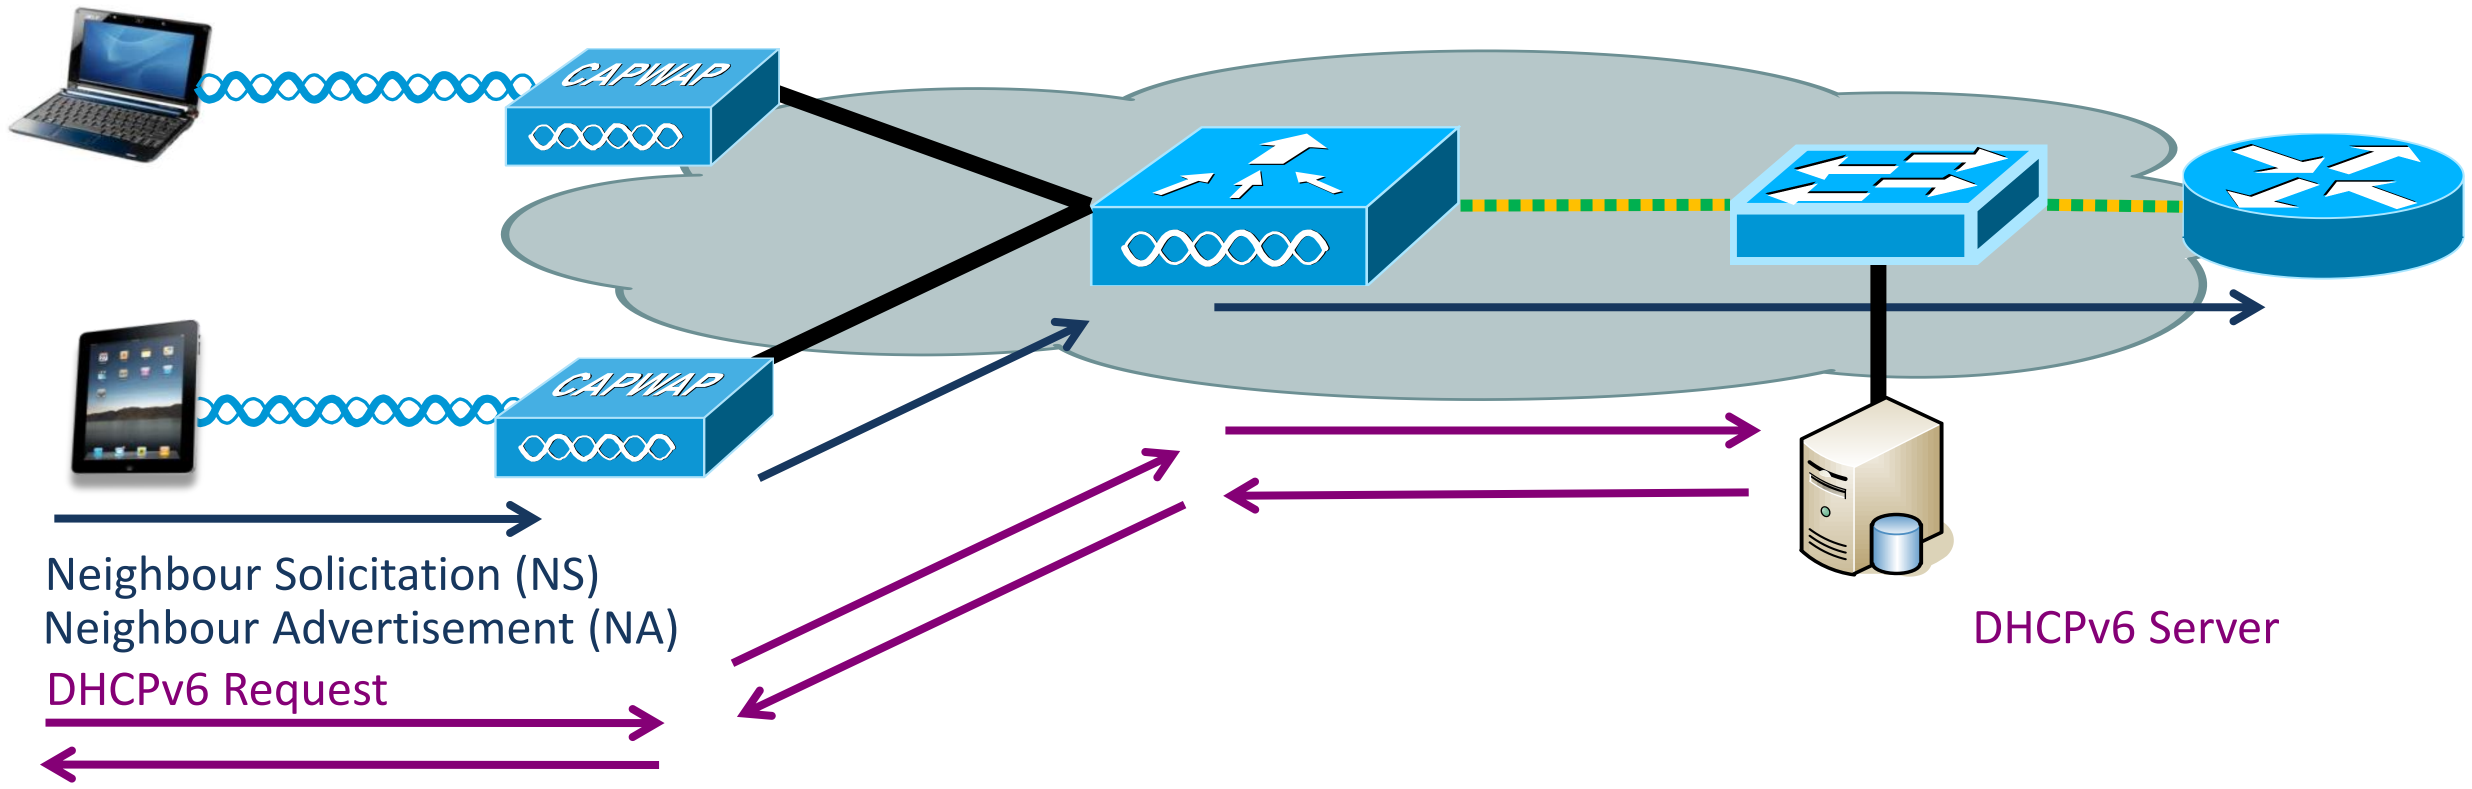

- DHCPv6 – A "managed" mode of IPv6 address assignment
  – Not present, or enabled by default on most IPv6-capable clients, yet
- Enabled on Windows 7 (SP1), Apple OSX "Lion" 10.7, Apple iOS 4.3
- DHCPv6 can also used to provide DNS, Domain Name and other options when SLACC is used
  – This mode is called "Other Config" mode

Cisco Supports Many IPv6 Addresses Per Client

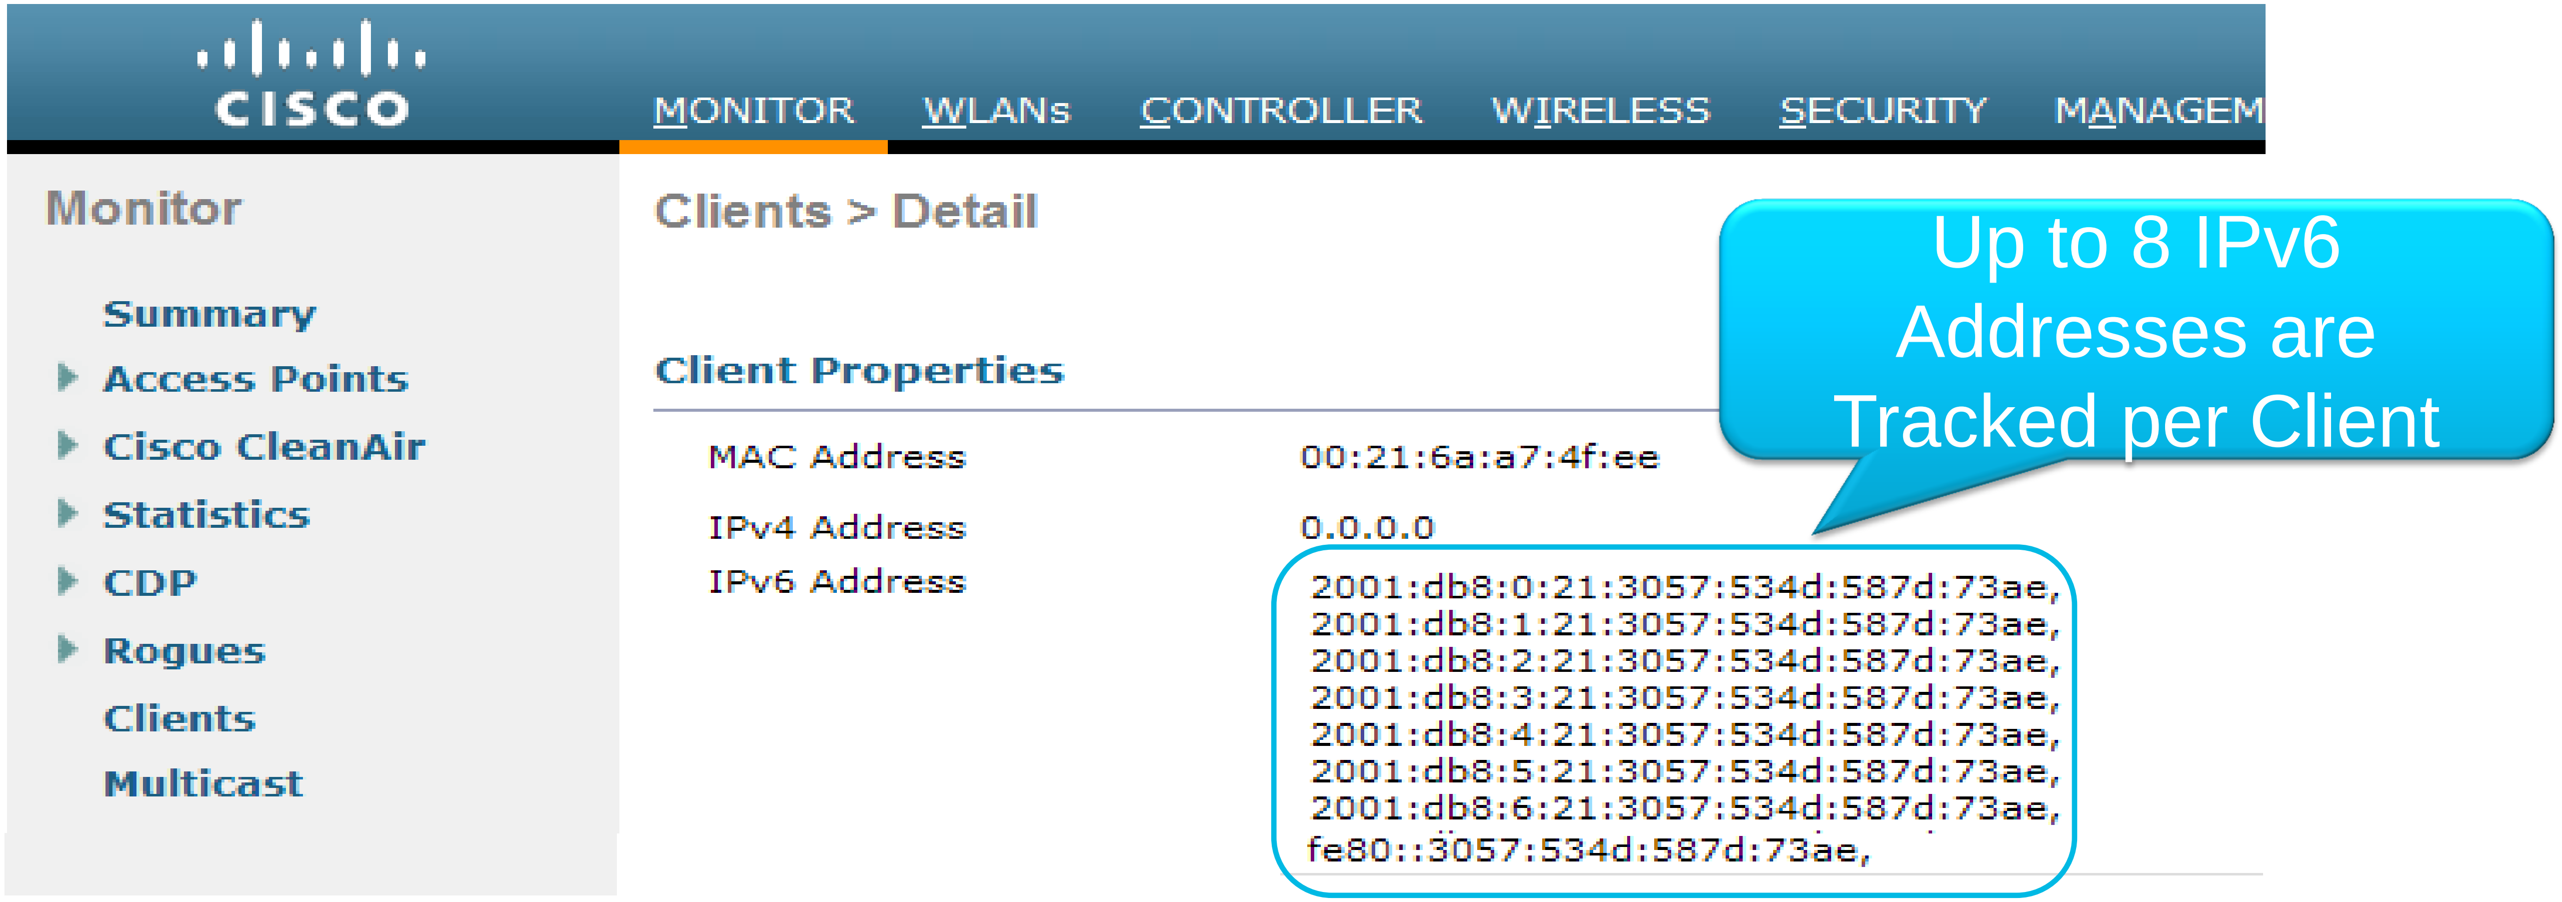

- Support for many IPv6 addresses per client is necessary because:
  - Clients can have multiple address types per interface
  - Clients can be assigned addresses via multiple methods such as SLAAC and DHCPv6
  - Most clients automatically generate a temporary address in addition to assigned addresses

Why is it Important to Support Many IPv6 Addresses?

Clients can have multiple IPv6 addresses per interface which can be static, SLAAC or DHCPv6 assigned

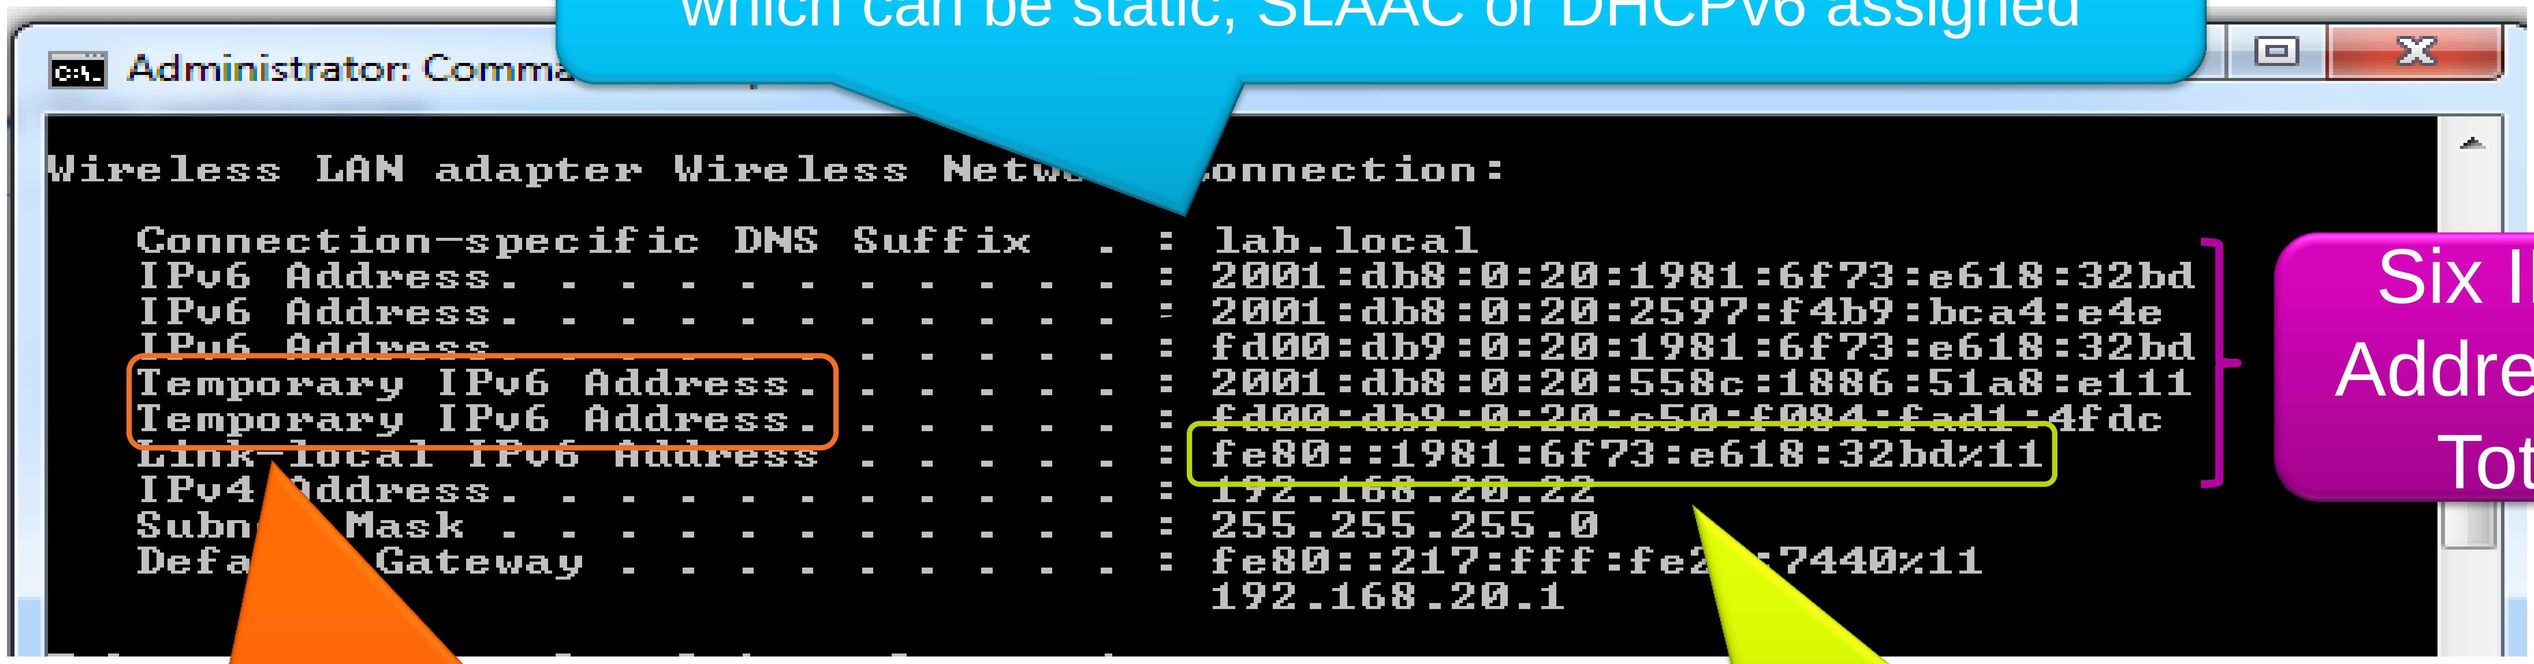

Six IPv6 Addresses Total

Most operating systems use a temporary address by default to mask their tracks on the global IPv6 Internet

All clients have a Link-Local address in addition to Global and/or Unique-Local addresses

Use Case #1: Mobility

Challenge	Dropped connections when roaming to a different network
Solution	Intelligent IPv6 packets processing RA follows roaming clients through mobility tunnel
Benefits	Reliable connectivity while roaming
Differentiator	Seamless layer-3 mobility for IPv6 clients

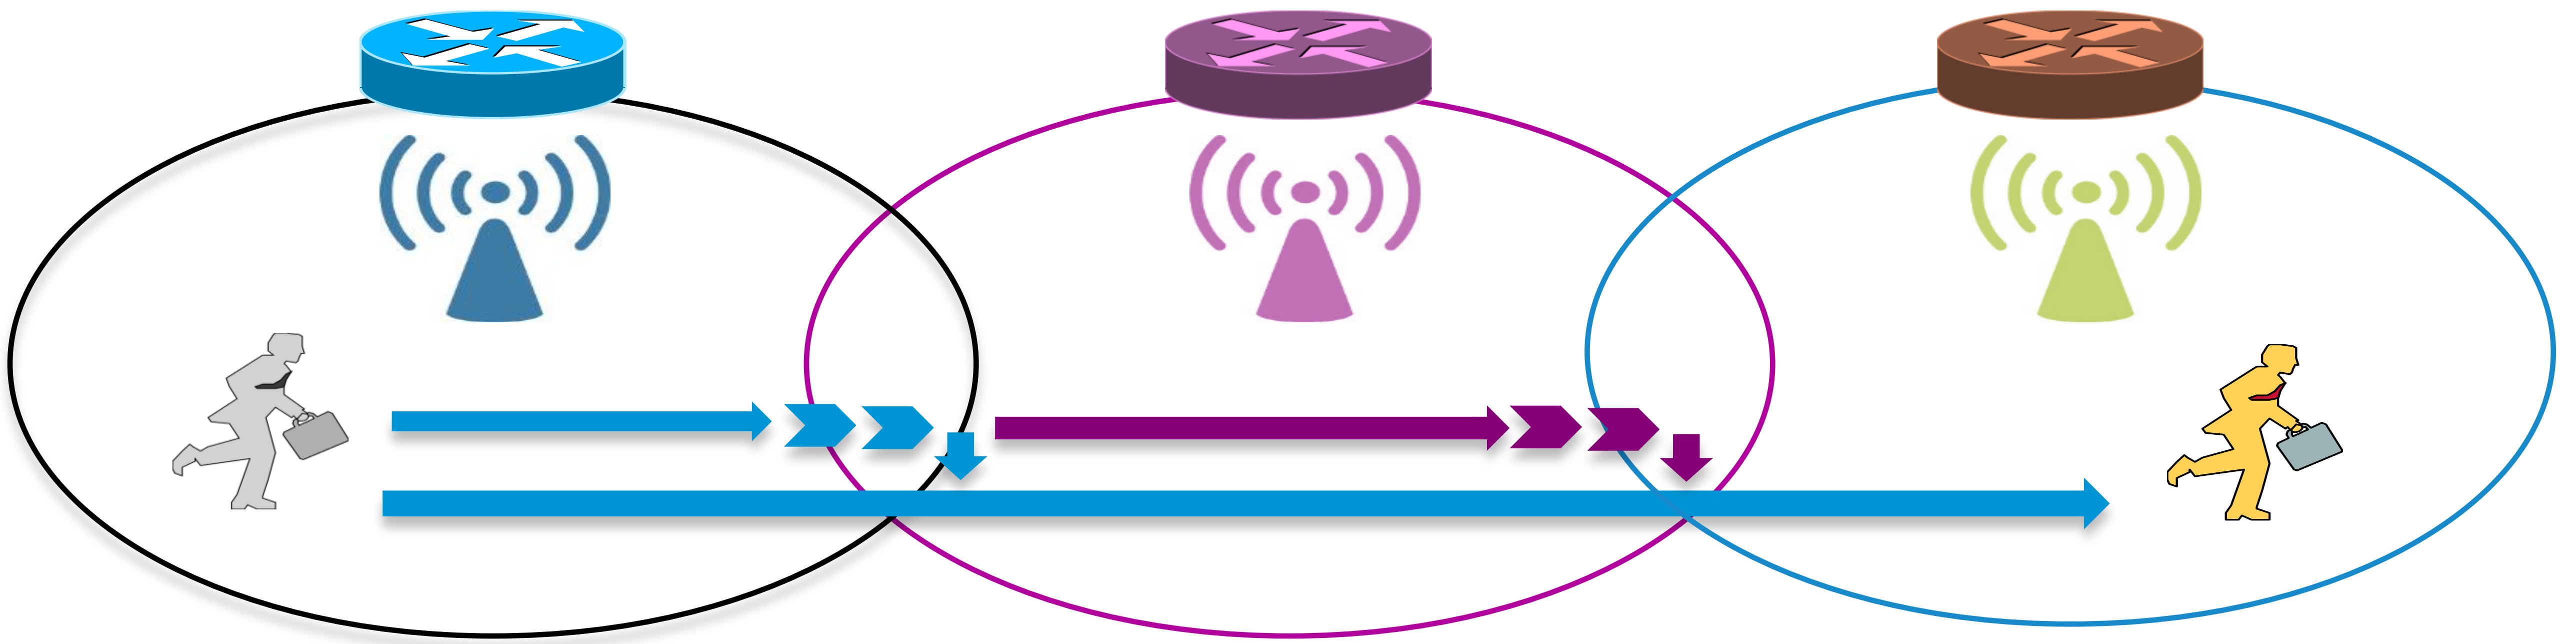

How Does Cisco Solve IPv6 Mobility?

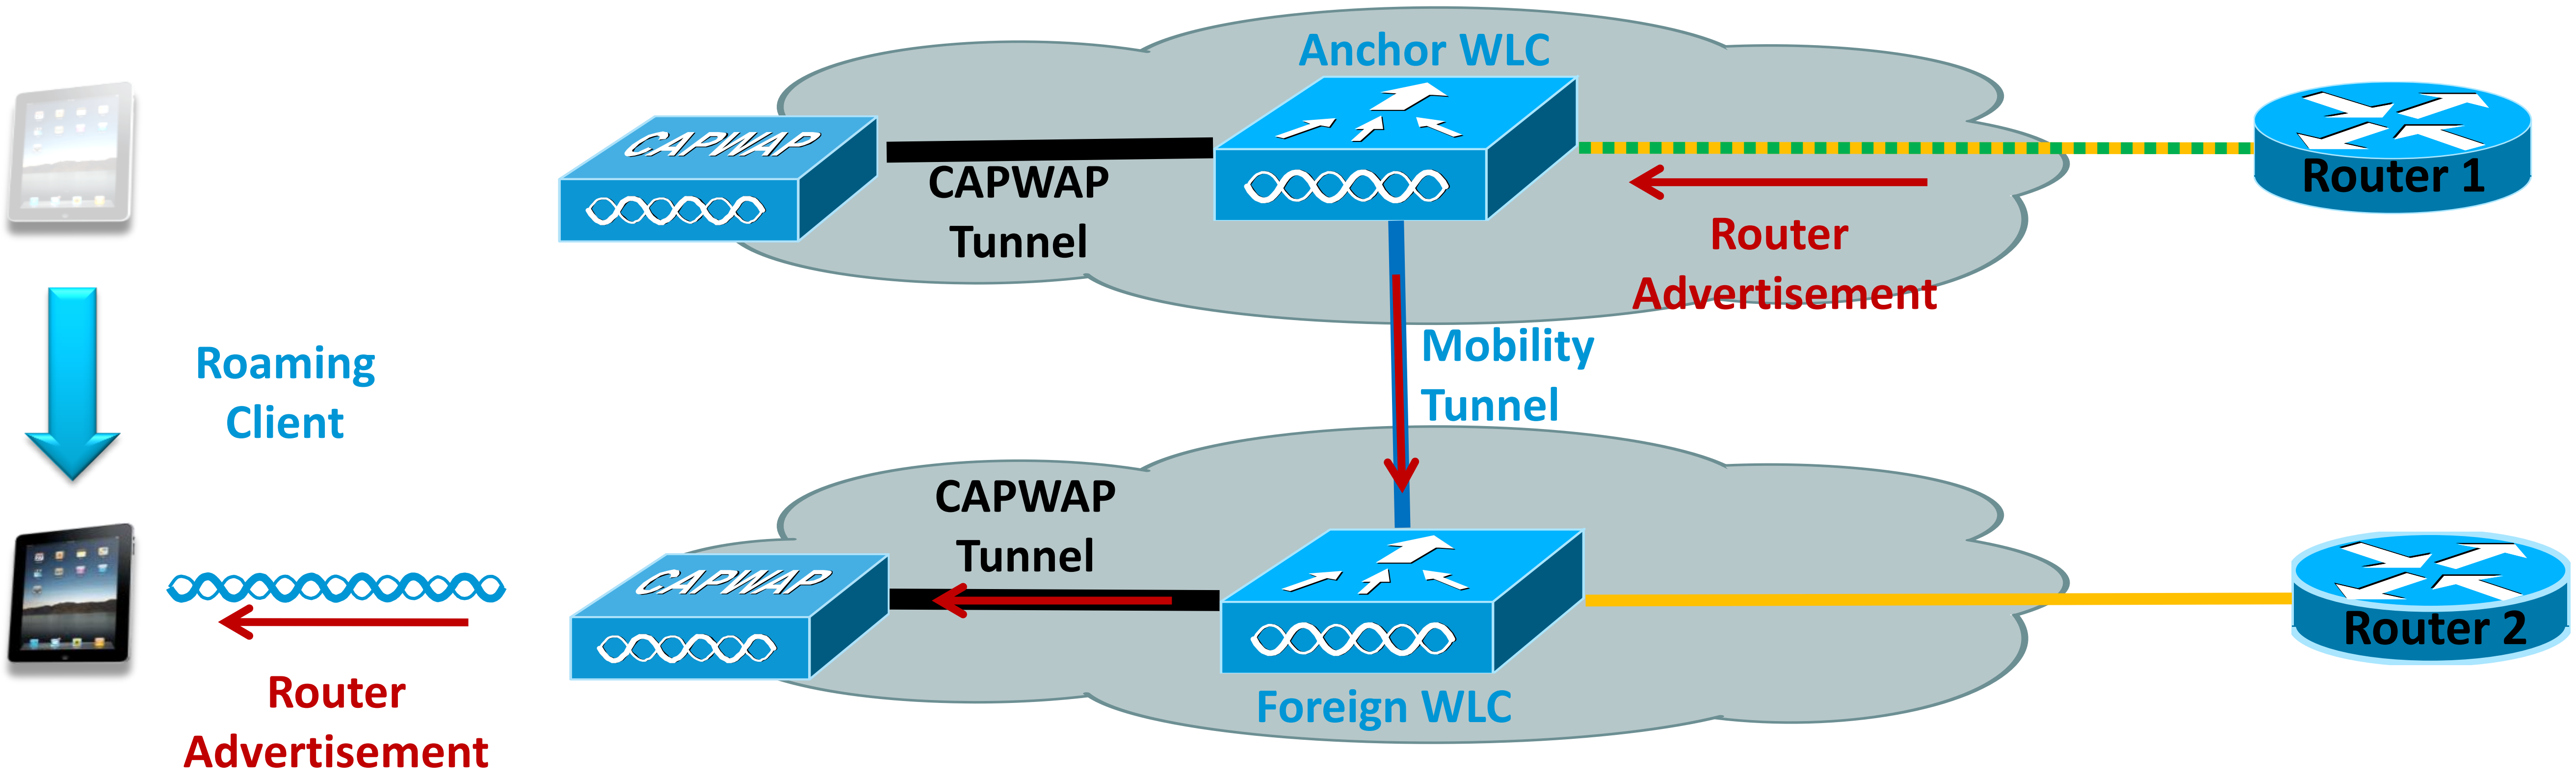

- To address this issue, the roaming client must be able to receive the original router advertisement
- The anchor controller sends the RA to the foreign in the mobility tunnel
- When the Access Point receives the RA, it will convert the multicast RA to unicast (MC2UC) and send RA to each client individually

New IPv6 Addresses Learning with Mobility

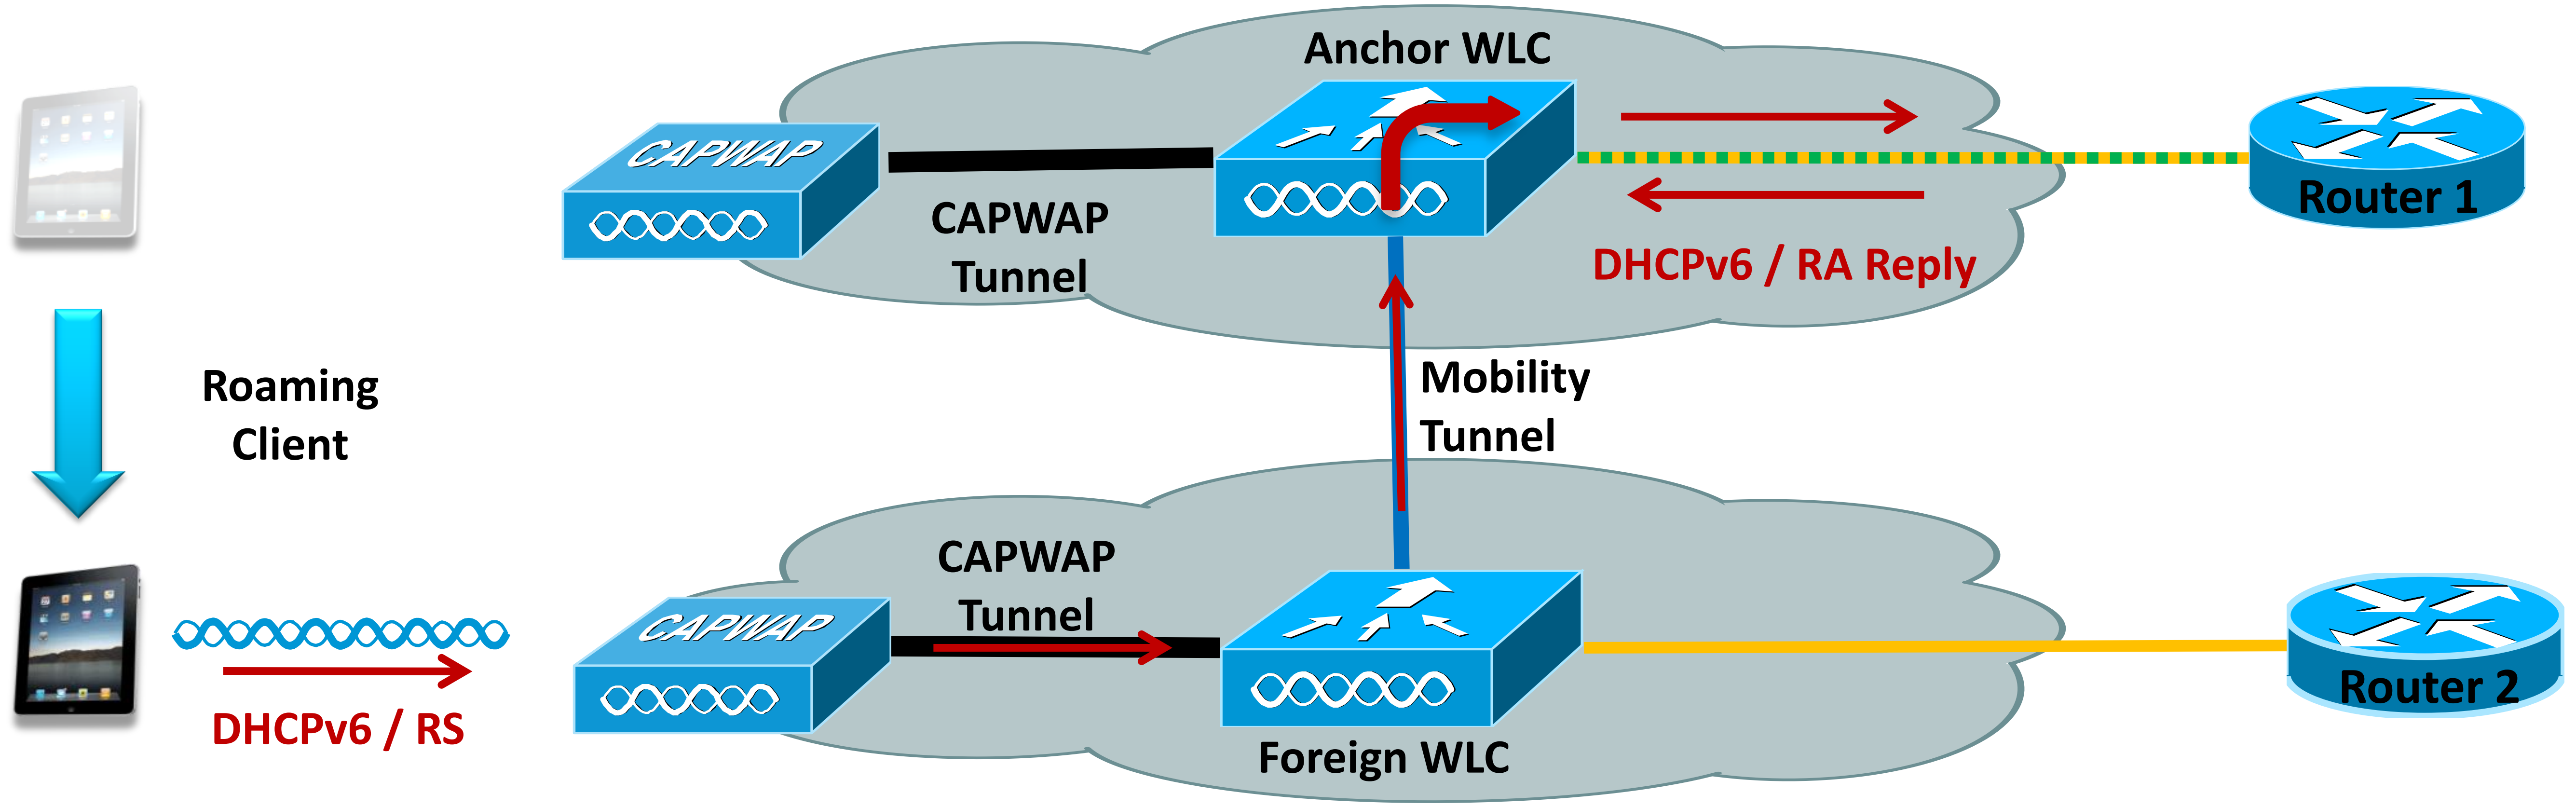

- IPv6 address is always learned at the anchor either through DHCPv6 or NDP
- DHCPv6 packets from a roamed client at the foreign controller will be tunneled to the anchor controller, which will learn the IPv6 address from the DHCPv6 replies
- Similarly NDP messages for a roamed client are processed at the anchor controller
- Whenever a new IPv6 address is learned at the anchor the new address is sent in a mobility message to the foreign controller

Testing IPv6-Only Client Mobility with Cisco

Clients – Windows 7 (SP1) with only IPv6 Enabled

Before Roaming

```
Wireless LAN adapter Wireless Network Connection:

   Connection-specific DNS Suffix  . : lab.local
   IPv6 Address. . . . . . . . . . . : 2001:db8:0:20:3057    d:587d:73ae
   IPv6 Address. . . . . . . . . . . : 2001:db8:0:20:7476    ec22:8a80
   Link-local IPv6 Address . . . . . : fe80::3057:   d:587    11
   Default Gateway . . . . . . . . . : fe80::217:fff:fe29:7440%1
```

During Roaming

```
Reply from 2001:db8:0:113::200: time=1ms
Reply from 2001:db8:0:113::200: time=1ms
Reply from 2001:db8:0:113::200: time=87ms
Reply from 2001:db8:0:113::200: time=13ms
Reply from 2001:db8:0:113::200: time=14ms
Reply from 2001:db8:0:113::200: time=15ms
```

IPv6 PING Continues

Client Keeps IPv6 Address and Connectivity is Seamless

After Roaming

```
Wireless LAN adapter Wireless Network Connection:

   Connection-specific DNS Suffix  . : lab.local
   IPv6 Address. . . . . . . . . . . : 2001:db8:0:20:3057    87d:73ae
   IPv6 Address. . . . . . . . . . . : 2001:db8:0:20:7476    ec22:8a80
   Link-local IPv6 Address . . . . . : fe80::3057:   d:58 d:73ae%11
   Default Gateway . . . . . . . . . : fe80::217:fff:fe29:7440%11
```

Use Case #2: Security

7.2

Challenge	Vulnerabilities originated from client community
Solution	First Hop Security blocks rogue announcements IPv6 ACLs provides IPv6 traffic control
Benefits	Increased network availability and reliability Lower operational cost
Differentiator	Proactively block known threats from wireless side

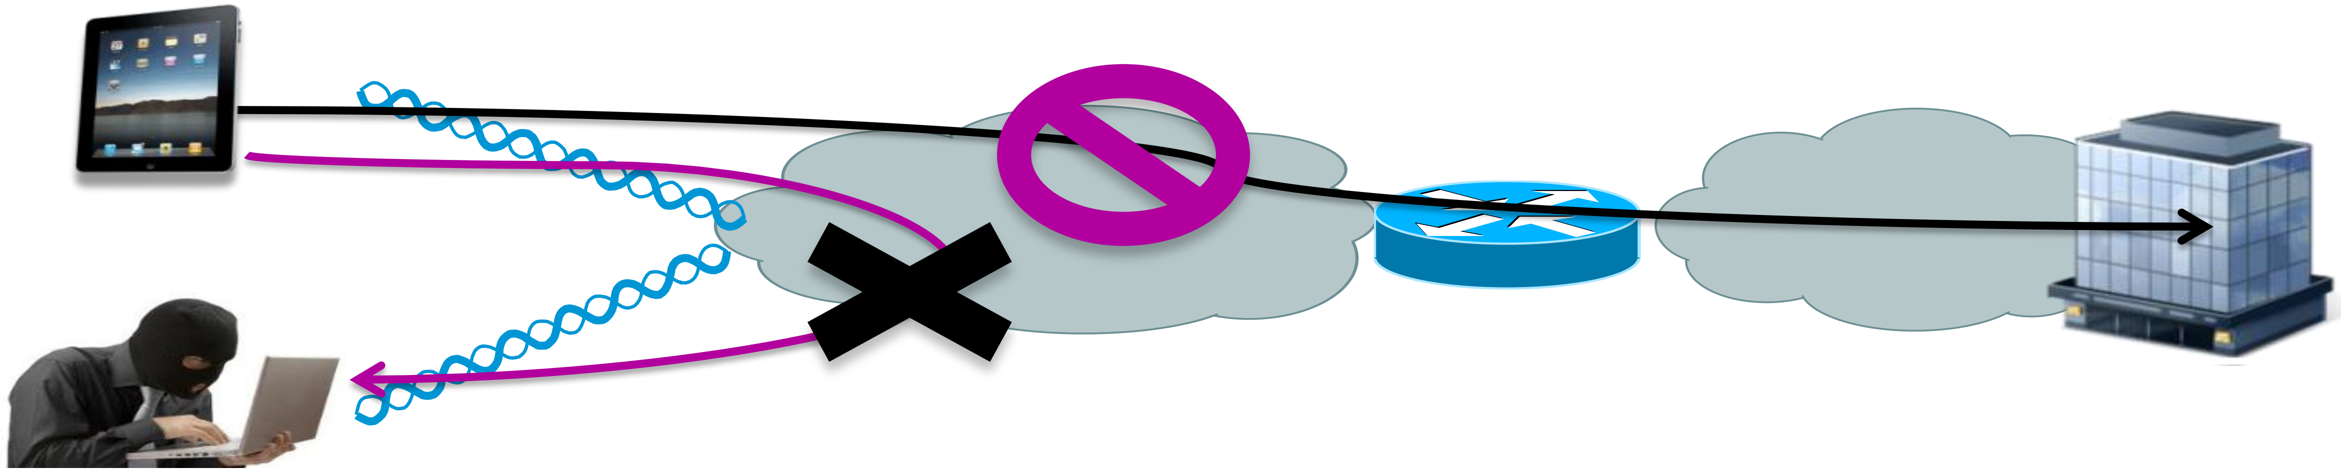

Rogue Router Announcement(RA)

Rogue DHCP Server

First Hop Security for Wireless IPv6 Clients

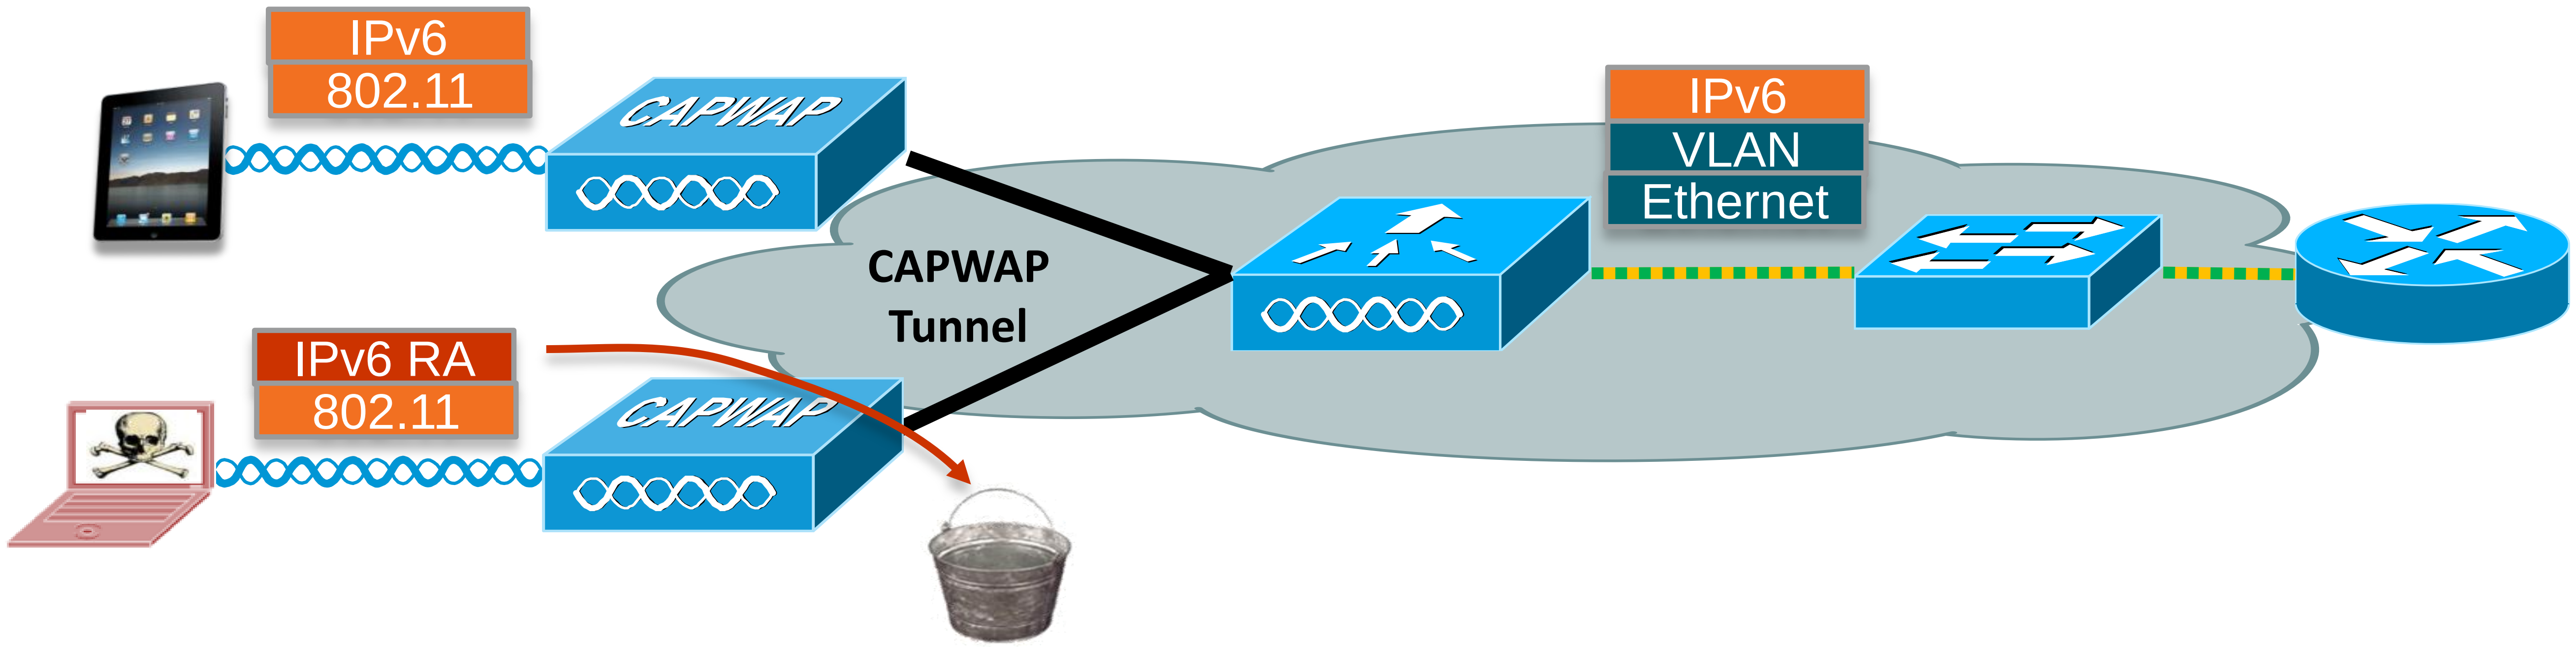

Router Advertisement Guard → RA From Client Dropped at the Access Point (Local and FlexConnect modes)

DHCPv6 Server Guard → DHCPv6 Advertisement Blocked at the Controller.

Undesired IPv6 Addresses/Prefix → IPv6 Source Guard Drops Undesired Packets at Controller

IPv6 ACL Support

Up to 128 ACL (64 for IPv4 and 64 for IPv6) supported

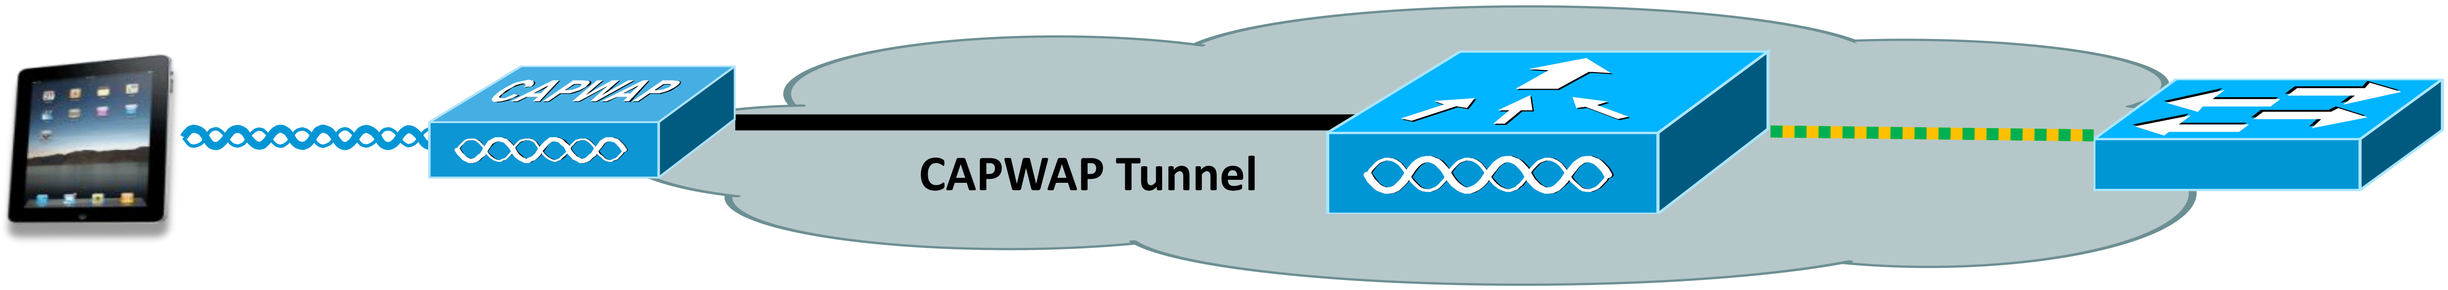

- Two ACL profiles (one for IPv4 and one for IPv6) are supported per dual stack client
- ACL profiles for wireless clients can be configured on Wireless Controller or provided by AAA Server
  - AAA server can send both IPv4 and IPv6 ACL attributes for dual stack clients after successful user authentication
- Counters are maintained on ACL matches for operational/maintenance purposes

IPv6 ACL Configuration

A separate IPv4 and IPv6 ACL can be applied on a per-WLAN basis.

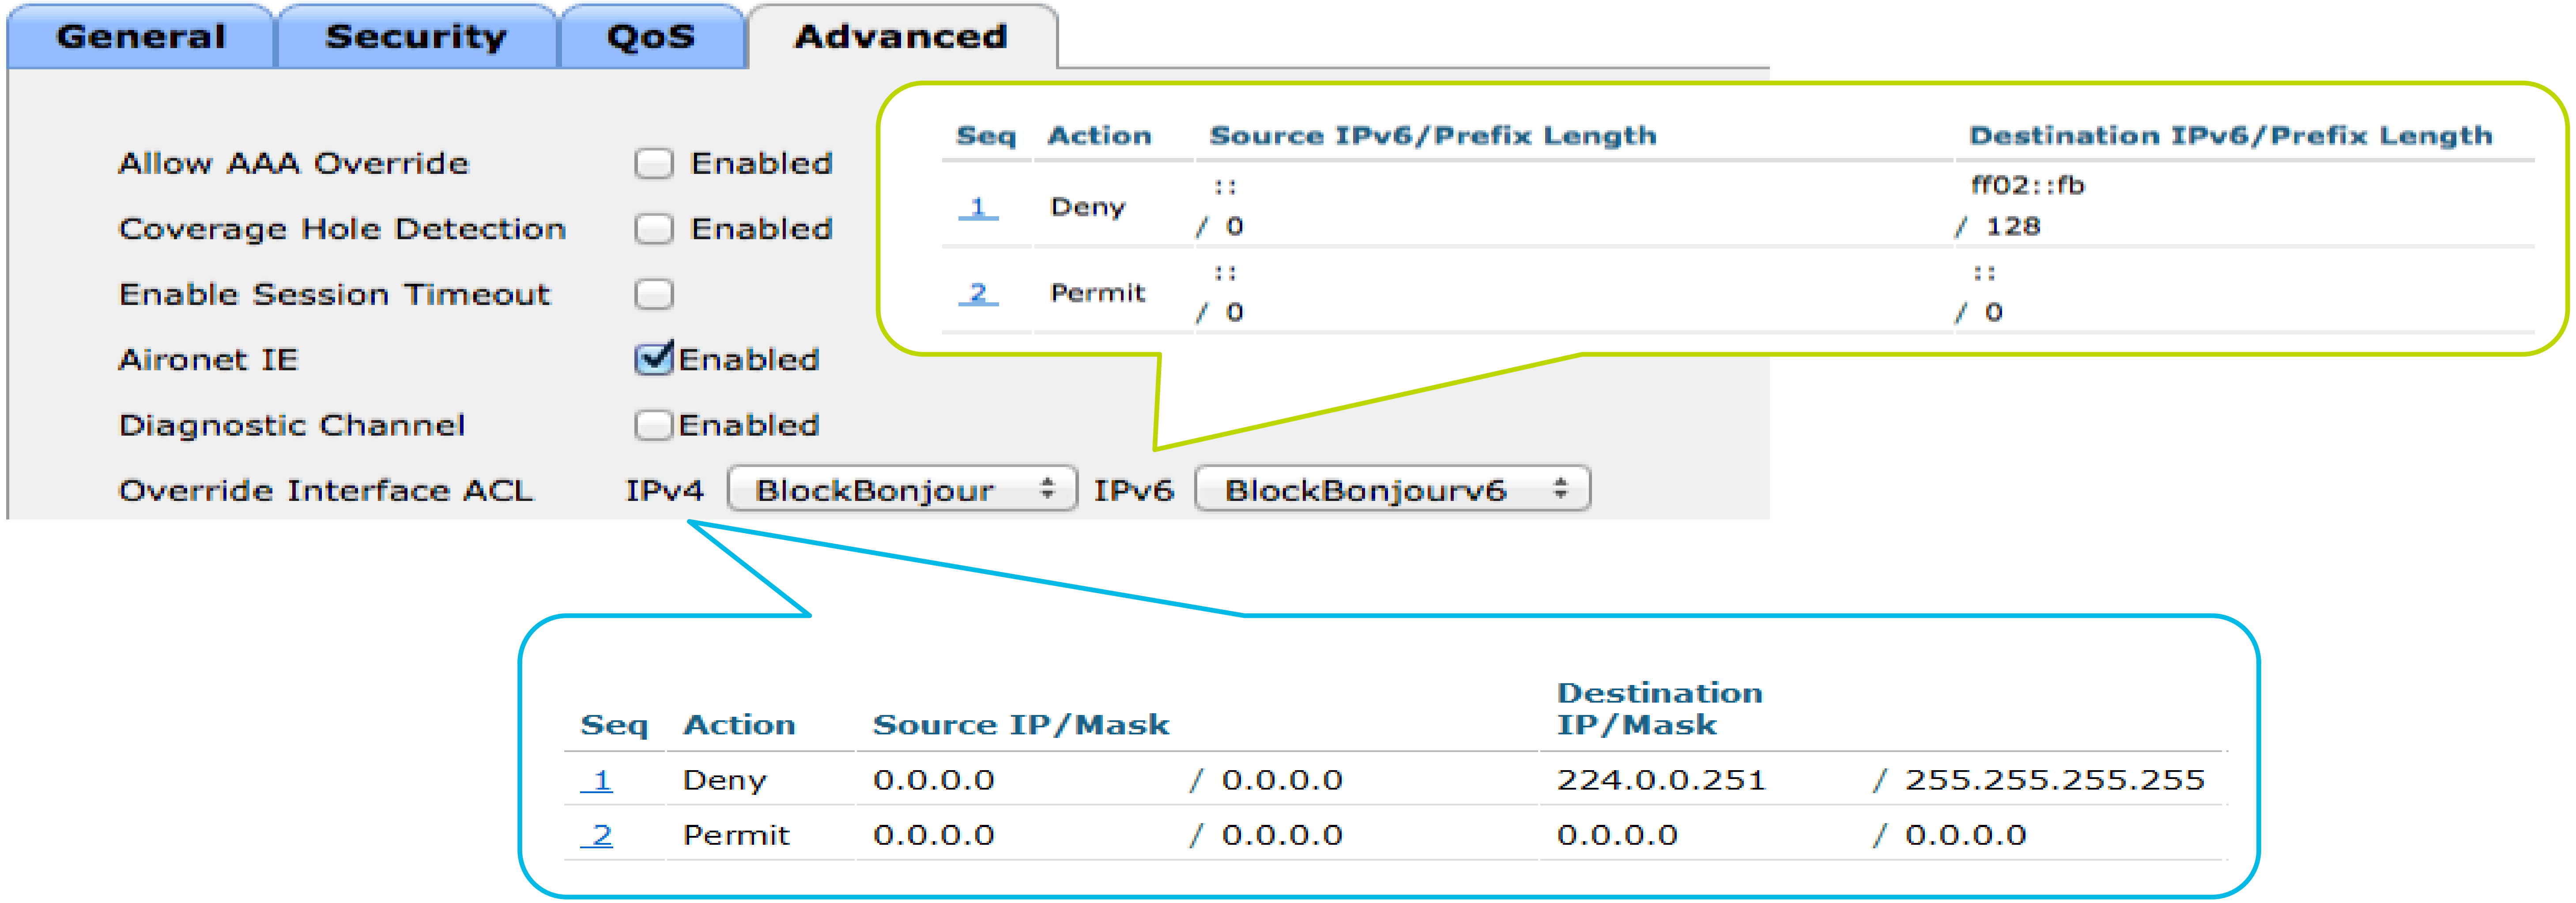

7.2

Use Case #3: Efficiency

Challenge	Chatty IPv6 packets, busy network, high CPU
Solution	Intelligent processing of IPv6 packets with proxy and rate limit
Benefits	Increase radio efficiency, decrease processing load on router
Differentiator	50% NDP reduction on wireless and 25% on wired side

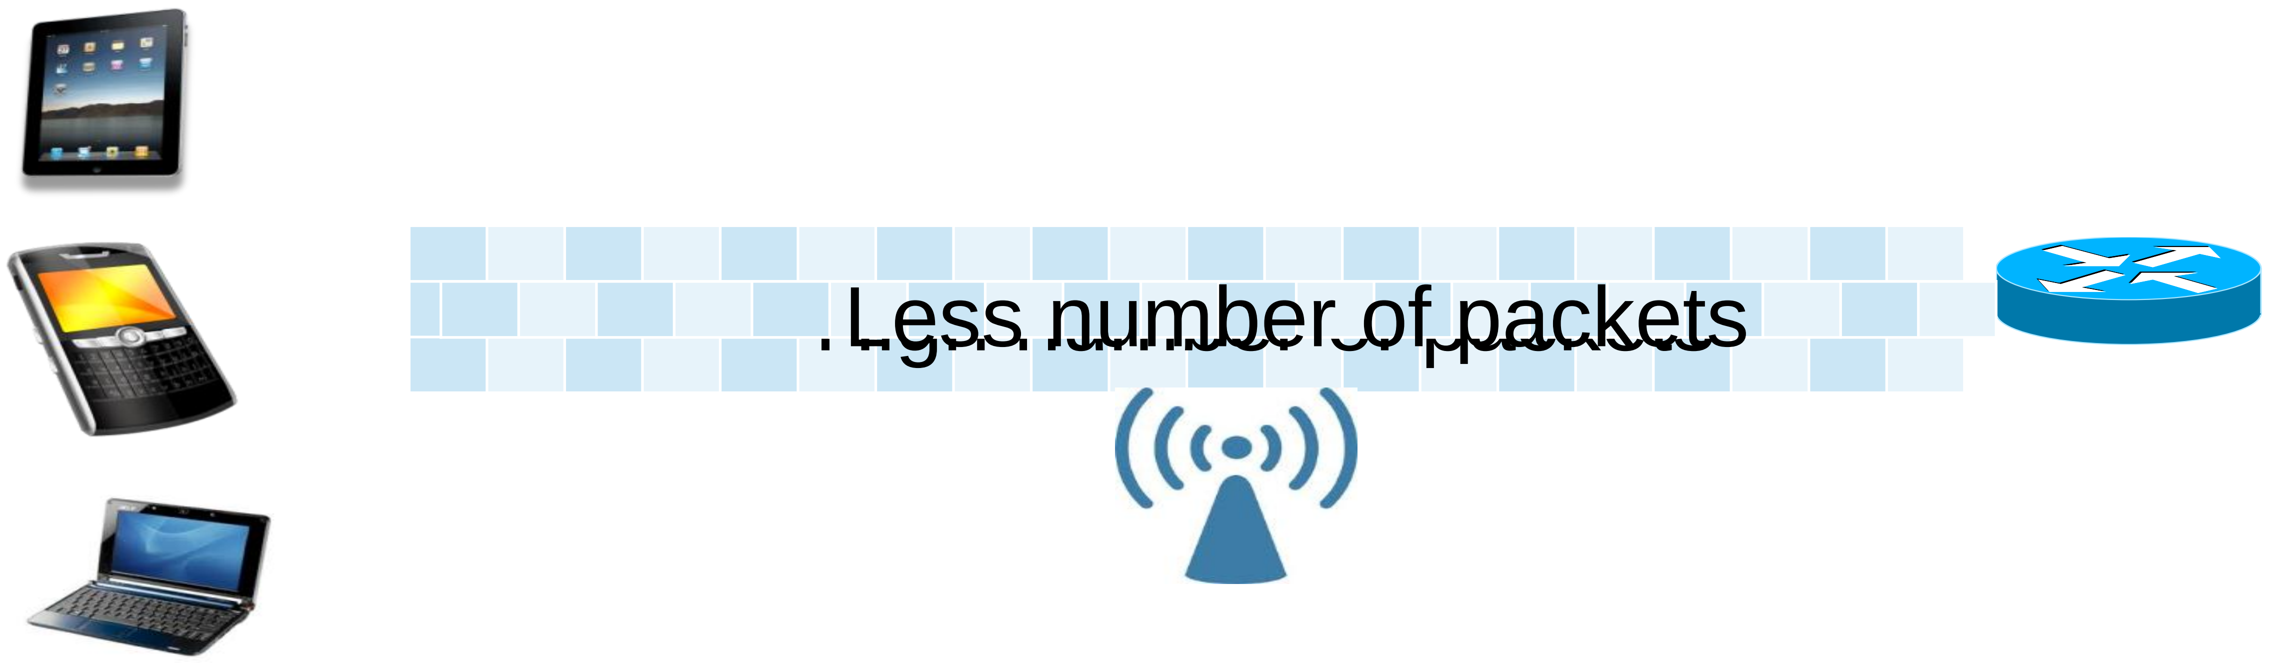

First Hop Optimisation for Wireless IPv6 Clients

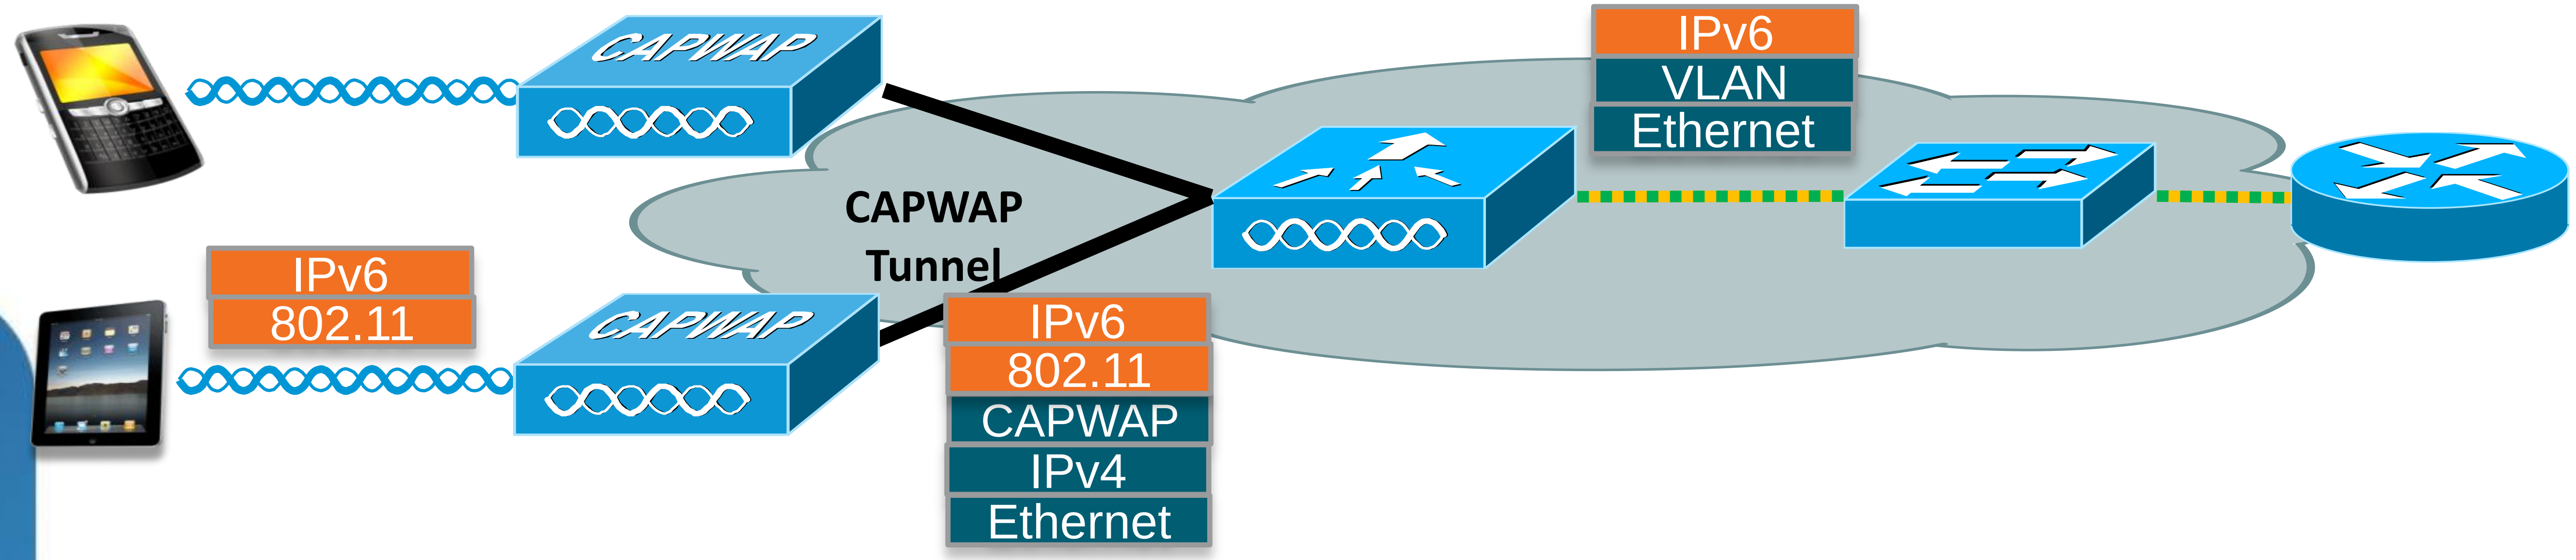

Rate Limiting/ Throttling — Router Advertisement (Periodic)

Neighbour Discovery Caching — Neighbour Solicitation → ; ← Proxy Neighbour Advertisement

Neighbour Solicitation (NS) Suppression
Respond to NS with cache binding table entry.

IPv6 Neighbour Discovery Caching

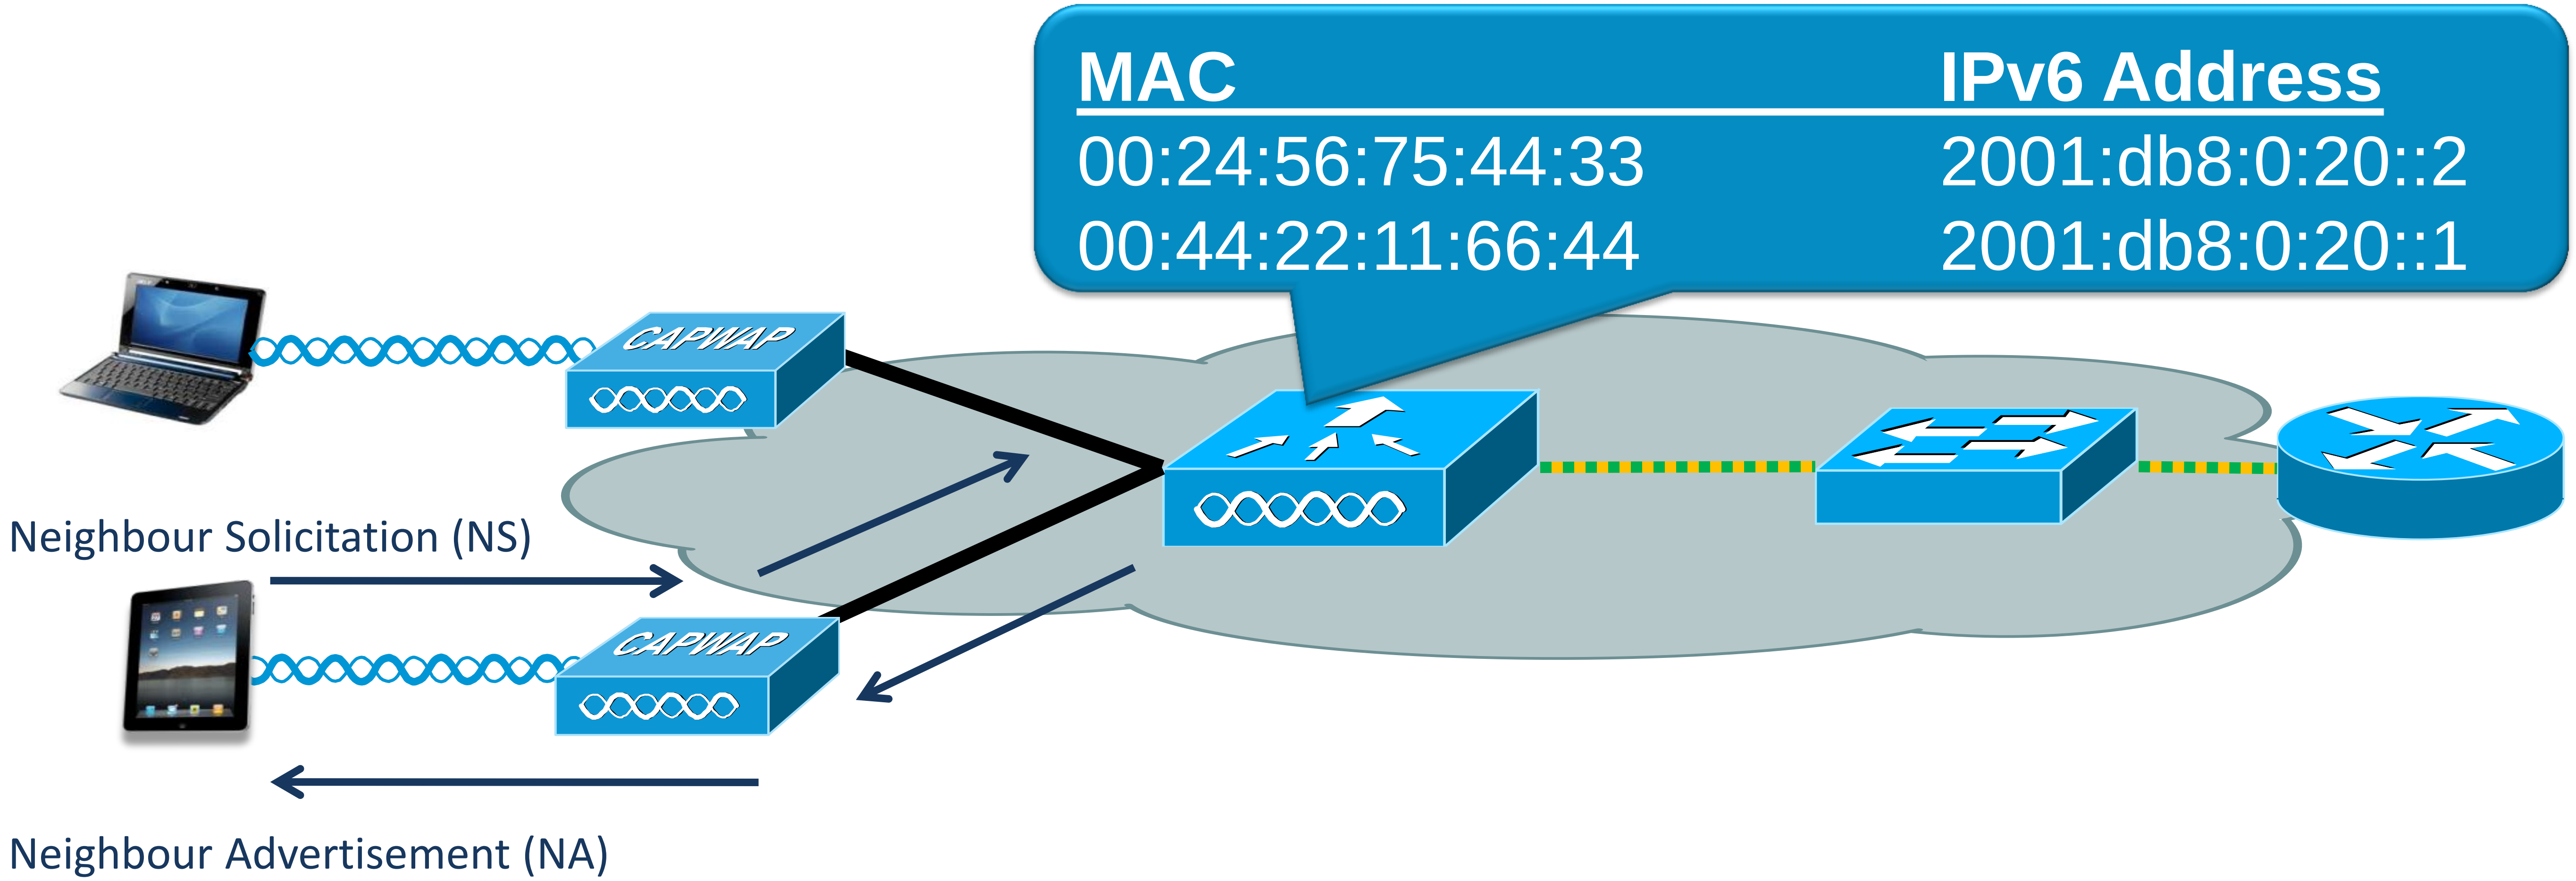

- The controller will respond to IPv6 neighbour solicitation messages by first checking it’s local cache for a match
- Neighbour Advertisements for the request are sent back via L2 unicast

First Hop Security / Efficiency Configuration

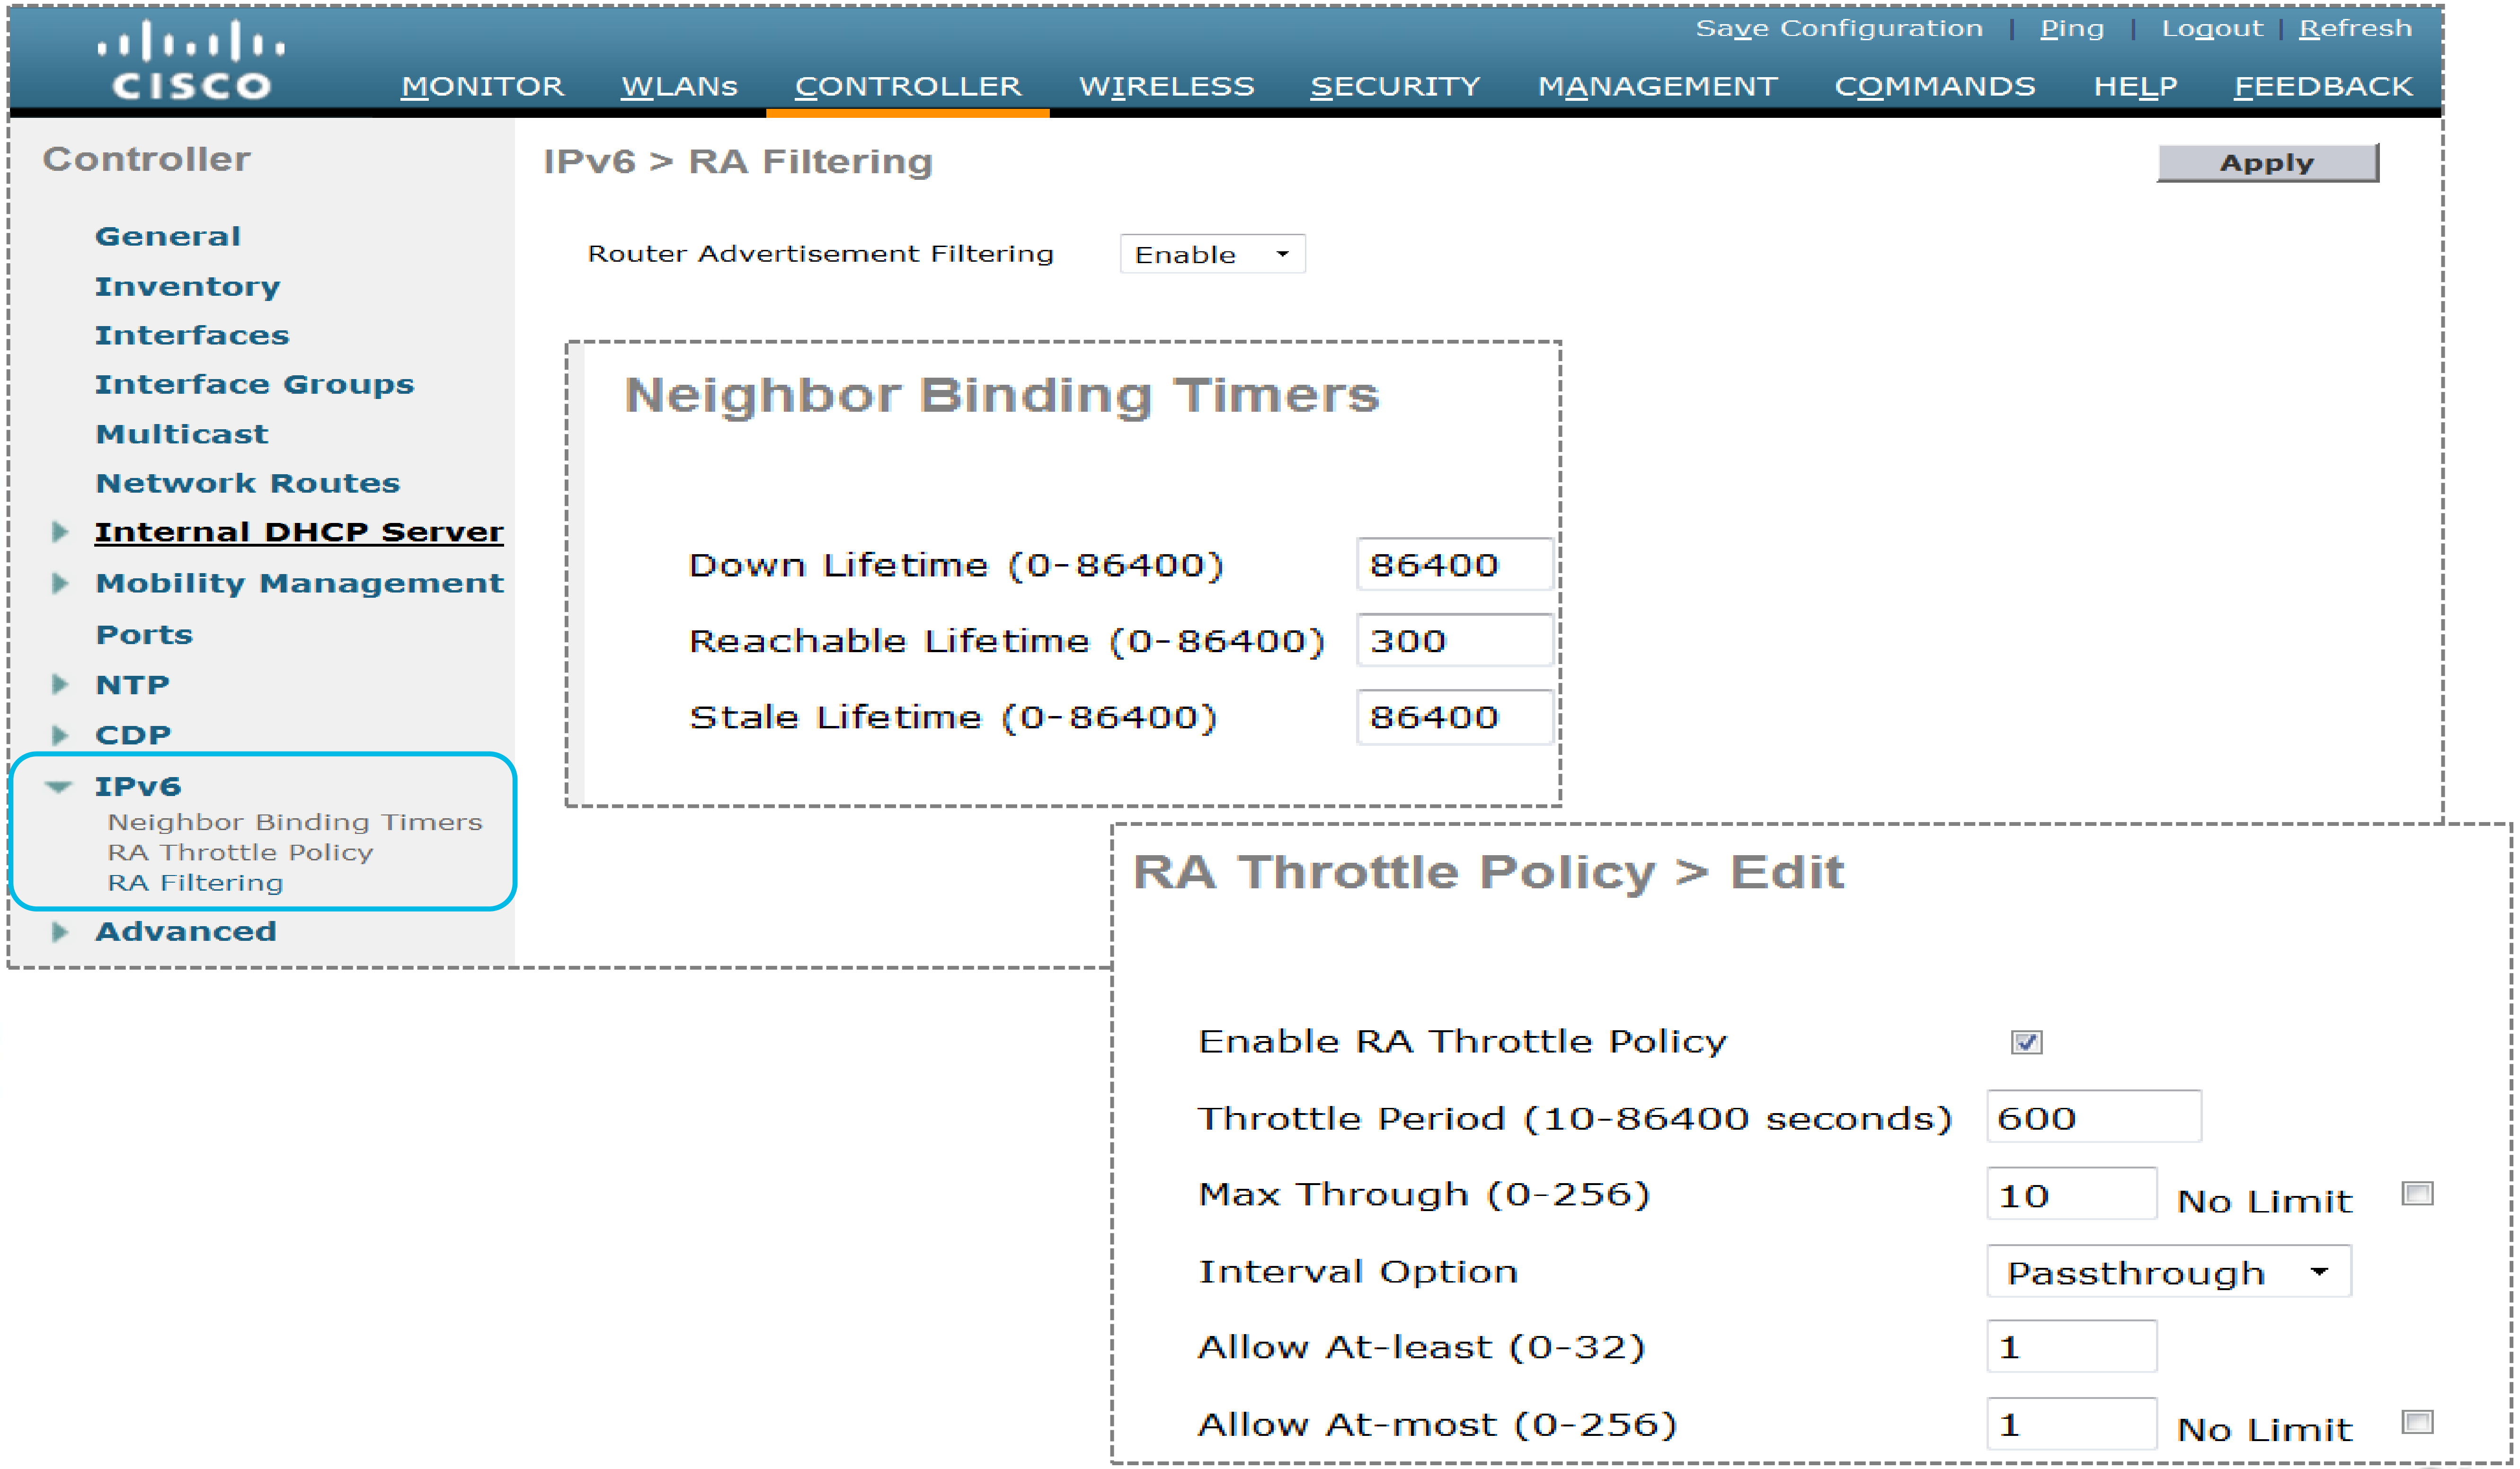

Guest Access for Dual Stack Clients

Web based authentication supported on wireless controllers:

- Web Pass-through
- Authentication
  - Local DB
  - RADIUS
  - LDAP
- Redirection
  - Splash/Conditional redirect
  - ISE integration redirects ***
  - External Webauth URL redirects

Dual-stack clients will need to authenticate only once via IPv4 or IPv6 to begin the guest session

Virtual Interface configured with IPv4 address only (IPv6 is auto-generated)

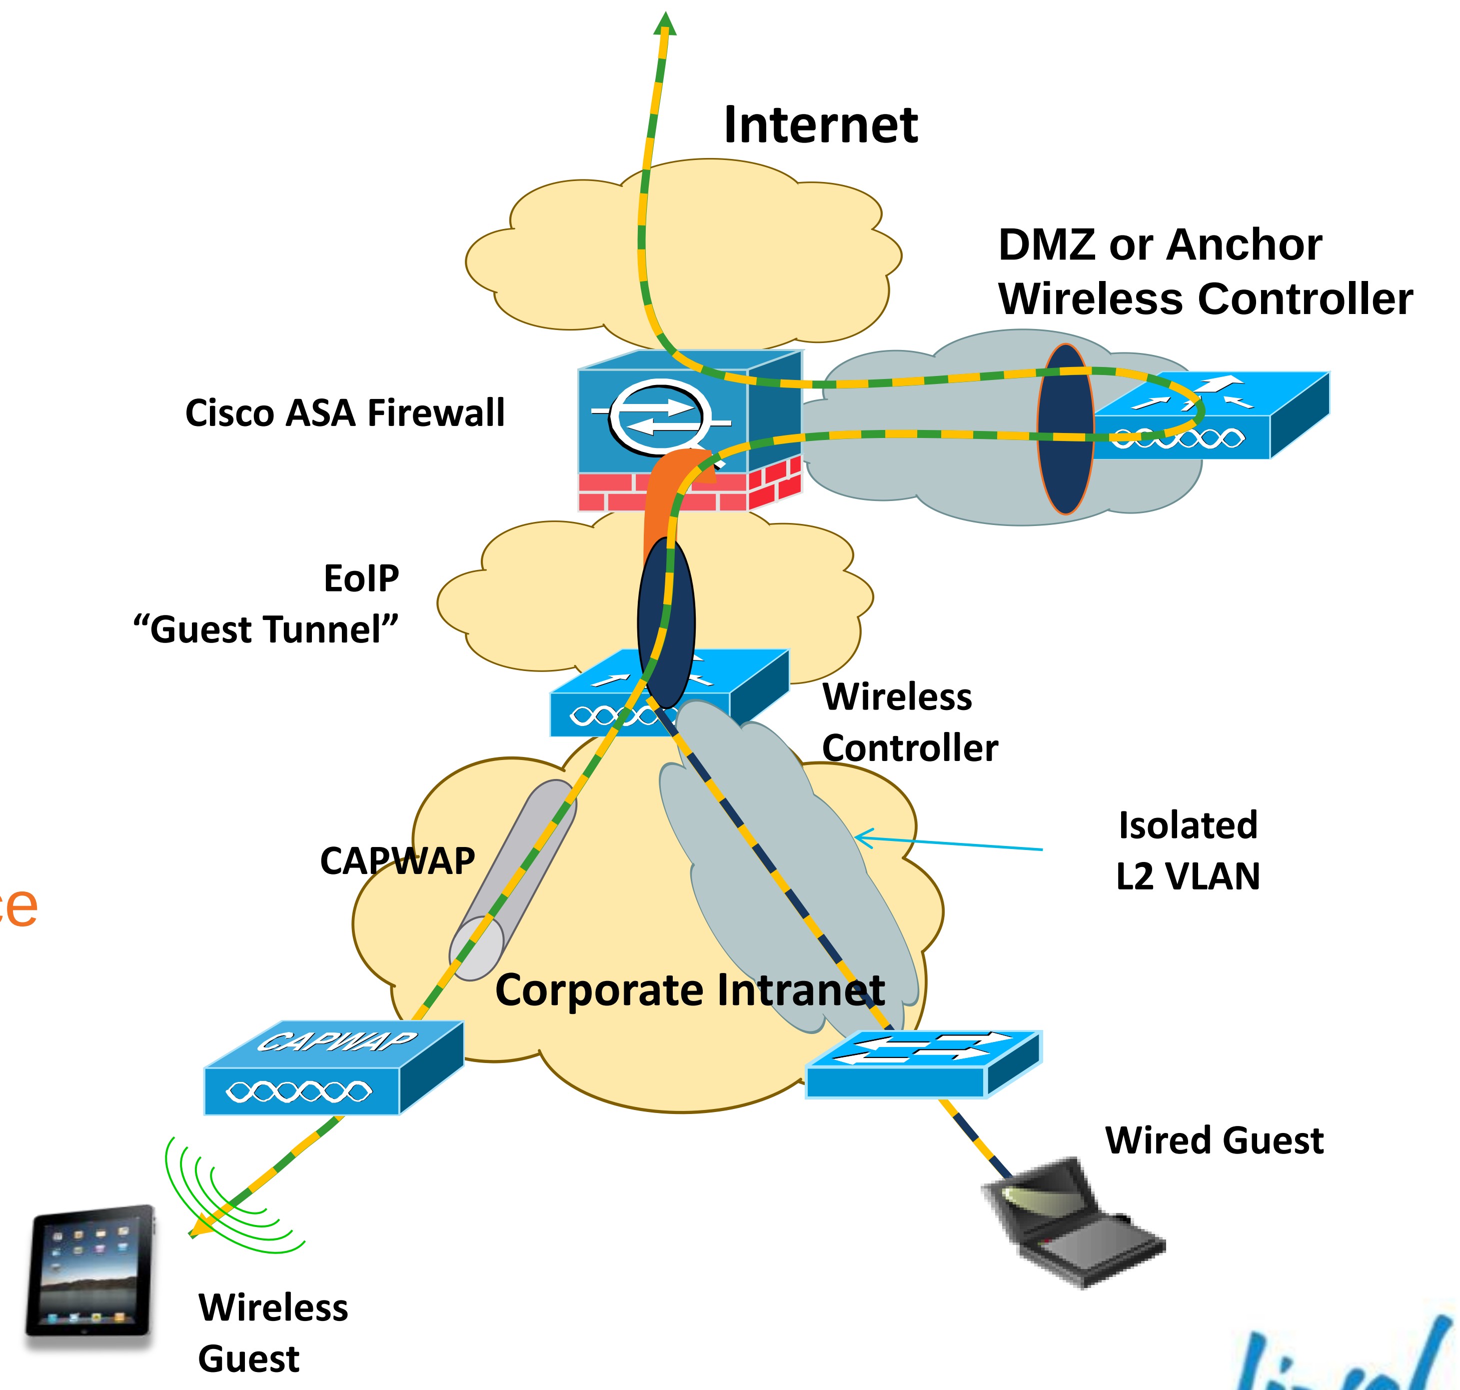

*** ISE integration will work with dual stack clients.

IPv6-Only Client Captive Portal

IPv6 Multicast Support

- Dual stack clients are supported through both IGMP for IPv4 and MLD for IPv6
- The controller supports MLDv1 and acts as a proxy for IPv6 clients
- When replying to upstream IPv6 router reports, the controller uses a native IPv6 link-local address derived from the MAC address

Multicast

Enable Global Multicast Mode	☑
Enable IGMP Snooping	☑
IGMP Timeout (seconds)	60
IGMP Query Interval (seconds)	20
Enable MLD Snooping	☑
MLD Timeout (seconds)	60
MLD Query Interval (seconds)	20

Multicast Groups

Layer3 MGID(Multicast Group ID) Mapping

Group address	Vlan	MGID	IGMP/MLD
224.0.0.251	20	555	IGMP
224.0.0.252	20	551	IGMP
239.255.255.250	20	553	IGMP
ff02::c	20	556	MLD
ff02::fb	20	550	MLD
ff02::1:3	20	552	MLD
ff02::2:fbb5:a199	20	554	MLD

Only Cisco Supports IPv6 Video Over Wireless

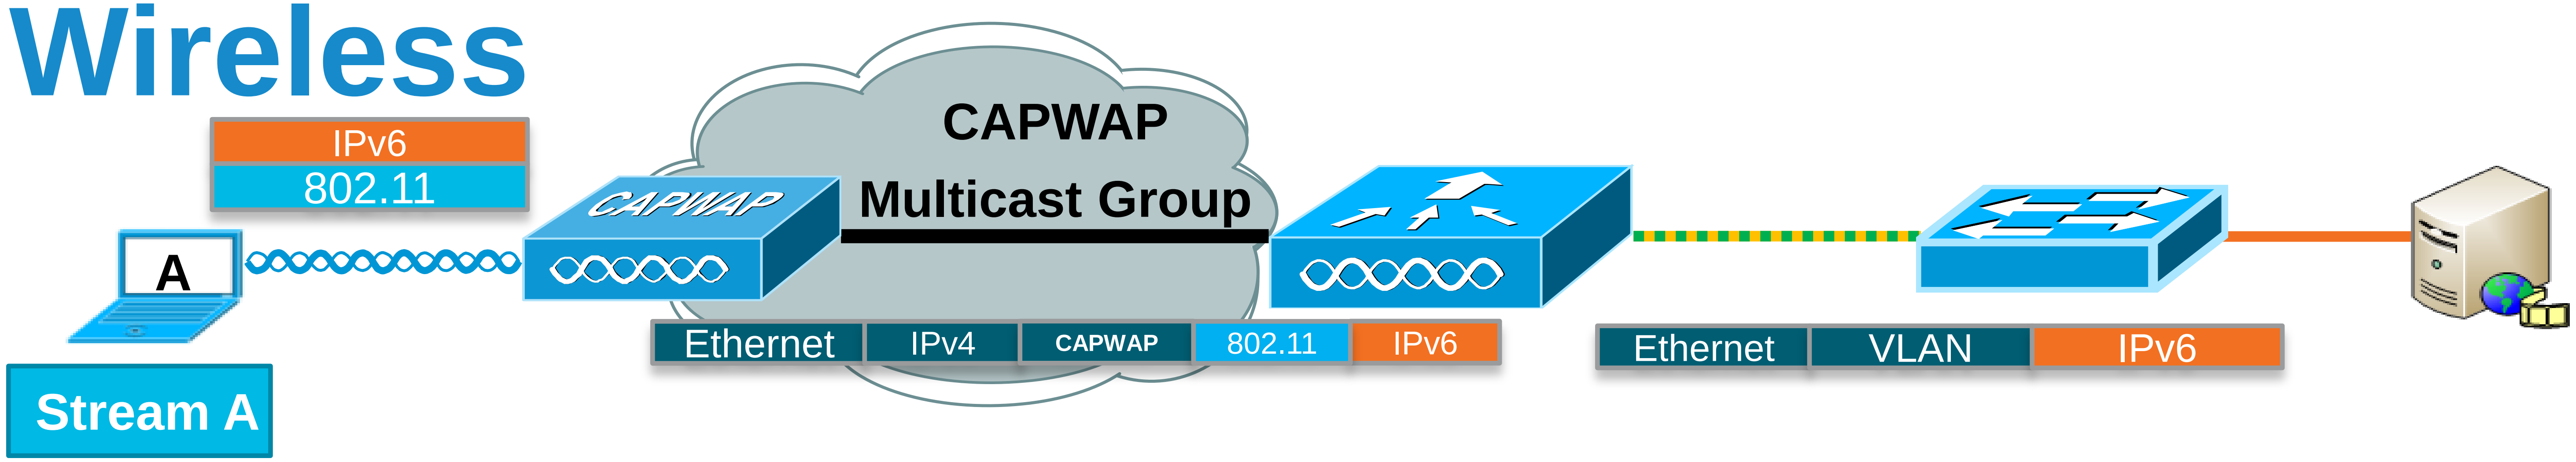

Media Streams

Stream Name	Start IP Address	End IP Address	Operation Status
VideoStream-IPv4	239.100.1.20	239.100.1.22	Multicast Direct
VideoStream-IPv6	ff00:0:2::20	ff00:0:2::22	Multicast Direct

- Existing VideoStream support using MC2UC (Multicast to Unicast) for IPv4 works the same for IPv6 multicast streams
- The multicast to unicast conversion occurs at the Access Point for efficiency and scalability

IPv6 Quality of Service Mapping

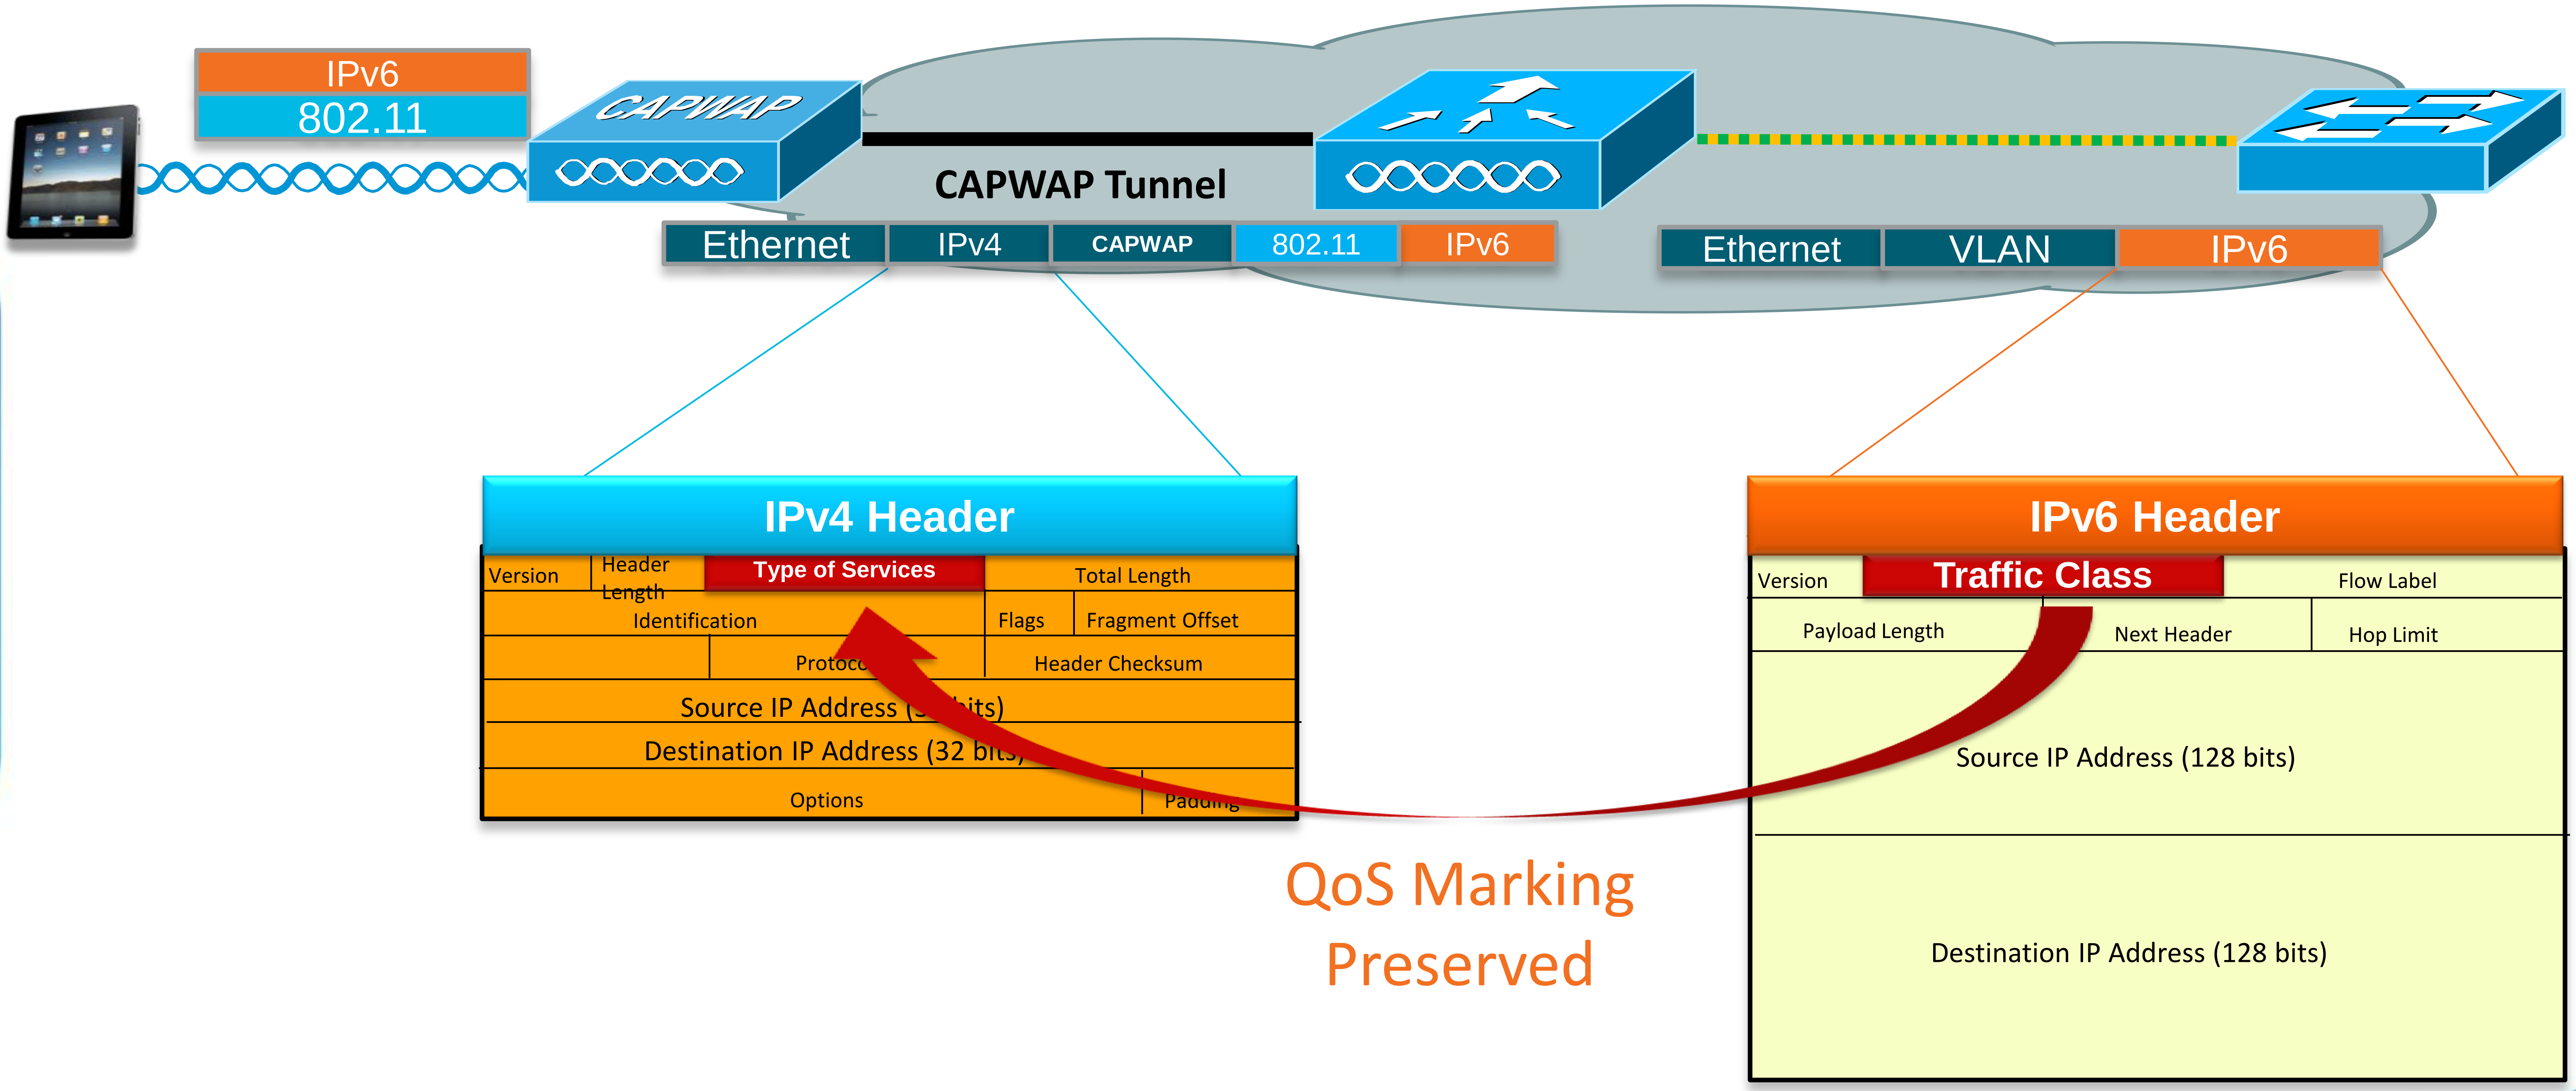

FlexConnect and IPv6

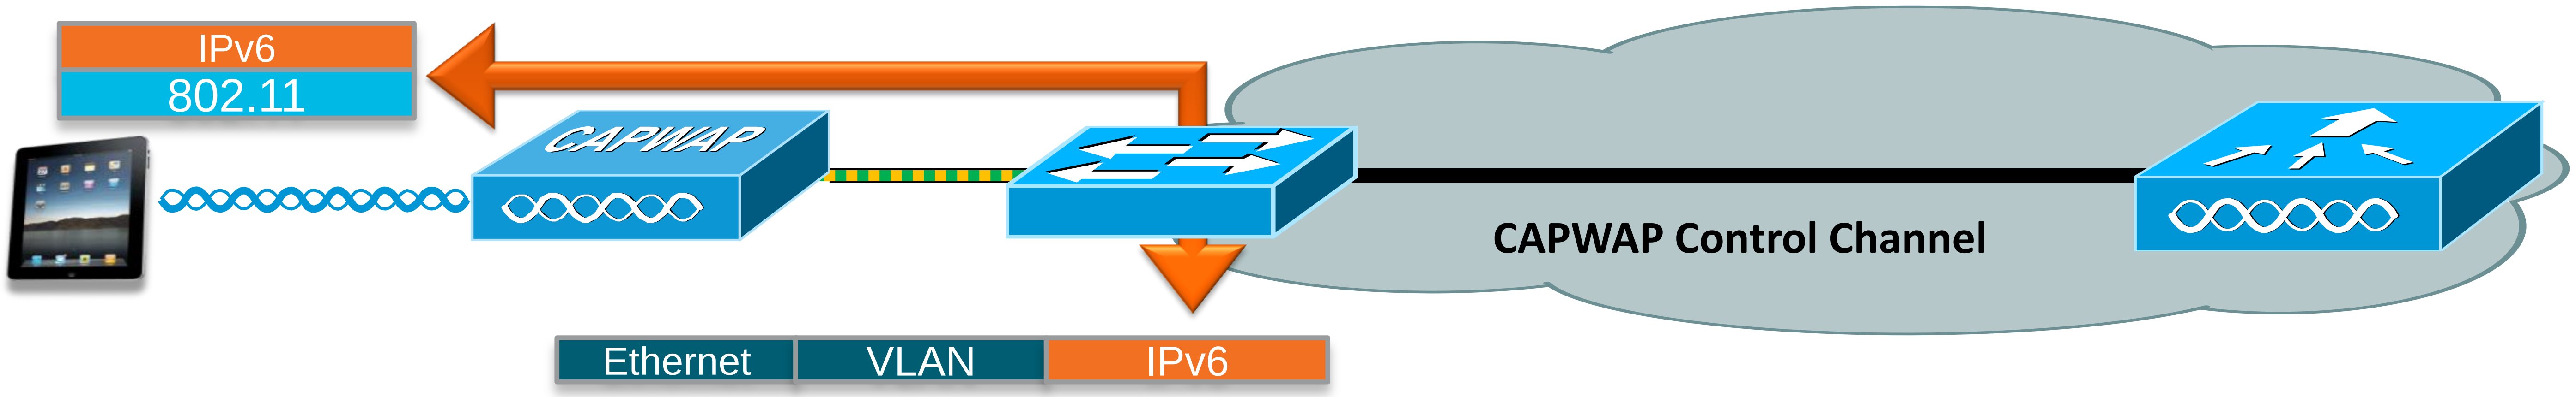

- Dual Stack and IPv6 clients will support bridging of IPv6 packets in local switching
  - No IPv6 ACLs supported on AP
  - No L3 mobility supported for locally switched-WLANs as is the case with IPv4
  - RA Guard is still active for Local Switched WLANs
  - IPv6 address of locally switched clients will not be displayed as the AP does not snoop the IPv6 NDP packets
- FlexConnect Central Switching WLANs do not support IPv6 (this is dependent on supporting MC2UC in FlexConnect)

Use Case #4: Client Management

7.2

Challenge	Client visibility is vital for Troubleshooting, Planning & Security
Solution	NCS tracks IPv6 client addresses, client IP version distribution and trending; MSE tracks IPv6 client locations
Benefits	Prepare admin for IPv6 troubleshooting, address planning Provide client traceability
Differentiator	Management system for wired + wireless, IPv4 + IPv6

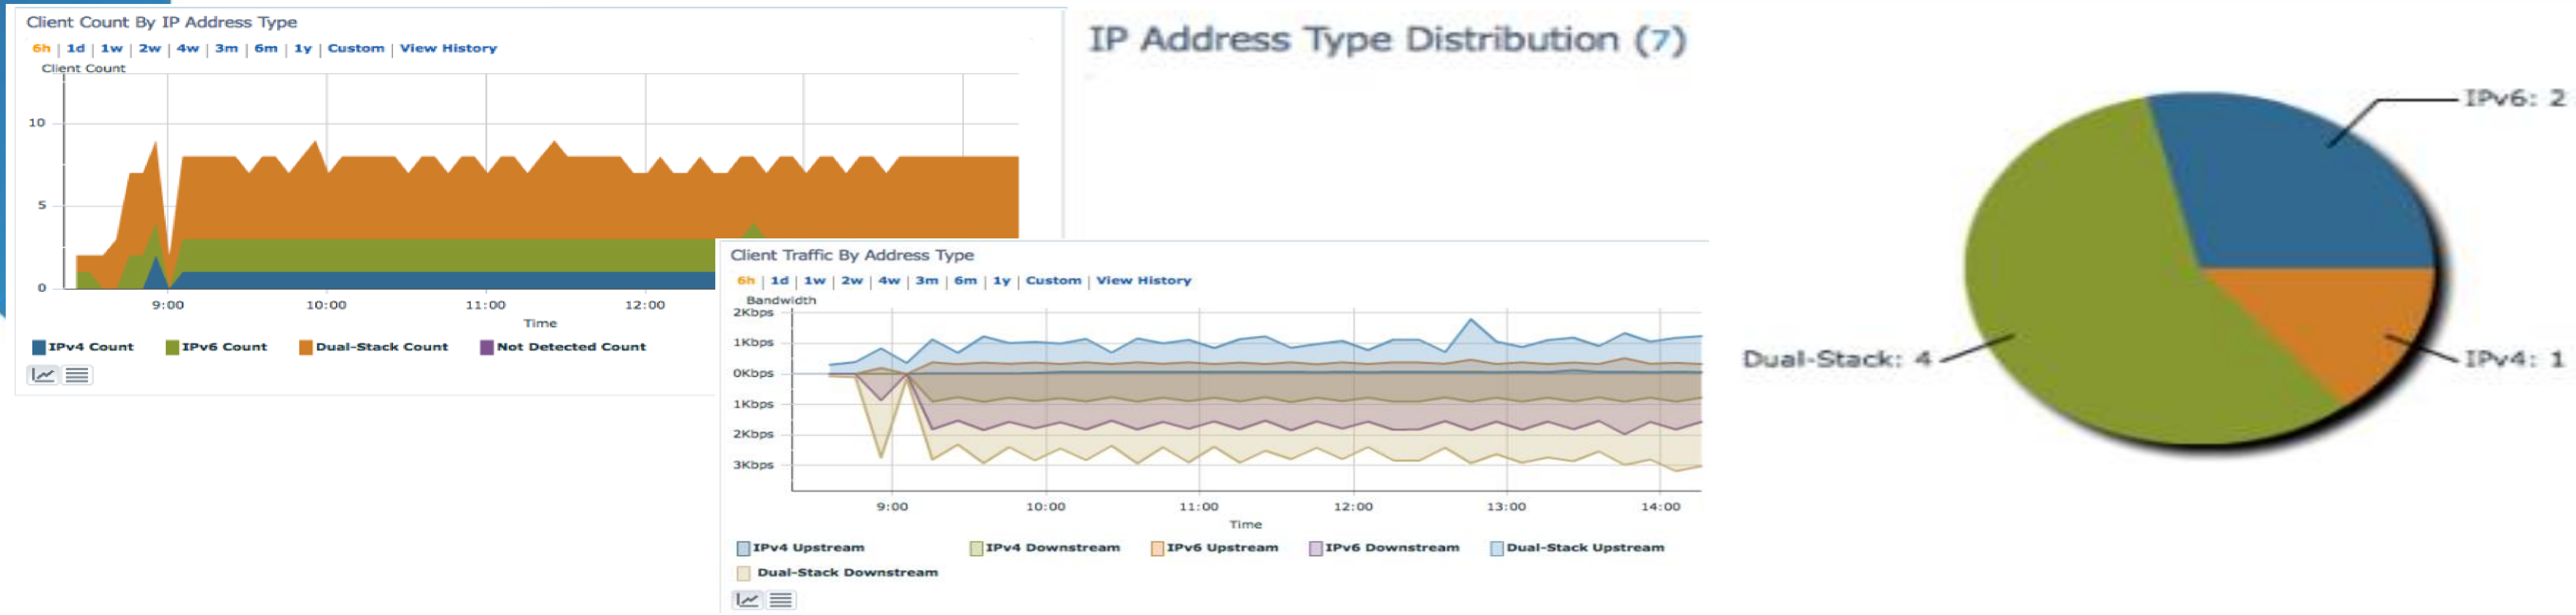

Cisco NCS 1.1 Provides Comprehensive IPv6 Client Visibility and Monitoring

Cisco Prime
Network Control System

Home | Monitor | Configure | Services | Rep

Log Out

Visibility – Recognition of IPv6 Global and Link Local Addresses

Clients and Users

Troubleshoot | Test | Disable | Remove | More | Track Clients | Identi… nknown Users

	MAC Address	Vendor	IP Address	IP Type	Link Local	Router Advertisements Dropped
○	00:21:6a:a7:4f:ee	Intel	2001:db8:0:20:3057:534d:587d:73ae	IPv6	fe80::3057:534d:587d:73ae	0
○	00:21:6a:a7:54:88	Intel	192.168.20.21	Dual-Stack	fe80::5dda:a8e0:a969:fde6	0
○	00:24:d7:99:97:08	Intel	192.168.20.23	Dual-Stack	fe80::224:d7ff:fe99:9708	70
○	00:21:6a:5a:86:70	Intel	192.168.20.30	Dual-Stack	fe80::221:6aff:fe5a:8670	0
○	00:21:6a:67:31:48	Intel	192.168.20.25	Dual-Stack	fe80::acec:d514:2a14:ca7d	0
○	00:21:6a:a7:54:4e	Intel	192.168.20.22	Dual-Stack	fe80::1981:6f73:e618:32bd	0
○	f8:1e:df:e5:5b:03	Apple	192.168.20.29	Dual-Stack	fe80::fa1e:dfff:fee5:5b03	0
○	f8:1e:df:e3:0a:76	Apple	192.168.20.28	Dual-Stack	fe80::fa1e:dfff:fee3:a76	0
○	00:21:6a:a7:78:64	Intel	192.168.20.27	Dual-Stack	fe80::b5ba:eb3d:848d:ab6a	0

Insight – Identification of IPv4, Dual-Stack or IPv6-Only Client Types

Security – Identification of Clients Acting as IPv6 Routers

Cisco NCS 1.1 Provides a Rich IPv6 Session History

Client IPv6 Addresses for: 00:21:6a:a7:4f:ee

Total 3

IP Address	Scope	Assignment	Discovery Time
2001:db8:0:20:3057:534d:587d:73ae	Global Unique	NDP	2011-Nov-04, 15:45:35 UTC
2001:db8:0:20:b3ac:f21d:3da6:833f	Global Unique	DHCP	2011-Nov-04, 15:45:35 UTC
fe80::3057:534d:587d:73ae	Link Local	NDP	2011-Nov-04, 15:45:35 UTC

Association History

Association Time	Duration	AP Name	IP Address Type	IP Address
2011-Nov-01, 16:36:37 UTC	1 hrs 15 min 1 sec	3500-1	IPv6	2001:db8:0:20:7476:cf0c:ec22:8a80
2011-Nov-01, 16:26:37 UTC	5 min 0 sec	3500-1	IPv6	2001:db8:0:21:3057:534d:587d:73ae
2011-Oct-31, 21:36:05 UTC	1 hrs 30 min 1 sec	3500-1	IPv6	2001:db8:0:21:d57f:63fc:34ac:663f
2011-Oct-31, 04:02:52 UTC	17 hrs 28 min 12 sec	3500-1	IPv6	2001:db8:1:21:3057:534d:587d:73ae

NCS Logs both current and past IPv6 Addresses

- Since IPv6 clients can change addresses so often (sometimes 1 per day with temporary addresses), they need to be tracked over time
- This is needed for tracking down attacks or copyright infringement violations that need to be audited all the way back to the user

Distribution Layer: Example with ULA and General Prefix Feature

```
ipv6 general-prefix ULA-CORE FD9C:58ED:7D73::/53
ipv6 general-prefix ULA-ACC FD9C:58ED:7D73:1000::/53
ipv6 unicast-routing
!
interface GigabitEthernet1/0/1
 description To 6k-core-right
 ipv6 address ULA-CORE ::3:0:0:0:D63/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
 ipv6 summary-address eigrp 10 FD9C:58ED:7D73:1000::/53
!
interface GigabitEthernet1/0/2
 description To 6k-core-left
 ipv6 address ULA-CORE ::C:0:0:0:D63/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
 ipv6 summary-address eigrp 10 FD9C:58ED:7D73:1000::/53
```

```
interface Vlan4
 description Data VLAN for Access
 ipv6 address ULA-ACC ::D63/64
 ipv6 eigrp 10
 standby version 2
 standby 2 ipv6 autoconfig
 standby 2 timers msec 250 msec 750
 standby 2 priority 110
 standby 2 preempt delay minimum 180
 standby 2 authentication ese
!
ipv6 router eigrp 10
 no shutdown
 router-id 10.122.10.10
 passive-interface Vlan4
 passive-interface Loopback0
```

Access Layer: Dual Stack (Routed Access)

```
ipv6 unicast-routing
ipv6 cef
!
interface GigabitEthernet1/0/25
 description To 6k-dist-1
 ipv6 address 2001:DB8:CAFE:1100::CAC1:3750/64
 ipv6 ospf network point-to-point
 ipv6 ospf 1 area 2
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
!
interface GigabitEthernet1/0/26
 description To 6k-dist-2
 ipv6 address 2001:DB8:CAFE:1101::CAC1:3750/64
 ipv6 ospf network point-to-point
 ipv6 ospf 1 area 2
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
```

```
interface Vlan2
 description Data VLAN for Access
 ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64
 ipv6 ospf 1 area 2
 ipv6 cef
!
ipv6 router ospf 1
 router-id 10.120.2.1
 log-adjacency-changes
 auto-cost reference-bandwidth 10000
 area 2 stub no-summary
 passive-interface Vlan2
 timers spf 1 5
```

Distribution Layer: Dual Stack (Routed Access)

```
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef distributed
!
interface GigabitEthernet3/1
 description To 3750-acc-1
 ipv6 address 2001:DB8:CAFE:1100::A001:1010/64
 ipv6 ospf network point-to-point
 ipv6 ospf 1 area 2
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
!
interface GigabitEthernet1/2
 description To 3750-acc-2
 ipv6 address 2001:DB8:CAFE:1103::A001:1010/64
 ipv6 ospf network point-to-point
 ipv6 ospf 1 area 2
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
```

```
ipv6 router ospf 1
 auto-cost reference-bandwidth 10000
 router-id 10.122.0.25
 log-adjacency-changes
 area 2 stub no-summary
 passive-interface Vlan2
 area 2 range 2001:DB8:CAFE:xxxx::/xx
 timers spf 1 5
```

Campus IPv6 Deployment Options

Hybrid Model

- Plan "B" if Layer 3 device can't support IPv6 but you have to get IPv6 over it
- Offers IPv6 connectivity via multiple options
  - Dual-stack
  - Configured tunnels—L3-to-L3
  - ISATAP—Host-to-L3
- Leverages existing network
- Offers natural progression to full dual-stack design
- May require tunnelling to less-than-optimal layers (i.e. core layer)
- Any sizable deployment will be an operational management challenge
- ISATAP creates a flat network (all hosts on same tunnel are peers)
- Provides basic HA of ISATAP tunnels via old Anycast-RP idea

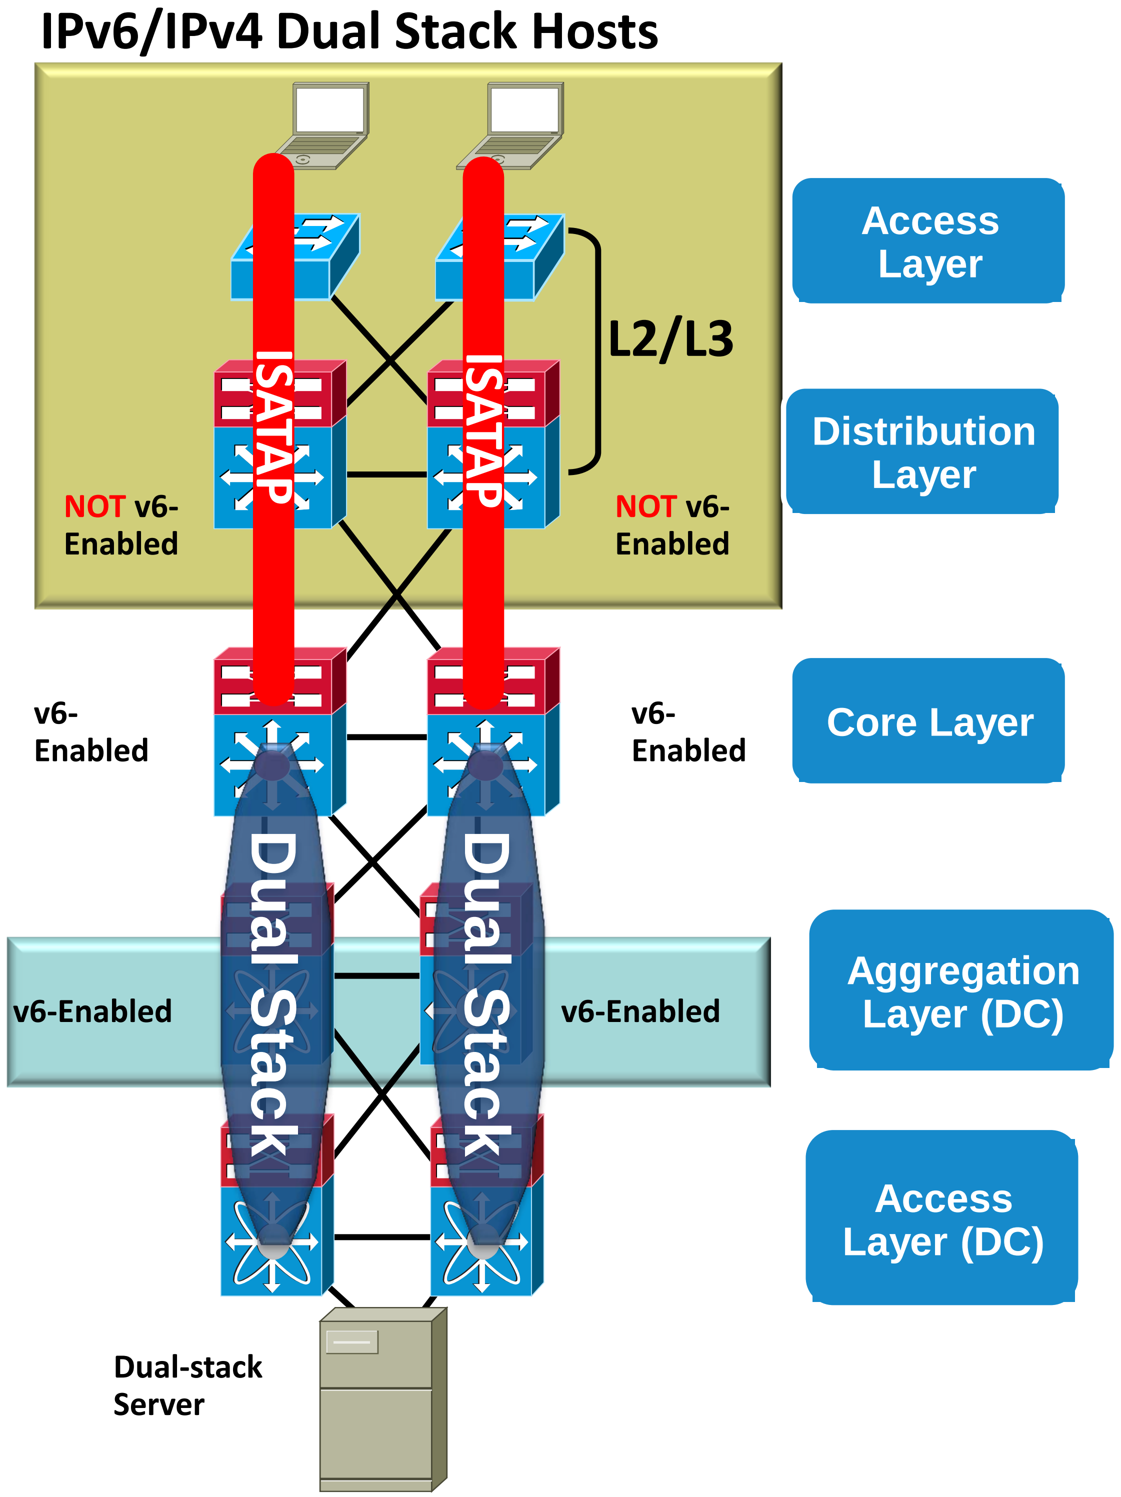

Campus Hybrid Model 1

QoS

1. Classification and marking of IPv6 is done on the egress interfaces on the core layer switches because packets have been tunneled until this point—QoS policies for classification and marking cannot be applied to the ISATAP tunnels on ingress
2. The classified and marked IPv6 packets can now be examined by upstream switches (e.g. aggregation layer switches) and the appropriate QoS policies can be applied on ingress. These polices may include trust (ingress), policing (ingress) and queuing (egress)

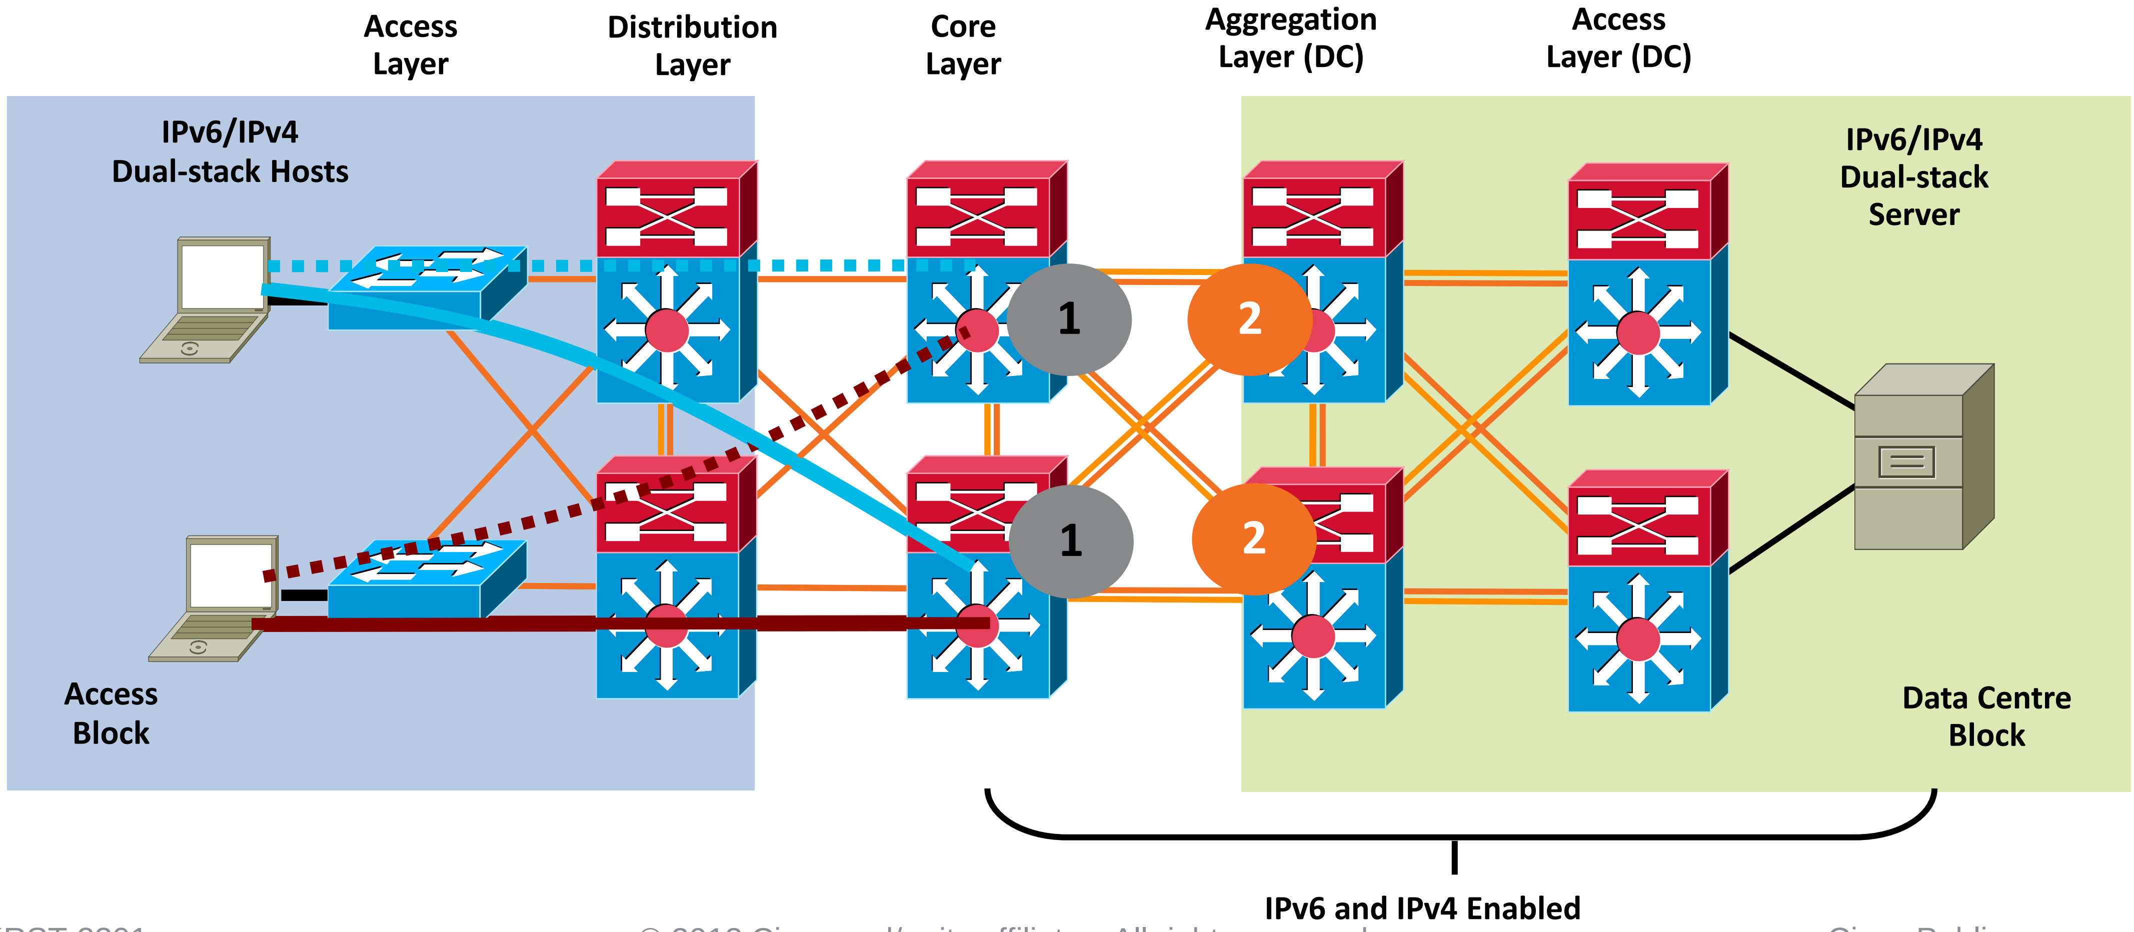

IPv6 ISATAP Implementation

ISATAP Host Considerations

- ISATAP is available on Windows XP, Windows 2003, Vista/W7/W8/Server 2008/2012, port for Linux
- If Windows host does not detect IPv6 capabilities on the physical interface then an effort to use ISATAP is started
- Can learn of ISATAP routers via DNS “A” record lookup “isatap” or via static configuration
  - If DNS is used then Host/Subnet mapping to certain tunnels cannot be accomplished due to the lack of naming flexibility in ISATAP
  - Two or more ISATAP routers can be added to DNS and ISATAP will determine which one to use and also fail to the other one upon failure of first entry
  - If DNS zoning is used within the enterprise then ISATAP entries for different routers can be used in each zone
- In the presented design the static configuration option is used to ensure each host is associated with the correct ISATAP tunnel
- Can conditionally set the ISATAP router per host based on subnet, userid, department and possibly other parameters such as role

Highly Available ISATAP Design

Topology

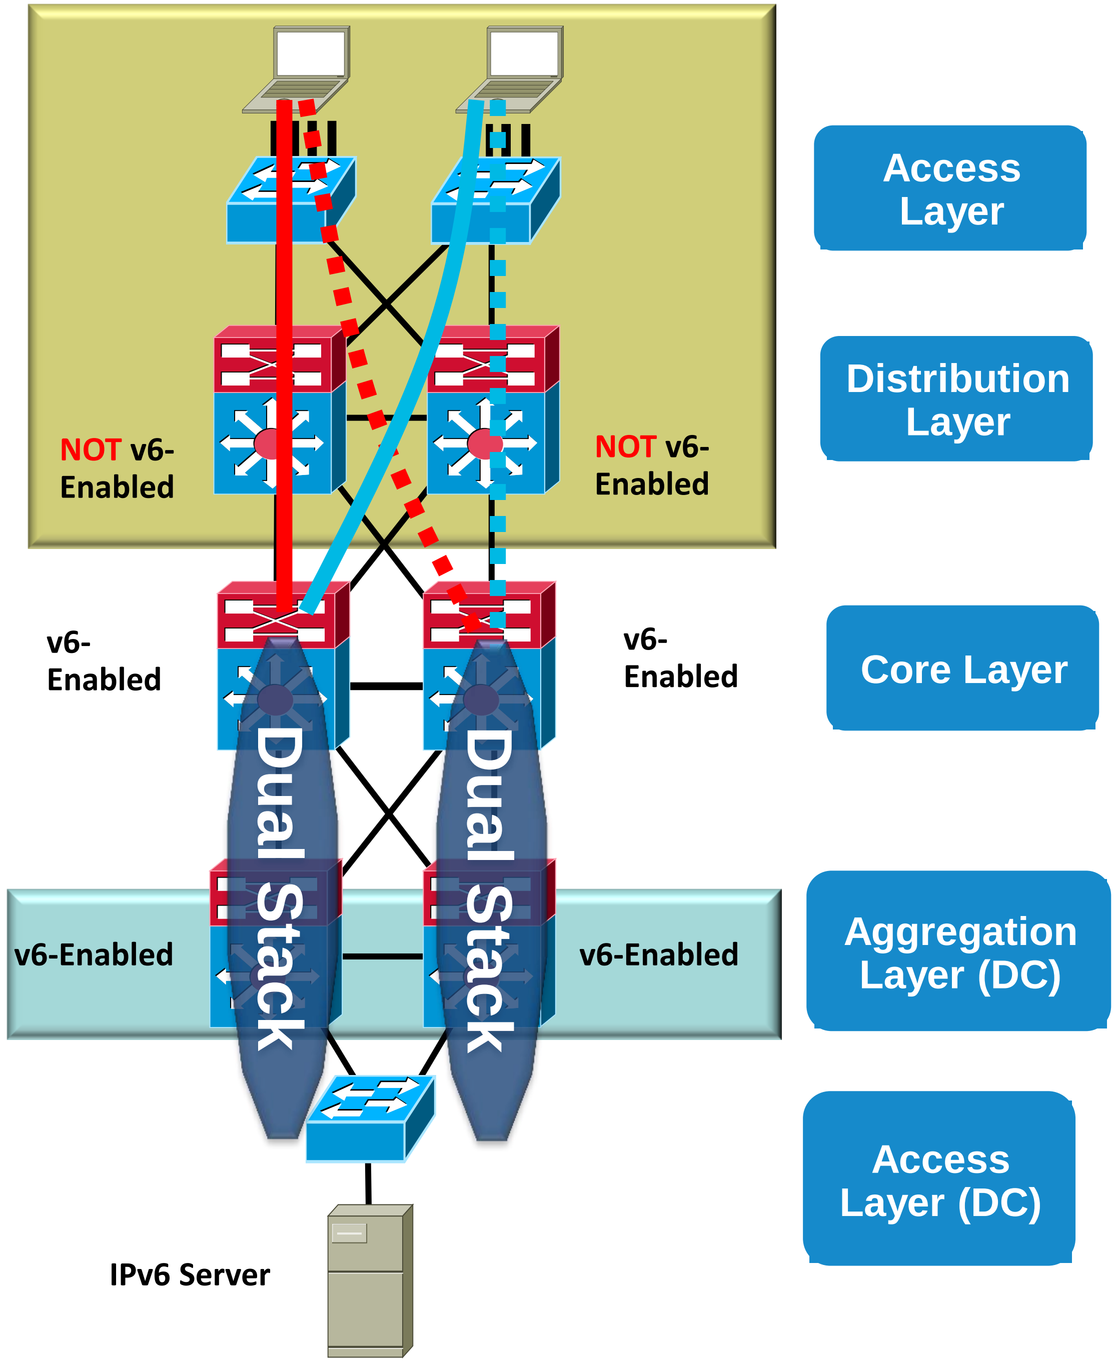

- ISATAP tunnels from PCs in access layer to core switches
- Redundant tunnels to core or service block
- Use IGP to prefer one core switch over another (both v4 and v6 routes)—deterministic
- Preference is important due to the requirement to have traffic (IPv4/IPv6) route to the same interface (tunnel)
- Works like Anycast-RP with IPmc ☺

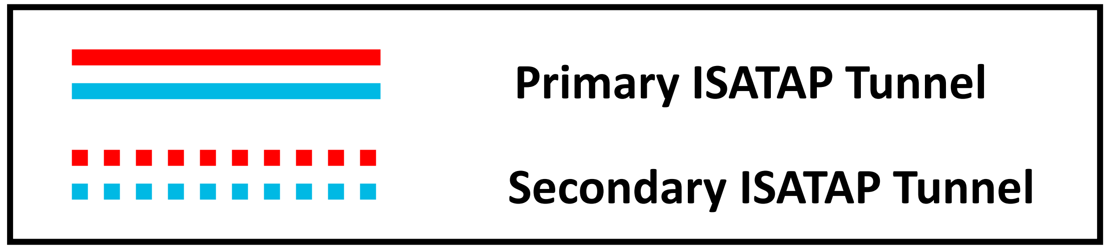

IPv6 Campus ISATAP Configuration

Redundant Tunnels

ISATAP Primary

```
interface Tunnel2
 ipv6 address 2001:DB8:CAFE:2::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 tunnel source Loopback2
 tunnel mode ipv6ip isatap
!
interface Tunnel3
 ipv6 address 2001:DB8:CAFE:3::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 tunnel source Loopback3
 tunnel mode ipv6ip isatap
!
interface Loopback2
 description Tunnel source for ISATAP-VLAN2
 ip address 10.122.10.102 255.255.255.255
!
interface Loopback3
 description Tunnel source for ISATAP-VLAN3
 ip address 10.122.10.103 255.255.255.255
```

ISATAP Secondary

```
interface Tunnel2
 ipv6 address 2001:DB8:CAFE:2::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 ipv6 ospf cost 10
 tunnel source Loopback2
 tunnel mode ipv6ip isatap
!
interface Tunnel3
 ipv6 address 2001:DB8:CAFE:3::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 ipv6 ospf cost 10
 tunnel source Loopback3
 tunnel mode ipv6ip isatap
!
interface Loopback2
 ip address 10.122.10.102 255.255.255.255
 delay 1000
!
interface Loopback3
 ip address 10.122.10.103 255.255.255.255
 delay 1000
```

IPv6 Campus ISATAP Configuration

IPv4 and IPv6 Routing—Options

ISATAP Secondary—Bandwidth adjustment

```
interface Loopback2
 ip address 10.122.10.102 255.255.255.255
 delay 1000
```

ISATAP Primary—Longest-match adjustment

```
interface Loopback2
 ip address 10.122.10.102 255.255.255.255
```

ISATAP Secondary—Longest-match adjustment

```
interface Loopback2
 ip address 10.122.10.102 255.255.255.254
```

IPv4—EIGRP

```
router eigrp 10
 eigrp router-id 10.122.10.3
```

IPv6—OSPFv3

```
ipv6 router ospf 1
 router-id 10.122.10.3
```

- To influence IPv4 routing to prefer one ISATAP tunnel source over another—alter delay/cost or mask length
- Lower timers (timers spf, hello/hold, dead) to reduce convergence times
- Use recommended summarisation and/or use of stubs to reduce routes and convergence times

Set RID to ensure redundant loopback addresses do not cause duplicate RID issues

Distribution Layer Routes

Primary/Secondary Paths to ISATAP Tunnel Sources

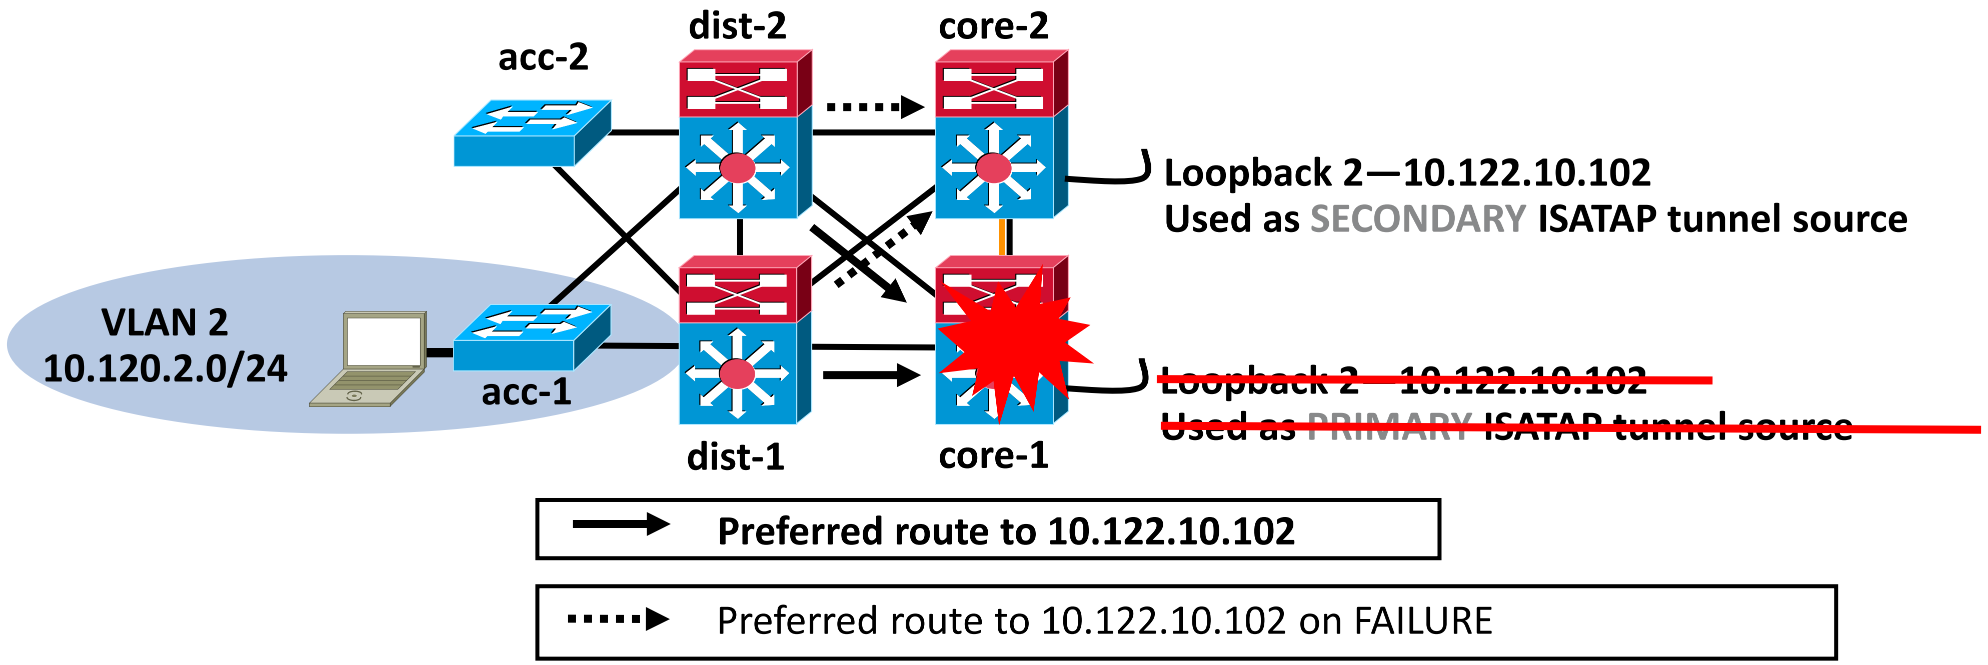

Before Failure

```
dist-1#show ip route | b 10.122.10.102/32
D       10.122.10.102/32 [90/130816] via 10.122.0.41, 00:09:23, GigabitEthernet1/0/27
```

After Failure

```
dist-1#show ip route | b 10.122.10.102/32
D       10.122.10.102/32 [90/258816] via 10.122.0.49, 00:00:08, GigabitEthernet1/0/28
```

IPv6 Campus ISATAP Configuration

ISATAP Client Configuration

Windows XP/Vista/W7/W8 Host

```
C:\>netsh int ipv6 isatap set router 10.122.10.103
Ok.
```

```
interface Tunnel3
 ipv6 address 2001:DB8:CAFE:3::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 eigrp 10
 tunnel source Loopback3
 tunnel mode ipv6ip isatap
!
interface Loopback3
 description Tunnel source for ISATAP-VLAN3
 ip address 10.122.10.103 255.255.255.255
```

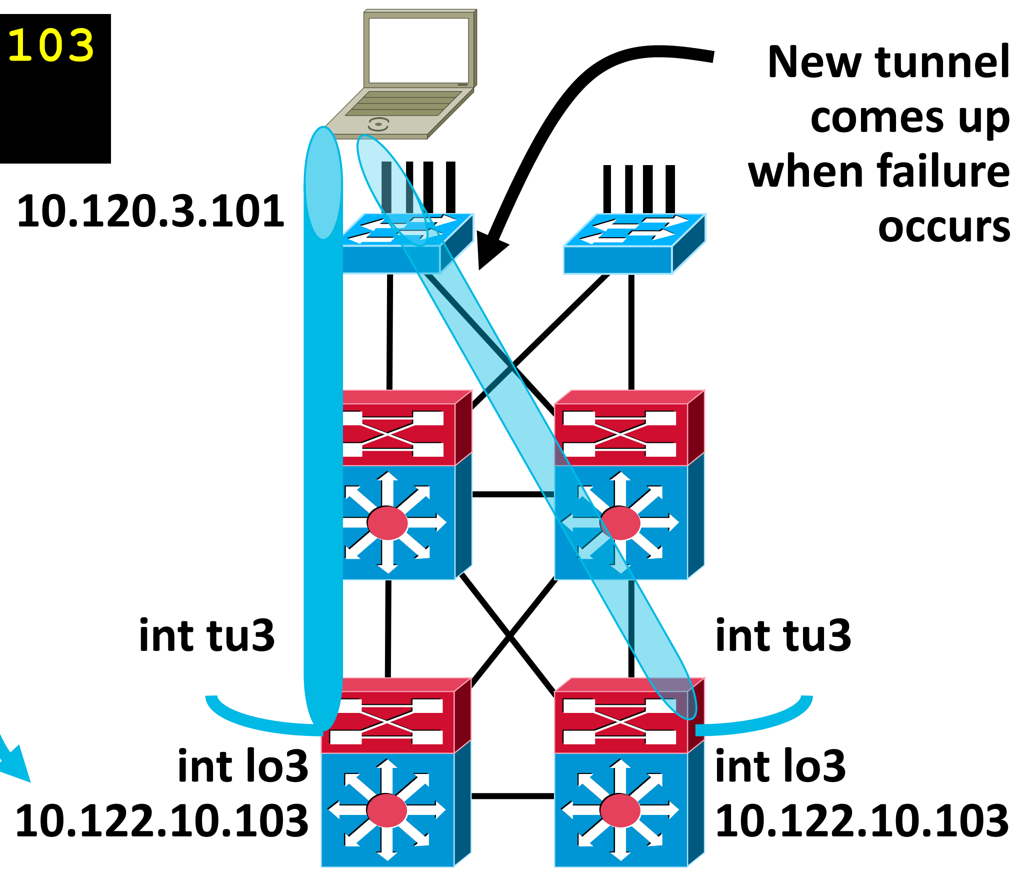


```
Tunnel adapter Automatic Tunneling Pseudo-Interface:
Connection-specific DNS Suffix  . :
IP Address. . . . . . . . . . . . : 2001:db8:cafe:3:0:5efe:10.120.3.101
IP Address. . . . . . . . . . . . : fe80::5efe:10.120.3.101%2
Default Gateway . . . . . . . . . : fe80::5efe:10.122.10.103%2
```

Campus Hybrid Model 1

QoS Configuration Sample—Core Layer

```
mls ipv6 acl compress address unicast
mls qos
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
   set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
   set dscp af21
  class class-default
   set dscp default
!
ipv6 access-list BULK-APPS
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
 permit tcp any any eq telnet
 permit tcp any any eq 22
```

```
ipv6 access-list BULK-APPS
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
 permit tcp any any eq telnet
 permit tcp any any eq 22
!
interface GigabitEthernet2/1
 description to 6k-agg-1
 mls qos trust dscp
 service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/2
 description to 6k-agg-2
 mls qos trust dscp
 service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/3
 description to 6k-core-1
 mls qos trust dscp
 service-policy output IPv6-ISATAP-MARK
```

Campus IPv6 Deployment Options

IPv6 Service Block—Rapid Deployment/Pilot

- Provides ability to rapidly deploy IPv6 services without touching existing network
- Provides tight control of where IPv6 is deployed and where the traffic flows (maintain separation of groups/locations)
- Get lots of operational experience with limited impact to existing environment – Ideal for Pilot
- Similar challenges as Hybrid Model – Lots of tunnelling
- Configurations are very similar to the Hybrid Model
  - ISATAP tunnels from PCs in access layer to service block switches (instead of core layer—Hybrid)
- 1) Leverage existing ISP block for both IPv4 and IPv6 access
- 2) Use dedicated ISP connection just for IPv6—Can use IOS Zone FW or ASA

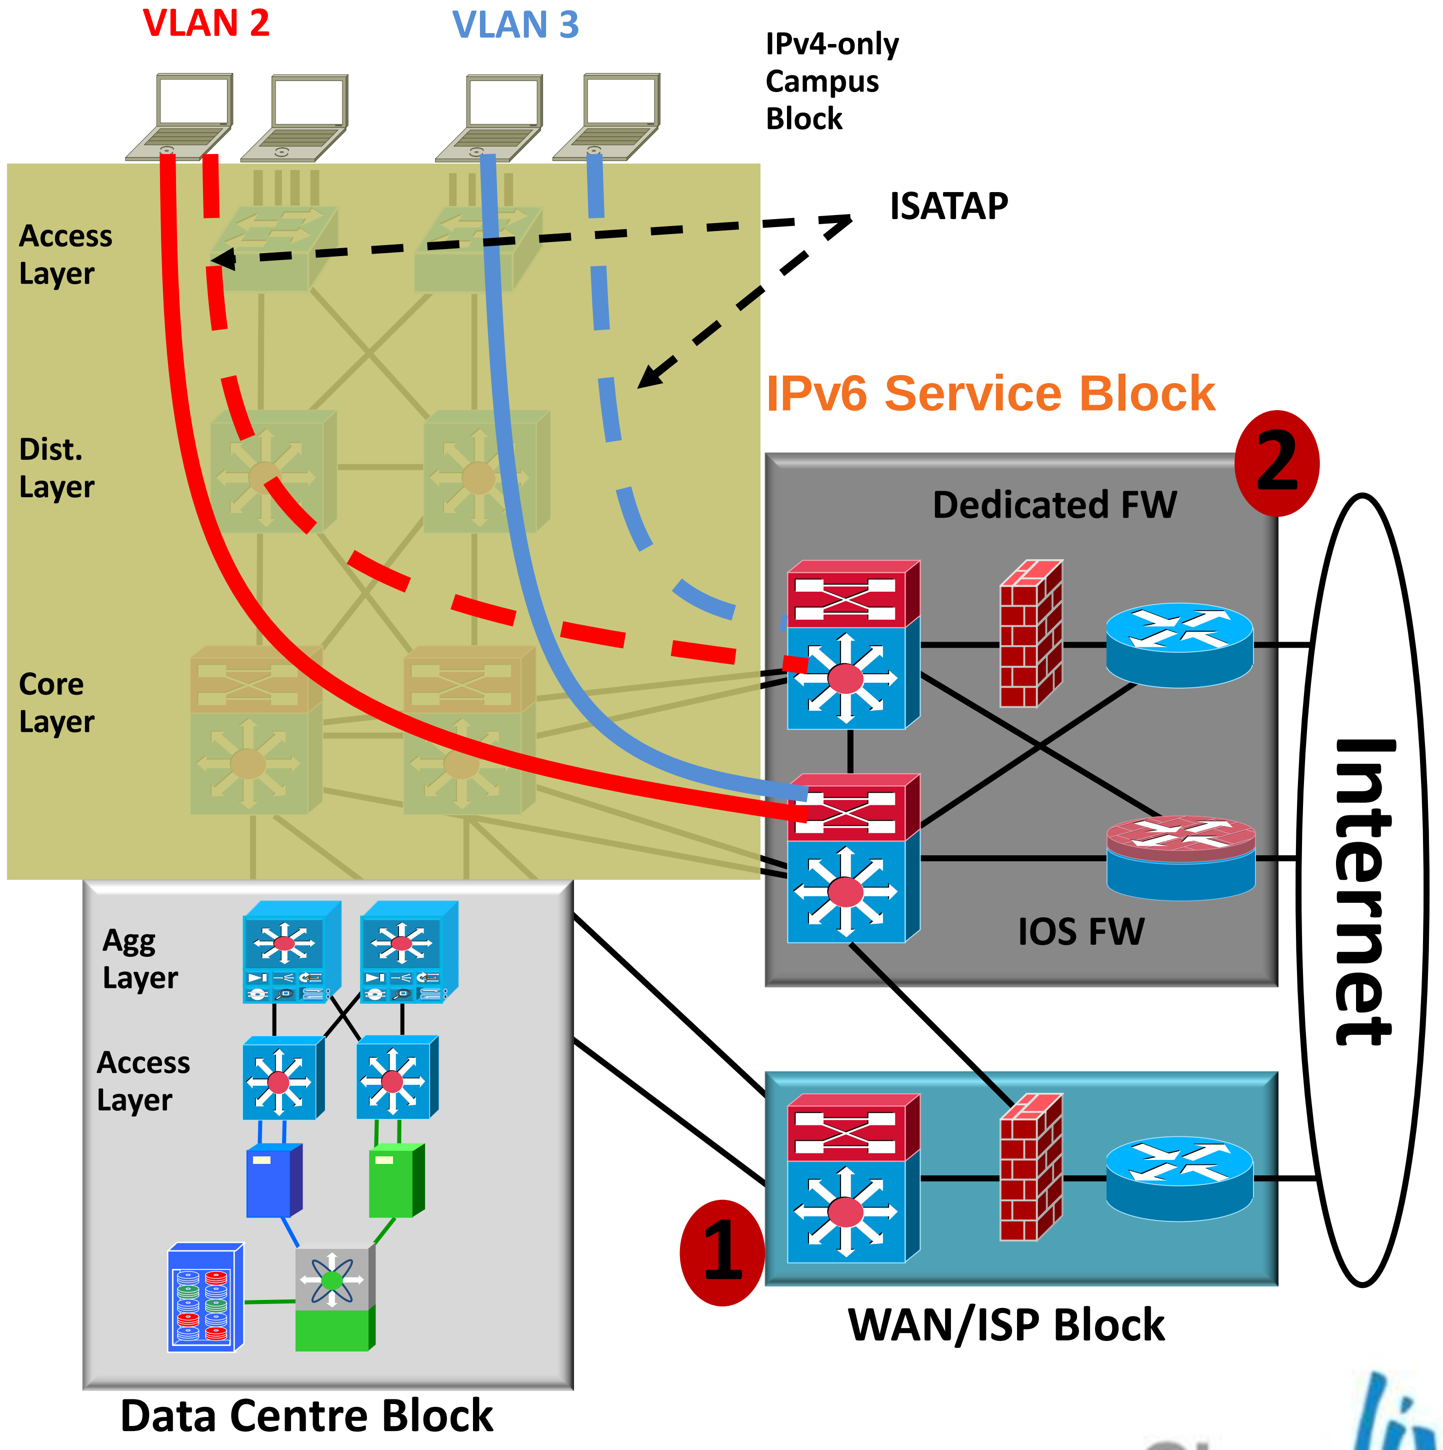

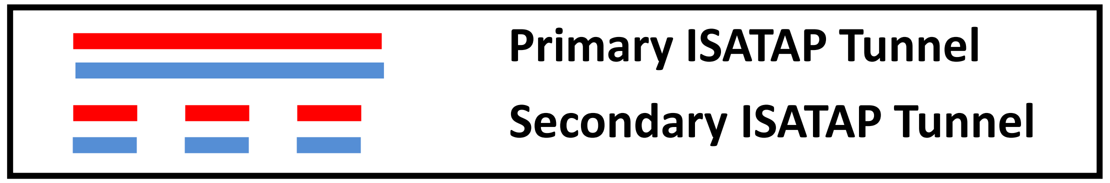

Data Centre & Internet Edge

- http://www.cisco.com/en/US/docs/solutions/Enterprise/Borderless_Networks/Internet_Edge/InternetEdgeIPv6.html

IPv6 Data Centre Integration

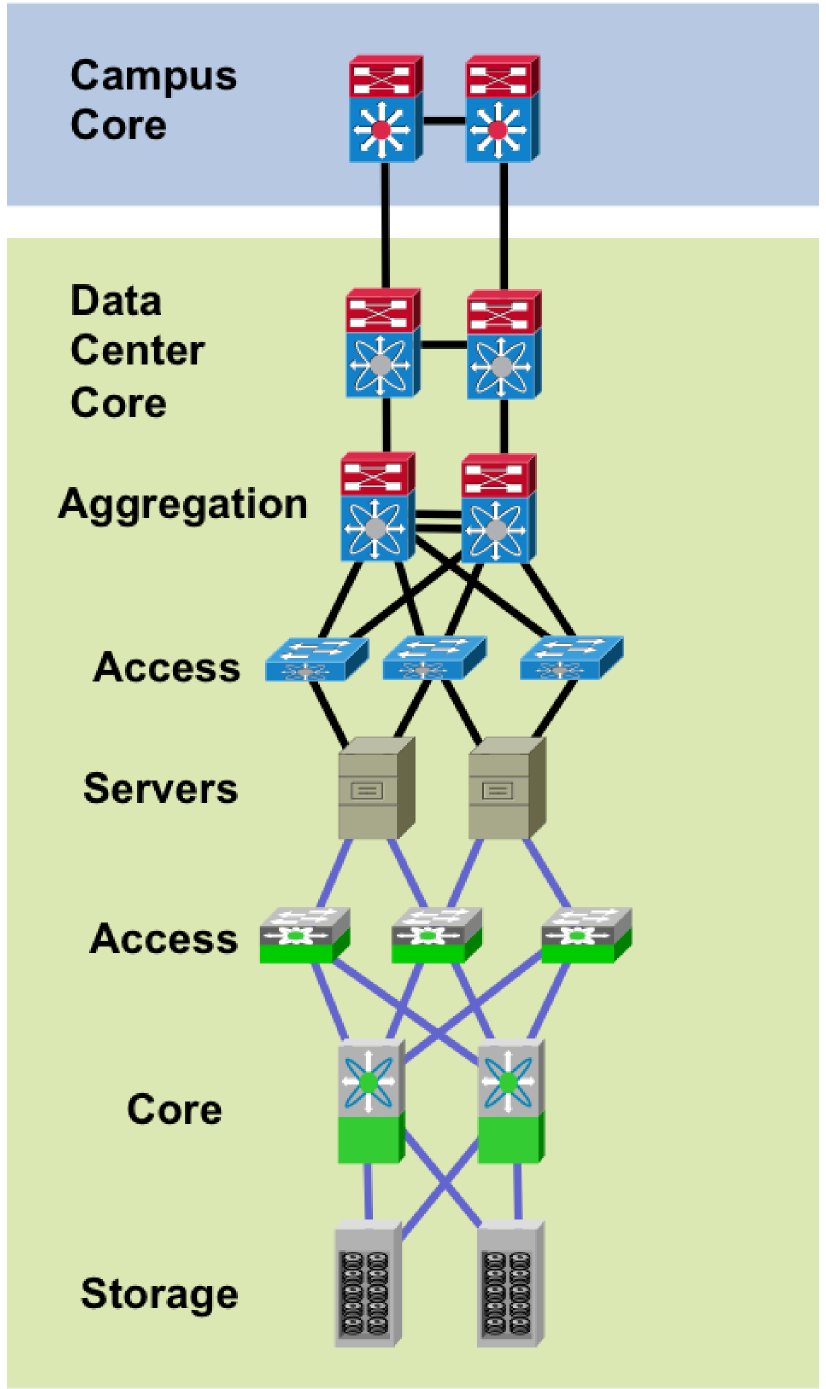

- Routing and switch concepts are no different with IPv6 in the DC
- More stuff going on in DC like: SAN, L4-7, every kind of virtualisation, multi-site, L2 extension, private/hybrid cloud, etc.. and L3)
- Internet-facing Data Centre = Internet Edge – many of the same design requirements, perhaps at different scale
- Large scale DC deployments (MSDC, SPDC, large enterprise) can die a horrible death using default IPv6 neighbour timers (think ARP meltdown)
- Do enough to get apps alive and available to IPv6 clients and expand from there

NDP Scaling Issues in the DC

- Large DCs with very dense hosts populations can cause severe performance problems on the control plane of switches due to IPv4 and IPv6 'control' traffic (i.e. ARP and Neighbour Discovery Protocol [NDP] – RFC4861)
- In IPv4 we have used ARP timers and other techniques to reduce the impact
- In IPv6 we have similar issues that need to be dealt with
- Lessons learned from production deployments from some of our largest accounts and excellent work by Swami Venkataraman & Ming Zhang (most of the NDP scaling content is based on their paper)

NDP Theory in 6 Bullets

Reference

- Neighbour Unreachability Detection (NUD) – Used by NDP to detect loss of nodes and/or gateways
- Neighbour Cache (== ARP cache) can be in one of 5 states: INCOMPLETE, REACHABLE, STALE, DELAY, PROBE
- Default gateway interface config has ‘AdvReachableTime’ of 0 (not specified) in RAs which indicates to host to use ‘BaseReachableTime’ which is 30 seconds
- Host sends NS to get link layer address of gateway(s) > Gateway sends NA back, marks cache for host as STALE, Host marks entry as REACH
- 5 seconds later, gateway sends NS for host > NA comes back and entry marked as REACH
- While traffic is being forwarded to host, NUD runs for AdvReachableTime [+/- 30 seconds] (In IPv4 ARP would refresh every 4 hours or based on adjusted timers)

*There is more to this than I listed and some configuration scenarios change this behaviour

NDP Scaling – Scenario 1

Reference

- Based on NDP and NUD mechanics, we have a lot of NS/NA going on between hosts and gateways to ensure connectivity
- During link failures and/or HSRP flaps on large host-populated networks, we can have SIGNIFICANT NDP activity
- Switch perspective: Switch 1 and Switch 2 (Agg layer DC switches) send NS to 2000 hosts every 30 seconds <> +/- 2000 hosts reply with NA == 4000 NS/NA sent/received by switch for one VLAN
- Host perspective: +/- 2000 hosts (just for one VLAN) send NS to Switch 1 and Switch 2 <> Both switches reply with NA == 4000 NS/NA sent/received by switch
- We are at 8000 NS/NA processed by control plane (CPU) by switches every 30 seconds
- **`ipv6 nd reachable-time`** *`value-in-milliseconds`*

NDP Scaling – Scenario 2

Reference

- Unsolicited NA – NA messages sent with Solicit Flag unset to indicate changes to its link-layer address – triggered post DAD
- Not used as a reliable means of updating cache
- Host A comes online, sends unsolicited NA to FF02::1 (all nodes) – Switch 1 and 2 receive NA, but by default, ignores NA
- Unsolicited NA Glean feature allows for creation/update of STALE entry AND /128 entry in HW by switch when unsolicited NA received
- This GREATLY reduces packet loss and CPU spike when failover happens as entry is already in HW
- **`ipv6 nd na glean`**

NDP Scaling – Scenario 3

Reference

- Switch will wait 4 hours to flush an entry after it goes into STALE
- Scavenge and Refresh timer can now be updated via the ND cache expiry
- Test with various ranges of time, but it is recommended that the expiration be at least 1 hour
- **`ipv6 nd cache expire`** *`expire-time-in-seconds`* **`[refresh]`**
  - **`ipv6 nd cache expire 3600`**

NDP Scaling – Scenario 4

Reference

- NUD defaults to send 3 NS packets every 1 second
- This may hurt with large boot storms or STP events
- Exponential NUD allows for tuning of retransmit rate of NUD (initial resolution is still 3 NS packets per second interval)
- **`ipv6 nd nud retry`** *`base interval-in-milliseconds maximum-attempts`*
- The retransmit interval is calculated with the following formula - tm^n
  – t = Time Interval
  – m = Base (1,2,3)
  – n = Current NS number (Where first NS is 0)
- ipv6 nd nud retry 1 1000 3 will give a fixed interval of 1 second and 3 retransmits
- ipv6 nd nud retry 2 1000 3 will give a retransmit interval of 1, 2, 4 seconds
- ipv6 nd nud retry 3 1000 4 will give a retransmit interval of 1, 3, 9, 27 seconds

NDP Scaling – Scenario 5

Reference

- Glean adjacency – Switch is missing MAC rewrite information for next-hop that is directly connected
- When MAC address is not present for destination, the packets will go to CPU to trigger an NDP request
- In a busy and volatile virtualised environment (like with large VDI farm with lots of VM coming and going), this CPU activity can hurt
- The MLS Glean Rate-Limiter can rate limit the number of packets that are sent to CPU
- **`mls rate-limit unicast cef glean <pps> <burst>`**
  - **`mls rate-limit unicast cef glean 200 10`**

NDP Scaling Commands - Summary

- You must test values of commands to find the right balance in your environment
- Not every deployment needs the same ‘aggressive’ changes
- There is no one-size-fits-all

- NUD Reachable Time: **ipv6 nd reachable-time** *time-in-milliseconds*
  #Some setting as high as 45 minutes (2700000 msec)

- Unsolicited NA Glean: **ipv6 nd na glean**

- Scavenge and Refresh Timer: **ipv6 nd cache expire** *time-in-seconds*
- NUD runs when cache timer expires: **ipv6 nd cache expire 3600 refresh**

- **ipv6 nd nud retry** *base interval-in-milliseconds maximum-attempts*

- Glean rate limiter: **mls rate-limit unicast cef glean** *pps burst*

IPv6 Configuration on Nexus

- IP-is-IP – minor syntax changes based on different platforms between campus & data centre
- Check for the features you need, platform support, performance capabilities
- Same stuff you do for any new platform you invest in

```
interface Vlan114
  no shutdown
  description Outside FW VLAN
  ipv6 address 2001:0db8:cafe:0114::0002/64
  hsrp version 2
  hsrp 114 ipv6
    preempt delay minimum 180
    timers  1  3
    ip autoconfig
```

iSCSI/VRRP for IPv6

- Same configuration requirements and operation as with IPv4
- Can use automatic preemption—configure VR address to be the same as physical interface of “primary”
- Host-side HA uses NIC teaming (see slides for NIC teaming)
- Support for iSCSI with IPsec

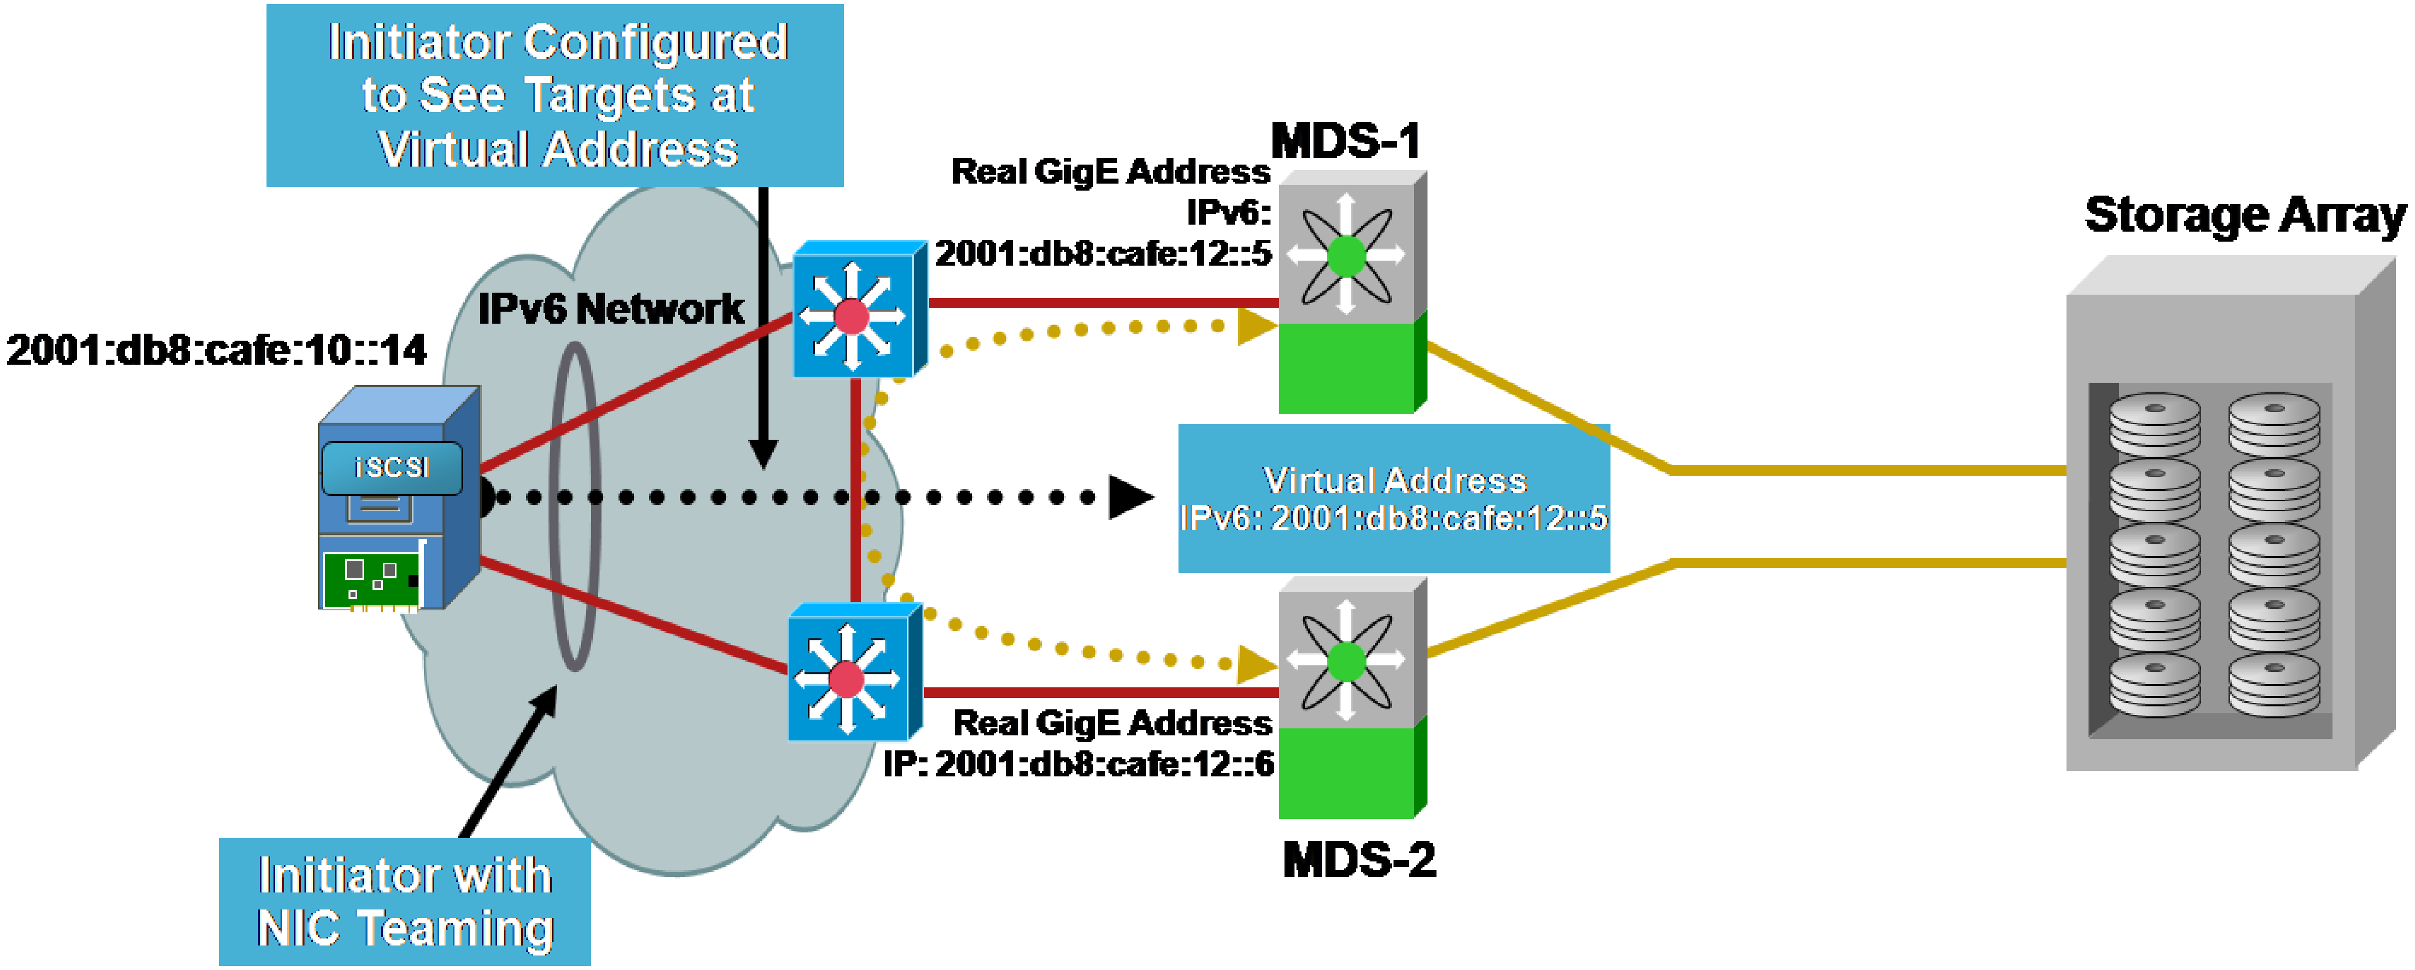

iSCSI IPv6 Example—MDS

Initiator/Target

```
iscsi virtual-target name iscsi-atto-target
  pWWN 21:00:00:10:86:10:46:9c
  initiator iqn.1991-05.com.microsoft:w2k8-svr-01.cisco.com permit
iscsi initiator name iqn.1991-05.com.microsoft:w2k8-svr-01.cisco.com
  static pWWN 24:01:00:0d:ec:24:7c:42
  vsan 1
zone default-zone permit vsan 1
zone name iscsi-zone vsan 1
    member symbolic-nodename iqn.1991-05.com.microsoft:w2k8-svr-01.cisco.com
    member pwwn 21:00:00:10:86:10:46:9c
    member pwwn 24:01:00:0d:ec:24:7c:42
    member symbolic-nodename iscsi-atto-target
zone name Generic vsan 1
    member pwwn 21:00:00:10:86:10:46:9c
zoneset name iscsi_zoneset vsan 1
    member iscsi-zone
zoneset name Generic vsan 1
    member Generic
```

iSCSI/VRRP IPv6 Example—MDS

Interface

MDS-1

```
interface GigabitEthernet2/1
  ipv6 address 2001:db8:cafe:12::5/64
  no shutdown
  vrrp ipv6 1
    address 2001:db8:cafe:12::5
    no shutdown
```

MDS-2

```
interface GigabitEthernet2/1
  ipv6 address 2001:db8:cafe:12::6/64
  no shutdown
  vrrp ipv6 1
    address 2001:db8:cafe:12::5
    no shutdown
```

```
mds-1# show vrrp ipv6 vr 1
Interface  VR IpVersion Pri   Time Pre State   VR IP addr
------------------------------------------------------------------
  GigE2/1   1   IPv6    255  100cs     master  2001:db8:cafe:12::5
```

```
mds-2# show vrrp ipv6 vr 1
Interface  VR IpVersion Pri   Time Pre State   VR IP addr
------------------------------------------------------------------
  GigE2/1   1   IPv6    100  100cs     backup  2001:db8:cafe:12::5
```

iSCSI Initiator Example—W2K8 IPv6

1


```
iscsi initiator name iqn.1991-05.com.microsoft:w2k8-svr-01.cisco.com
```

2

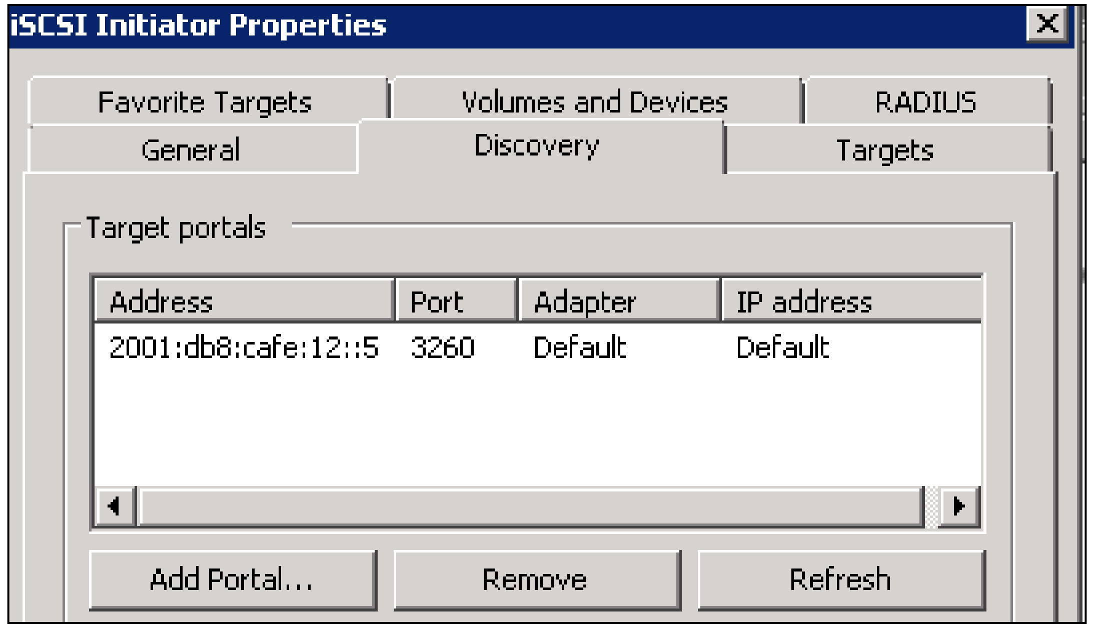


```
interface GigabitEthernet2/1
  ipv6 address 2001:db8:cafe:12::5/64
```

3

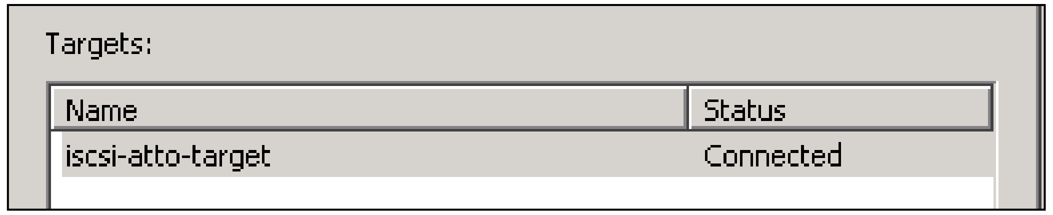


```
mds9216-1# show fcns database vsan 1
VSAN 1:
----------------------------------------------------------------------
FCID      TYPE  PWWN                    (VENDOR)  FC4-TYPE:FEATURE
----------------------------------------------------------------------
0x670400  N     21:00:00:10:86:10:46:9c           scsi-fcp:target
0x670405  N     24:01:00:0d:ec:24:7c:42 (Cisco)   scsi-fcp:init isc..w
```

SAN-OS 3.x—FCIP(v6)

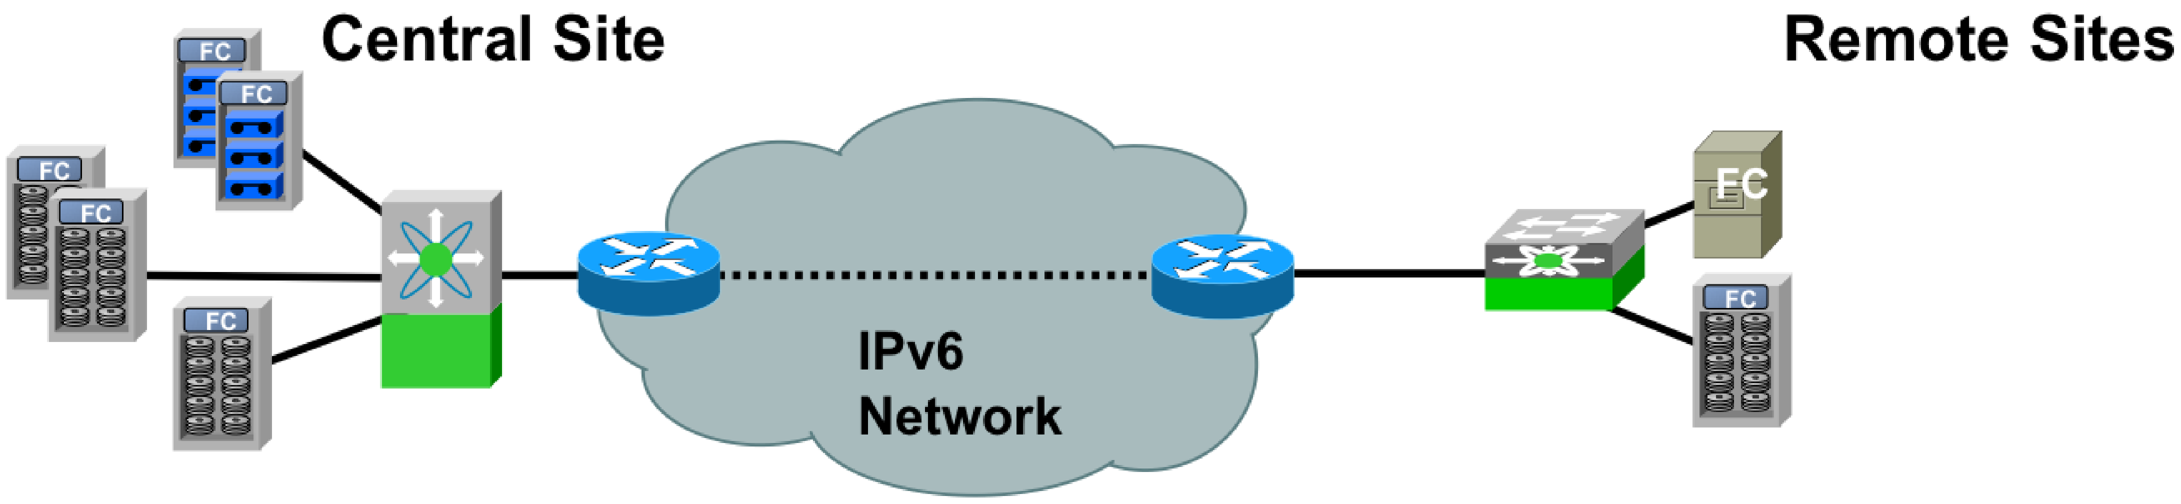


```
fcip profile 100
  ip address 2001:db8:cafe:50::1
  tcp max-bandwidth-mbps 800 min-available-
bandwidth-mbps 500  round-trip-time-us 84
!
interface fcip100
  use-profile 100
  peer-info ipaddr 2001:db8:cafe:50::2
!
interface GigabitEthernet2/2
  ipv6 address 2001:db8:cafe:50::1/64
```

```
fcip profile 100
  ip address 2001:db8:cafe:50::2
  tcp max-bandwidth-mbps 800 min-available-
bandwidth-mbps 500  round-trip-time-us 84
!
interface fcip100
  use-profile 100
  peer-info ipaddr 2001:db8:cafe:50::1
!
interface GigabitEthernet2/2
  ipv6 address 2001:db8:cafe:50::2/64
```

Data Centre NIC Teaming Issue

What Happens If IPv6 Is Unsupported?

Auto-configuration

```
Interface 10: Local Area Connection #VIRTUAL TEAM INTERFACE

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Public     Preferred  29d23h58m41s  6d23h58m41 2001:db8:cafe:10:20d:9dff:fe93:b25d
```

Static configuration

```
netsh interface ipv6> add address "Local Area Connection" 2001:db8:cafe:10::7
Ok.

netsh interface ipv6>sh add
Querying active state...
Interface 10: Local Area Connection
Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Manual     Duplicate      infinite     infinite 2001:db8:cafe:10::7
Public     Preferred  29d23h59m21s  6d23h59m21s 2001:db8:cafe:10:20d:9dff:fe93:b25d
```

Note: Same Issue Applies to Linux

Intel ANS NIC Teaming for IPv6

- Intel IPv6 NIC Q&A—Product support
- http://www.intel.com/support/network/sb/cs-009090.htm
- Intel now supports IPv6 with Express, ALB, and AFT deployments
- Check for Broadcom/HP and other NIC vendors support – most have it now

Interim Hack for Unsupported NICs

- Main issue for NICs with no IPv6 teaming support is DAD—Causes duplicate checks on Team and Physical even though the physical is not used for addressing
- Set DAD on Team interface to "0"—Understand what you are doing ☺
- Microsoft Vista/W7/Server 2008 allows for a command line change to reduce the "DAD transmits" value from 1 to 0
  - `netsh interface ipv6 set interface 19 dadtransmits=0`
- Microsoft Windows 2003—Value is changed via a creation in the registry
  - `\\HKLM\System\CurrentControlSet\Services\Tcpip6\Parameters\Interfaces\(InterfaceGUID)\DupAddrDetectTransmits` - **Value "0"**
- Linux
  - `# sysctl -w net/ipv6/conf/bond0/dad_transmits=0`
  - `net.ipv6.conf.eth0.dad_transmits = 0`

Intel NIC Teaming—IPv6 (Pre Team)

```
Ethernet adapter Local Area Connection 3:
   Connection-specific DNS Suffix  . :
   Autoconfiguration IP Address. . . : 169.254.25.192
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   IP Address. . . . . . . . . . . . : fe80::204:23ff:fec7:b0d7%11
   Default Gateway . . . . . . . . . : fe80::212:d9ff:fe92:de76%11

Ethernet adapter LAN:
   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.89.4.230
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : 2001:db8:cafe:1::2
   IP Address. . . . . . . . . . . . : fe80::204:23ff:fec7:b0d6%12
   Default Gateway . . . . . . . . . : fe80::212:d9ff:fe92:de76%12
```

Intel NIC Teaming—IPv6 (Post Team)

```
Ethernet adapter TEAM-1:
   Connection-specific DNS Suffix  . :
   IP Address. . . . . . . . . . . . : 10.89.4.230
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   IP Address. . . . . . . . . . . . : 2001:db8:cafe:1::2
   IP Address. . . . . . . . . . . . : fe80::204:23ff:fec7:b0d6%13
   Default Gateway . . . . . . . . . : fe80::212:d9ff:fe92:de76%13
```

```
Interface 13: TEAM-1

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Public     Preferred   m11s         4m11s        2001:db8:cafe:1::2
Link       Preferred   infinite     infinite     fe80::204:23ff:fec7:b0d6
```

IPv6 in the Enterprise Data Centre

Biggest Challenges Today

- Application support for IPv6 – Know what you don't know
  - If an application is protocol centric (IPv4):
  - Needs to be rewritten – Probably not going to happen
  - Needs to be translated until it is replaced – We will talk about this next
  - Wait and pressure vendors to move to protocol agnostic framework
- Deployment of translation
  - SLB66 or SLB64
  - Stateful NAT64
  - Apache Reverse Proxy
  - Windows Port Proxy
  - 3rd party proxy solutions
- Network services above L3
  - SLB, SSL-Offload, application monitoring (probes)
  - Application Optimisation
  - High-speed security inspection/perimeter protection

Commonly Deployed IPv6-enabled OS/Apps

Operating Systems

- Windows 7
- Windows Server 2008/R2
- SUSE
- Red Hat
- Ubuntu
- The list goes on

Virtualisation & Applications

- VMware vSphere 4.1
- Microsoft Hyper-V
- Microsoft Exchange 2007 SP1/2010
- Apache/IIS Web Services
- Windows Media Services
- Multiple Line of Business apps

Most commercial applications won't be your problem – it will be the custom/home-grown apps

Common Deployment Models for Data

Pure Dual Stack | Conditional Dual Stack | Translation as a Service

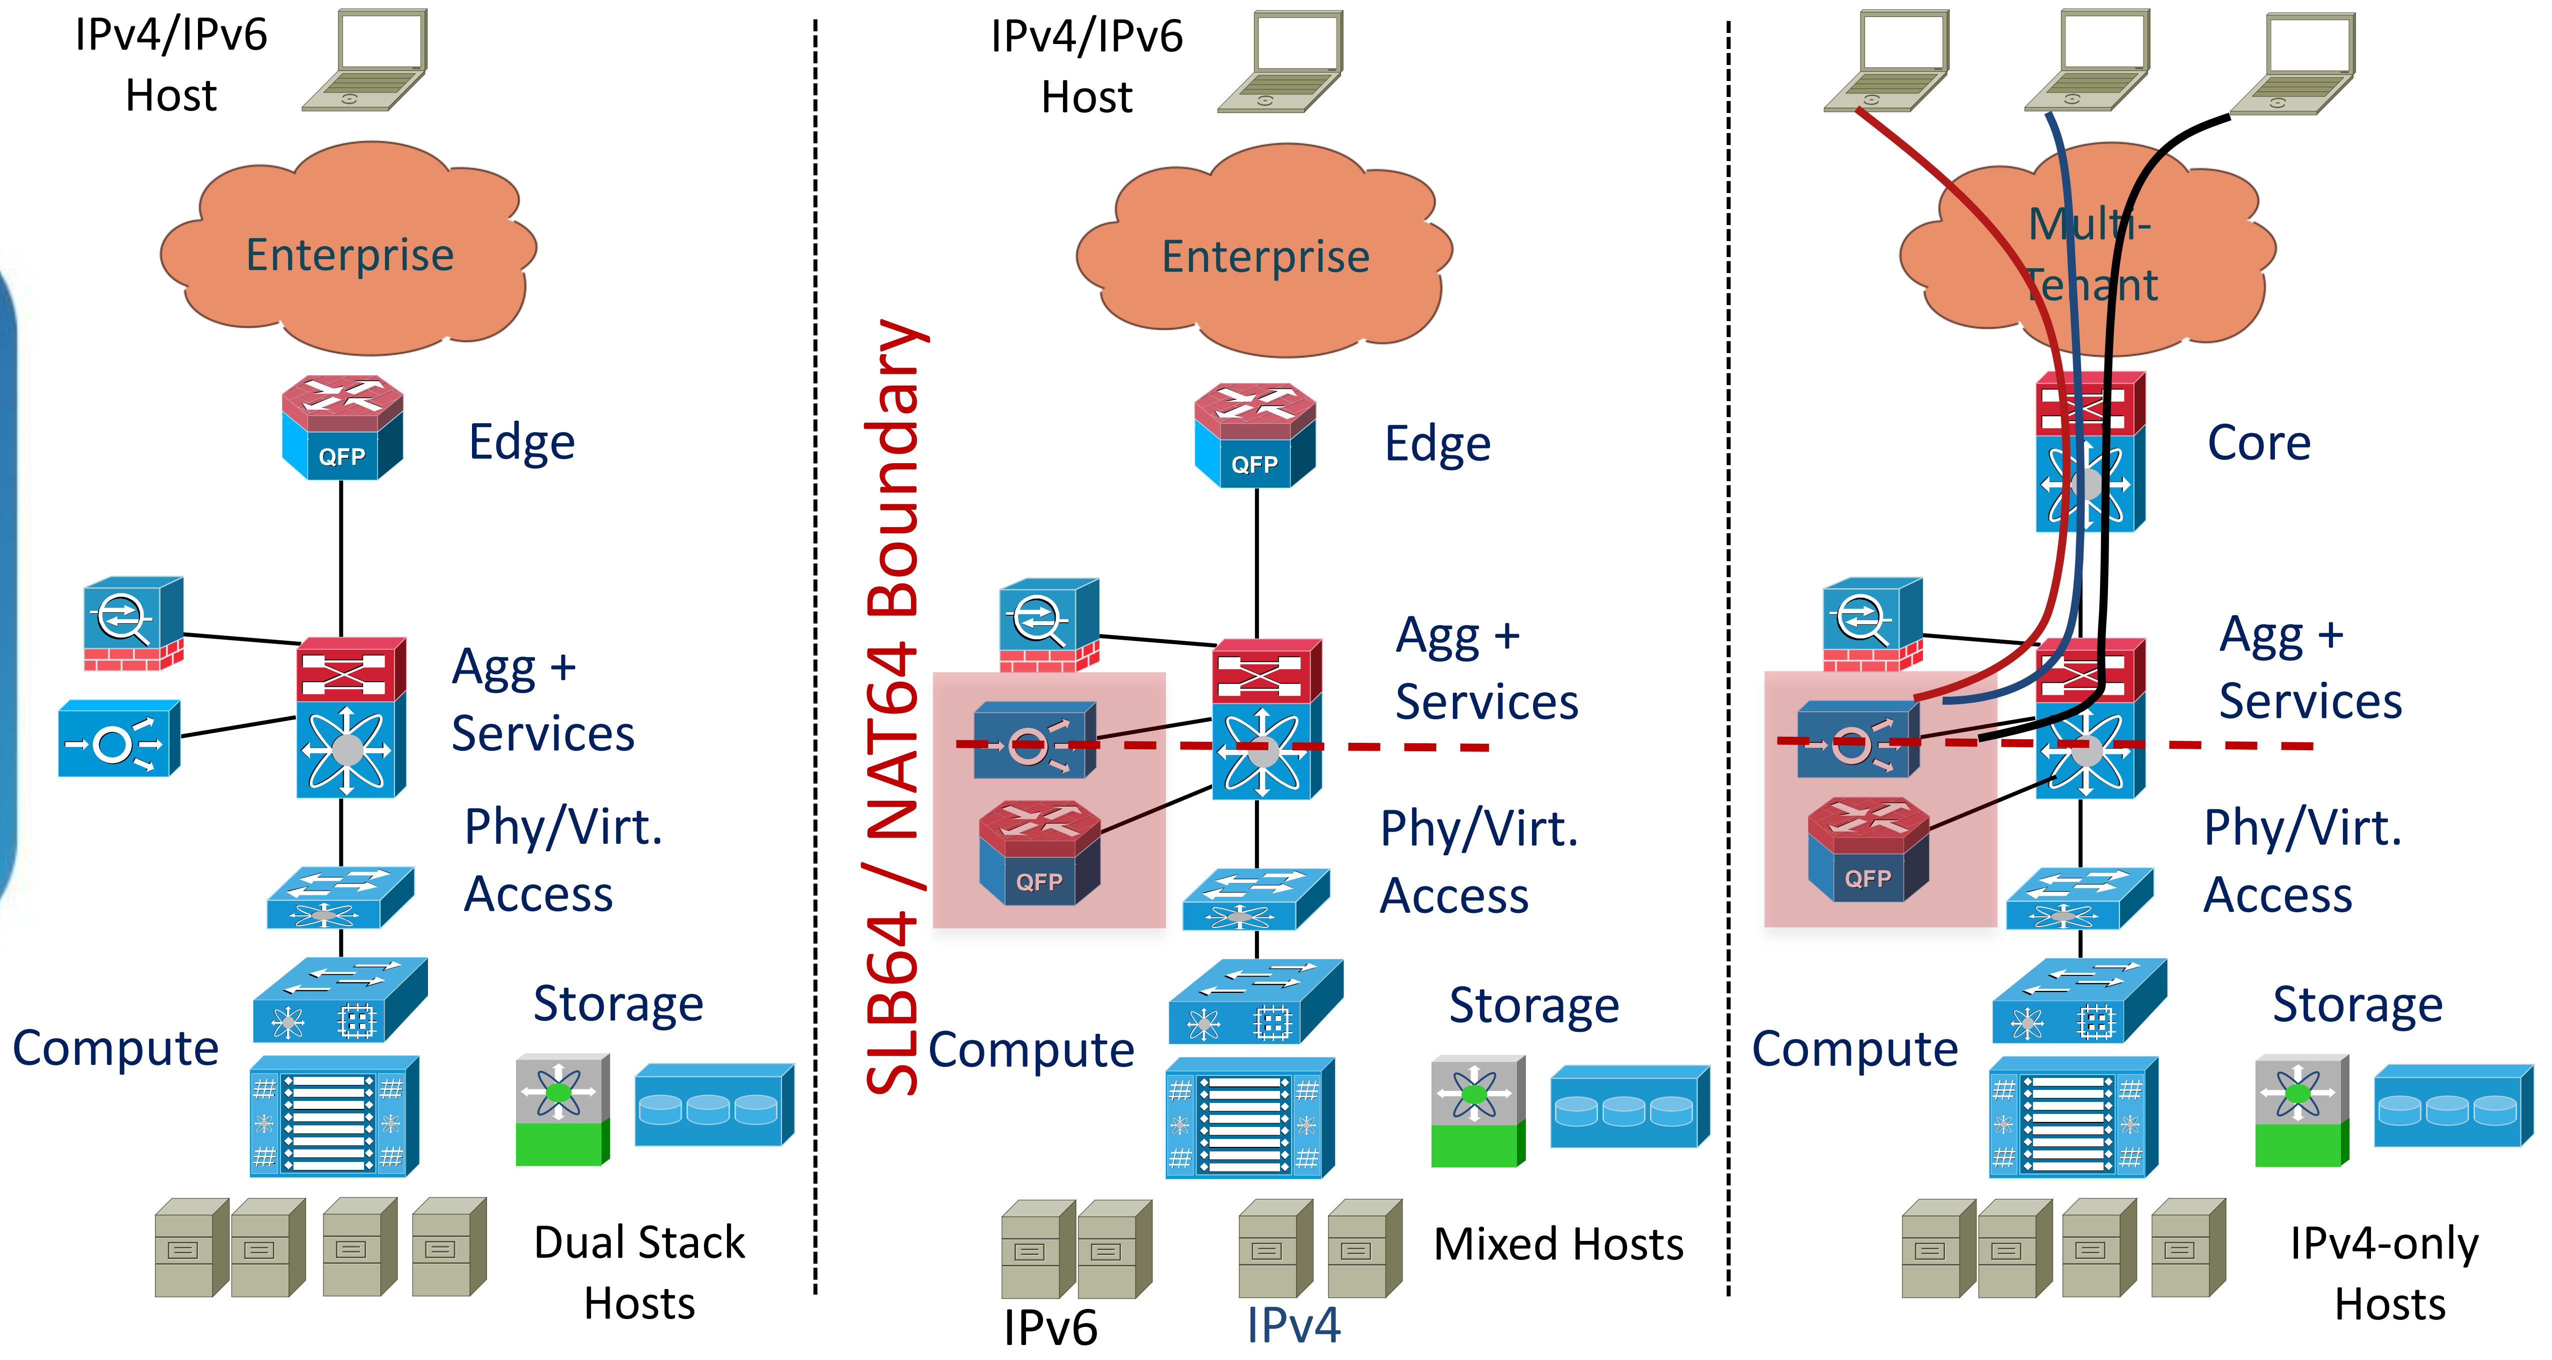

Application/OS Support Drives DC Design Options

Exchange 2003/2007/W2K3

No App IPv6 Support / Limited OS Support

- Leave on IPv4
- Translation won't work – no ALG support for MAPI/RPC, etc…

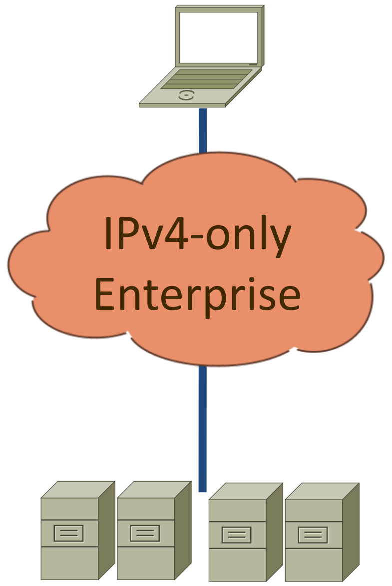

IPv4-only Host

Exchange 2007 SP1/W2K8

Most of App Supports IPv6 / Full OS Support

- Dual stack what you can
- IPv4 legacy components (i.e. MSFT UC)
- Lazy man's method – Translate HTTP/S components

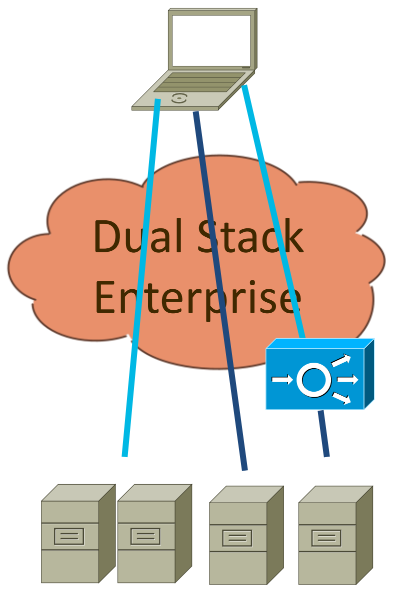

IPv6 IPv4

Exchange 2010/W2K8

Full App Support / Full OS Support

- Dual stack everything
- Lazy man's method – Translate HTTP/S components

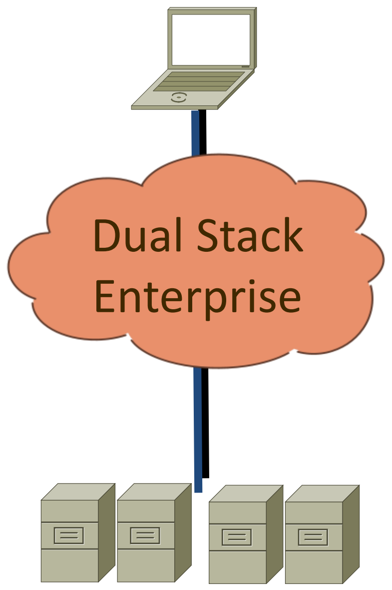

Dual Stack Hosts

IPv6/IPv4 Translation = Crack Addiction

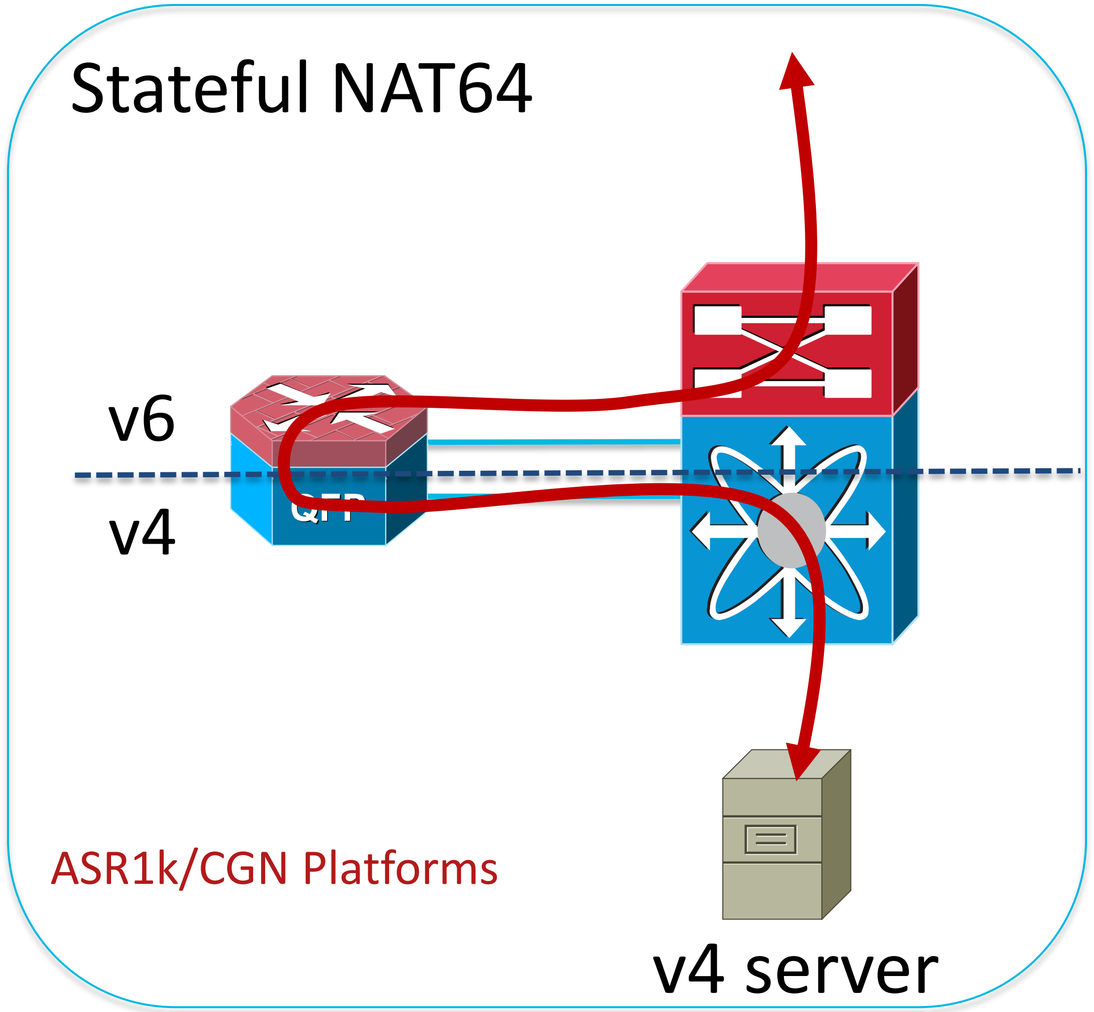

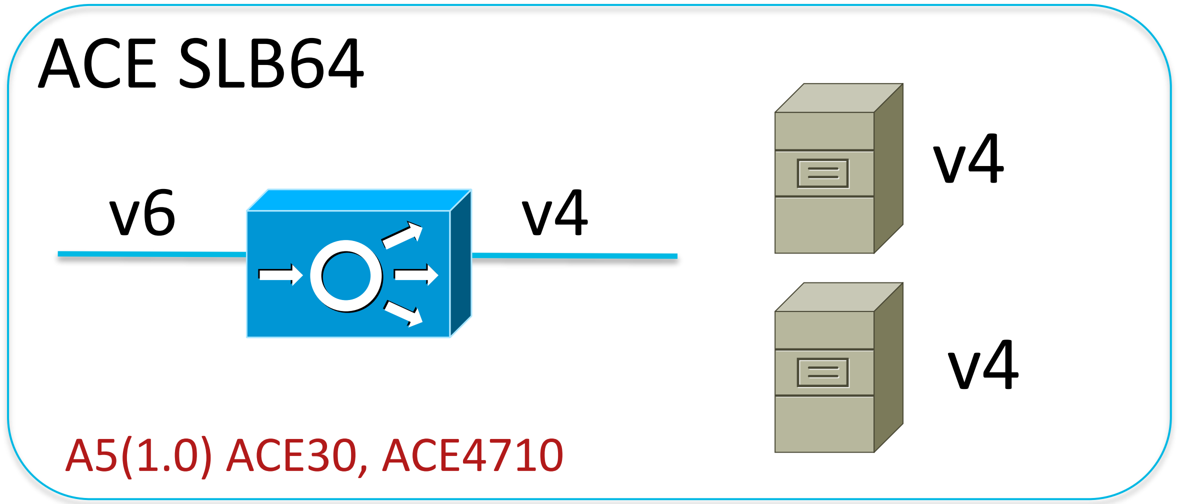

- Others like Apache2 Reverse Proxy, Windows PortProxy, Citrix NetScaler, etc..
- Needed for specific use cases in specific place, NOT a permanent solution
- Put translation as deep into DC/IE as possible (get full visibility of IPv6)

Apache2 Reverse Proxy

Netstat - Client

```
TCP   [2001:db8:beef:10::16]:54640  [2001:db8:cafe:12::5]:80  ESTABLISHED
TCP   [2001:db8:beef:10::16]:54641  [2001:db8:cafe:12::5]:80  ESTABLISHED
```

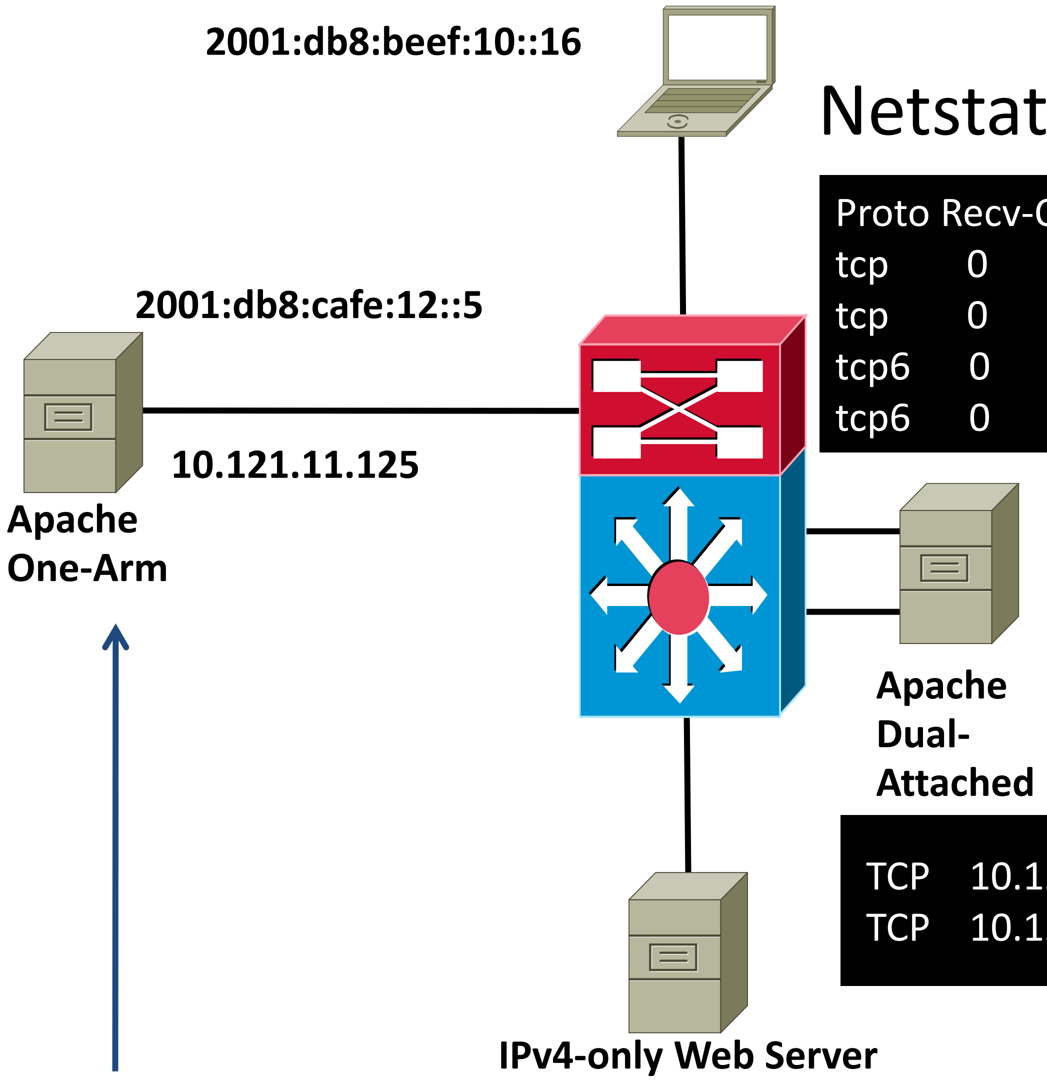

Netstat - Proxy

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	0	10.121.11.125:40475	10.121.11.60:80	ESTABLISHED
tcp	0	0	10.121.11.125:40476	10.121.11.60:80	ESTABLISHED
tcp6	0	0	2001:db8:cafe:12::5:80	2001:db8:beef:10::16:54640	ESTABLISHED
tcp6	0	0	2001:db8:cafe:12::5:80	2001:db8:beef:10::16:54641	ESTABLISHED

Netstat - Server

```
TCP   10.121.11.60:80       10.121.11.125:40475   ESTABLISHED
TCP   10.121.11.60:80       10.121.11.125:40476   ESTABLISHED
```

```
<VirtualHost *:80>
      ProxyPass / http://10.121.11.60:80/
      ProxyPassReverse / http://10.121.11.60:80/
```

Microsoft Windows PortProxy

- Can be treated like an appliance
  - One-arm
  - Dual-attached (better perf)
- Outside traffic comes in on IPv6—PortProxy to v4 (VIP address on ACE)
- Traffic is IPv4 to server

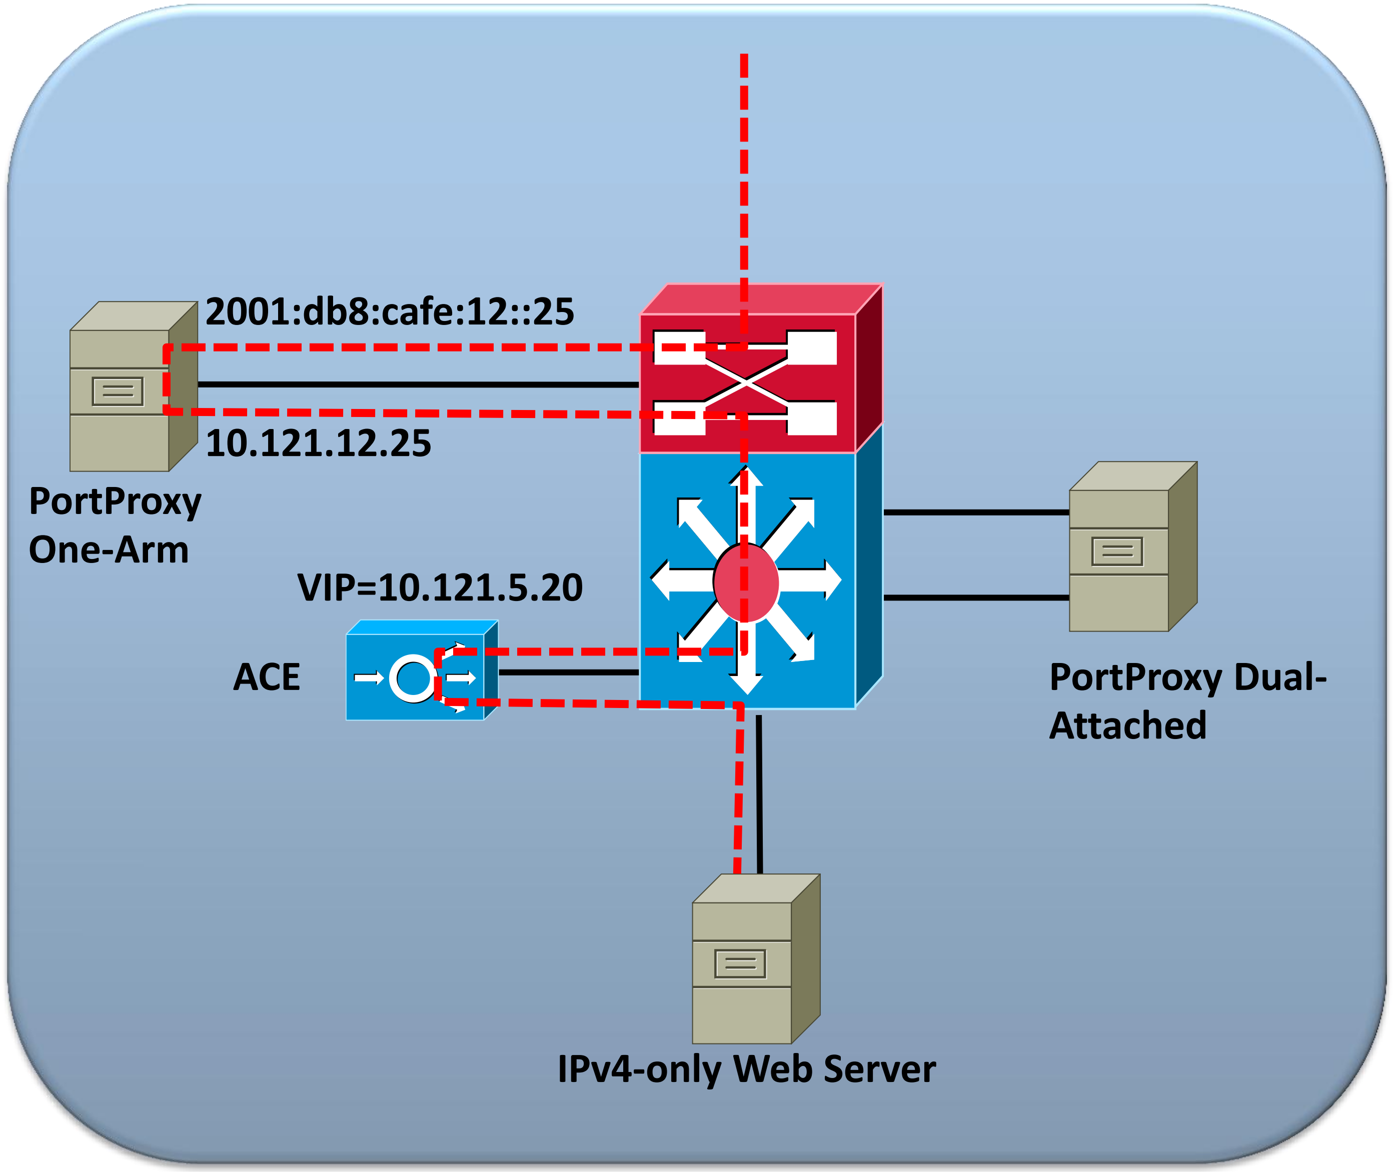

PortProxy Configuration/Monitoring

adsf

```
netsh interface portproxy>sh all
Listen on ipv6:             Connect to ipv4:
Address          Port       Address          Port
--------------- ----------  --------------- ----------
2001:db8:cafe:12::25 80     10.121.5.20      80
```

```
Active Connections
  Proto  Local Address            Foreign Address         State
  TCP    10.121.12.25:58141       10.121.5.20:http        ESTABLISHED
  TCP    [2001:db8:cafe:12::25]:80  [2001:db8:cafe:10::17]:52047  ESTABLISHED
```

```
conn-id    np dir proto vlan source                  destination            state
----------+--+---+-----+----+----------------------+----------------------+-----
-+
14         1  in  TCP   5    10.121.12.25:58573     10.121.5.20:80         ESTAB
13         1  out TCP   5    10.121.14.15:80        10.121.5.12:1062       ESTAB
```

Dual Stack the Internet Edge

- Dual stack the same network you have
- If not, do just enough IPv6-only to get you going
- Most design elements should be the same as with IPv4 (minus pure NAT/PAT)
- You may have to embrace SLB64/Proxy/NAT64 for IPv4-only apps
- LISP (Locator/ID Separation Protocol)

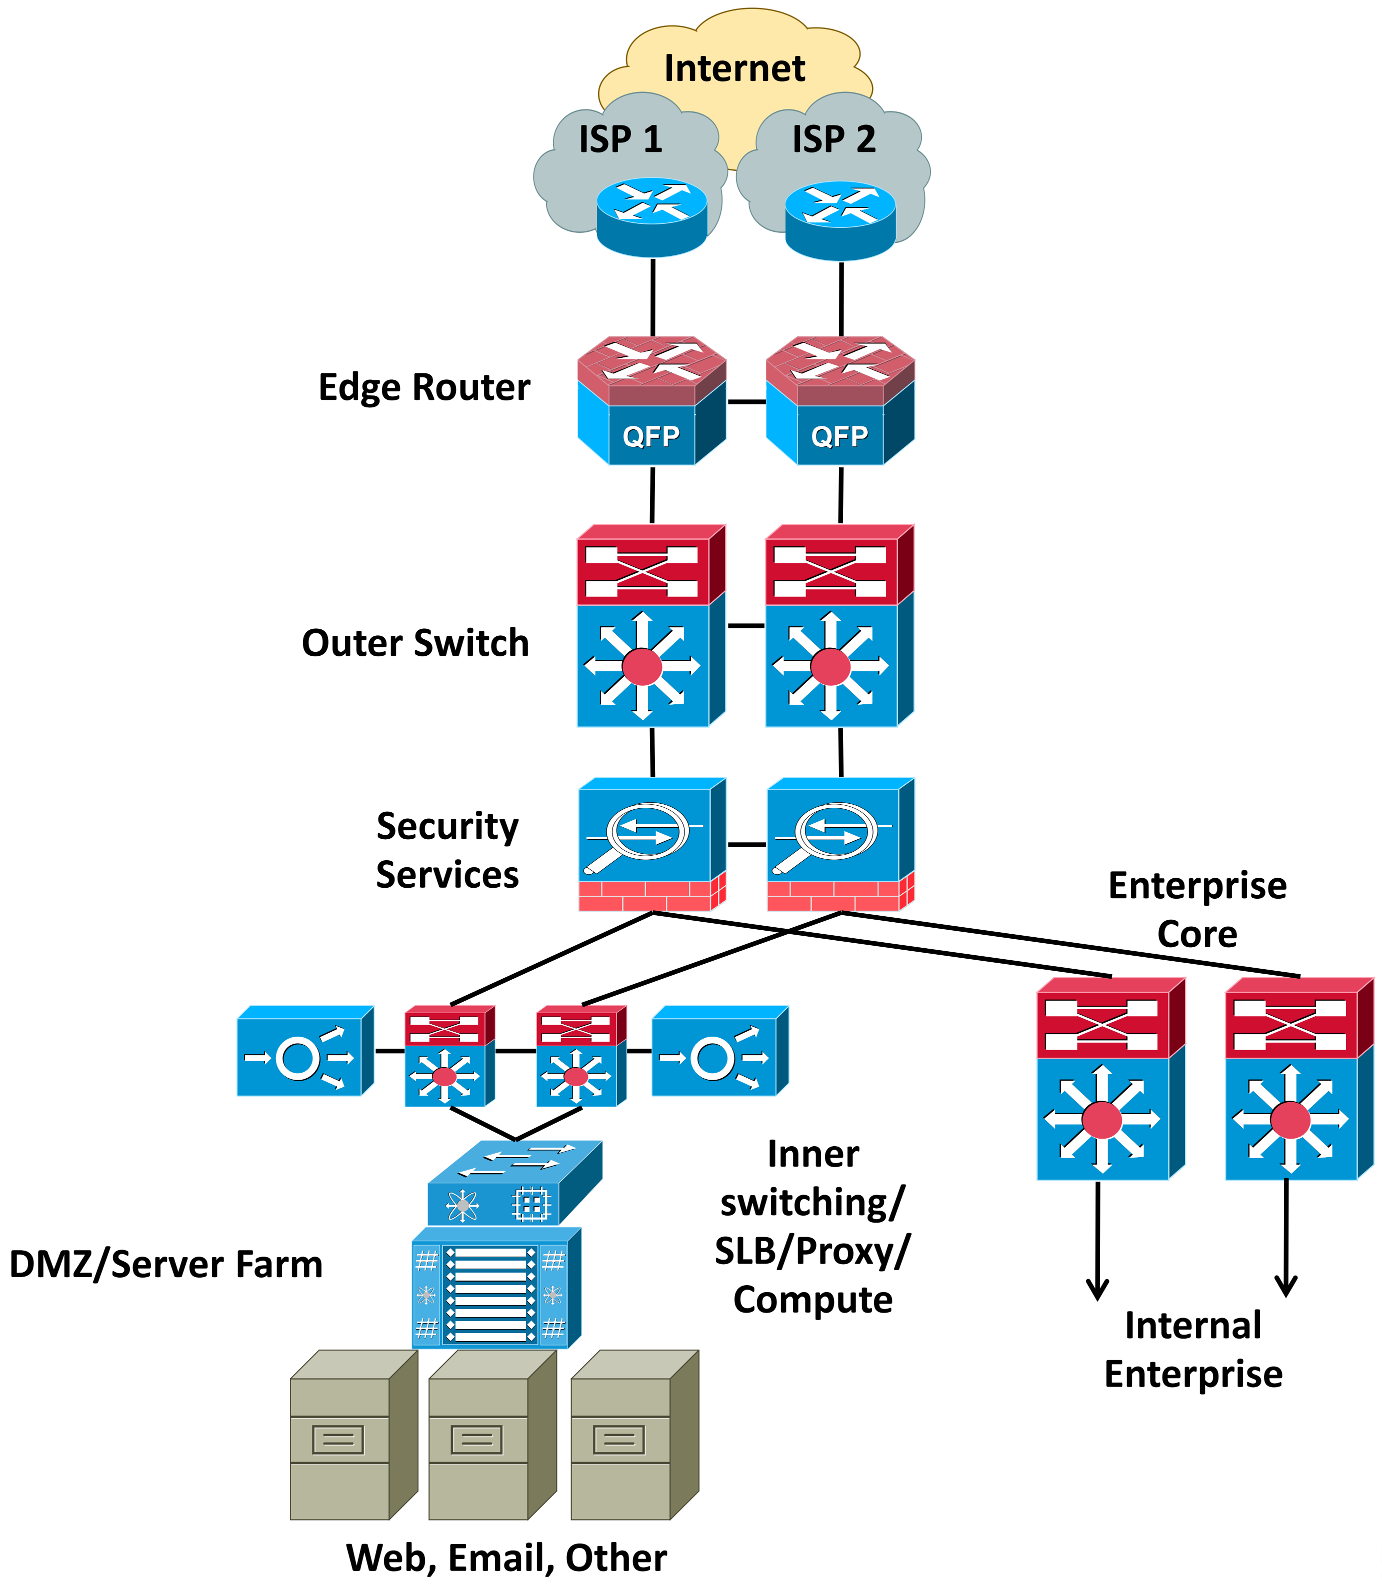

Top Enterprise Careabouts for Internet Edge

- Dual Stack peering
- Tunnel brokers as backup plan
- Address plans/Prefix-lengths
- To translate or not
- User experience – All things being equal – IPv6 wins and that may be a bad thing

Global Addressing Dilemma

- Today, many do NAT44 and 'hide' their RFC1918 space allowing for easier m
- If you are all PA or all PI and peering in multiple regions, then what?
  - PI from one region and run it everywhere?
  - PI each region and tune routing?
  - Will ISP in one region all of the sudden reject PI block from another?
  - What about translation?
- NPTv6 – Translating your prefix for the sake of multi-homing
  - RFC6296 – IPv6-to-IPv6 Network Prefix Translation
  - Make sure you understand the "Prefix" part well and what it really does
  - Internal PI, PA, ULA
  - STUN, TURN, ICE will all be used like with IPv4
- draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-xx

Some enterprises are getting a prefix per RIR, but deploy only one (building plan to use all as backup)

Not available on most shipping hardware – we will come back to this in the future

Internet Edge – to – ISP

Boatloads of Options

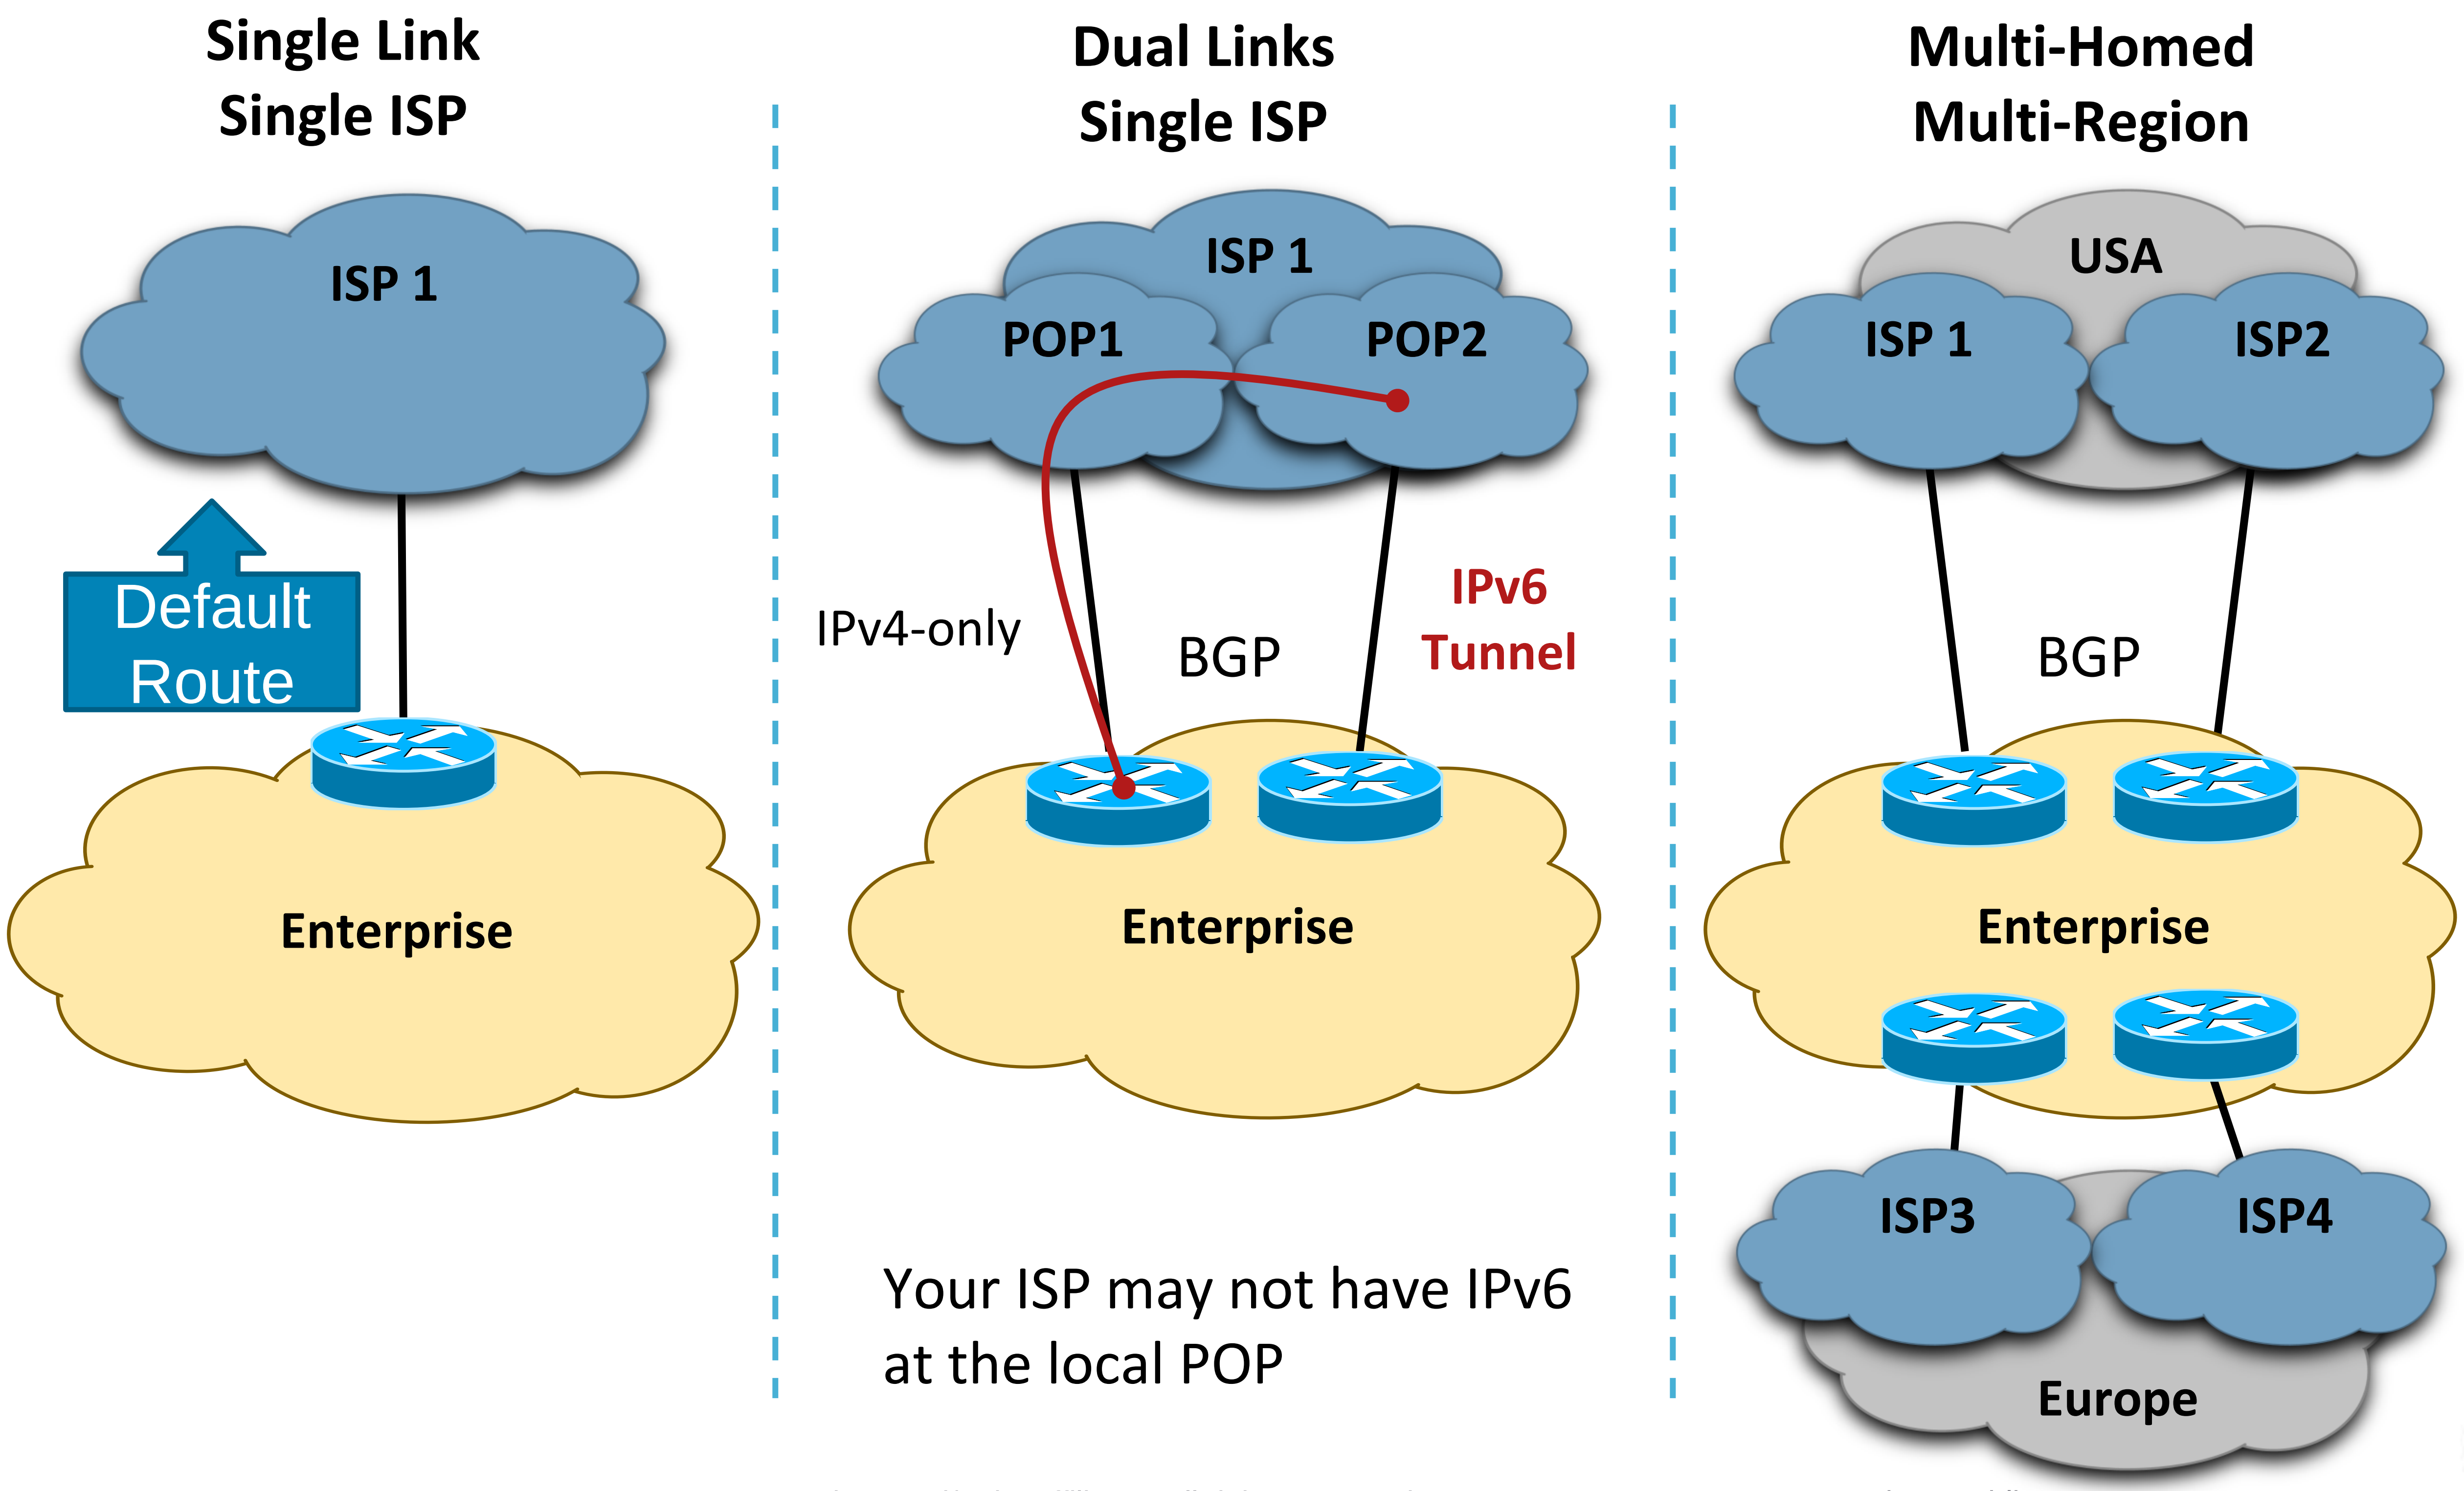

Single Homed

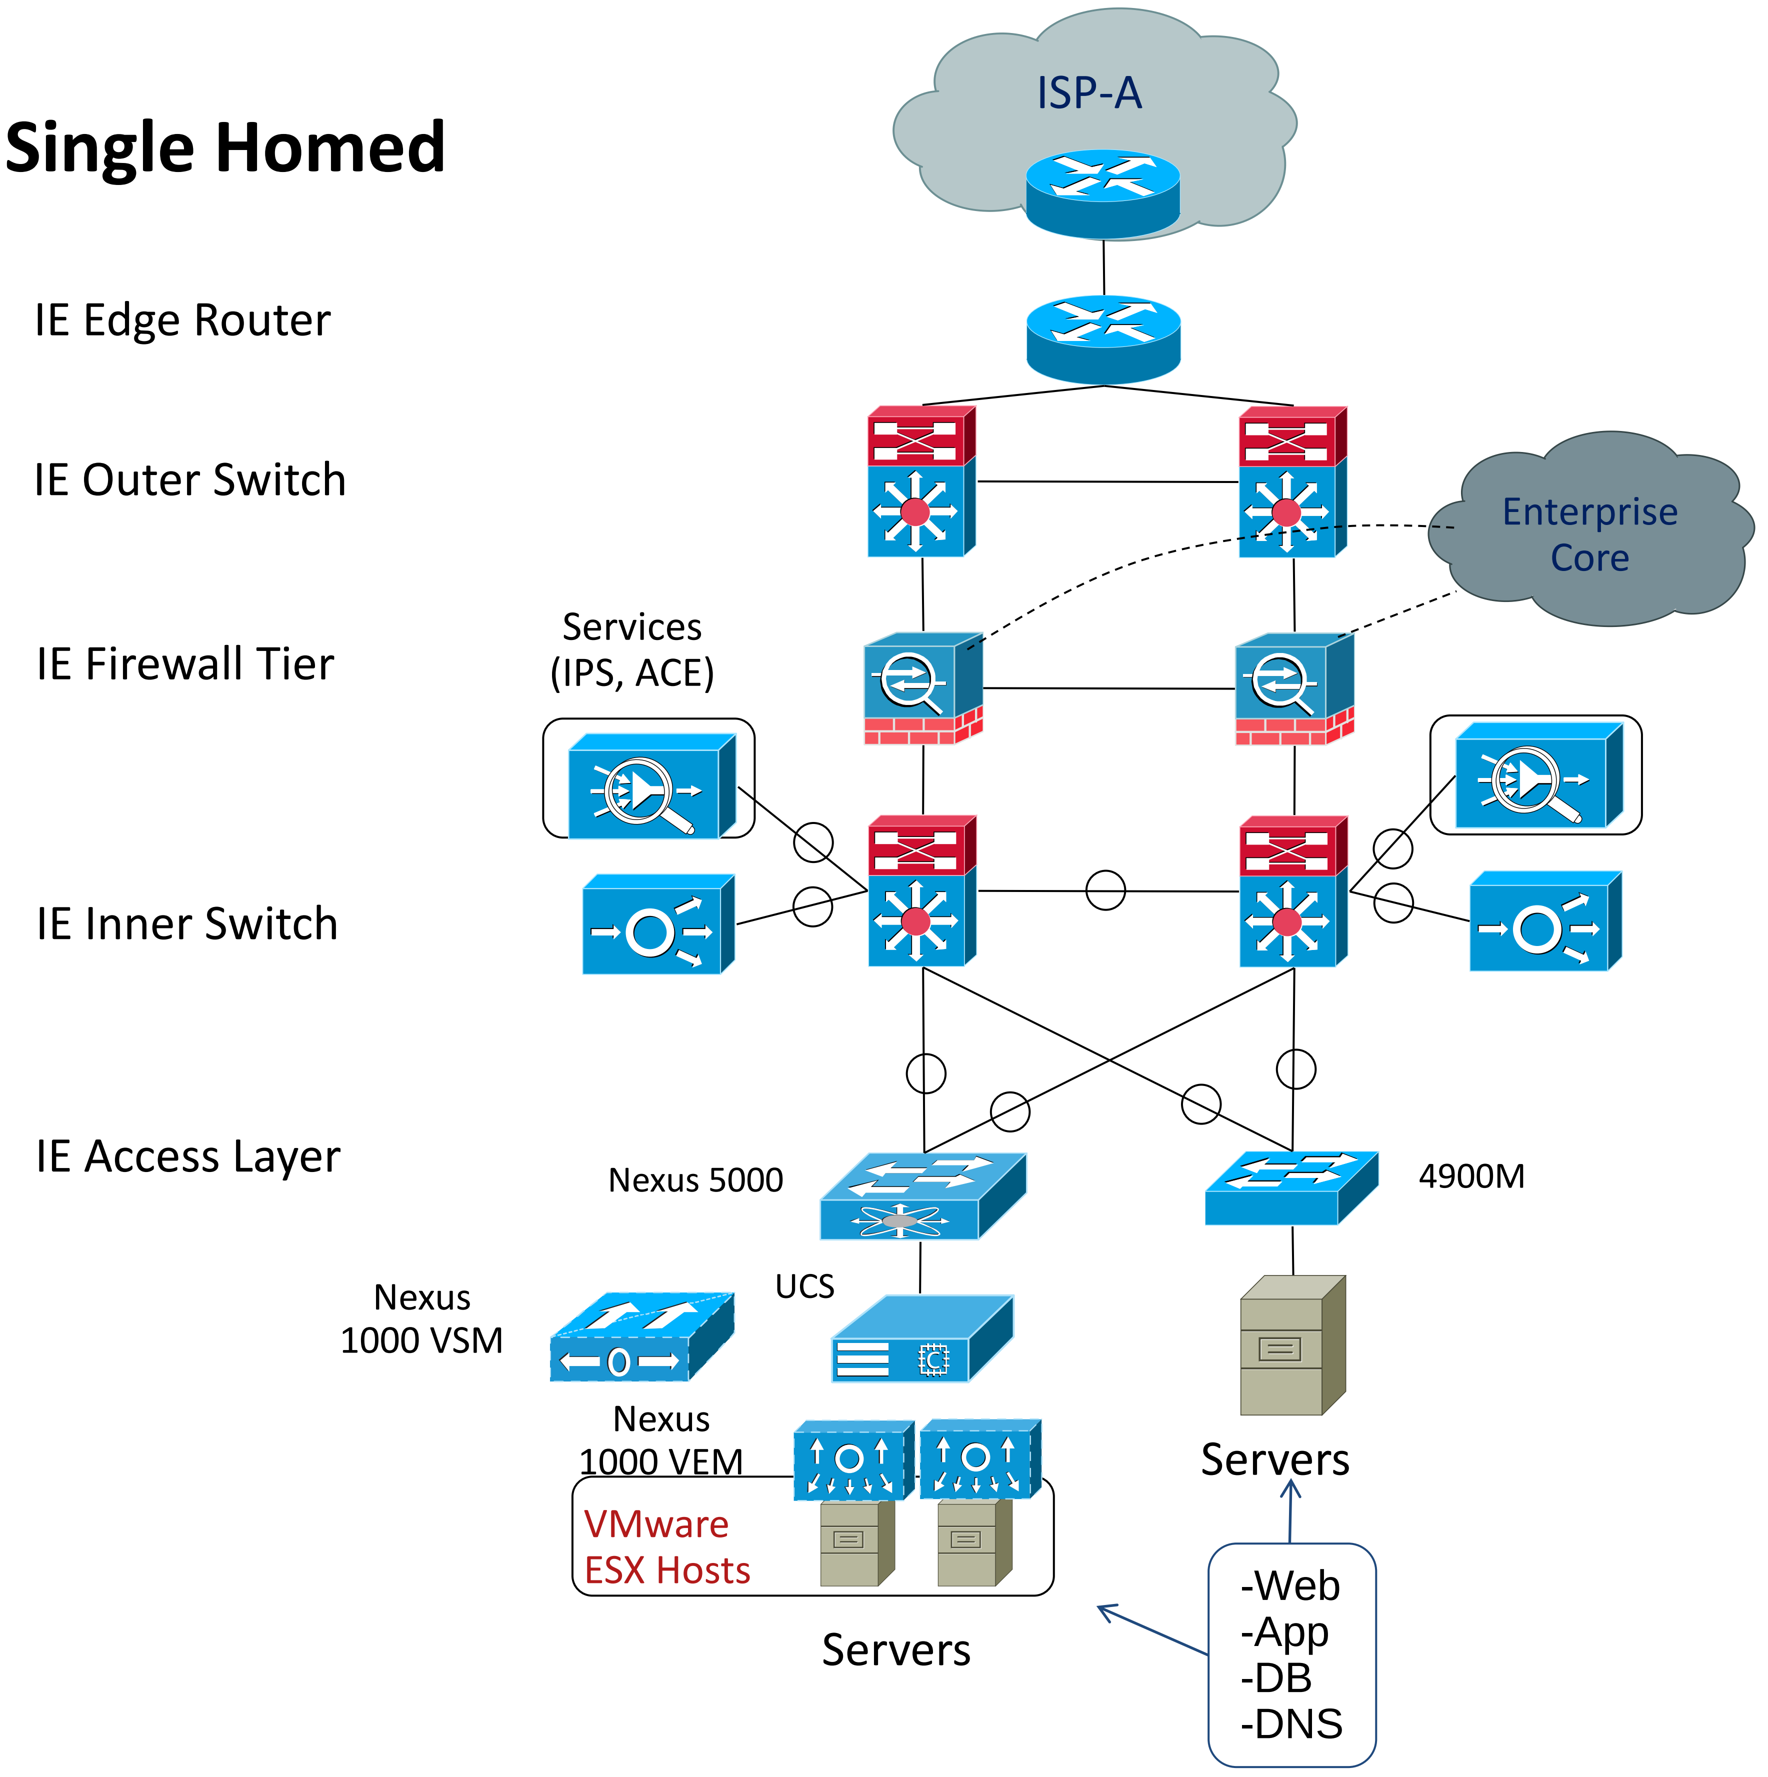

- L2 or L3 outer and inner switch designs
- Small DMZ to large DC-sized DMZ
- Baremetal or fully virtualized
- Dedicated IPv6 SF or dual stack

Multihomed – DS
(single ISP or multiple)

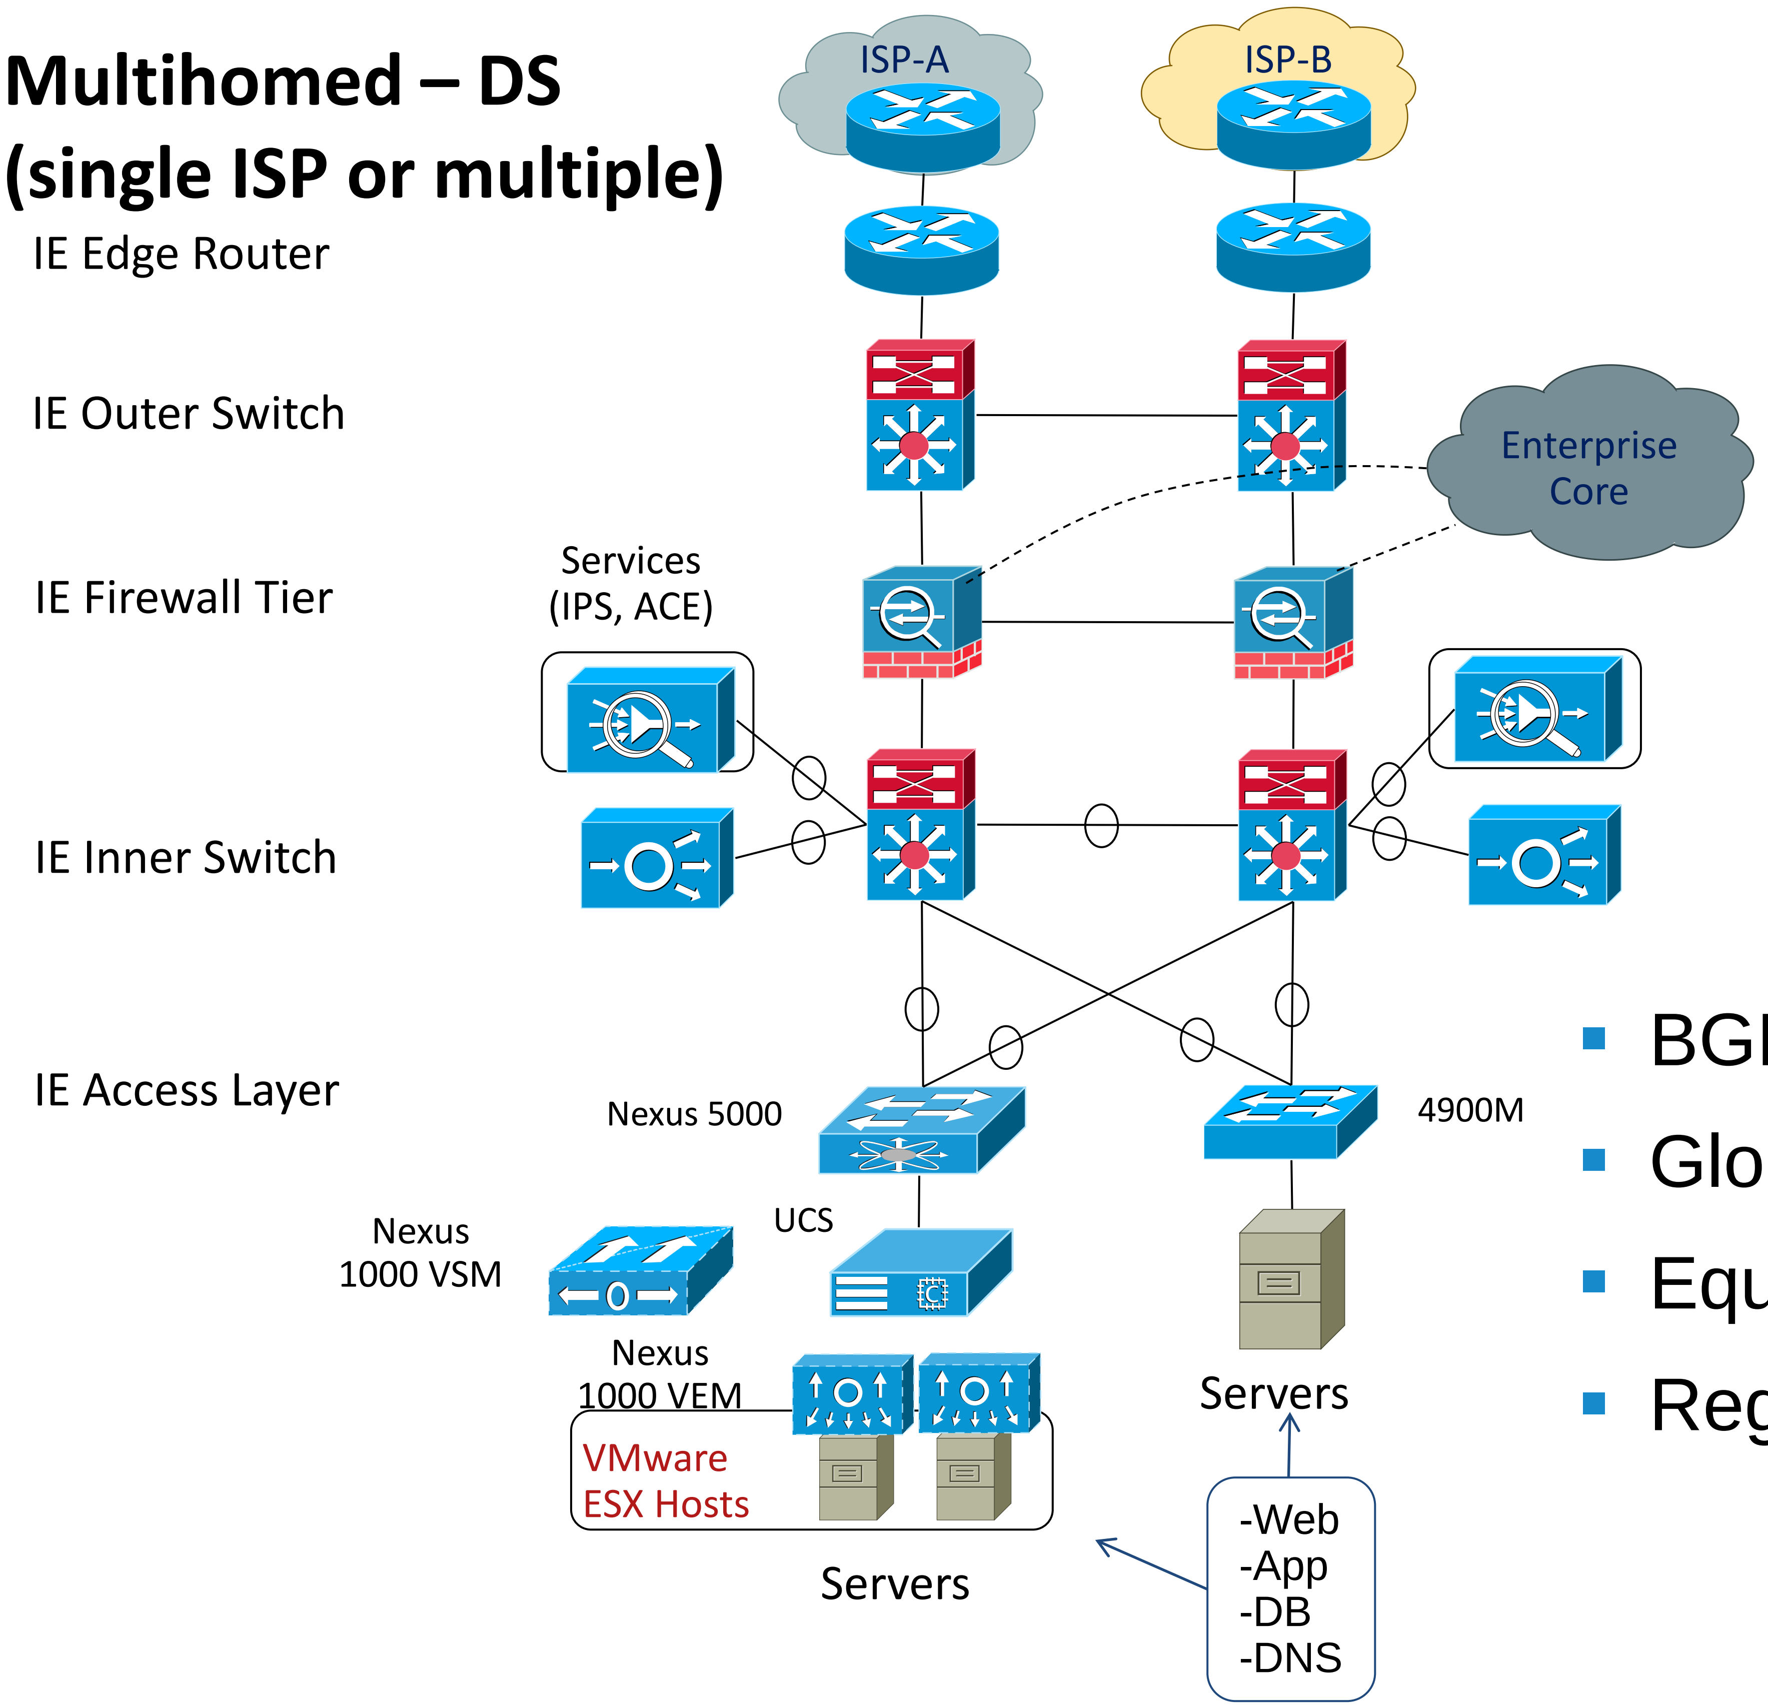

- BGP and address policy stuff
- Global Load Balancing
- Equal cost
- Regional egress/ingress policies

Multihomed – SLB64

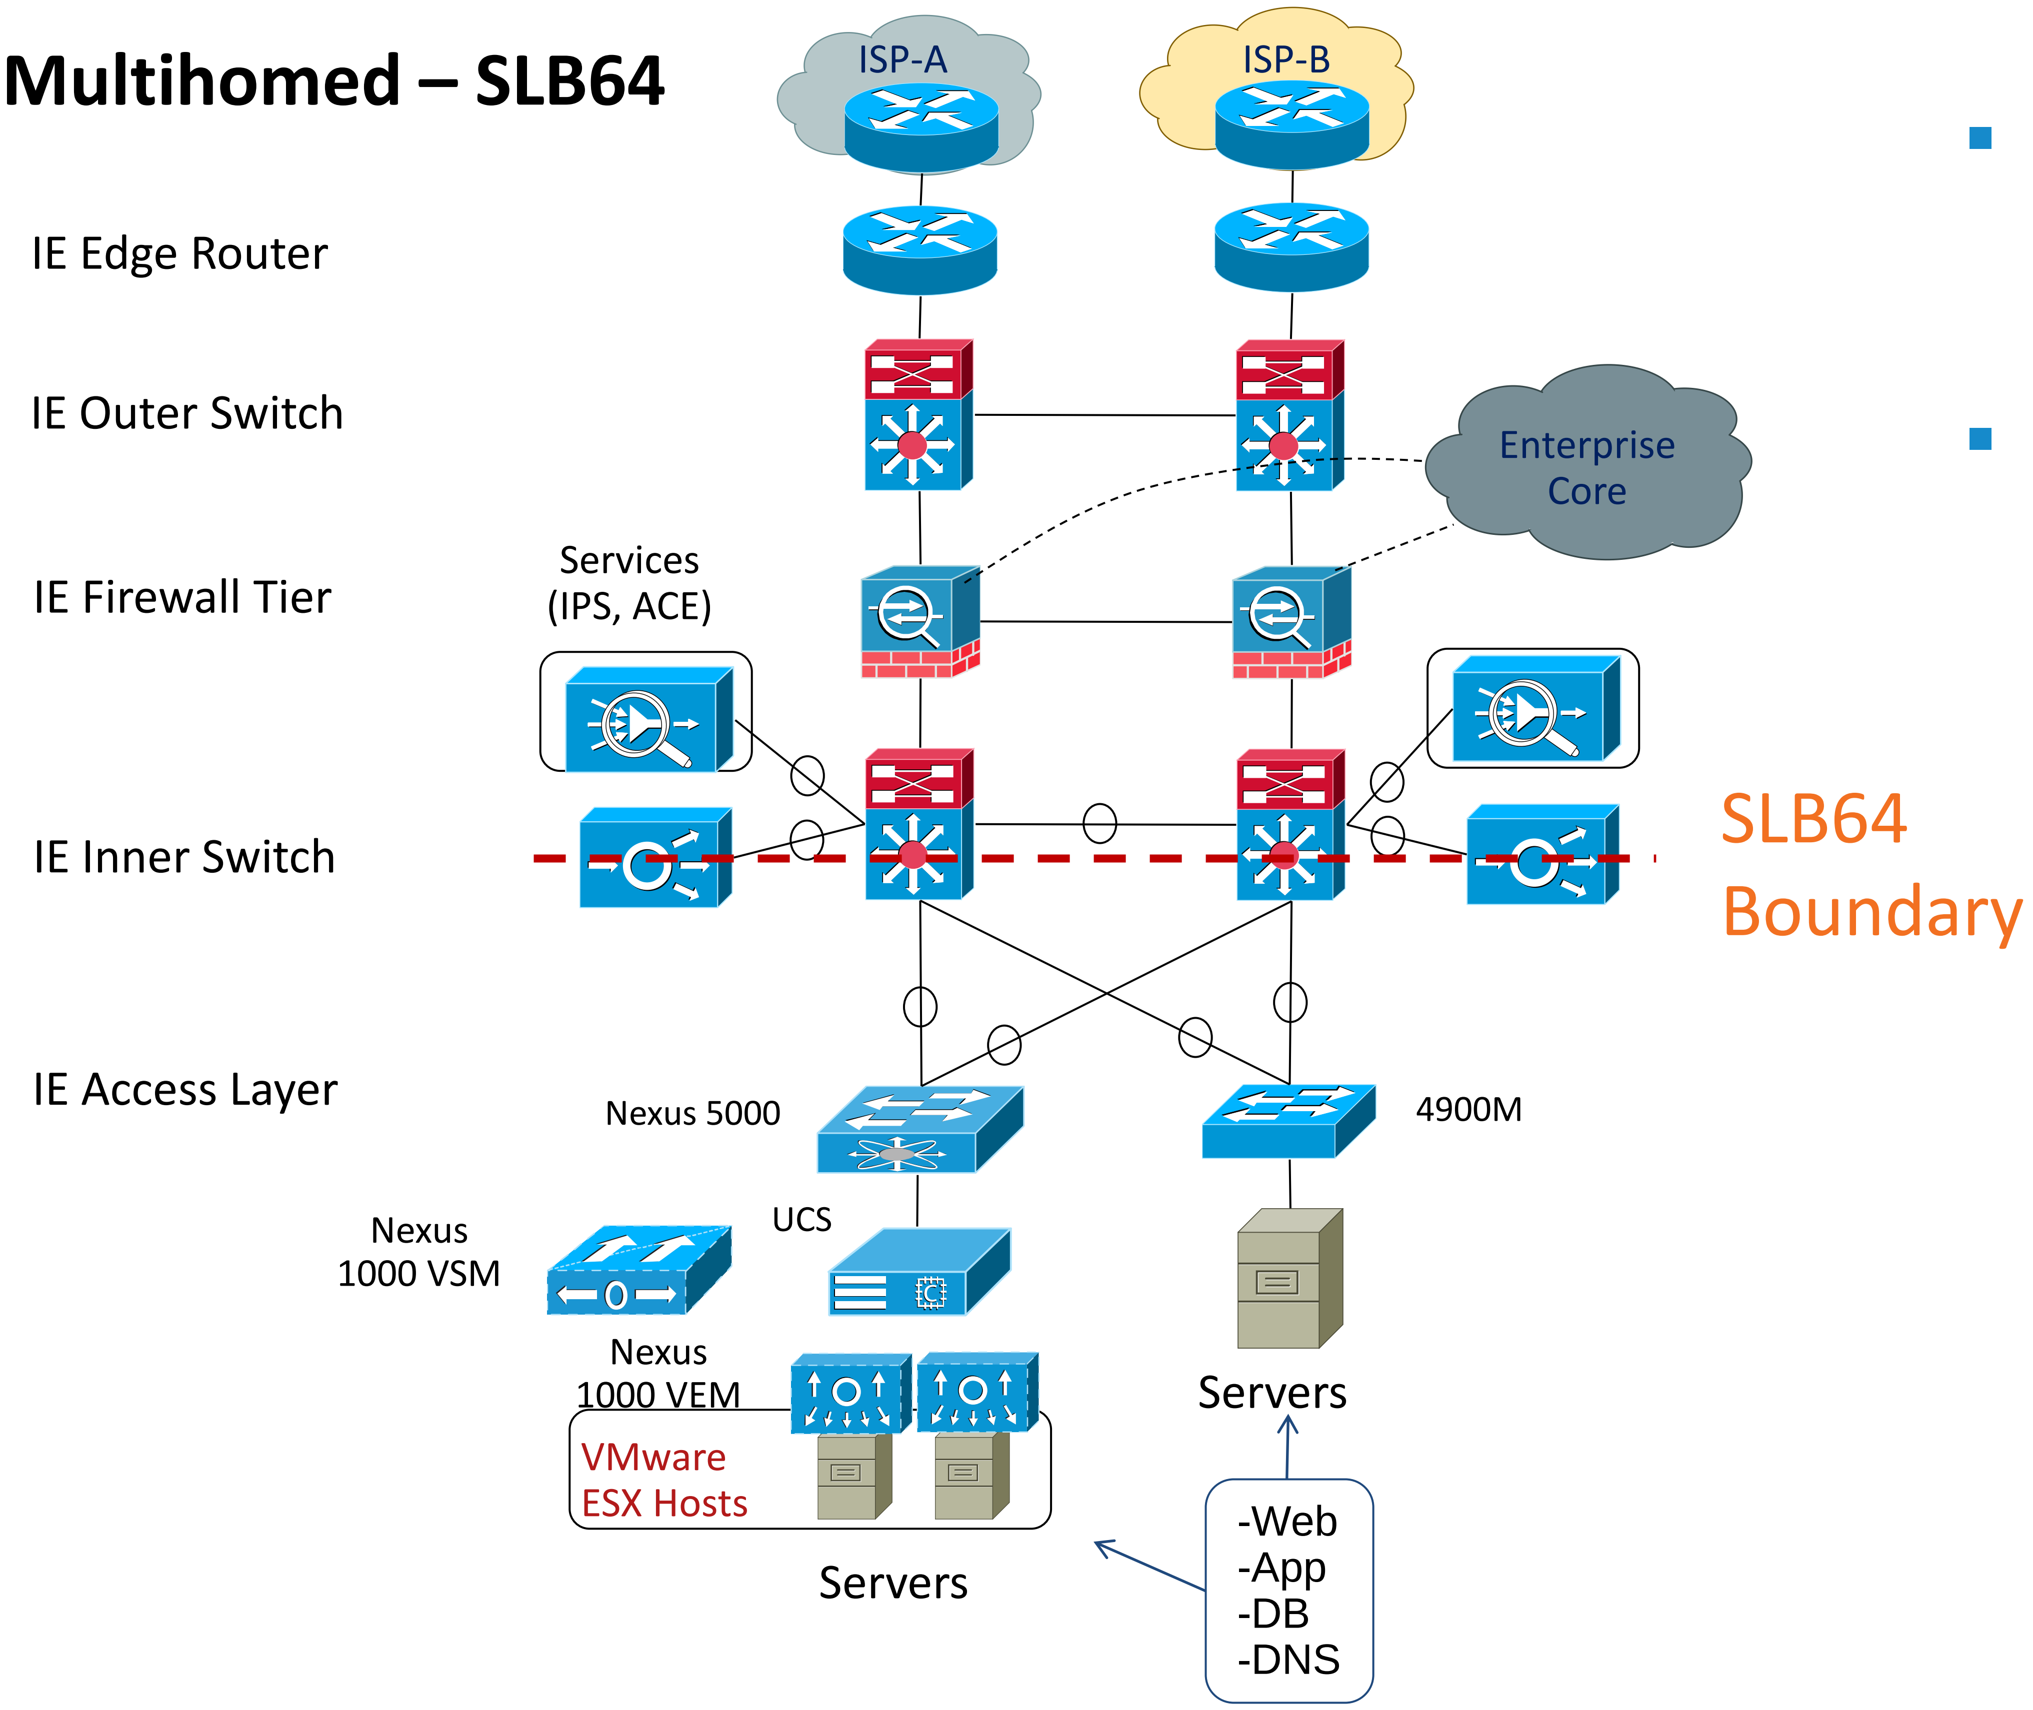

- OS and/or App will dictate design
- Operation issues also contribute to design choice

Multihomed – NAT64

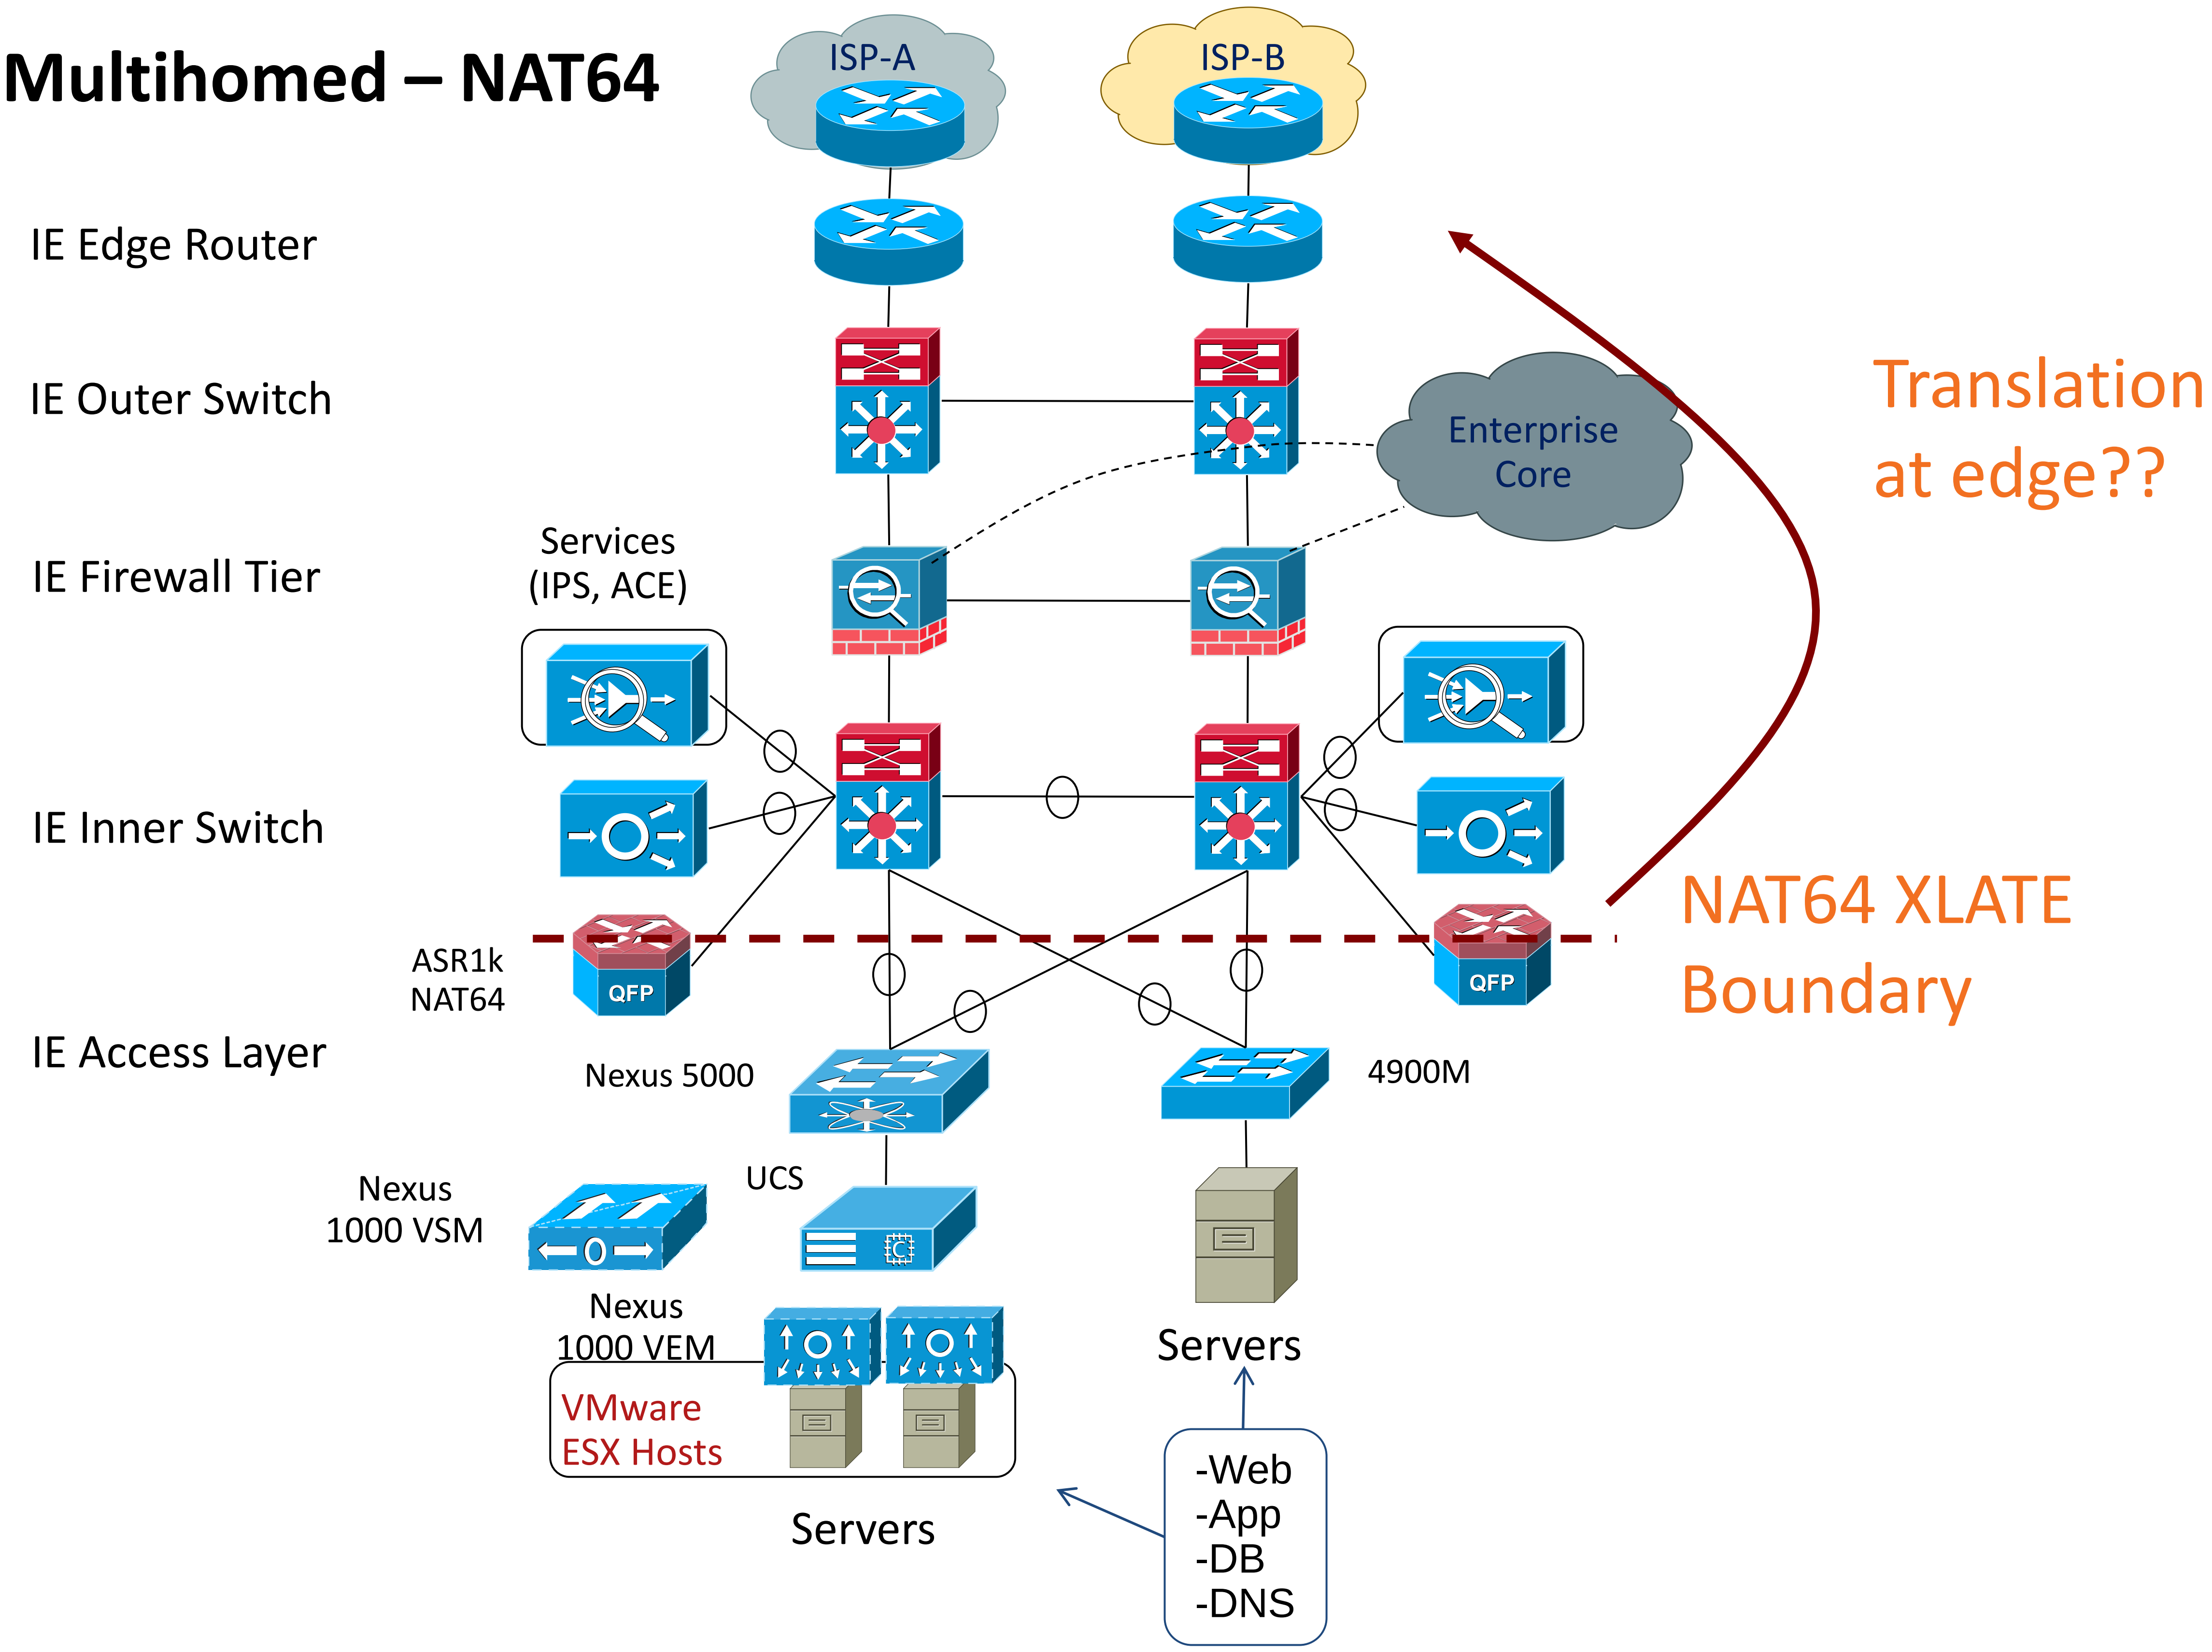

- 'Bolt-on' design
- Additional cost if not done in SLB
- Logging Src IP?

Multi-homed – Dual Stack

Single ISP – Multi-Peer - DS

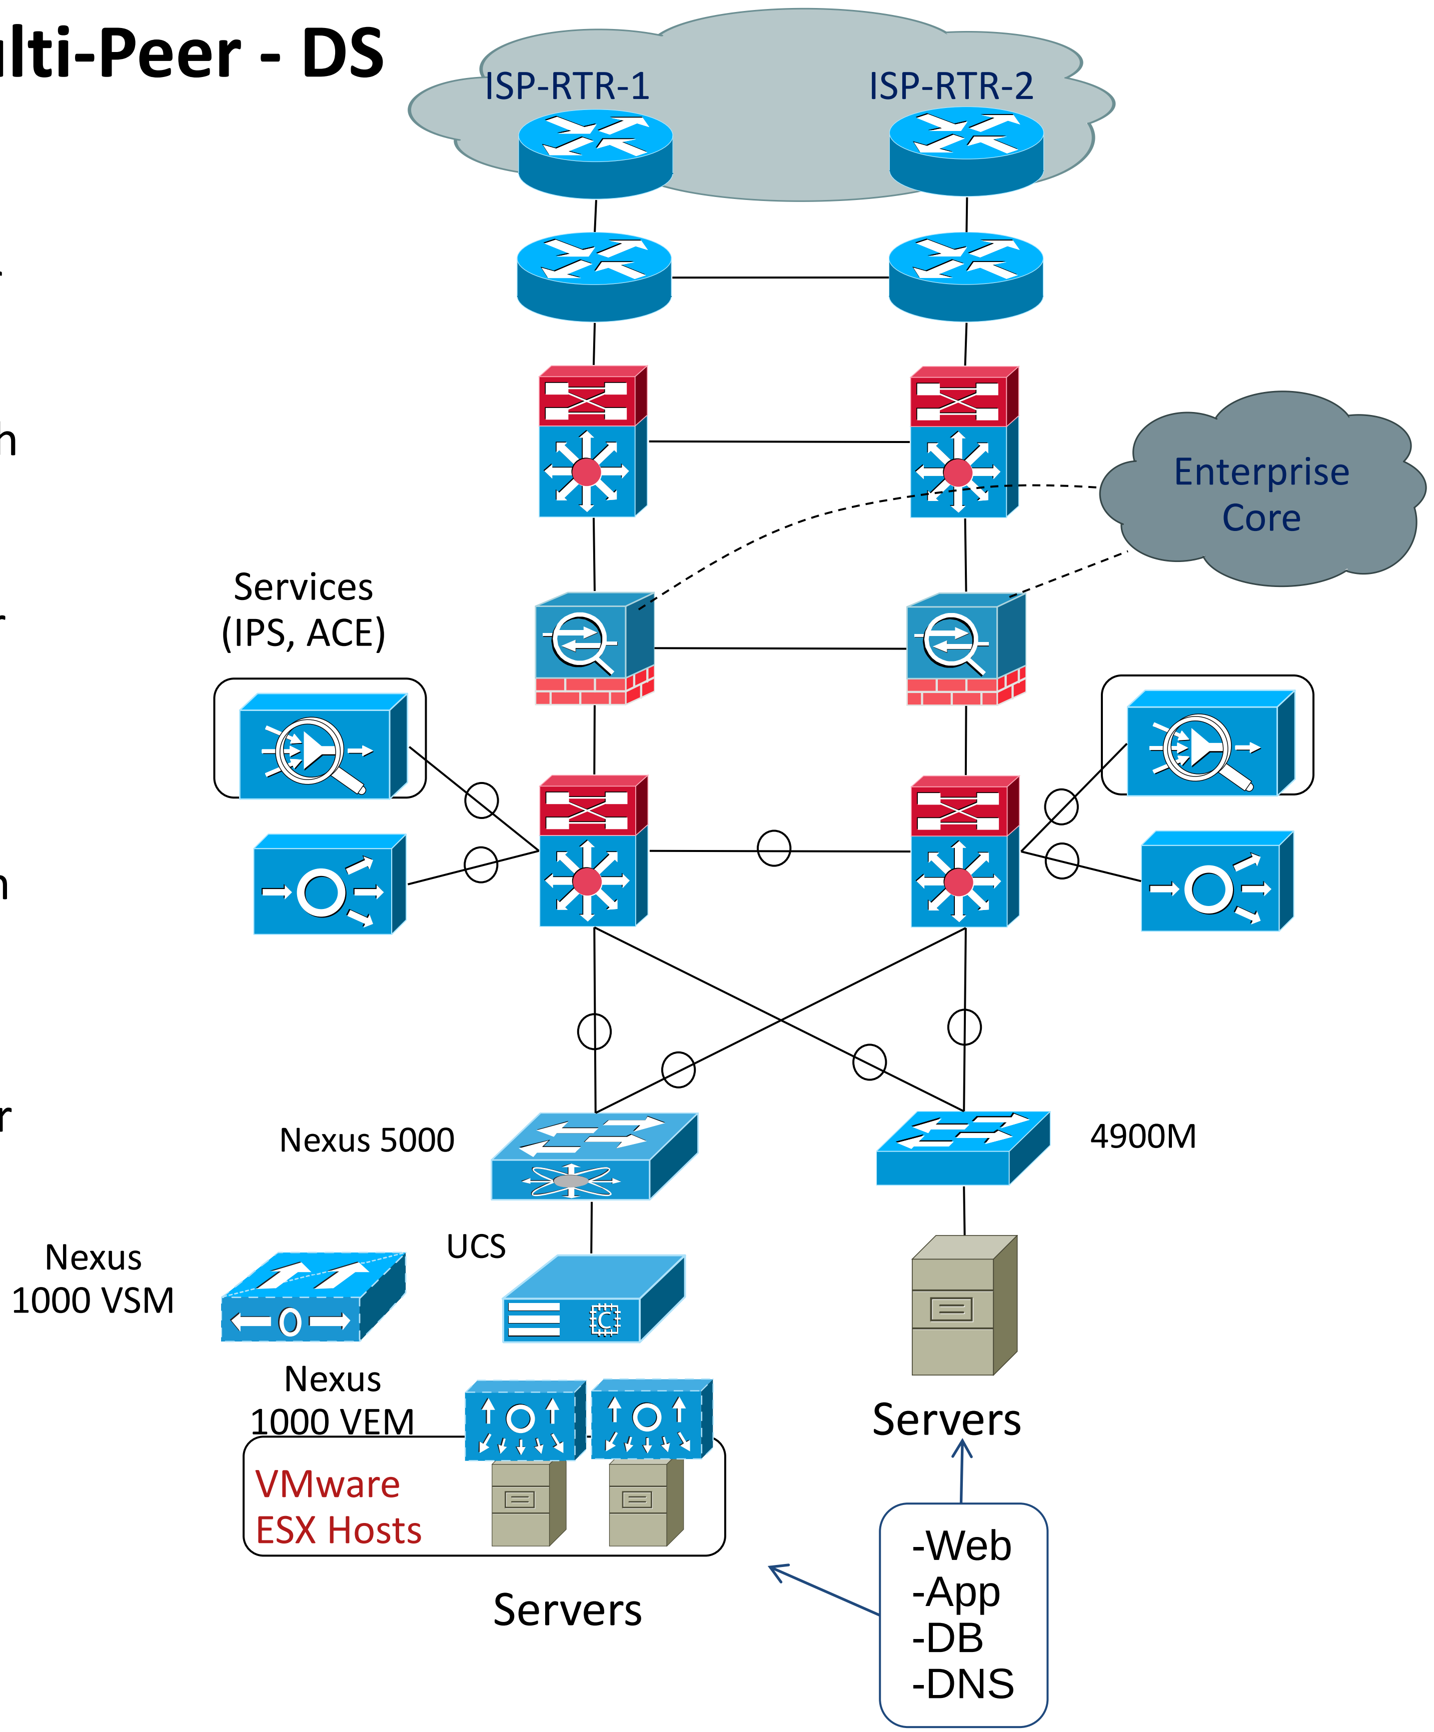

- Single ISP or multi-ISP will change BGP slightly
- No translation in this design
- Dual stack along the traffic flow from client-to-server (app)
- Keep a careful eye out on limitations in SW/HW and/or security details

Routing at the Edge

- Many, many different peering, HA and routing scenarios
  - eBGP to single ISP or multiple ISPs
  - IGP internally between edge routers and ASA or L3 switch
  - Equal cost routing or primary/secondary with FHRP
  - Dynamic or static
  - Etc…
- Our scenario is:
  - eBGP peering to single ISP but different ISP routers
  - iBGP between edge routers for re-routing during link failures
  - HSRP on edge-to-ASA links
  - Primary/Secondary routing preference with BGP
  - Inject default route from ISP

Edge Peering

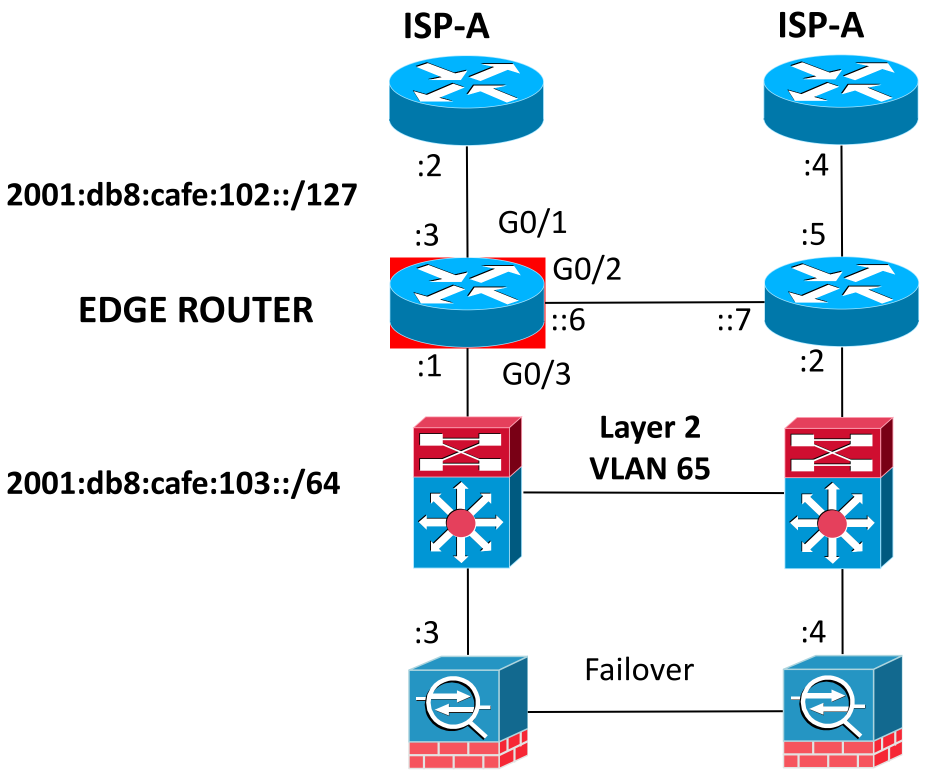

- Basic IP/Interface of left edge router
- /127s used on P2P
- /64 on ‘shared’ links (more than 2 network devices)
- May need special ACLs (like for HbH protection)

```
ipv6 unicast-routing
no ipv6 source-route
ipv6 cef
!
interface GigabitEthernet0/1
 description to ISPA (7604-1)
 ipv6 address 2001:DB8:CAFE:102::3/127
 ipv6 verify unicast reverse-path
!
interface GigabitEthernet0/2
 description LINK to EDGE-2
 ipv6 address 2001:DB8:CAFE:102::6/127
!
interface GigabitEthernet0/3
 description to ASA
 ipv6 address 2001:DB8:CAFE:103::1/64
 standby version 2
 standby 2 ipv6 autoconfig
 standby 2 priority 110
 standby 2 preempt delay minimum 300 reload 300
 standby 2 authentication CISCO
 standby 2 track GigabitEthernet0/1 20
!
ipv6 route 2001:DB8:CAFE::/48 2001:DB8:CAFE:103::3
```

Check HW Dependency

Apply Appropriate ACLs/CoPP

- Protect infrastructure that can be hurt by control plane processing
- HbH, RH0 (http://tools.ietf.org/html/rfc5095), etc. …

```
ipv6 access-list HBH
 deny hbh any any
 deny ipv6 any any routing-type 0
 permit icmp any any
 permit ipv6 any any
```

- Check that all networking vendors can handle /127 and/or protect against ICMP “ping pong” attacks

BGP - Edge Router

Reference

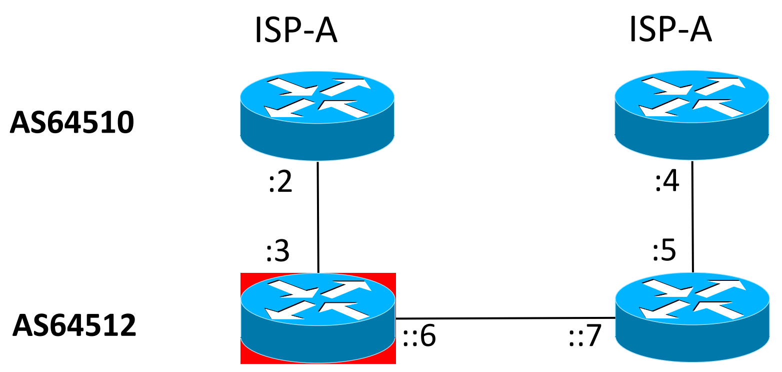

- Private AS Used (removed by ISP)
- eBGP to ISP
- iBGP to local edge router

```
router bgp 64512
 bgp router-id 192.168.1.33
 no bgp default ipv4-unicast
 bgp log-neighbour-changes
 neighbor 2001:DB8:CAFE:102::2 remote-as 64510
 neighbor 2001:DB8:CAFE:102::2 description IPv6_PEER_ISP
 neighbor 2001:DB8:CAFE:102::2 password CISCO
 neighbor 2001:DB8:CAFE:102::7 remote-as 64512
 neighbor 2001:DB8:CAFE:102::7 description EDGE_RTR_2
 neighbor 2001:DB8:CAFE:102::7 password CISCO
!
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
  neighbor 2001:DB8:CAFE:102::2 activate
  neighbor 2001:DB8:CAFE:102::7 activate
  neighbor 2001:DB8:CAFE:102::7 next-hop-self
  network 2001:DB8:CAFE::/48
  no synchronization
 exit-address-family
```

Reference

BGP Filters

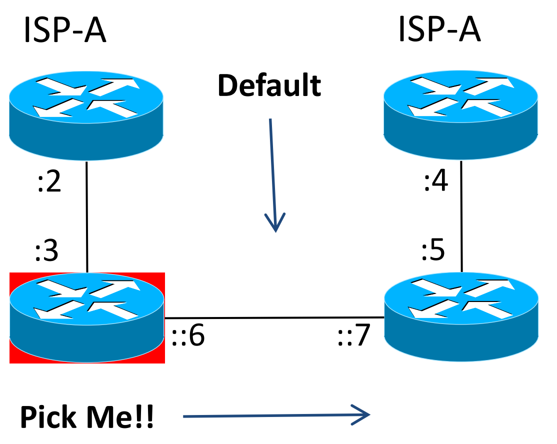

- Accepting default only
- Setting higher local pref
- ACLs for BGP

```
 address-family ipv6
  neighbor 2001:DB8:CAFE:102::2 prefix-list v6Default-Only in
  neighbor 2001:DB8:CAFE:102::2 route-map LOCAL in
 exit-address-family
!
ipv6 prefix-list v6Default-Only seq 5 permit ::/0
!
route-map LOCAL permit 10
 set local-preference 200
!
ipv6 access-list BGP
 permit tcp host 2001:DB8:CAFE:102::3 host 2001:DB8:CAFE:102::2 eq bgp
 deny tcp any any eq bgp
 permit ipv6 any any
!
ipv6 access-list IBGP
 permit tcp host 2001:DB8:CAFE:102::6 host 2001:DB8:CAFE:102::7 eq bgp
 deny tcp any any eq bgp
 permit ipv6 any any
!
interface GigabitEthernet0/1
 ipv6 traffic-filter BGP in
!
interface GigabitEthernet0/2
 ipv6 traffic-filter IBGP in
```

BGP Filters - Secondary

Reference

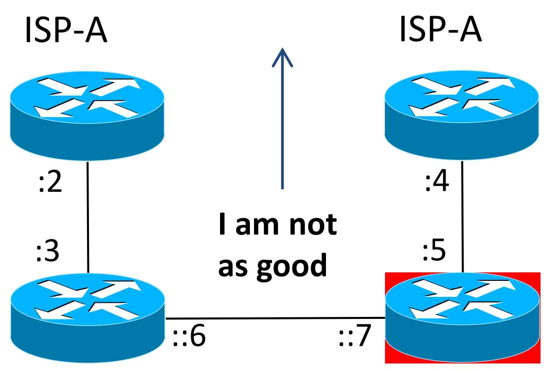

- Accepting default only
- AS PATH Prepend
- ACLs for BGP

```
 address-family ipv6
  neighbor 2001:DB8:CAFE:102::4 activate
  neighbor 2001:DB8:CAFE:102::4 prefix-list v6Default-Only in
  neighbor 2001:DB8:CAFE:102::4 route-map AS_PATH_PREPEND out
  neighbor 2001:DB8:CAFE:102::6 activate
  neighbor 2001:DB8:CAFE:102::6 next-hop-self
  network 2001:DB8:CAFE::/48
  no synchronization
 exit-address-family
!
route-map AS_PATH_PREPEND permit 10
 set as-path prepend 64512
```

Reference

Routing at Edge

Primary Edge Router

```
B    ::/0 [20/0]
      via FE80::216:9CFF:FE6D:5980, GigabitEthernet0/1
S    2001:DB8:CAFE::/48 [1/0]
      via 2001:DB8:CAFE:103::3
```

Default from ISP

Static towards ASA

Secondary Edge Router

```
B    ::/0 [200/0]
      via 2001:DB8:CAFE:102::6
S    2001:DB8:CAFE::/48 [1/0]
      via 2001:DB8:CAFE:103::3
```

Local Pref makes IBGP peer Favorable

ASA Interfaces

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ipv6 address 2001:db8:cafe:103::3/64 standby 2001:db8:cafe:103::4
!
interface GigabitEthernet0/1.19
 vlan 19
 nameif WEB
 security-level 50
 ipv6 address 2001:db8:cafe:115::3/64 standby 2001:db8:cafe:115::4
!
interface GigabitEthernet0/1.22
 vlan 22
 nameif DNS
 security-level 50
 ipv6 address 2001:db8:cafe:118::3/64 standby 2001:db8:cafe:118::4
!
interface Management0/0
 nameif management
 security-level 100
 ipv6 address 2001:db8:cafe:11a::10/64 standby 2001:db8:cafe:11a::11
 management-only
!
ipv6 route outside ::/0 fe80::5:73ff:fea0:2
```

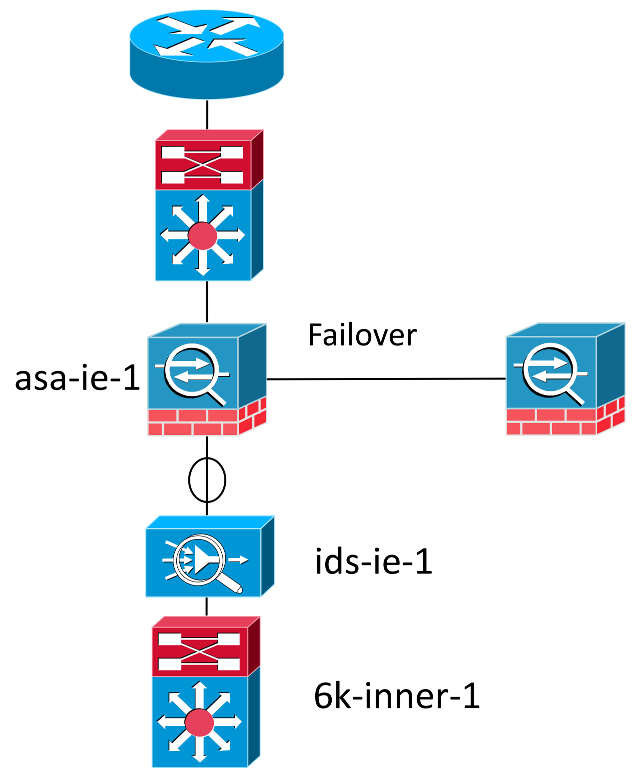

- VLANs on ASA or on inside switch
- L2 or L3 sandwich does not impact much

ASA HA/Failover

- Configuring Failover on the ASA is an either/or setup
- State for both protocols will be synced over a single failover configuration

```
interface GigabitEthernet0/3
 description LAN/STATE Failover Interface
!
failover
failover lan unit primary
failover lan interface fail GigabitEthernet0/3
failover polltime unit msec 200 holdtime msec 800
failover polltime interface msec 500 holdtime 5
failover key *****
failover replication http
failover link fail GigabitEthernet0/3
failover interface ip fail 10.140.3.1 255.255.255.252 standby 10.140.3.2
monitor-interface WEB
monitor-interface DNS
```

```
failover interface ip fail 2001:db8:cafe:fa11::2/127 standby 2001:db8:cafe:fa11::3
```

One or the other

ASA Object/ACL Configuration

```
object network IE-V6-WEB-VIP
 host 2001:db8:cafe:115::a
 description ACE IPv6 VIP address for Web Farm
object network ie-v6-dns
 host 2001:db8:cafe:118::a
object-group protocol TCPUDP
 protocol-object udp
 protocol-object tcp
!
ipv6 access-list outside_access_ipv6_in permit object-group TCPUDP any object ie-v6-dns eq domain
ipv6 access-list outside_access_ipv6_in permit tcp any object IE-V6-WEB-VIP eq www
!
access-group outside_access_ipv6_in in interface outside
```

HTTP or HTTPS?

ASA Device Manager

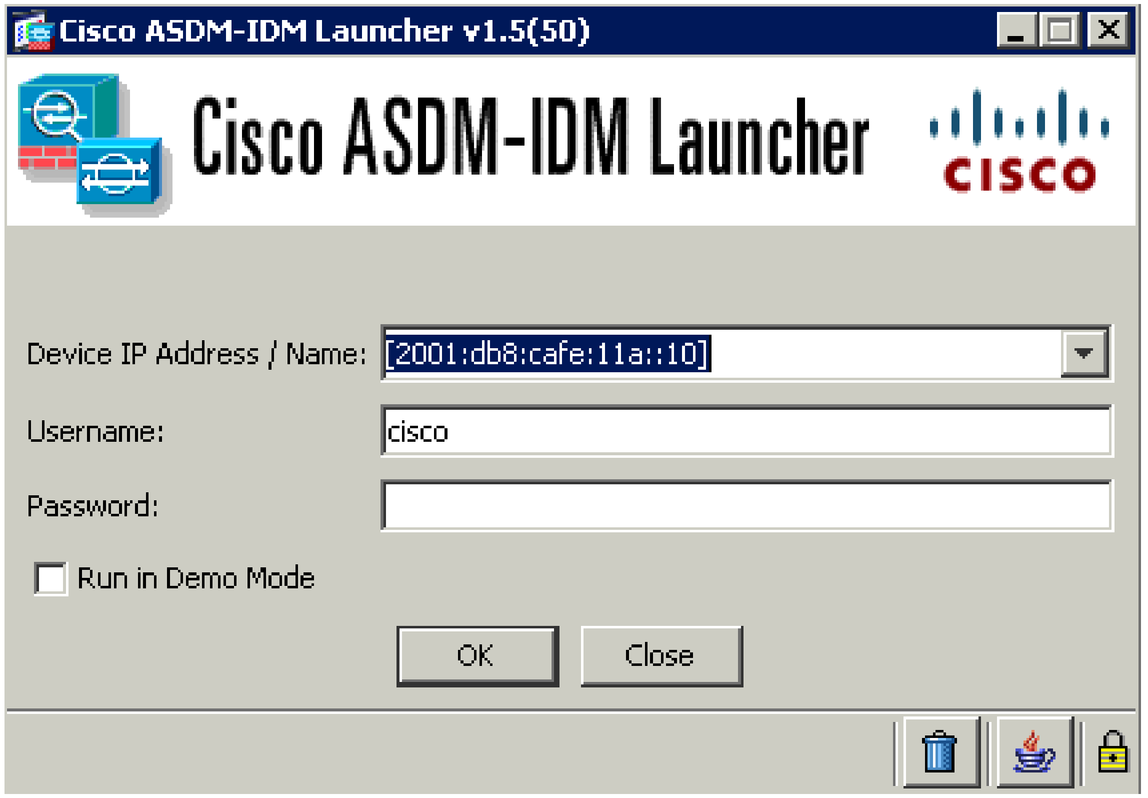


```
http server enable
http 2001:db8:cafe::/48 management
```

#	Enabled	Source	User	Destination	Service	Action
outside IPv6 (2 incoming rules)						
1	☑	any		ie-v6-dns	domain	Permit
2	☑	any		IE-V6-WEB-VIP	http	Permit
Global IPv6 (1 implicit rule)						
1		any		any	ip	Deny

IDS/IPS

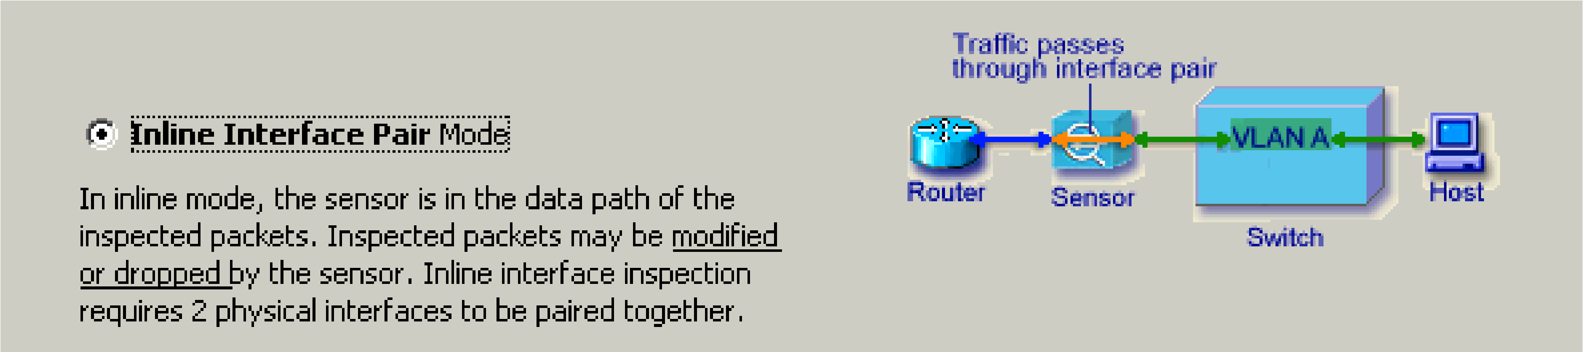

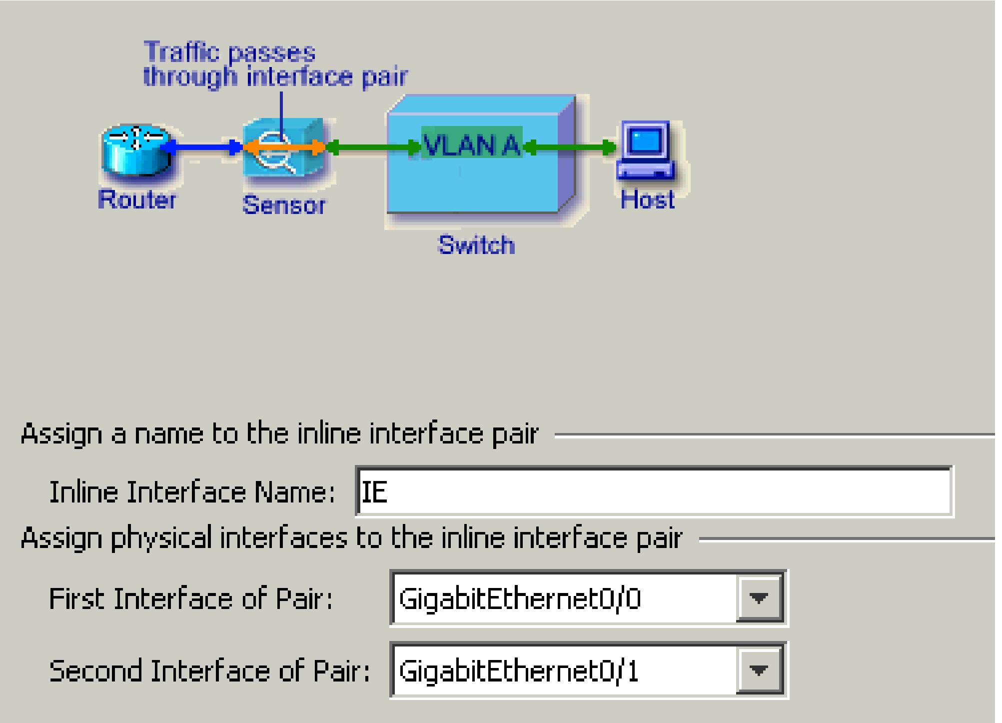

Severity	Date	Time	Device	Sig. Name	Sig. ID	Attacker IP	Victim IP	Vicitm ...	Threa...
high	09/27/2011	12:24:02	ids-ie-1	WWW WinNT cmd.exe Access	5081/0	172.16.99.100	10.140.19.10	80	90
high	09/27/2011	12:24:42	ids-ie-1	WWW WinNT cmd.exe Access	5081/0	2001:db8:ea5e:1:b878:ef18:e055:6476	2001:db8:cafe:115:0:0:0:a	80	90
high	09/27/2011	12:24:44	ids-ie-1	WWW WinNT cmd.exe Access	5081/0	2001:db8:ea5e:1:b878:ef18:e055:6476	2001:db8:cafe:115:0:0:0:a	80	90

Connecting the Inside

- L2 or L3 – Pick your HA/ECMP design
- It is no different than IPv4

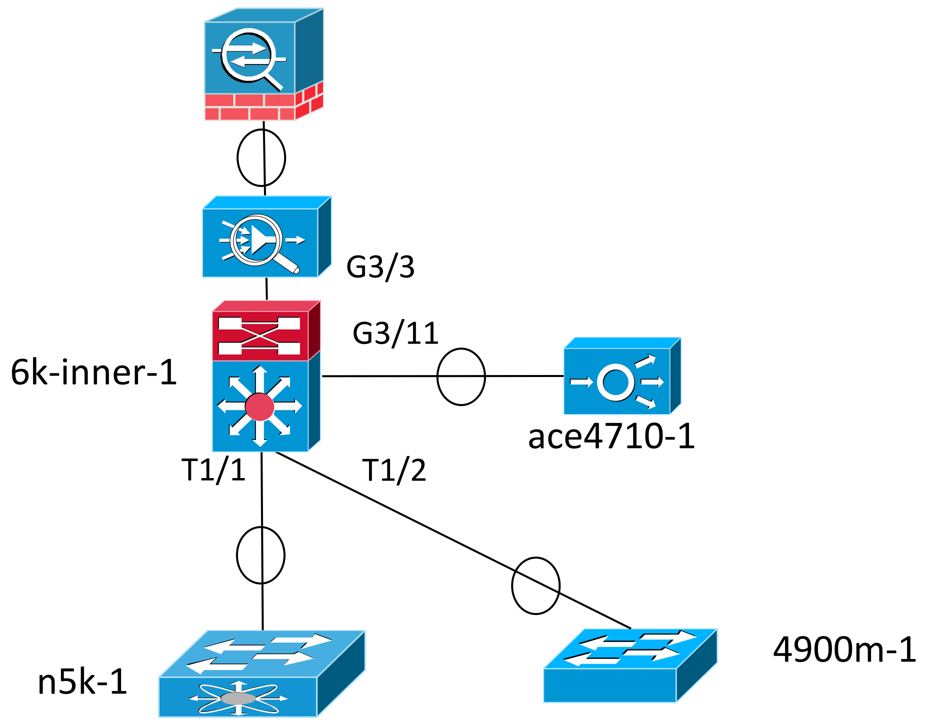


```
interface TenGigabitEthernet1/1
 description to N5k-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 18-25
 switchport mode trunk
 switchport nonegotiate
 spanning-tree guard root
!
interface TenGigabitEthernet1/2
 description to 4900m-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 18-25
 switchport mode trunk
 switchport nonegotiate
 spanning-tree guard root
```

```
interface GigabitEthernet3/3
 description to L2-IDS-ASA
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 18-25
 switchport mode trunk
!
interface GigabitEthernet3/11
 description to ACE4710 1-arm
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 19,24
 switchport mode trunk
```


Reference

Cisco ACE – Context Definition

Trunked Interface – One-arm Mode

```
interface gigabitEthernet 1/3
  description to IE-Trunk
  switchport trunk allowed vlan 19-22,24
  no shutdown
```

VLAN for Management

```
interface vlan 24
  ipv6 enable
  ip address 2001:db8:cafe:11a::b/64
  alias 2001:db8:cafe:11a::d/64
  peer ip address 2001:db8:cafe:11a::c/64
  access-group input ALL
  service-policy input remote_mgmt_allow_policy
  no shutdown
```

This will bring on the Mayan prediction if left off

Define Context

```
context IE-WEB
  allocate-interface vlan 19
```

Reference

Cisco ACE – Fault Tolerance (over IPv4)

FT Interface or just trunk with Port Channel

```
interface port-channel 1
  ft-port vlan 132
  no shutdown
```

```
ft interface vlan 132
  ip address 10.140.132.1 255.255.255.0
  peer ip address 10.140.132.2 255.255.255.0
  no shutdown

ft peer 1
  heartbeat interval 300
  heartbeat count 10
  ft-interface vlan 132
  query-interface vlan 19
ft group 2
  peer 1
  priority 110
  associate-context IE-WEB
  inservice
```

Reference

IE-WEB Context - MGMT

```
class-map type management match-any MGMT-CM
  2 match protocol xml-https any
  3 match protocol https any
  4 match protocol ssh any
  5 match protocol snmp any
  6 match protocol icmp any
  7 match protocol http any
  8 match protocol telnet any
class-map type management match-any MGMT-CM-v6
  2 match protocol icmpv6 anyv6

policy-map type management first-match MGMT
  class MGMT-CM
    permit
  class MGMT-CM-v6
    permit
interface vlan 19
  service-policy input MGMT
```

IP Access through the Cisco ACE

```
access-list EVERYONE line 10 extended permit icmp any any
access-list EVERYONE line 20 extended permit ip any any
access-list EVERYONE-v6 line 8 extended permit icmpv6 anyv6 anyv6
access-list EVERYONE-v6 line 16 extended permit ip anyv6 anyv6
interface vlan 19
 access-group input EVERYONE
 access-group input EVERYONE-v6
```

IE-WEB SLB66 Context Specific Configurations

```
probe http WEB_V6_PROBE
  interval 15
  passdetect interval 5
  request method get url /probe.html
  expect status 200 200
  open 1

rserver host WEB_V6_1
  ip address 2001:db8:cafe:115::10
  inservice
rserver host WEB_V6_2
  ip address 2001:db8:cafe:115::11
  inservice
serverfarm host WEB_V6_SF
  predictor leastconns slowstart 300
  probe WEB_V6_PROBE
  rserver WEB_V6_1 80
    inservice
  rserver WEB_V6_2 80
    inservice
```

```
class-map match-all WEB_V6_VIP
  2 match virtual-address 2001:db8:cafe:115::a tcp eq www

policy-map type loadbalance first-match WEB_V6_SLB
  class class-default
    serverfarm WEB_V6_SF
    insert-http x-forward header-value "%is"
policy-map multi-match WEB_V6_POL
  class WEB_V6_VIP
    loadbalance vip inservice
    loadbalance policy WEB_V6_SLB
    loadbalance vip icmp-reply active
    nat dynamic 1 vlan 19
interface vlan 19
  ipv6 enable
  ip address 2001:db8:cafe:115::d/64
  alias 2001:db8:cafe:115::f/64
  peer ip address 2001:db8:cafe:115::e/64
  access-group input EVERYONE-v6
  nat-pool 1 2001:db8:cafe:115::ace
2001:db8:cafe:115::ace/128 pat
  service-policy input MGMT
  service-policy input WEB_V6_POL
ip route ::/0 2001:db8:cafe:115::3
```

More on this later (insert-http x-forward header-value)

Don't screw this up (ipv6 enable)

1. Define probe
2. Define real servers (IPv6)
3. Serverfarm
4. Define VIP (IPv6)
5. Glue together with policy maps
6. One arm with SNAT example

Reference

SSL Offload

```
class-map match-all WEB_V6_VIP
  2 match virtual-address 2001:db8:cafe:115::a tcp eq https

ssl-proxy service SSL_PROXY_WEB
  key cisco-sample-key
  cert cisco-sample-cert

policy-map multi-match WEB_V6_POL
  class WEB_V6_VIP
    loadbalance vip inservice
    loadbalance policy WEB_V6_SLB
    loadbalance vip icmp-reply active
    nat dynamic 1 vlan 19
    ssl-proxy server SSL_PROXY_WEB
```

Health Monitoring (Probes) - HTTP

```
ace4710-1/IE-WEB# show probe
 probe          : WEB_V6_PROBE
 type           : HTTP
 state          : ACTIVE
----------------------------------------------
   port       : 80            address    : 0.0.0.0
   addr type  : -             interval   : 15      pass intvl   : 5
   pass count: 3              fail count: 3        recv timeout: 10
                ------------------ probe results ------------------
   associations     ip-address         port porttype probes failed passed health
   ------------ ----------------------+----+--------+------+------+------+------
   serverfarm   : WEB_V6_SF
     real       : WEB_V6_1[80]
                  2001:db8:cafe:115::10   80 REAL     7000   11     6989   SUCCESS
     real       : WEB_V6_2[80]
                  2001:db8:cafe:115::11   80 REAL     7623   942    6681   SUCCESS
```

Application Networking Manager 5.1

- Full Monitoring
- Configure all elements of policies

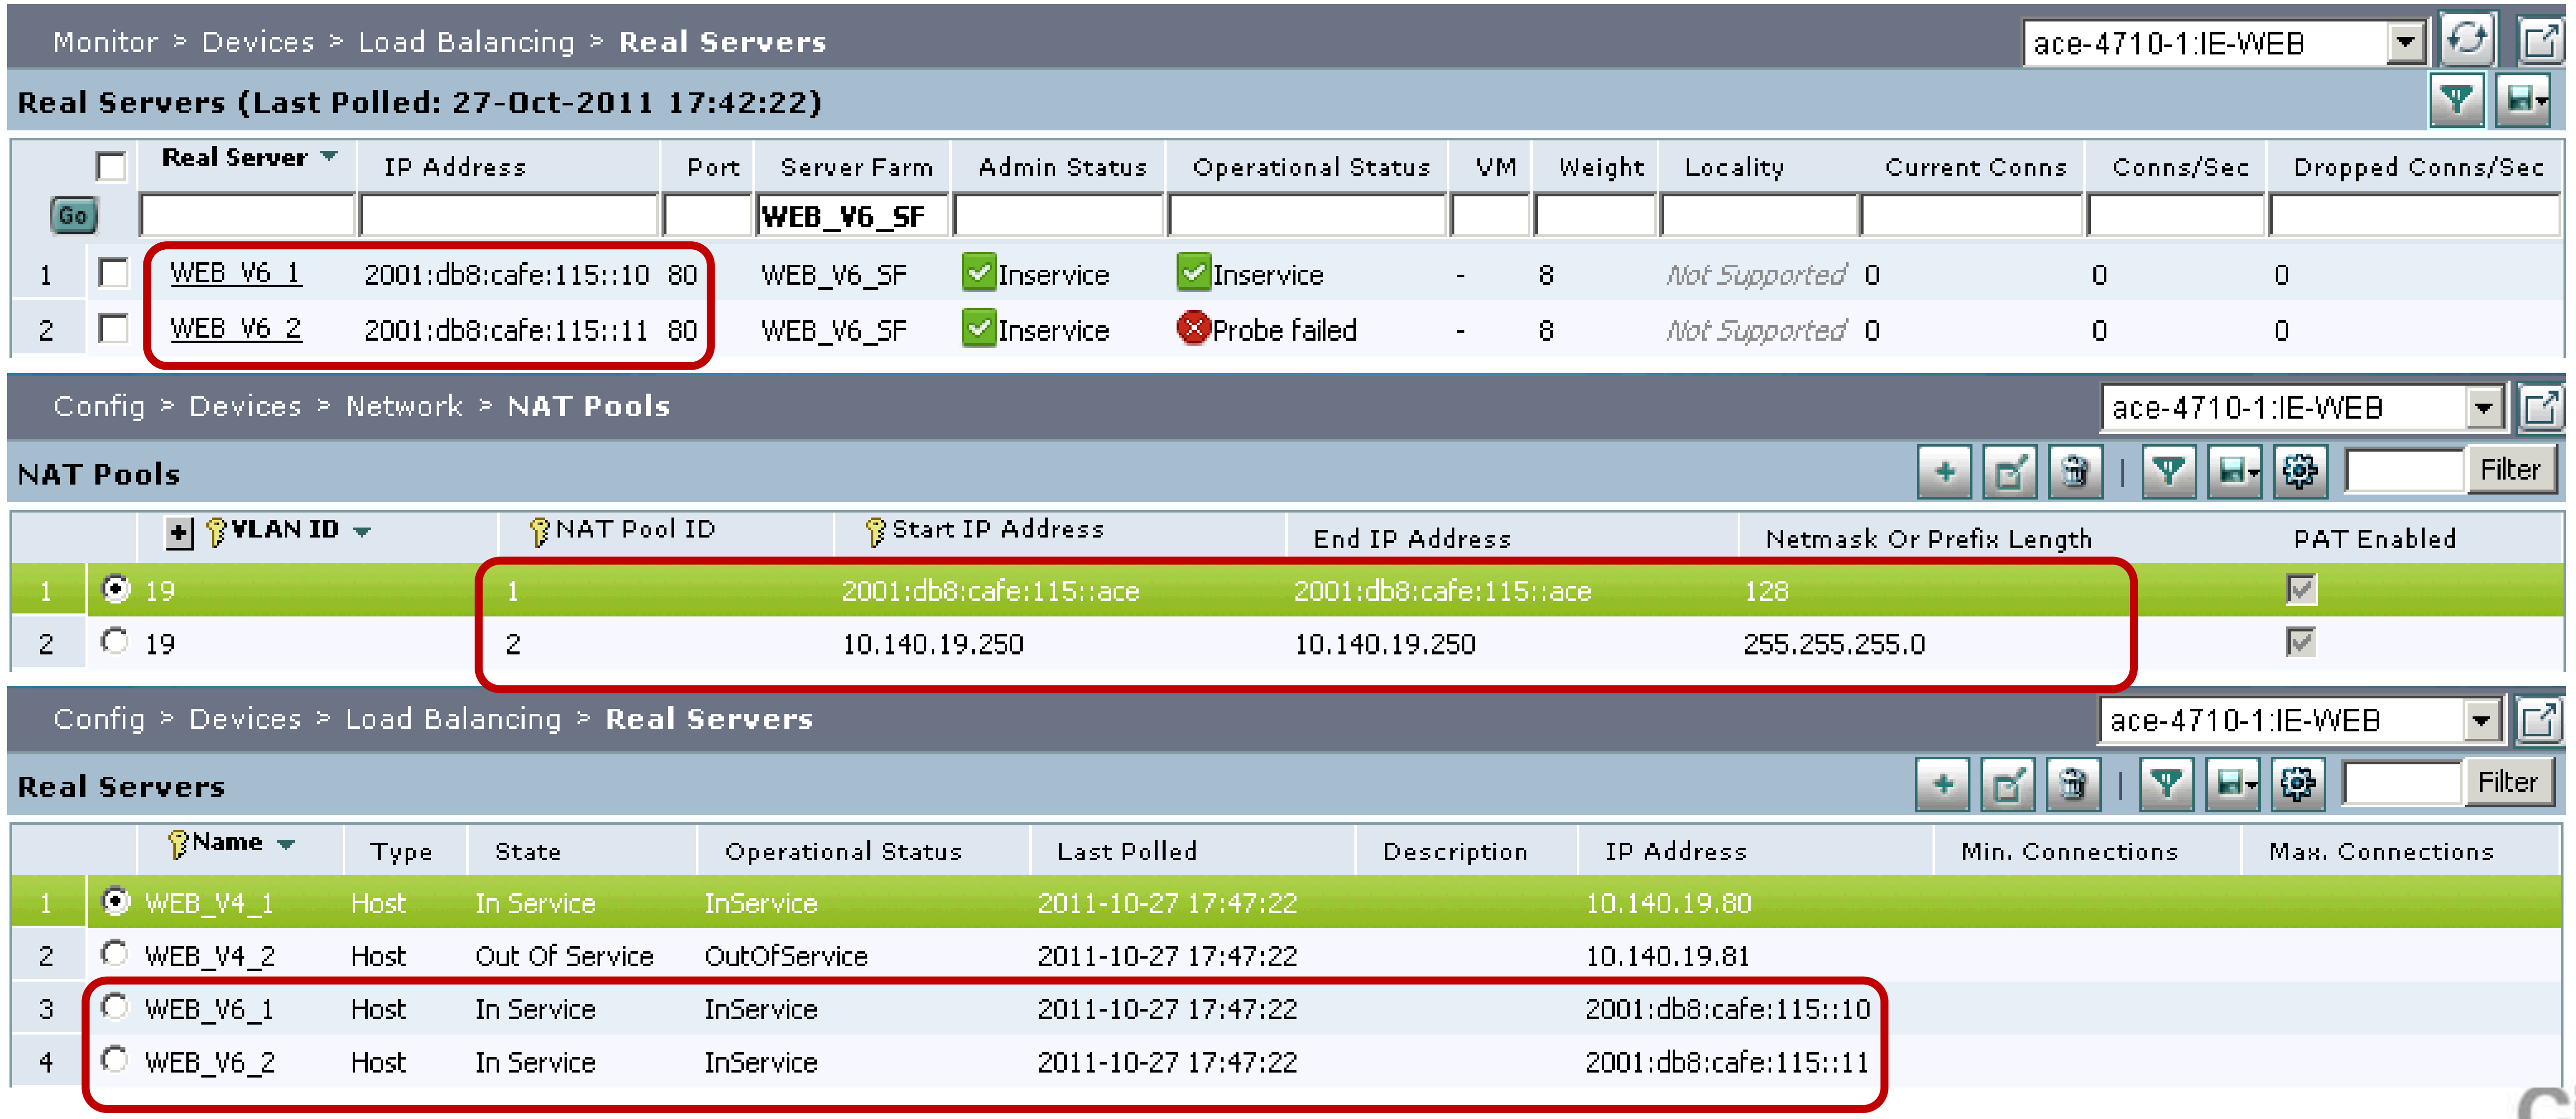

Monitor > Devices > Load Balancing > **Real Servers** — ace-4710-1:IE-WEB

Real Servers (Last Polled: 27-Oct-2011 17:42:22)

	Real Server	IP Address	Port	Server Farm	Admin Status	Operational Status	VM	Weight	Locality	Current Conns	Conns/Sec	Dropped Conns/Sec
Go				WEB_V6_SF								
1	WEB_V6_1	2001:db8:cafe:115::10	80	WEB_V6_SF	Inservice	Inservice	-	8	Not Supported	0	0	0
2	WEB_V6_2	2001:db8:cafe:115::11	80	WEB_V6_SF	Inservice	Probe failed	-	8	Not Supported	0	0	0

Config > Devices > Network > **NAT Pools** — ace-4710-1:IE-WEB

NAT Pools

	VLAN ID	NAT Pool ID	Start IP Address	End IP Address	Netmask Or Prefix Length	PAT Enabled
1	19	1	2001:db8:cafe:115::ace	2001:db8:cafe:115::ace	128	☑
2	19	2	10.140.19.250	10.140.19.250	255.255.255.0	☑

Config > Devices > Load Balancing > **Real Servers** — ace-4710-1:IE-WEB

Real Servers

	Name	Type	State	Operational Status	Last Polled	Description	IP Address	Min. Connections	Max. Connections
1	WEB_V4_1	Host	In Service	InService	2011-10-27 17:47:22		10.140.19.80		
2	WEB_V4_2	Host	Out Of Service	OutOfService	2011-10-27 17:47:22		10.140.19.81		
3	WEB_V6_1	Host	In Service	InService	2011-10-27 17:47:22		2001:db8:cafe:115::10		
4	WEB_V6_2	Host	In Service	InService	2011-10-27 17:47:22		2001:db8:cafe:115::11		

Access Layer

Reference

Nexus 5000 – We are doing basic management access

```
vrf context management
 ipv6 route 0::/0 fe80::0005:73ff:fea0:0002 mgmt0

interface mgmt0
 ipv6 address 2001:0db8:cafe:011a::0030/64
```

Catalyst 4900M

```
interface Vlan24
 ipv6 address 2001:DB8:CAFE:11A::12/64
!
ipv6 route ::/0 Vlan24 FE80::5:73FF:FEA0:2
```

Nexus 1000v

```
interface mgmt0
  ipv6 address 2001:0db8:cafe:011a::0013/64
!
ipv6 route 0::/0 fe80::0005:73ff:fea0:0002 mgmt0
```

VMware ESXi – IPv6 (1)

Host/vCenter do not need IPv6 for Guest VMs to use IPv6!

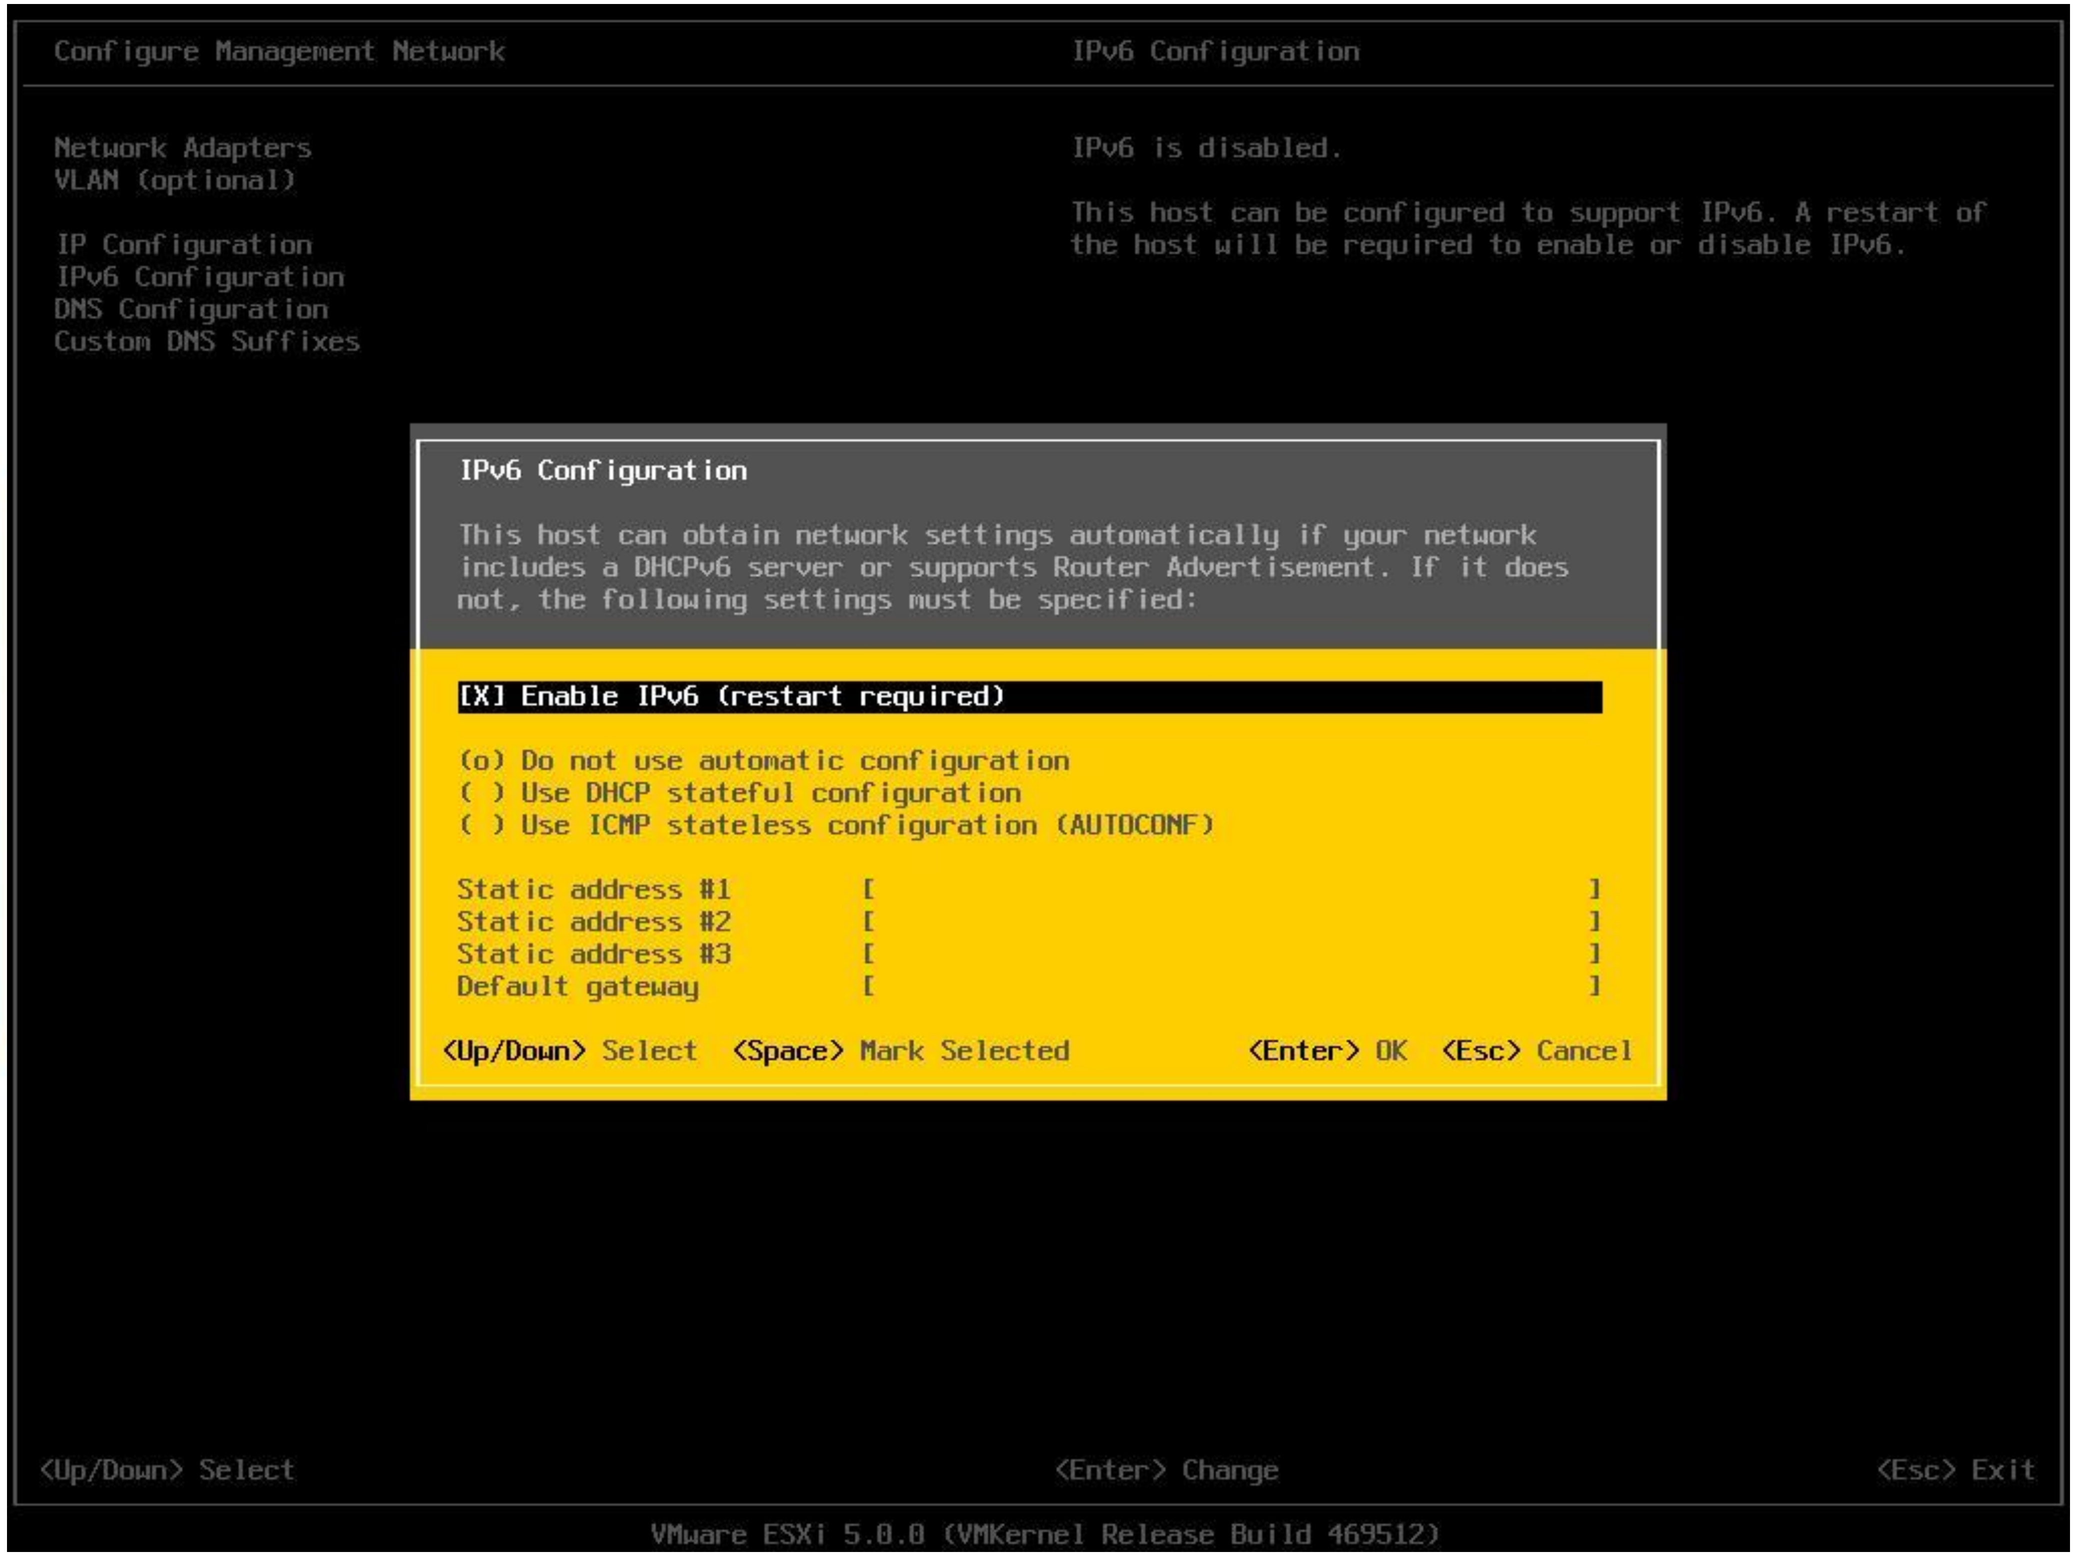

VMware ESXi – IPv6 (2)

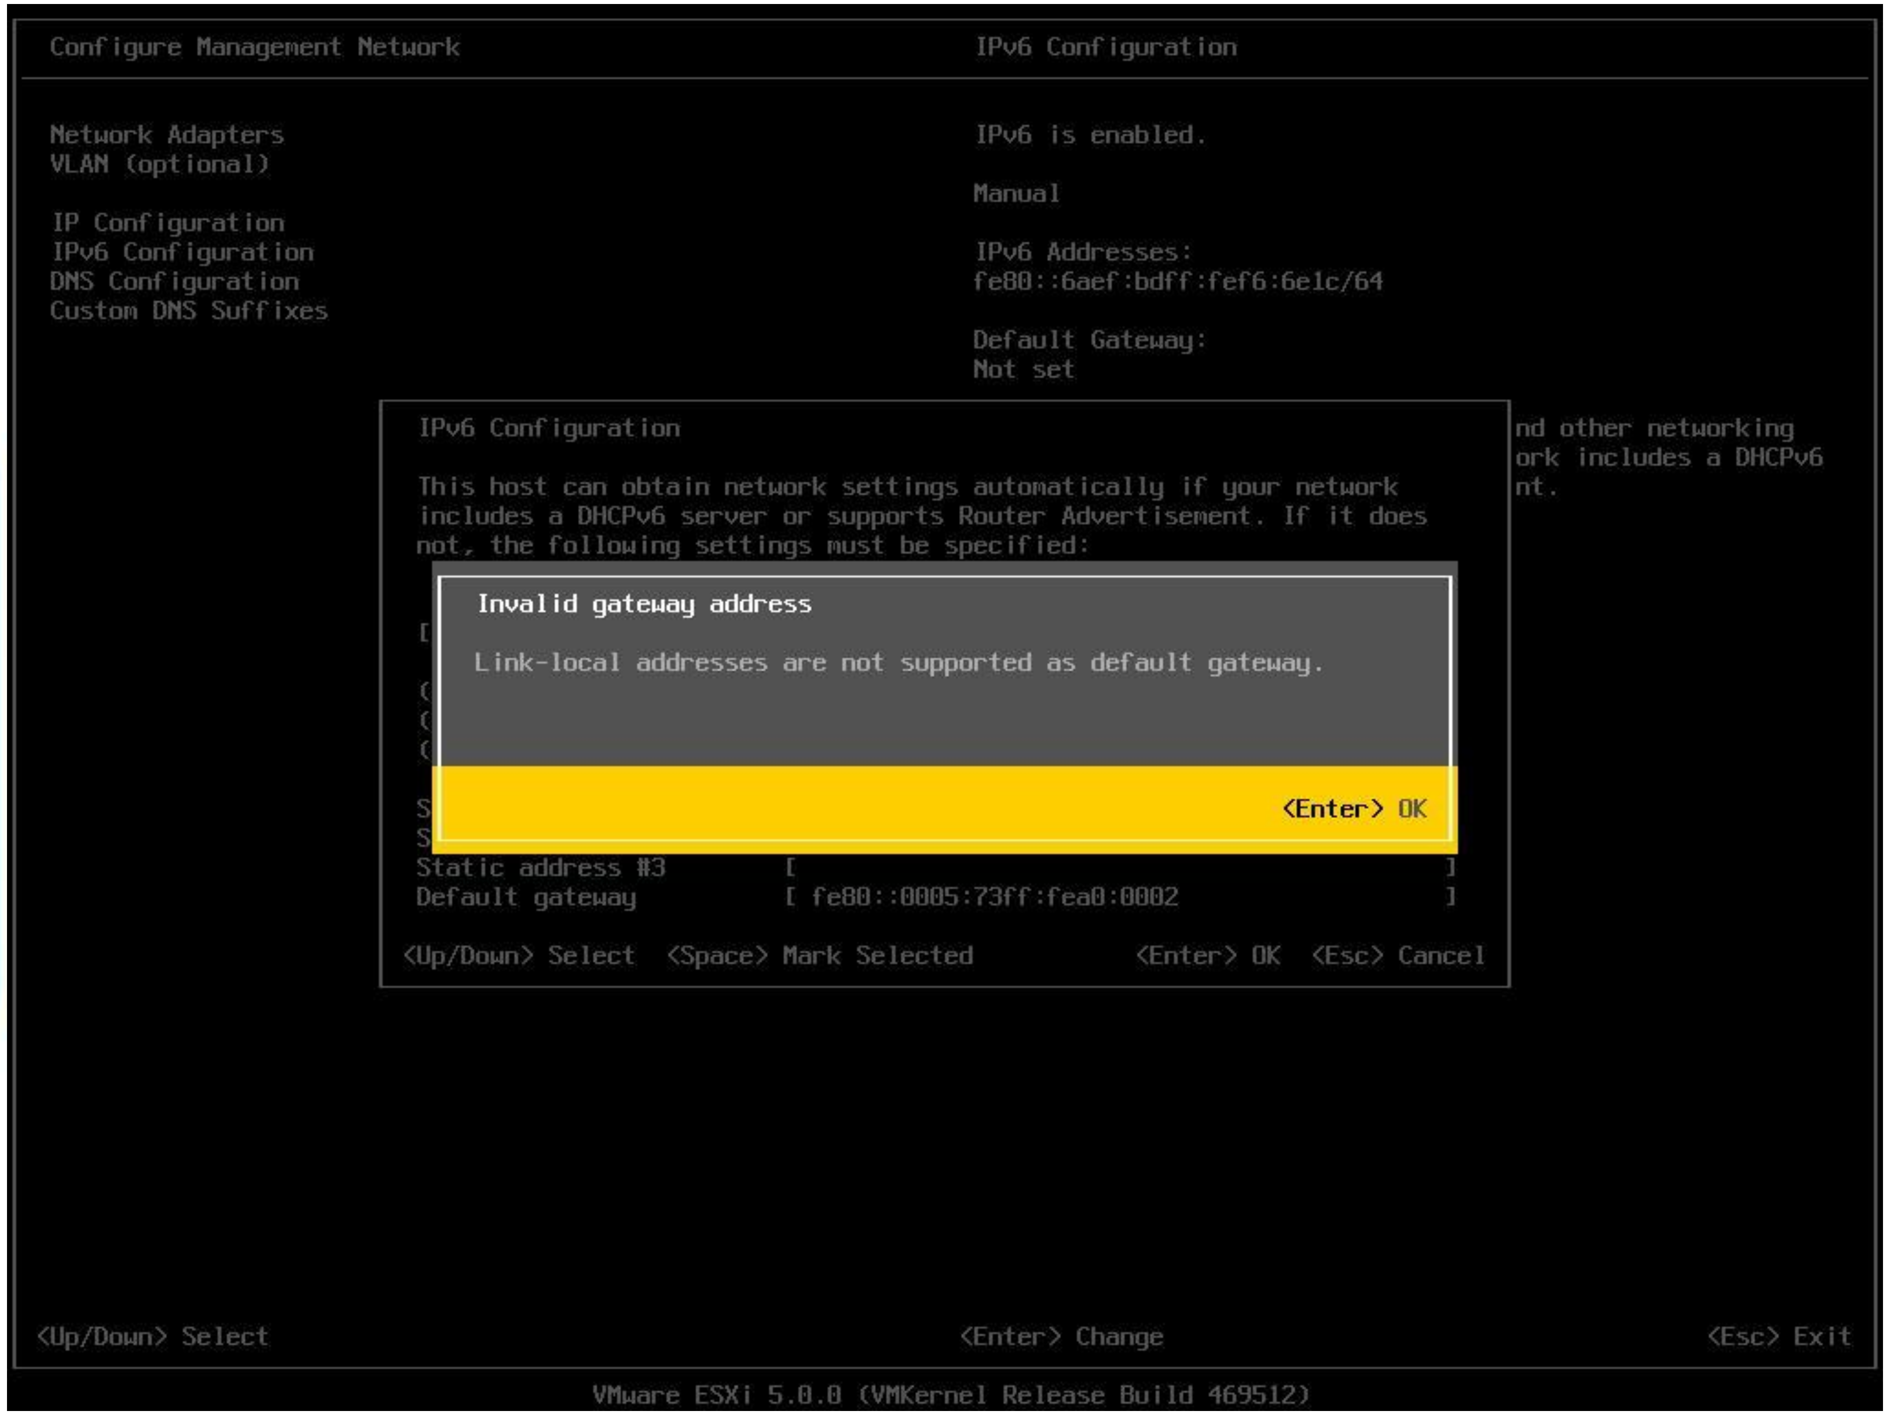

- As of ESX 5 you cannot set a LL address as a gateway
- VERY BAD
- Global or let it learn via RA

VMware ESXi – IPv6 (3)

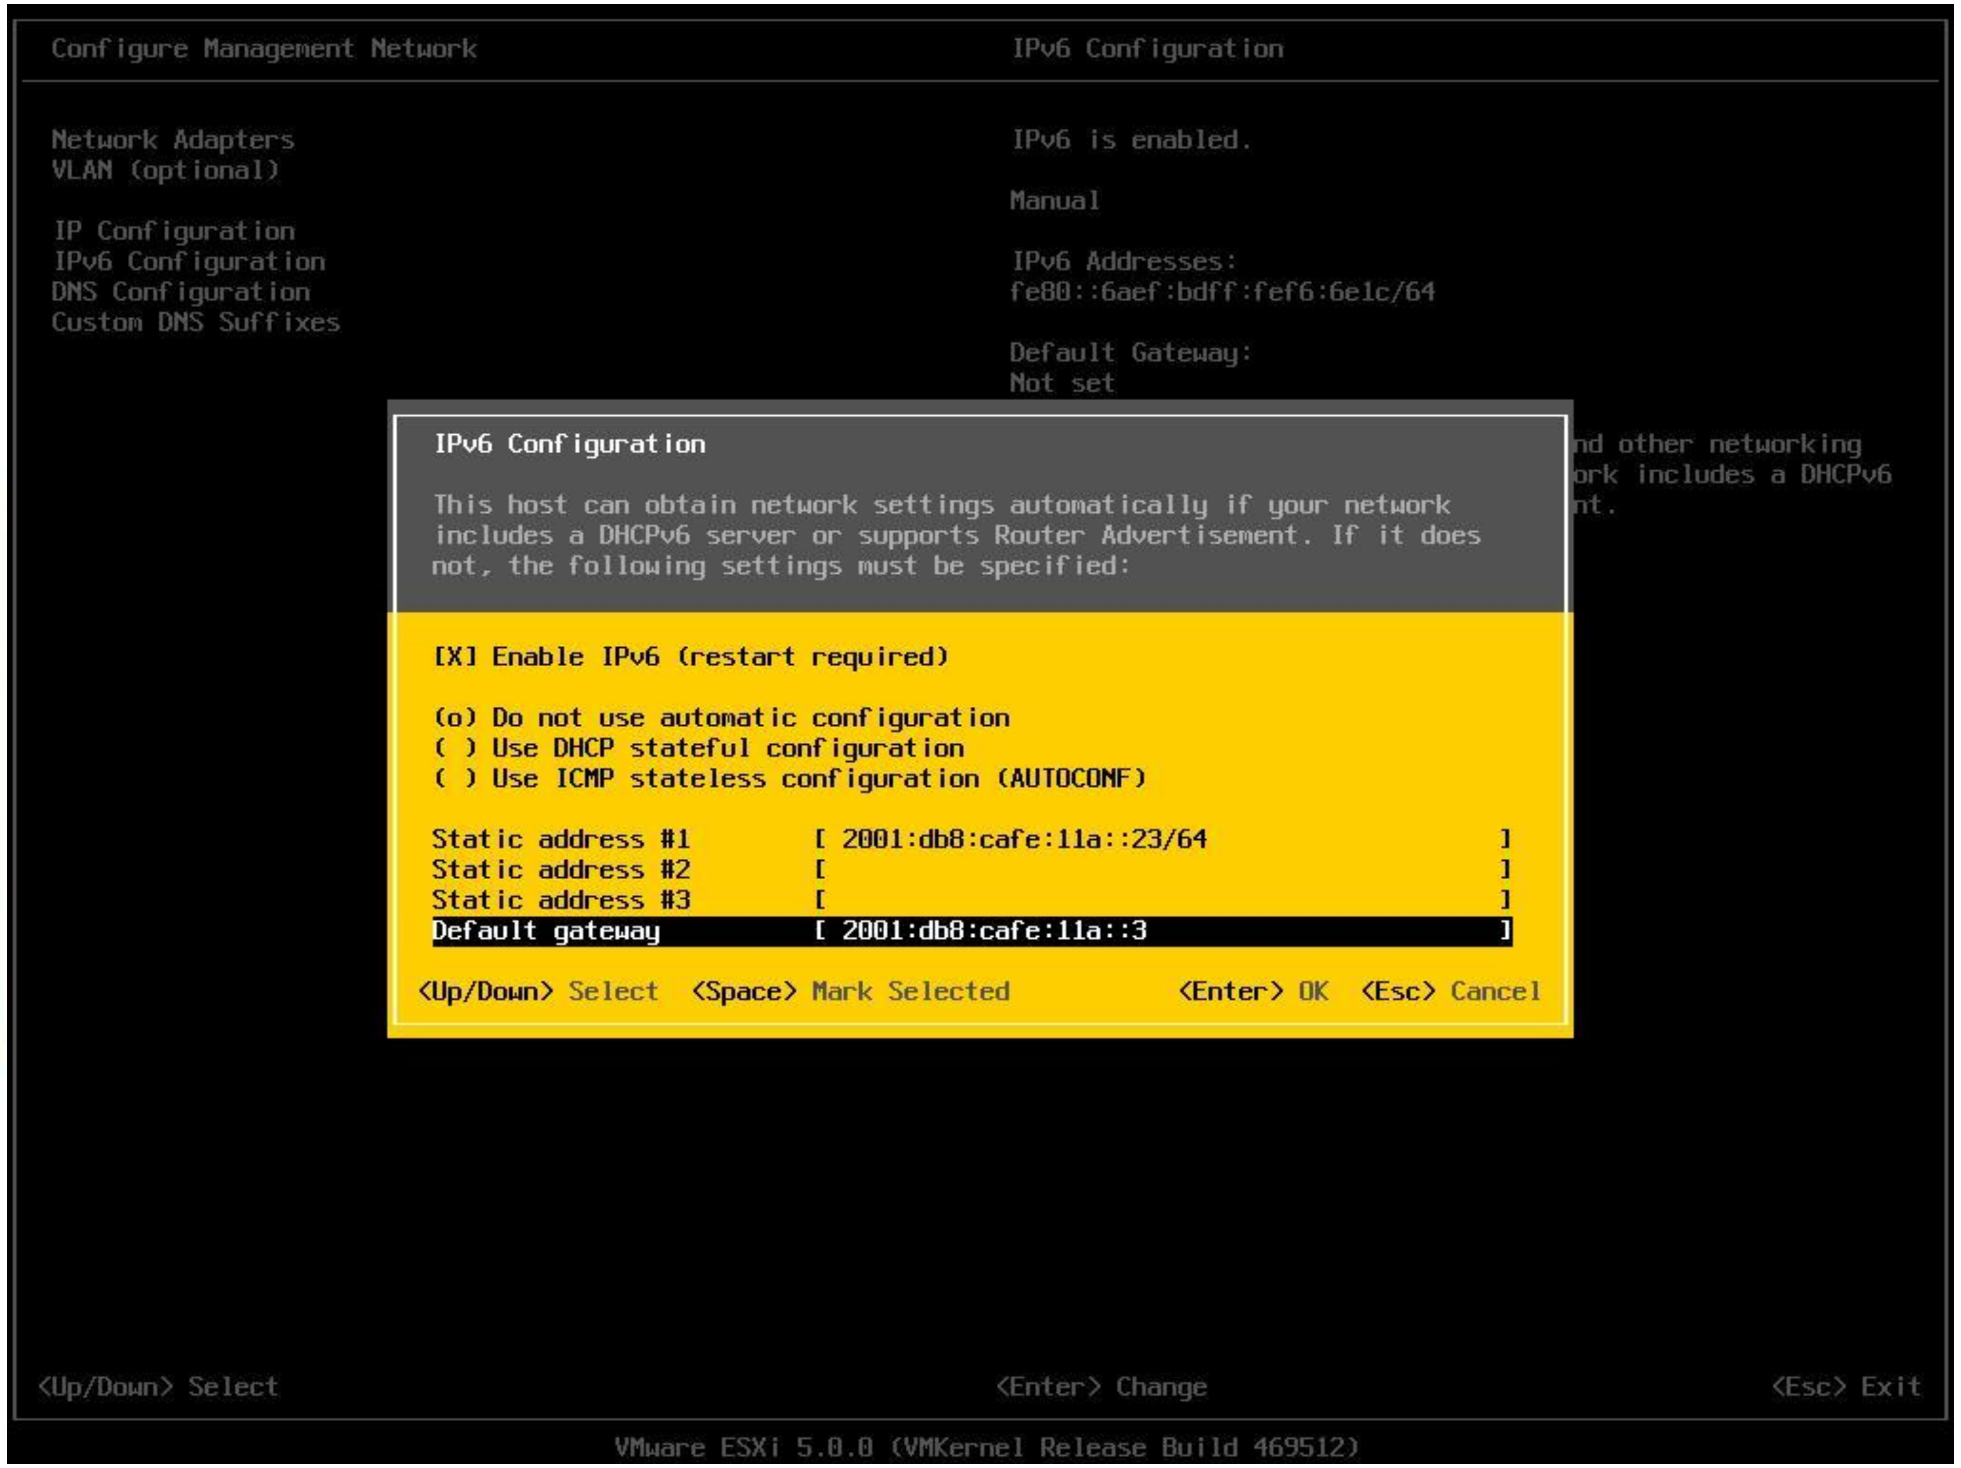

- Single GW or if GW can support FHRP on Global = OK
- If not, let host learn GW via RA (Test this!!)

VMware vCenter – IPv6 (1)

- Configure vCenter OS with IPv6 address, GW, DNS
- Login to vCenter host by name or address

VMware vCenter – IPv6 (2)

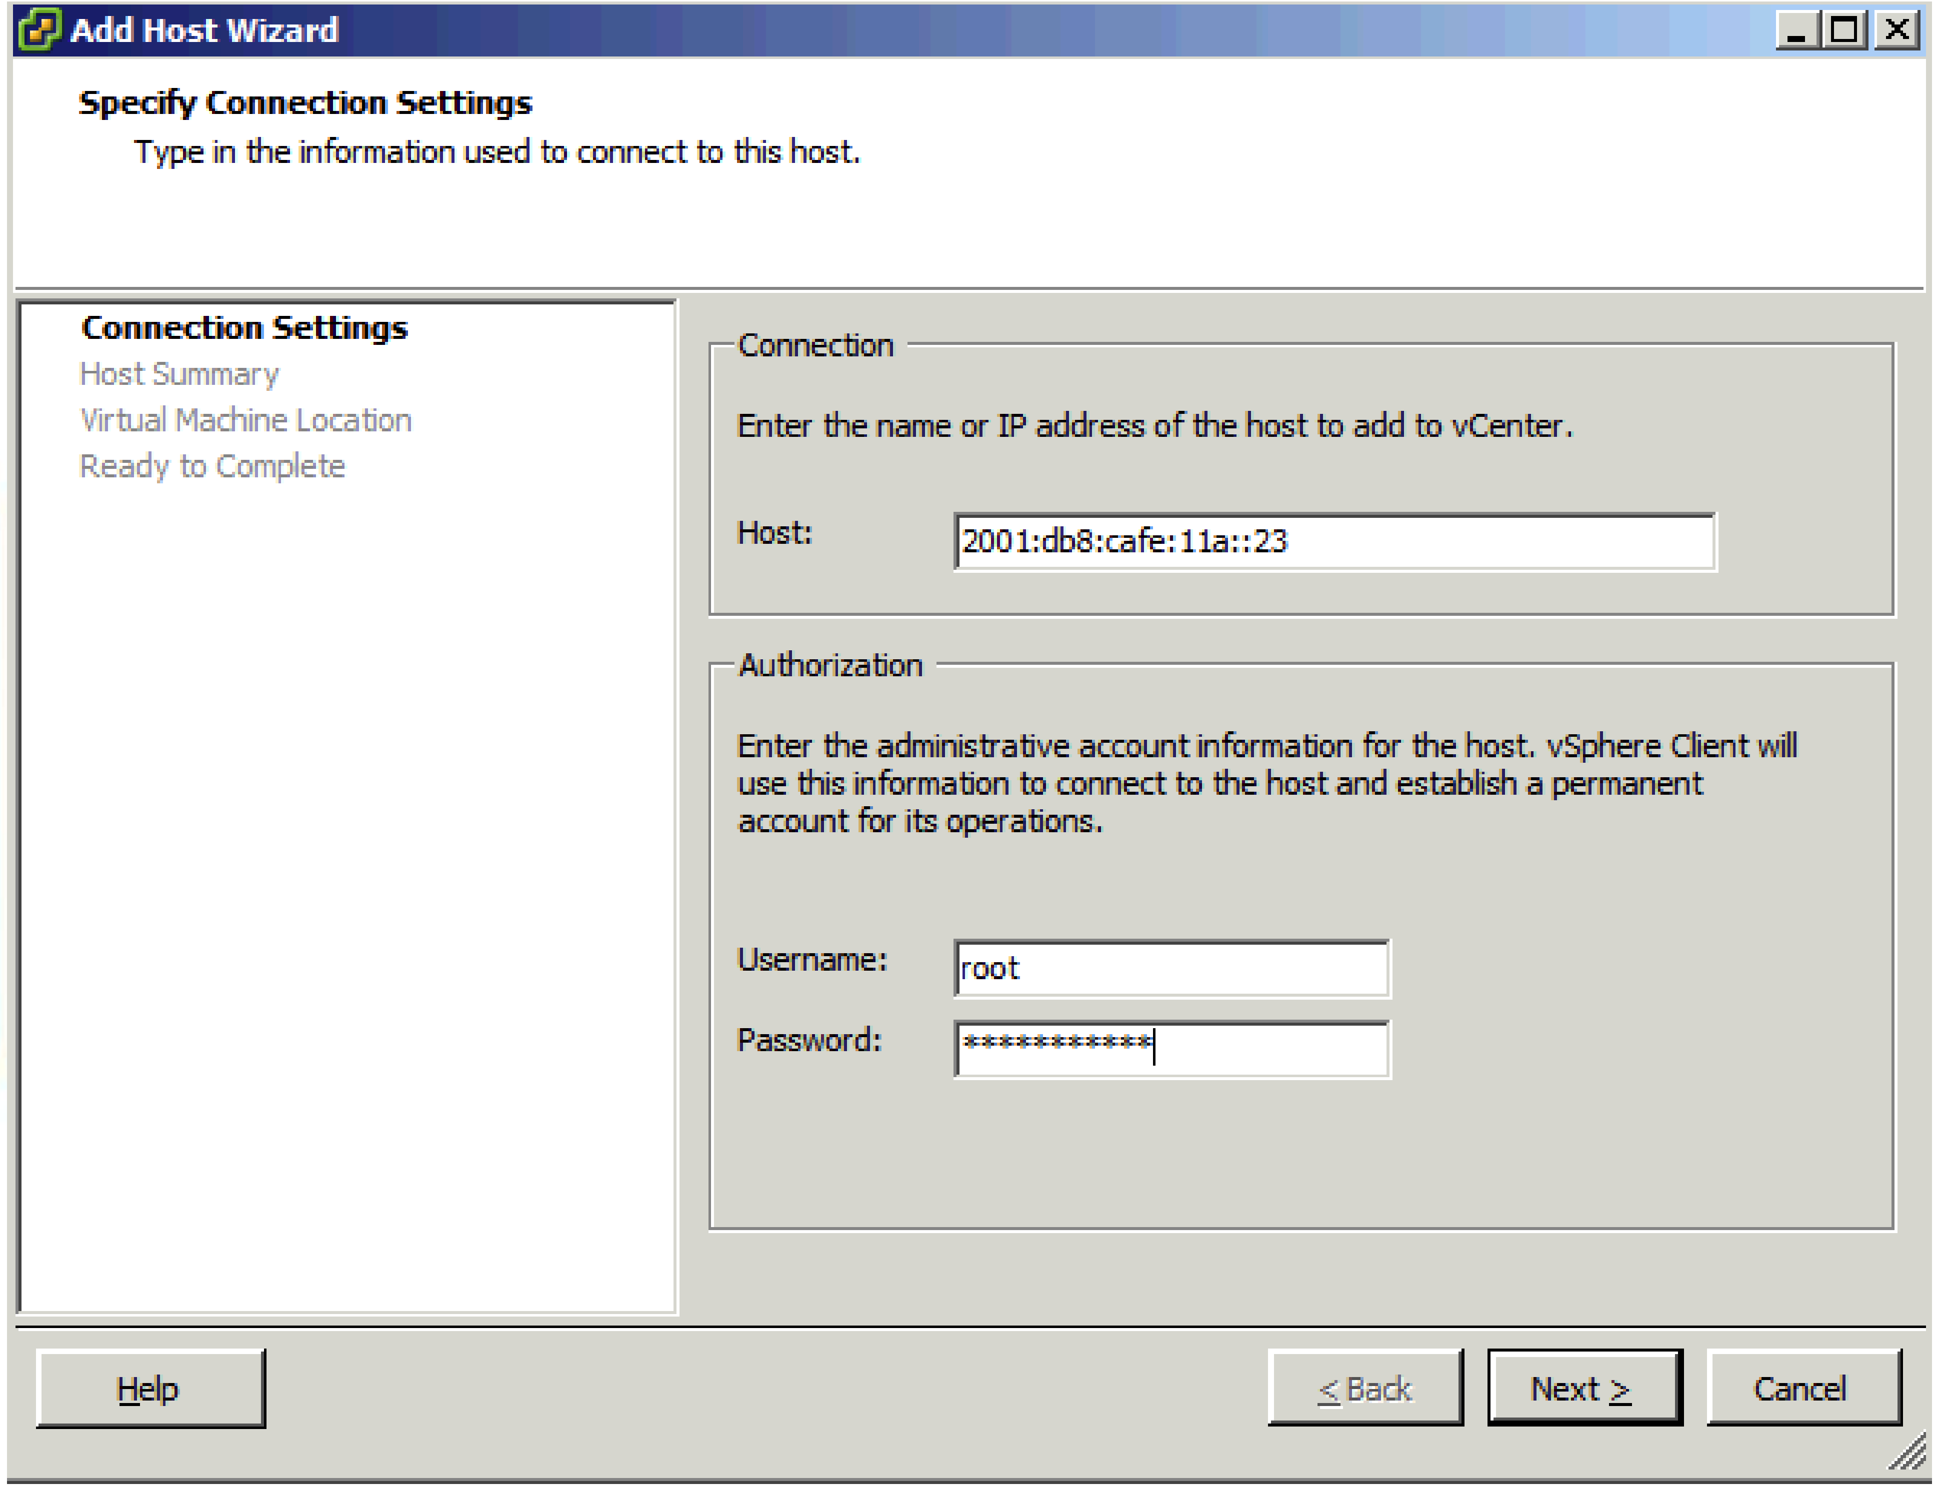

- Add host using name or IPv4 or IPv6 address

Other Stuff

- NetFlow
- Deep Packet Inspection
- Email, DNS, other apps
- More comprehensive security recommendations
  - Blocking routing type 0
  - Blocking HbH to protect against CPU spikes
  - uRPF – different capabilities based on platform
  - no ipv6 source-route – not on by default prior to 12.4(15)T
  - Normal bogon filters
  - Basically, all usual IPv4 stuff plus platform/code specific CLI or security-focused differences
  - Pick up copy of “IPv6 Security” by Eric Vyncke and Scott Hogg
  - Book MTE time with Eric Vyncke ;-)
- NPTv6 for single address space multi-homing configurations

BRKSEC-2003

Multi-homed – SLB64

Multihomed – SLB64

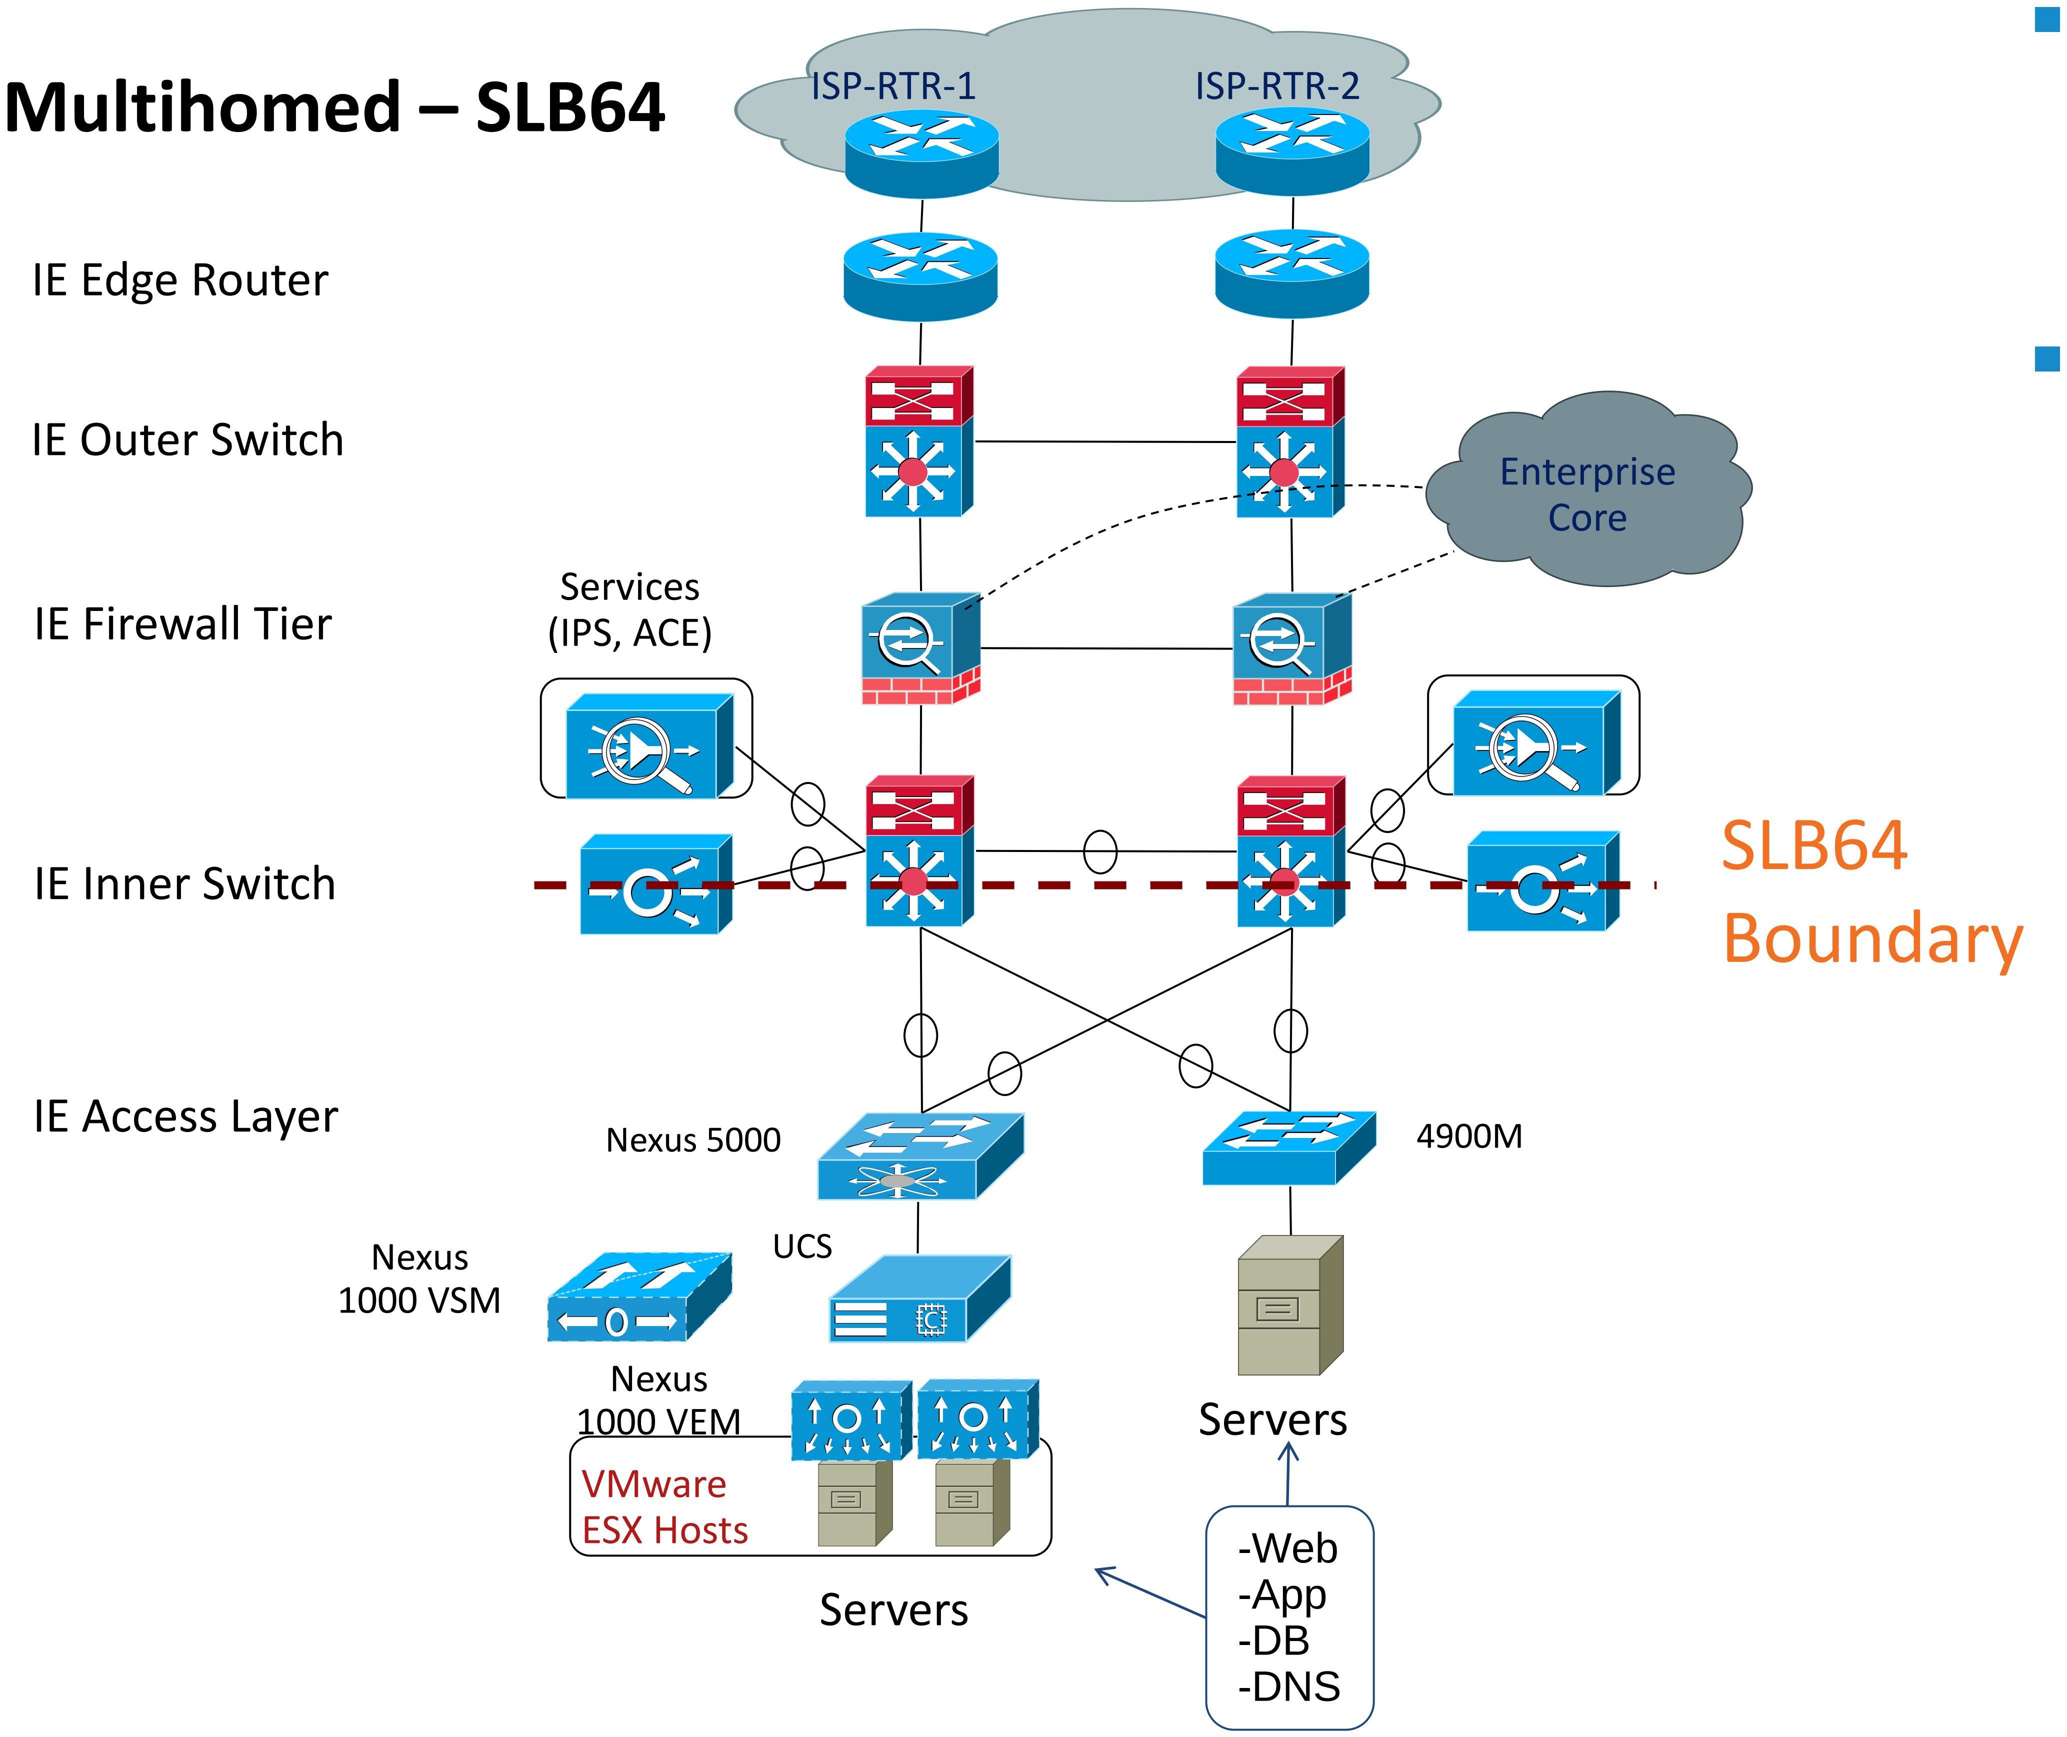

- Dual stack to the Cisco ACE
- IPv4-only South of Cisco ACE

Reference

Cisco ACE – Context Definition

Trunked Interface – One-arm Mode

```
interface gigabitEthernet 1/3
  description to IE-Trunk
  switchport trunk allowed vlan 19-22,24
  no shutdown
```

VLAN for Management

```
interface vlan 24
  ipv6 enable
  ip address 2001:db8:cafe:11a::b/64
  alias 2001:db8:cafe:11a::d/64
  peer ip address 2001:db8:cafe:11a::c/64
  access-group input ALL
  service-policy input remote_mgmt_allow_policy
  no shutdown
```

Define Context

```
context IE-WEB
  allocate-interface vlan 19
```

- Nothing changes from previous SLB66 example

Reference

Cisco ACE – Fault Tolerance (over IPv4)

FT Interface or just trunk with Port Channel

```
interface gigabitEthernet 1/4
  ft-port vlan 132
  no shutdown
```

```
ft interface vlan 132
  ip address 10.140.132.1 255.255.255.0
  peer ip address 10.140.132.2 255.255.255.0
  no shutdown

ft peer 1
  heartbeat interval 300
  heartbeat count 10
  ft-interface vlan 132
  query-interface vlan 19
ft group 2
  peer 1
  priority 110
  associate-context IE-WEB
  inservice
```

- As of ACE A5(1.0) FT is over IPv4

Reference

IE-WEB Context - MGMT

```
class-map type management match-any MGMT-CM
  2 match protocol xml-https any
  3 match protocol https any
  4 match protocol ssh any
  5 match protocol snmp any
  6 match protocol icmp any
  7 match protocol http any
  8 match protocol telnet any
class-map type management match-any MGMT-CM-v6
  2 match protocol icmpv6 anyv6

policy-map type management first-match MGMT
  class MGMT-CM
    permit
  class MGMT-CM-v6
    permit
interface vlan 19
  service-policy input MGMT
```

IP Access through the Cisco ACE

```
access-list EVERYONE line 10 extended permit icmp any any
access-list EVERYONE line 20 extended permit ip any any
access-list EVERYONE-v6 line 8 extended permit icmpv6 anyv6 anyv6
access-list EVERYONE-v6 line 16 extended permit ip anyv6 anyv6
interface vlan 19
 access-group input EVERYONE
 access-group input EVERYONE-v6
```

SLB64 Context Specific Configurations

```
probe http WEB_V4_PROBE
  interval 15
  passdetect interval 5
  request method get url /probe.html
  expect status 200 200
  open 1

rserver host WEB_V4_1
  ip address 10.140.19.80
  inservice
rserver host WEB_V4_2
  ip address 10.140.19.81
  inservice
serverfarm host WEB_V6_V4_SF
  predictor leastconns slowstart 300
  probe WEB_V4_PROBE
  rserver WEB_V4_1 80
    inservice
  rserver WEB_V4_2 80
    inservice
```

1 2 3

```
class-map match-all WEB_V6_V4_VIP
  2 match virtual-address 2001:db8:cafe:115::a tcp eq www

policy-map type loadbalance first-match WEB_V6_V4_SLB
  class class-default
    serverfarm WEB_V6_V4_SF
    nat dynamic 2 vlan 19 serverfarm primary
    insert-http x-forward header-value "%is"
policy-map multi-match WEB_V6_V4_POL
  class WEB_V6_V4_VIP
    loadbalance vip inservice
    loadbalance policy WEB_V6_V4_SLB
    loadbalance vip icmp-reply active
interface vlan 19
  ipv6 enable
  ip address 2001:db8:cafe:115::d/64
  ip address 10.140.19.13 255.255.255.0
  access-group input EVERYONE
  access-group input EVERYONE-v6
  nat-pool 2 10.140.19.250 10.140.19.250 netmask
255.255.255.0 pat
  service-policy input MGMT
  service-policy input WEB_V6_V4_POL
```

4 5 6

Reference

SSL Offload

```
class-map match-all WEB_V6_VIP
  2 match virtual-address 2001:db8:cafe:115::a tcp eq https

ssl-proxy service SSL_PROXY_WEB
  key cisco-sample-key
  cert cisco-sample-cert

policy-map multi-match WEB_V6_POL
  class WEB_V6_VIP
    loadbalance vip inservice
    loadbalance policy WEB_V6_SLB
    loadbalance vip icmp-reply active
    nat dynamic 1 vlan 19
    ssl-proxy server SSL_PROXY_WEB
```

- Nothing changes from previous SLB66 example
- ‘North’ bound VIP is still IPv6

Reference

Health Monitoring (Probes) - IPv4 Real Servers

```
ace-4710-1/IE-WEB# sh probe

 probe        : WEB_V4_PROBE
 type         : HTTP
 state        : ACTIVE
----------------------------------------------
   port       : 80           address    : 0.0.0.0
   addr type  : -            interval   : 15      pass intvl   : 5
   pass count: 3             fail count: 3       recv timeout: 10
                ------------------ probe results ------------------
   associations     ip-address        port porttype probes failed passed health
   ------------ ---------------------+----+--------+------+------+------+------
   serverfarm  : WEB_V6_V4_SF
     real      : WEB_V4_1[80]
                           10.140.19.80    80  REAL     32      0      32     SUCCESS
     real      : WEB_V4_2[80]
                           10.140.19.81    80  REAL     32      0      32     SUCCESS
```

Reference

Application Networking Manager 5.1

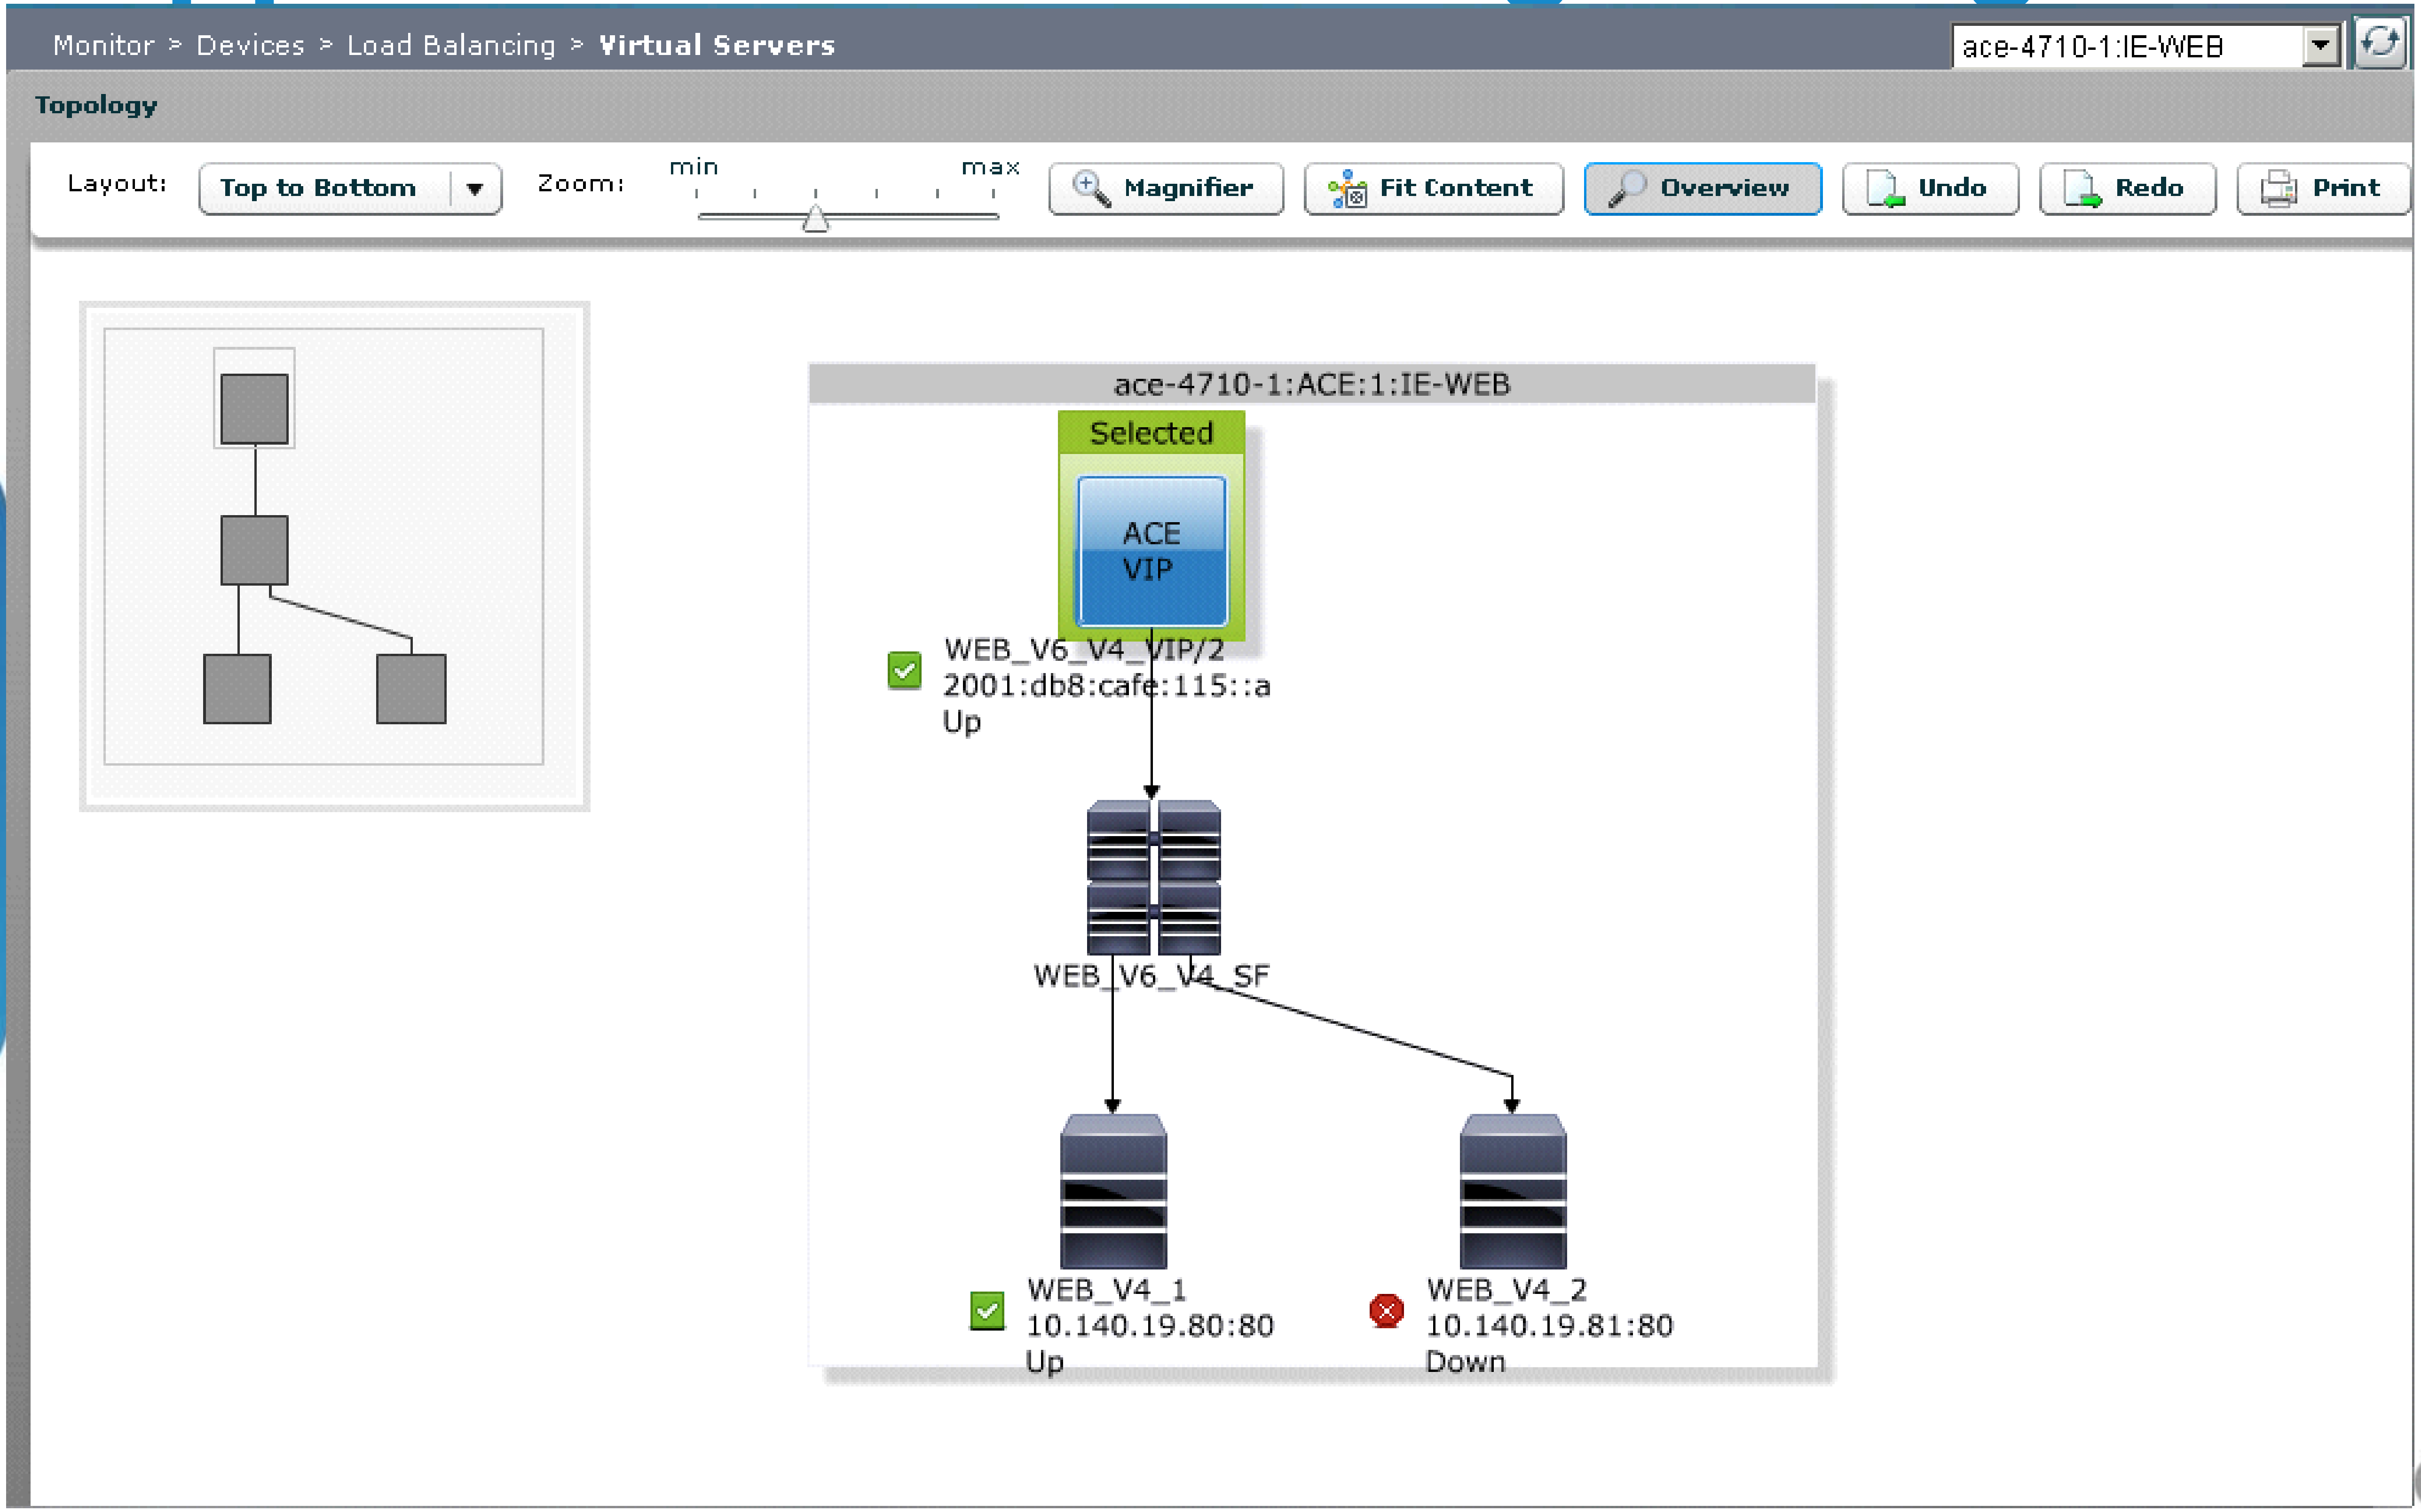

Validation of Connection

```
ace-4710-1/IE-WEB# show conn
conn-id    np dir proto source                                     sport state
                 vlan  destination                                dport
----------+--+---+-----+------------------------------------------+-----+------+
1640630    1  in  TCP   2001:db8:ea5e:1:49fa:b11a:aaf8:91a5         54911 ESTAB
                 19    2001:db8:cafe:115::a                         80
1647396    1  out TCP   10.140.19.80                                80    ESTAB
                 19    10.140.19.250                                1025
```

Client-2-VIP

Svr-2-SNAT

X-Forwarded-For

- By default the source IP of client requests that are logged will be the SNAT or other NAT’ed address
- You want to log the real source address – X-Forwarded-For (XFF) in HTTP
- e of XFF

```
cisco@ie-web-01:/$ tail -f /var/log/apache2/access.log
10.140.19.250 - - [25/Oct/2011:11:41:03 -0600] "GET / HTTP/1.1" 304 210
"-" "Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 6.1; Trident/4.0;
SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729;
Media Center PC 6.0; .NET4.0C)”
```

```
serverfarm WEB_V6_V4_SF
  insert-http x-forward header-value "%is"
```

ACE Policy Map – “is” = Source IP Address

```
Hypertext Transfer Protocol
    GET / HTTP/1.1\r\n
    x-forward: 2001:db8:ea5e:1:49fa:b11a:aaf8:91a5\r\n
```

Multi-homed – Stateful NAT64

Multihomed – NAT64

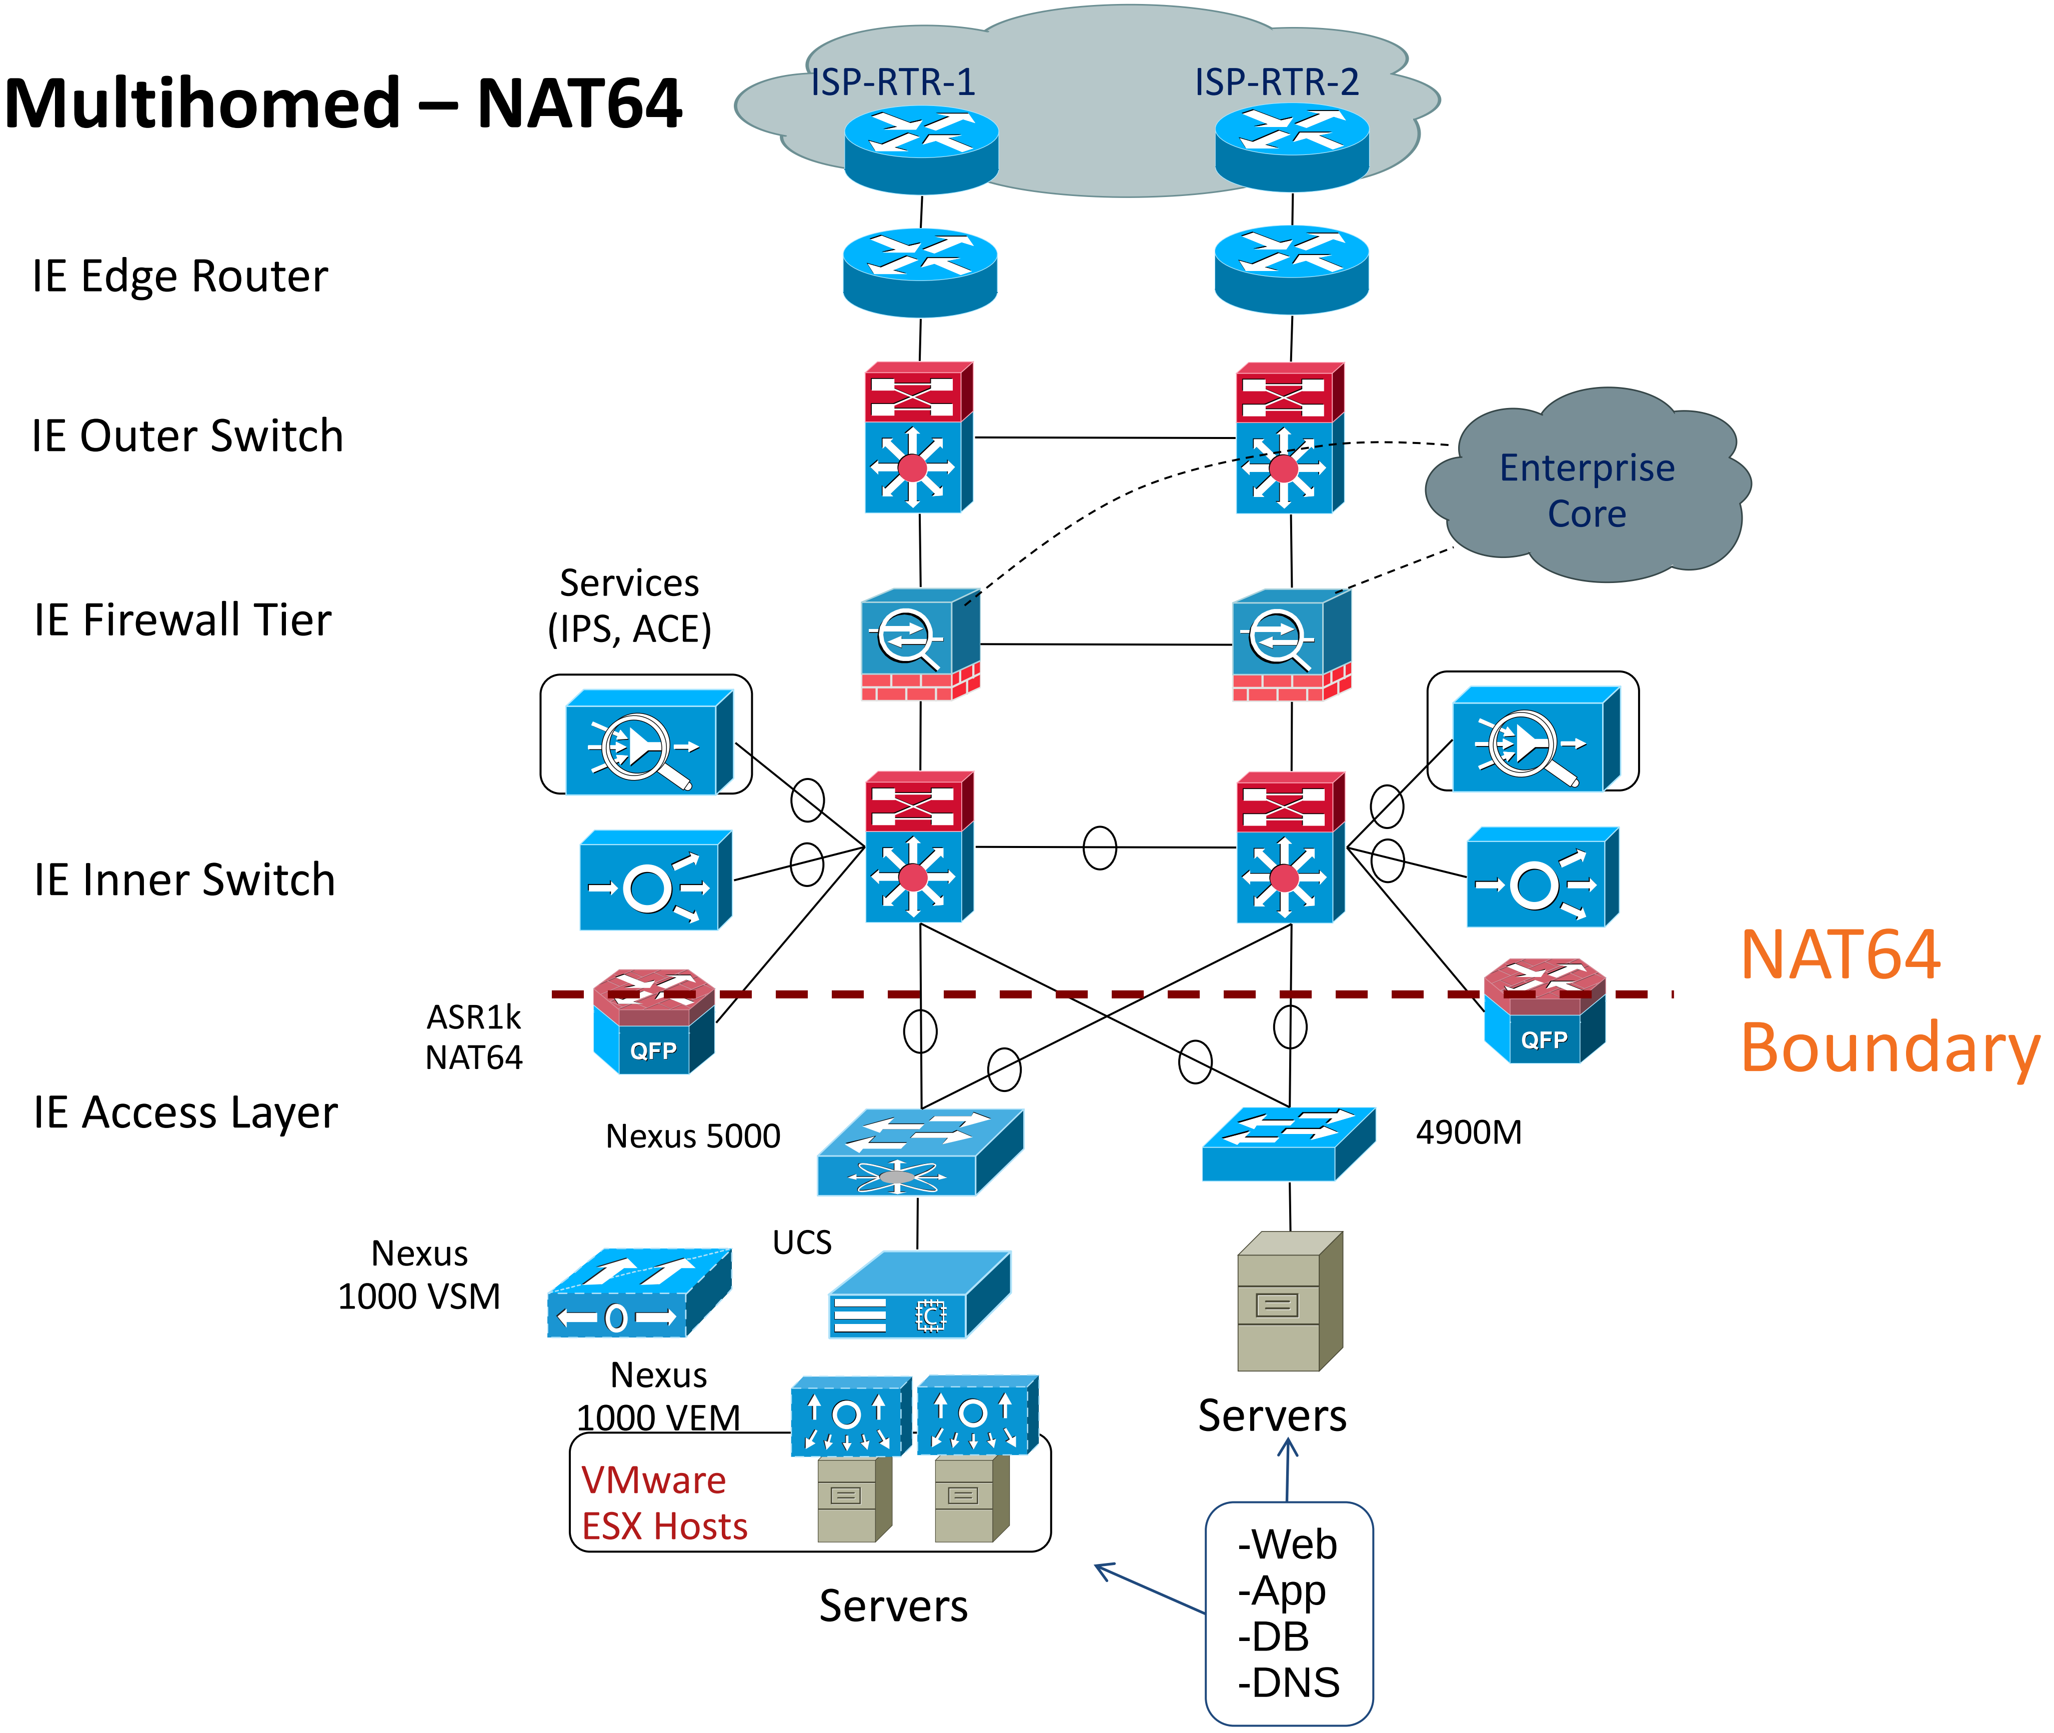

Reference

NAT64

- Lots of RFCs to check out:
  - RFC 6144 – Framework for IPv4/IPv6 Translation
  - RFC 6052 – IPv6 Addressing of IPv4/IPv6 Translators
  - RFC 6145 – IP/ICMP Translation Algorithm
  - RFC 6146 – Stateful NAT64
  - RFC 6147 – DNS64
- Stateless – Not your friend in the enterprise (corner case deployment)
  - 1:1 mapping between IPv6 and IPv4 addresses (i.e. 254 IPv6 hosts-to-254 IPv4 hosts)
  - Requires the IPv6-only hosts to use an “IPv4 translatable” address format
- Stateful – What we are after for translating IPv6-only hosts to IPv4-only host(s)
  - It is what it sounds like – keeps state between translated hosts
  - Several deployment models (PAT/Overload, Dynamic 1:1, Static, etc…)
  - This is what you will use to translate from IPv6 hosts (internal or Internet) to IPv4-only servers (internal DC or Internet Edge)
- Cisco NAT64 WP: http://bit.ly/poyOey

Stateful NAT64 – Example Topology

Static Example

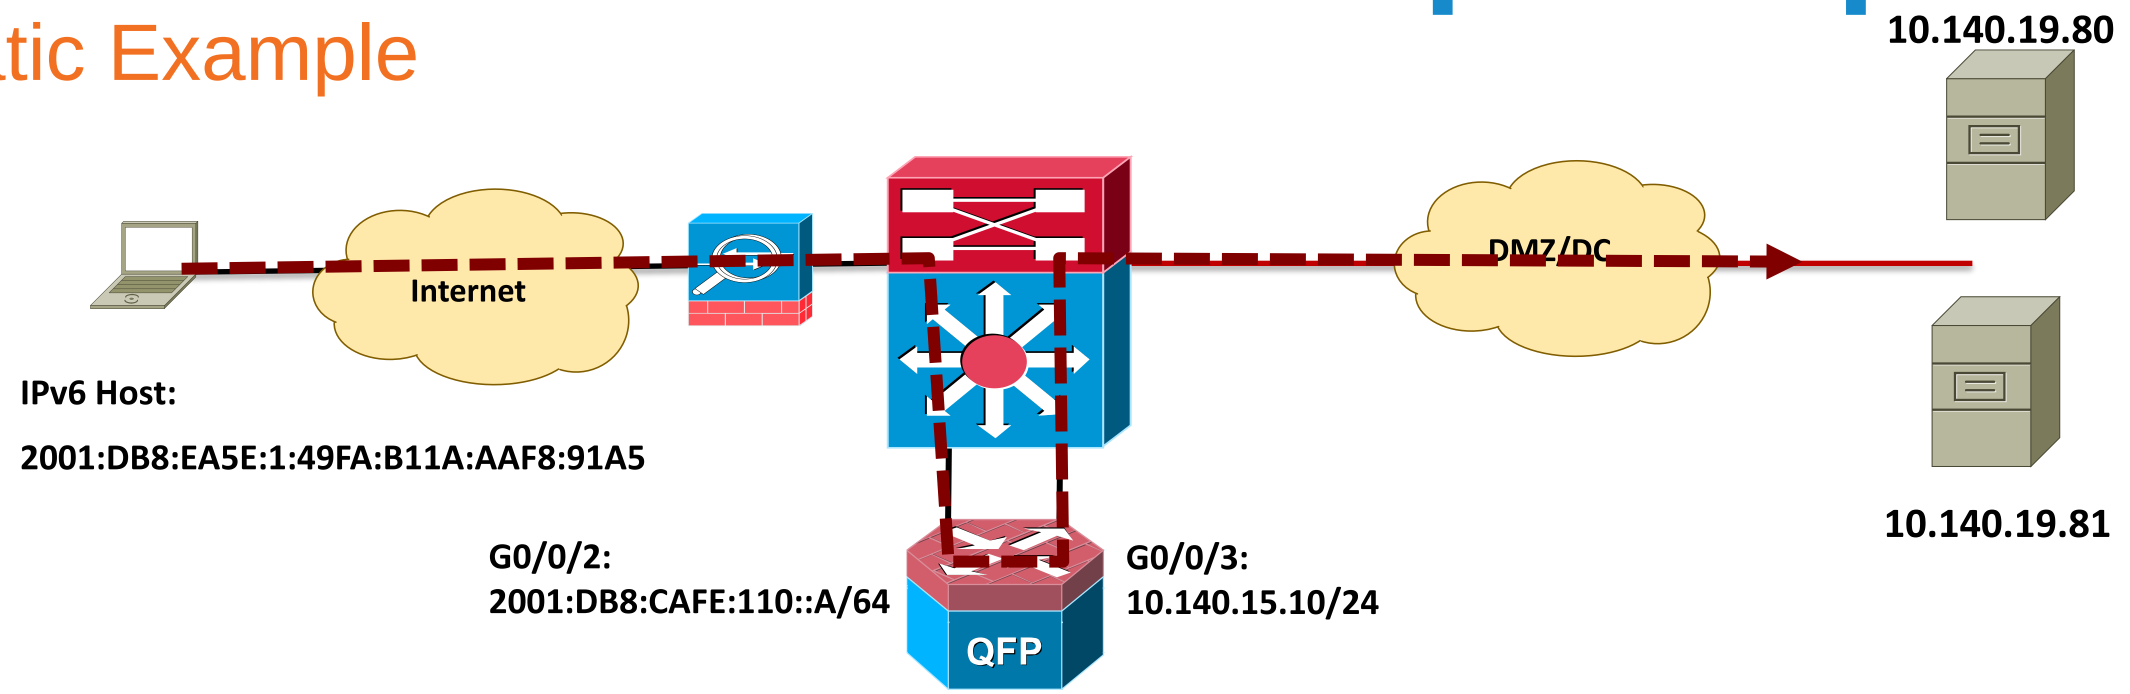


```
interface GigabitEthernet0/0/2
 description to 6k-inner-1 Outside
 no ip address
 ipv6 address 2001:DB8:CAFE:110::A/64
 nat64 enable
!
interface GigabitEthernet0/0/3
 description to 6k-inner-1 Inside
 ip address 10.140.15.10 255.255.255.0
 nat64 enable
```

```
ipv6 access-list EDGE_ACL
  permit ipv6 any host 2001:DB8:CAFE:BEEF::10
  permit ipv6 any host 2001:DB8:CAFE:BEEF::11
!
nat64 prefix stateful 2001:DB8:CAFE:BEEF::/96
nat64 v4 pool IE 10.140.15.20 10.140.15.20
nat64 v4v6 static 10.140.19.80 2001:DB8:CAFE:BEEF::10
nat64 v4v6 static 10.140.19.81 2001:DB8:CAFE:BEEF::11
nat64 v6v4 list EDGE_ACL pool IE overload
!
ipv6 route ::/0 2001:DB8:CAFE:110::10
router eigrp 10
 network 10.0.0.0
```

ASA Interfaces

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ipv6 address 2001:db8:cafe:103::3/64 standby 2001:db8:cafe:103::4
!
interface GigabitEthernet0/1.14
 vlan 14
 nameif nat64
 security-level 50
 ipv6 address 2001:db8:cafe:110::10/64 standby 2001:db8:cafe:110::11
 ipv6 enable
 ipv6 nd suppress-ra
!
ipv6 route outside ::/0 fe80::5:73ff:fea0:2
ipv6 route nat64 2001:db8:cafe:beef::/96 2001:db8:cafe:110::a
```

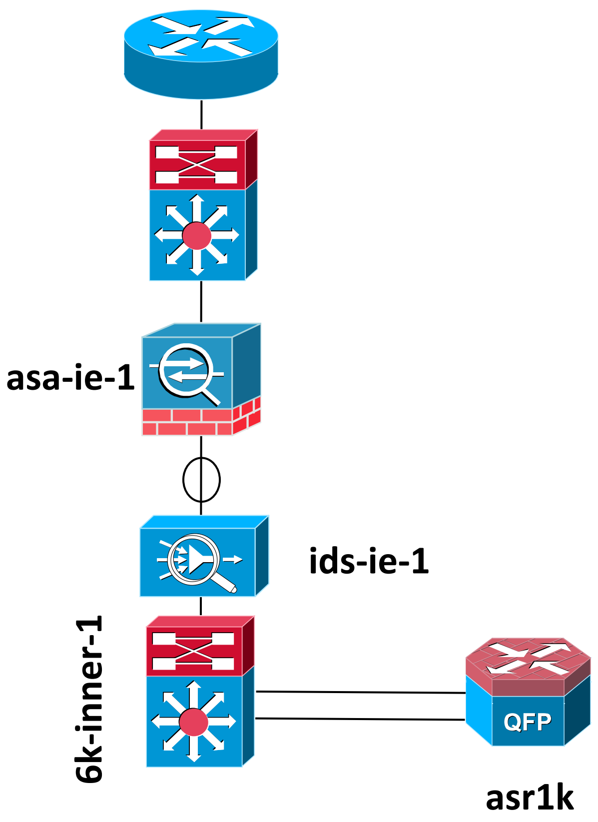

- Many connectivity types – Here, ASR is in VLAN14 that is trunked via 6k pair to the ASA pair
- If doing pure L3 P2P links to 6k then use IPv6 EIGRP to announce NAT64 prefix – here we have to do static route until ASA supports EIGRPv6 or OSPFv3

Reference

ASA Object/ACL Configuration

- External references are to the static NAT64 addresses from the “NAT64 Prefix”
- Object for each server
- ACL for L3/L4 stuff

```
object network NAT64-WEB-01
 host 2001:db8:cafe:beef::10
object network NAT64-WEB-02
 host 2001:db8:cafe:beef::11
!
ipv6 access-list outside_access_ipv6_in permit tcp any object NAT64-WEB-01 eq www
ipv6 access-list outside_access_ipv6_in permit tcp any object NAT64-WEB-02 eq www
!
access-group outside_access_ipv6_in in interface outside
```

NAT64 Translations

```
asr1k#show nat64 translations

Proto  Original IPv4        Translated IPv4
       Translated IPv6      Original IPv6
----------------------------------------------------------------------------

---    10.140.19.81         2001:db8:cafe:beef::11
       ---                  ---
---    10.140.19.80         2001:db8:cafe:beef::10
       ---                  ---
tcp    10.140.19.80:80      [2001:db8:cafe:beef::10]:80
       10.140.15.20:1024    [2001:db8:ea5e:1:49fa:b11a:aaf8:91a5]:57316
```

Static Entries

Dynamic Overloaded Entries

NAT64 Source NAT address

Outside Client Source Address

NAT64 Statistics

Reference

```
asr1k#sh nat64 statistics
Interface Statistics
GigabitEthernet0/0/2 (IPv4 not configured, IPv6 configured):
      Packets translated (IPv4 -> IPv6)
         Stateless: 0
         Stateful: 0
      Packets translated (IPv6 -> IPv4)
         Stateless: 0
         Stateful: 3
      Packets dropped: 0
   GigabitEthernet0/0/3 (IPv4 configured, IPv6 not configured):
      Packets translated (IPv4 -> IPv6)
         Stateless: 0
         Stateful: 3
      Packets translated (IPv6 -> IPv4)
         Stateless: 0
         Stateful: 0
      Packets dropped: 0
Dynamic Mapping Statistics
   v6v4
      access-list EDGE_ACL pool IE refcount 1
         pool IE:
            start 10.140.15.20 end 10.140.15.20
            total addresses 1, allocated 1 (100%)
            address exhaustion packet count 0
```

*Output reduced for clarity

NetFlow Export of Original Source IP

- In ACE example we used “x-forwarded-for” insertion to get original source IPv6 address
- With ASR1k we can use NetFlow to export original IPv6 Source address (flow record “ipv6 original-input”)
- You can export via IPv4 or IPv6 to your favorite collector
- THIS IS NOT A GREAT REPLACEMENT FOR XFF – Your existing web analytics and geolocation tools are probably XFF

Reference

NetFlow Record IPv6 Original-Input

```
asr1k#show flow record netflow ipv6 original-input
flow record netflow ipv6 original-input:
  Description:        Traditional IPv6 input NetFlow with ASs
  No. of users:       0
  Total field space:  97 bytes
  Fields:
    match ipv6 traffic-class
    match ipv6 flow-label
    match ipv6 protocol
    match ipv6 extension map
    match ipv6 source address
    match ipv6 destination address
    match transport source-port
    match transport destination-port
    match interface input
    match flow direction
    match flow sampler
    collect routing source as
    collect routing destination as
    collect routing next-hop address ipv6
    collect ipv6 source mask
    collect ipv6 destination mask
    collect transport tcp flags
    collect interface output
    collect counter bytes
    collect counter packets
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
```


Reference

NetFlow Export Example

- Normal NetFlow stuff
- Create a monitor
- Create an export destination
- Assign to interface

```
flow exporter EXPORT-IE
 destination 10.140.22.90
 transport udp 90
!
!
flow monitor NAT64
 record netflow ipv6 original-input
 exporter EXPORT-IE
 cache entries 200000
!
interface GigabitEthernet0/0/2
 description to 6k-inner-1 Outside
 ipv6 flow monitor NAT64 input
 ipv6 address 2001:DB8:CAFE:110::A/64
 nat64 enable
```

NetFlow Export Cache Output

```
asr1k#show flow monitor NAT64 cache
 . . . .
IPV6 FLOW LABEL:              0
IPV6 EXTENSION MAP:           0x00000000
IPV6 SOURCE ADDRESS:          2001:DB8:EA5E:1:49FA:B11A:AAF8:91A5
IPV6 DESTINATION ADDRESS:     2001:DB8:CAFE:BEEF::10
TRNS SOURCE PORT:             57227
TRNS DESTINATION PORT:        80
INTERFACE INPUT:              Gi0/0/2
FLOW DIRECTION:               Input
FLOW SAMPLER ID:              0
IP PROTOCOL:                  6
IP TOS:                       0x00
ip source as:                 0
ip destination as:            0
ipv6 next hop address:        ::100.0.0.1
ipv6 source mask:             /0
ipv6 destination mask:        /96
tcp flags:                    0x1A
interface output:             NV0
counter bytes:                661
counter packets:              4
timestamp first:              13:21:37.815
timestamp last:               13:21:38.039
```

Original Client Src IP

Outside IPv6 static host address

*Output reduced for clarity

LISP – For Edge IPv6 Support

LISP References

- **Sites you need to bookmark**
  - **http://lisp.cisco.com**
  - **http://www.lisp4.net** **http://www.lisp6.net**
- **Cisco LISP Mailer:**
  - **lisp-support@cisco.com**
- **The source of all goodness:**
  - **http://lisp.cisco.com/lisp_tech.html**

Reference

Definitions

- **ITR – Ingress Tunnel Router**: Receives packets from site-facing interfaces and encaps to remote LISP site or natively to non-LISP site
- **ETR – Egress Tunnel Router**: Receives packets from core-facing interfaces and de-caps and delivers to local EIDs at site
- **MR – Map-Resolver**: Receives Map-Requests from ITRs and forwards to authoritative Map-Server, or sends Negative-Map-Replies in response to Map-Requests for non-LISP sites
- **MS – Map-Server**: LISP ETRs register here, injects routes for LISP sites and forwards Map-Requests to registered ETRs
- **PITR – Proxy ITR**: Receives traffic from non-LISP sites, encapsulates traffic to LISP sites and advertises coarse-aggregate EID prefixes
- **PETR – Proxy ETR**: Allows IPv6 LISP sites with IPv4 RLOCs to reach Non-LISP IPv6 sites

The focus area for our discussion (ITR, ETR)

Provider will own these in most cases (unless you are doing internal LISP deployment) (MR, MS, PITR, PETR)

LISP Operations

LISP Components – Ingress/Egress Tunnel Router (xTR)

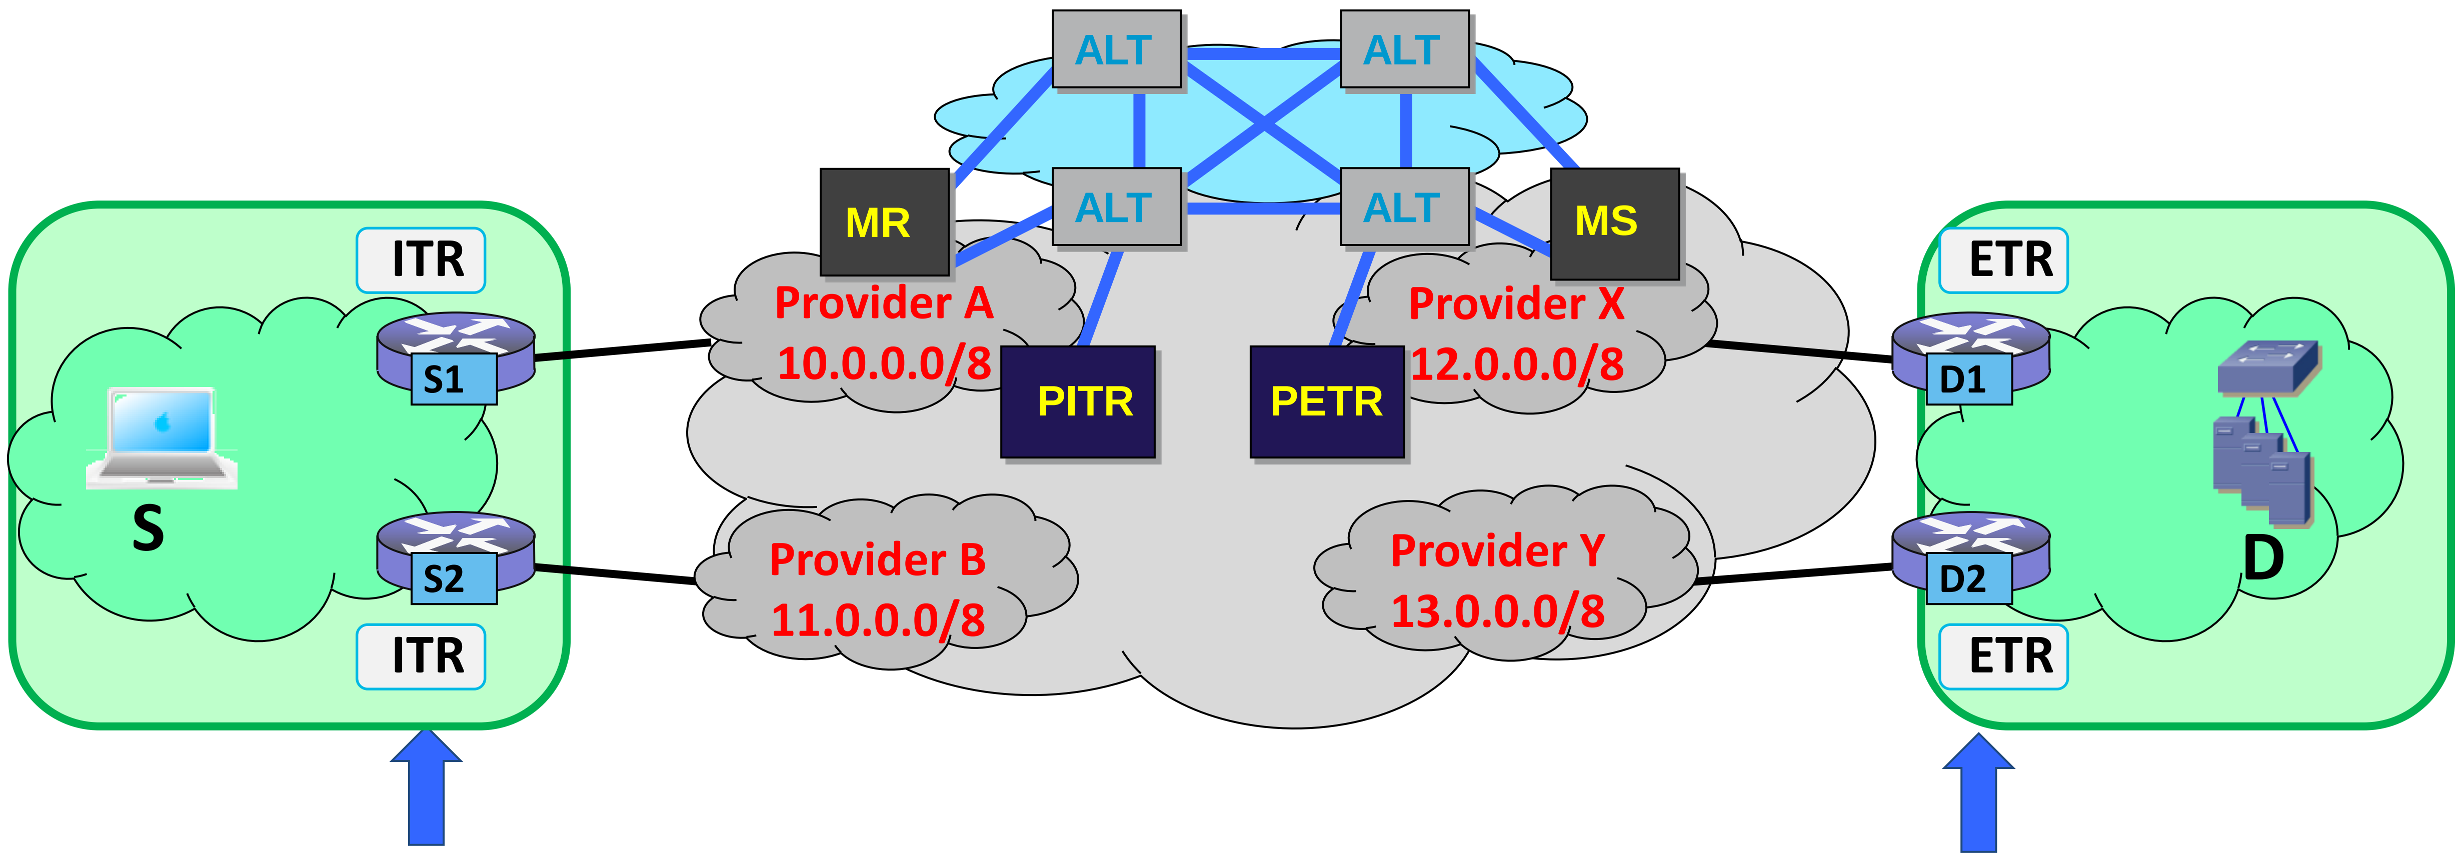

ITR – Ingress Tunnel Router
- Receives packets from site-facing interfaces
- Encaps to remote LISP site or natively forwards to non-LISP site

ETR – Egress Tunnel Router
- Receives packets from core-facing interfaces
- De-caps and delivers to local **EIDs** at the site

LISP Operations

LISP Components – Map-Server/Map-Resolver (MS/MR)

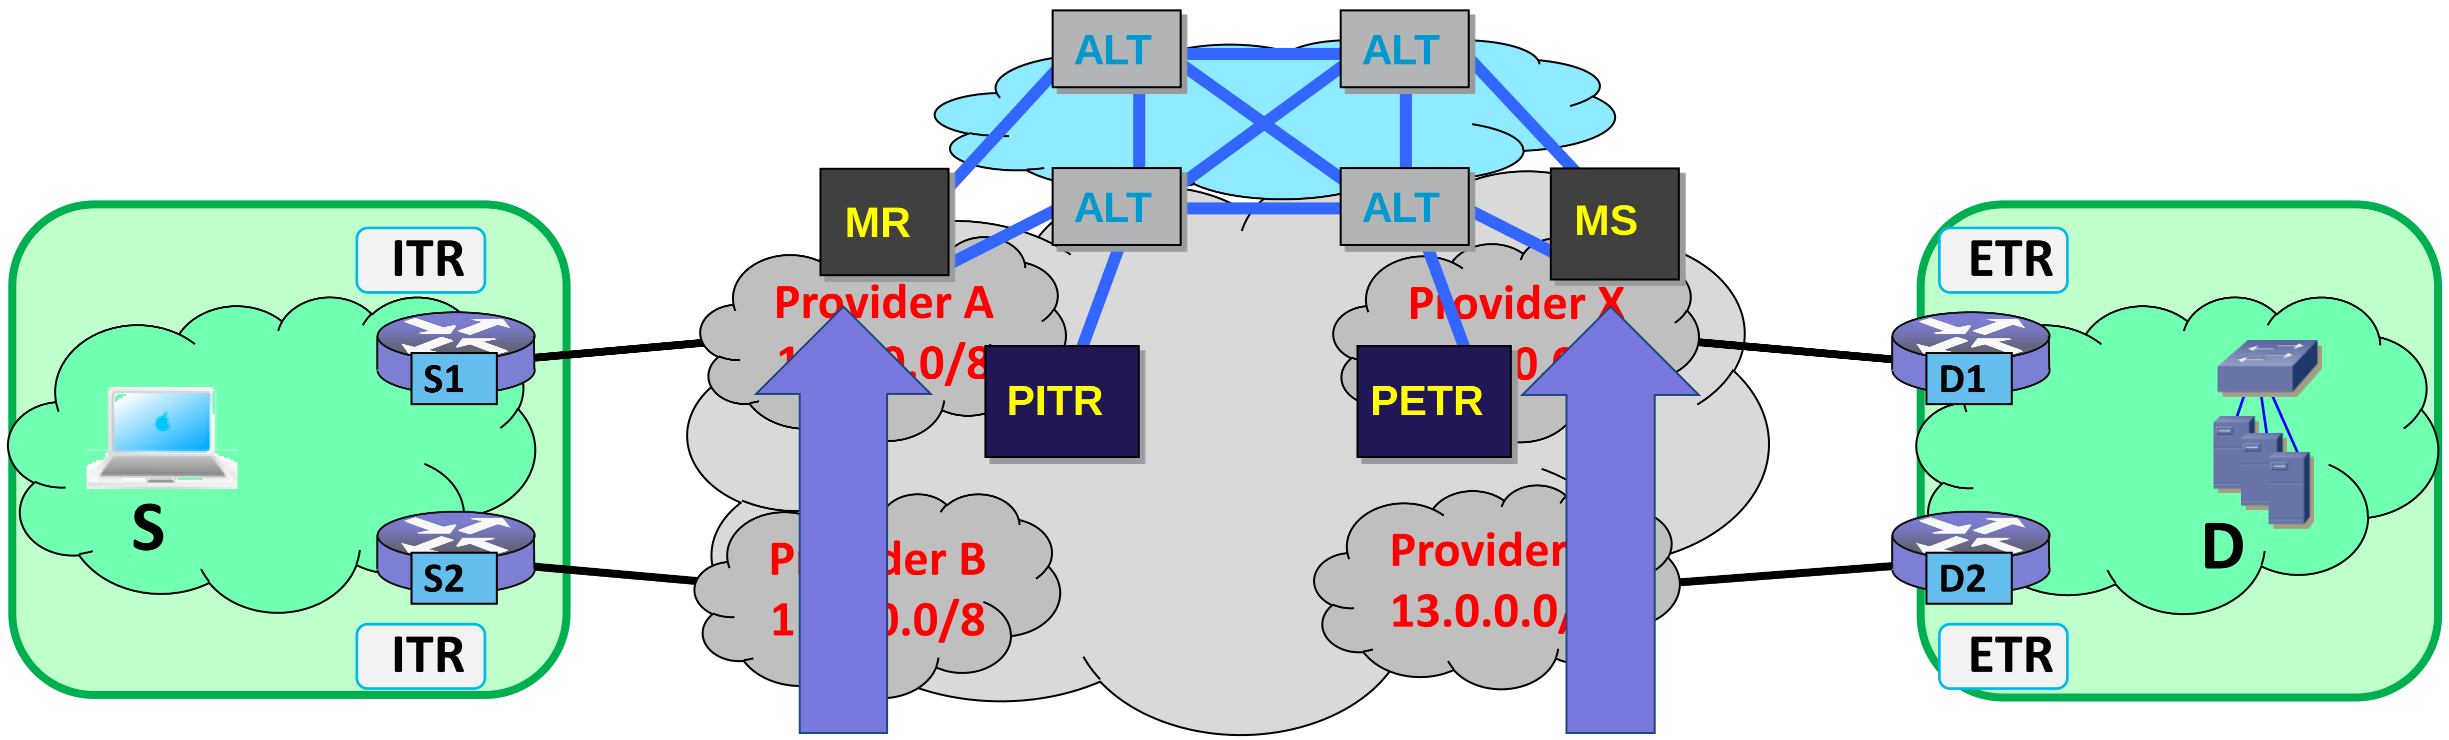

MR – Map-Resolver

- Receives Map-Request encapsulated from ITR
- De-caps Map-Request, forwards thru service interface onto the ALT topology
- Sends Negative Map-Replies in response to Map-Requests for non-LISP sites

MS – Map-Server

- LISP ETRs Register here; requires configured “lisp site” policy, key
- Injects routes for registered LISP sites into ALT thru ALT service interface
- Receives Map-Requests via ALT; en-caps Map-Requests to registered ETRs

LISP Operations

Interworking Mechanisms

- **Early Recognition – LISP will not be widely deployed day-one**
- **Interworking for:**
  - – LISP-capable sites to non-LISP sites (i.e. the rest of the Internet)
  - – non-LISP sites to LISP-capable sites
- **Two basic Techniques**
  - – LISP Network Address Translators (LISP-NAT)
  - – Proxy Ingress Tunnel Routers & Proxy Egress Tunnel Routers
- **Proxy-ITR/Proxy-ETR have the most promise**
  - – Infrastructure LISP network entity
  - – Creates a monetised service opportunity for infrastructure players

LISP

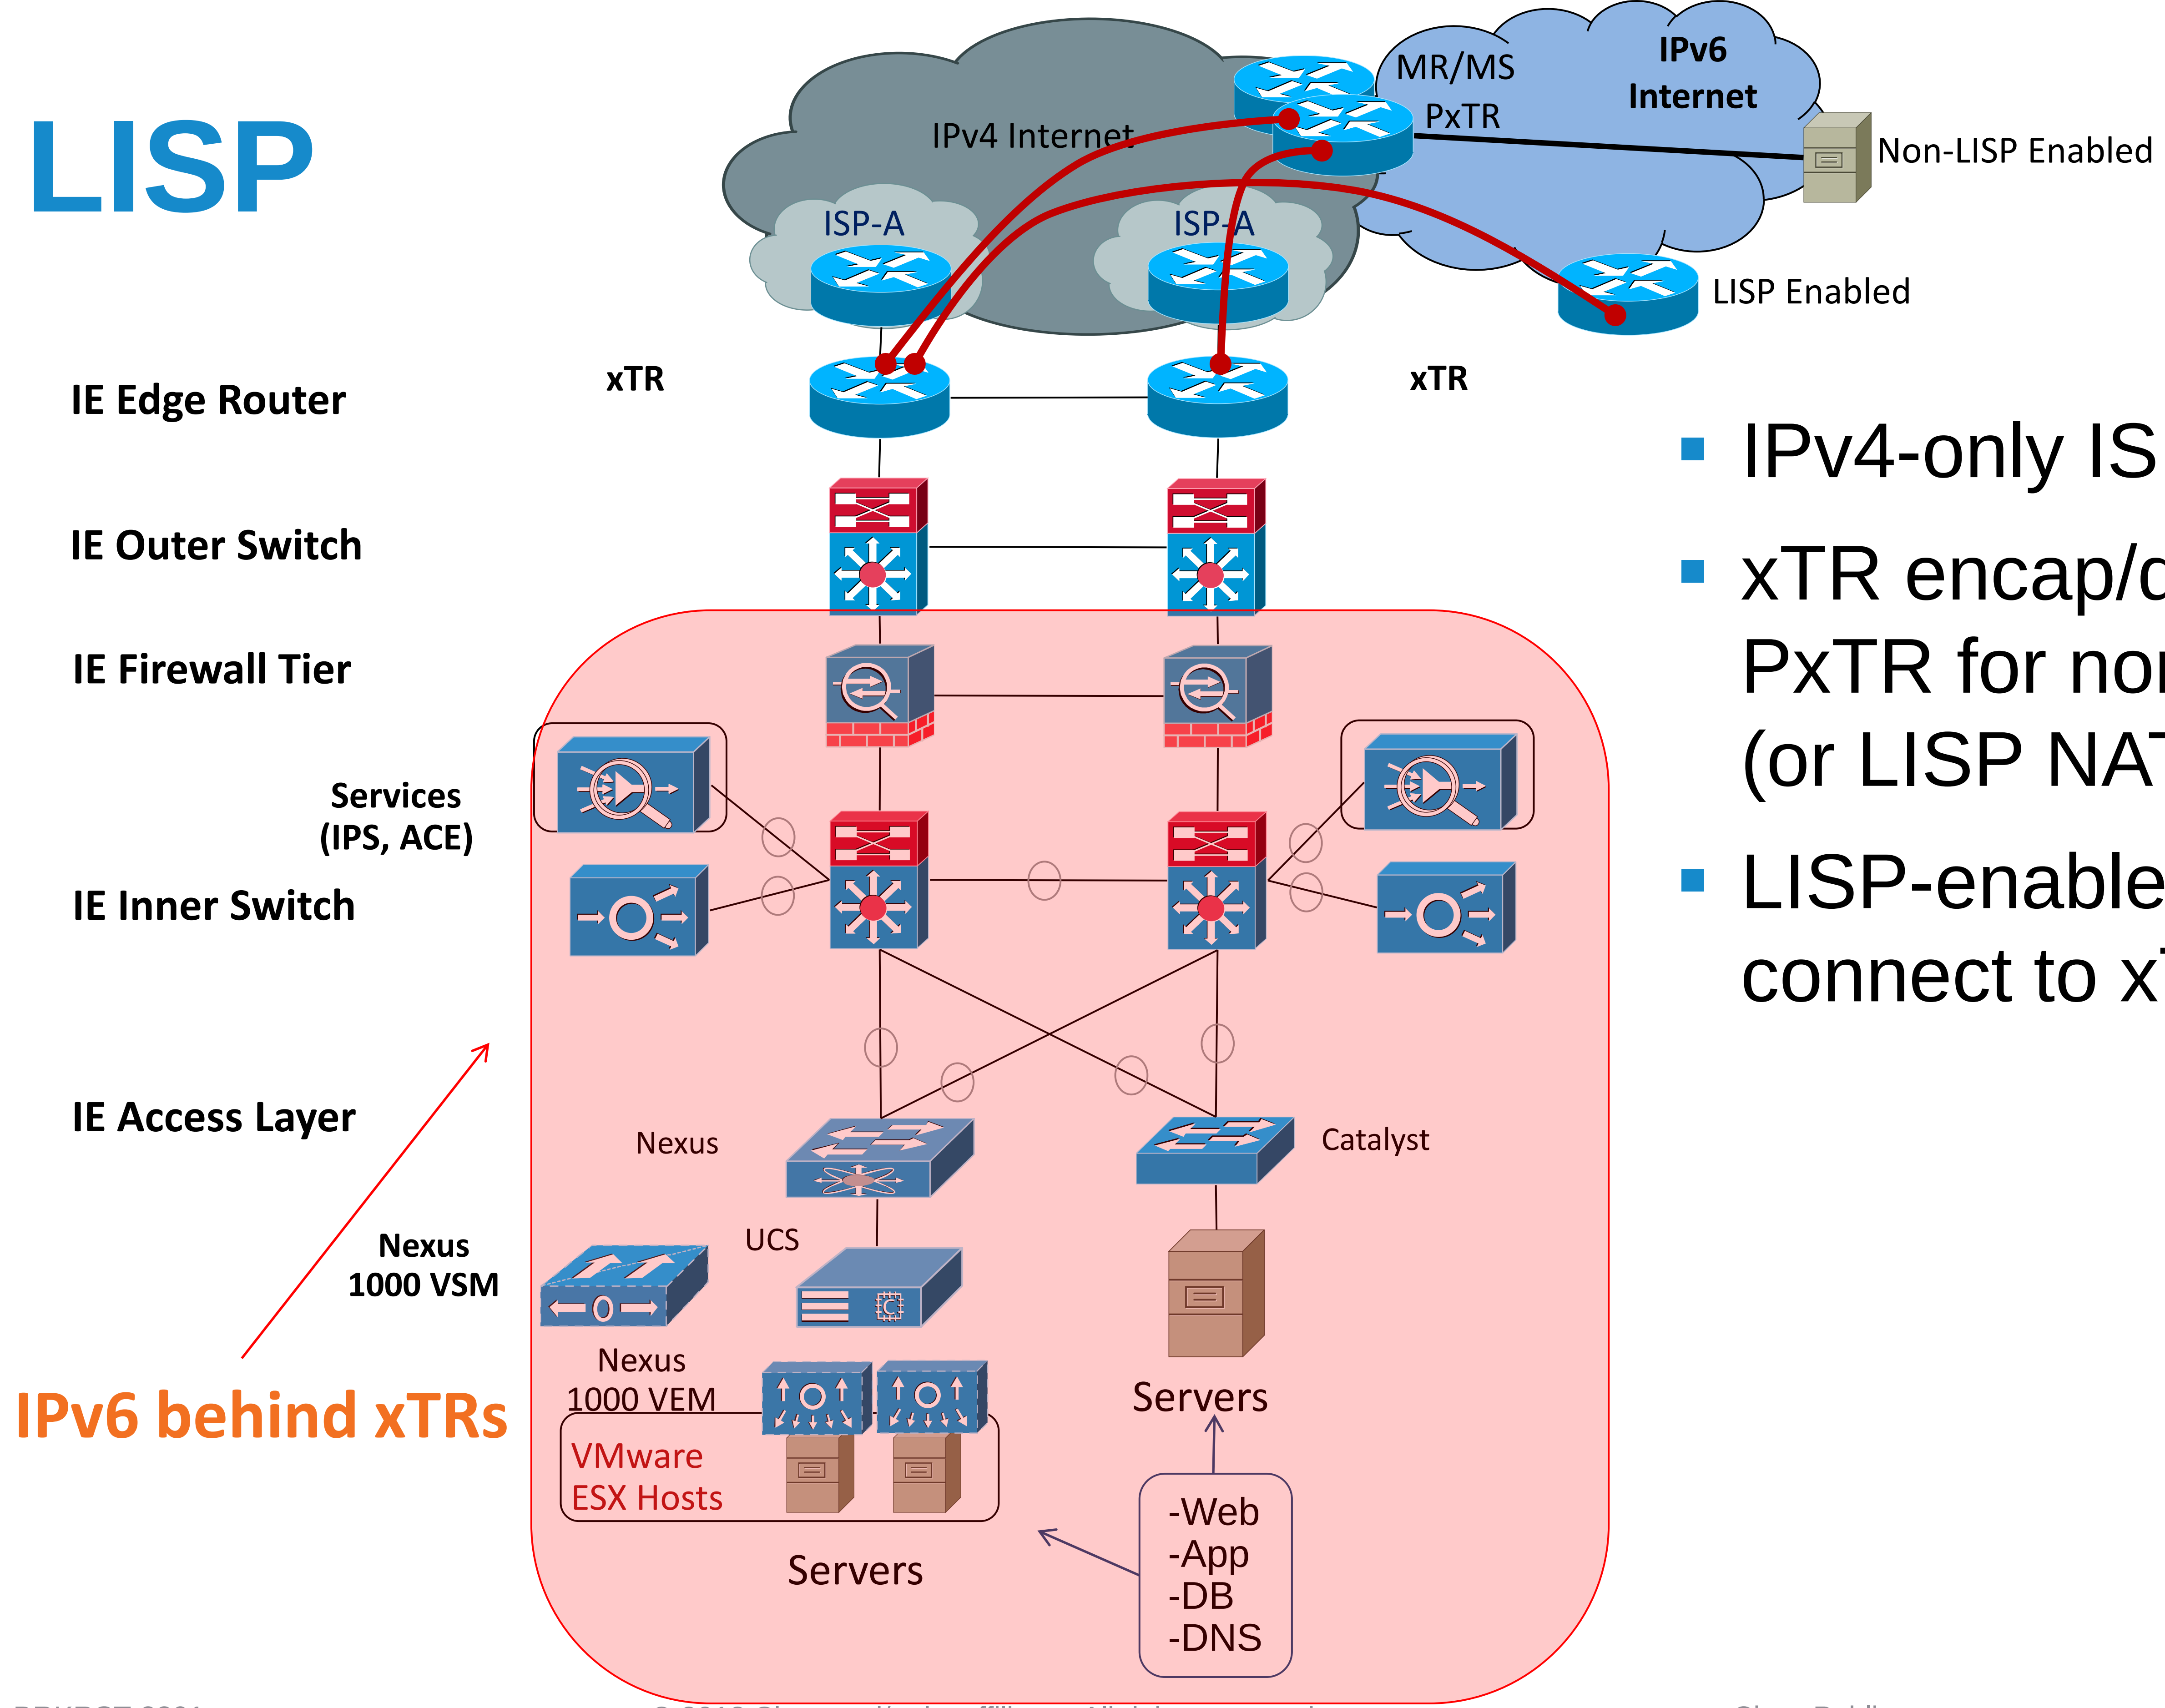

- IPv4-only ISP
- xTR encap/decap to PxTR for non-LISP sites (or LISP NAT)
- LISP-enable sites connect to xTR

Reference

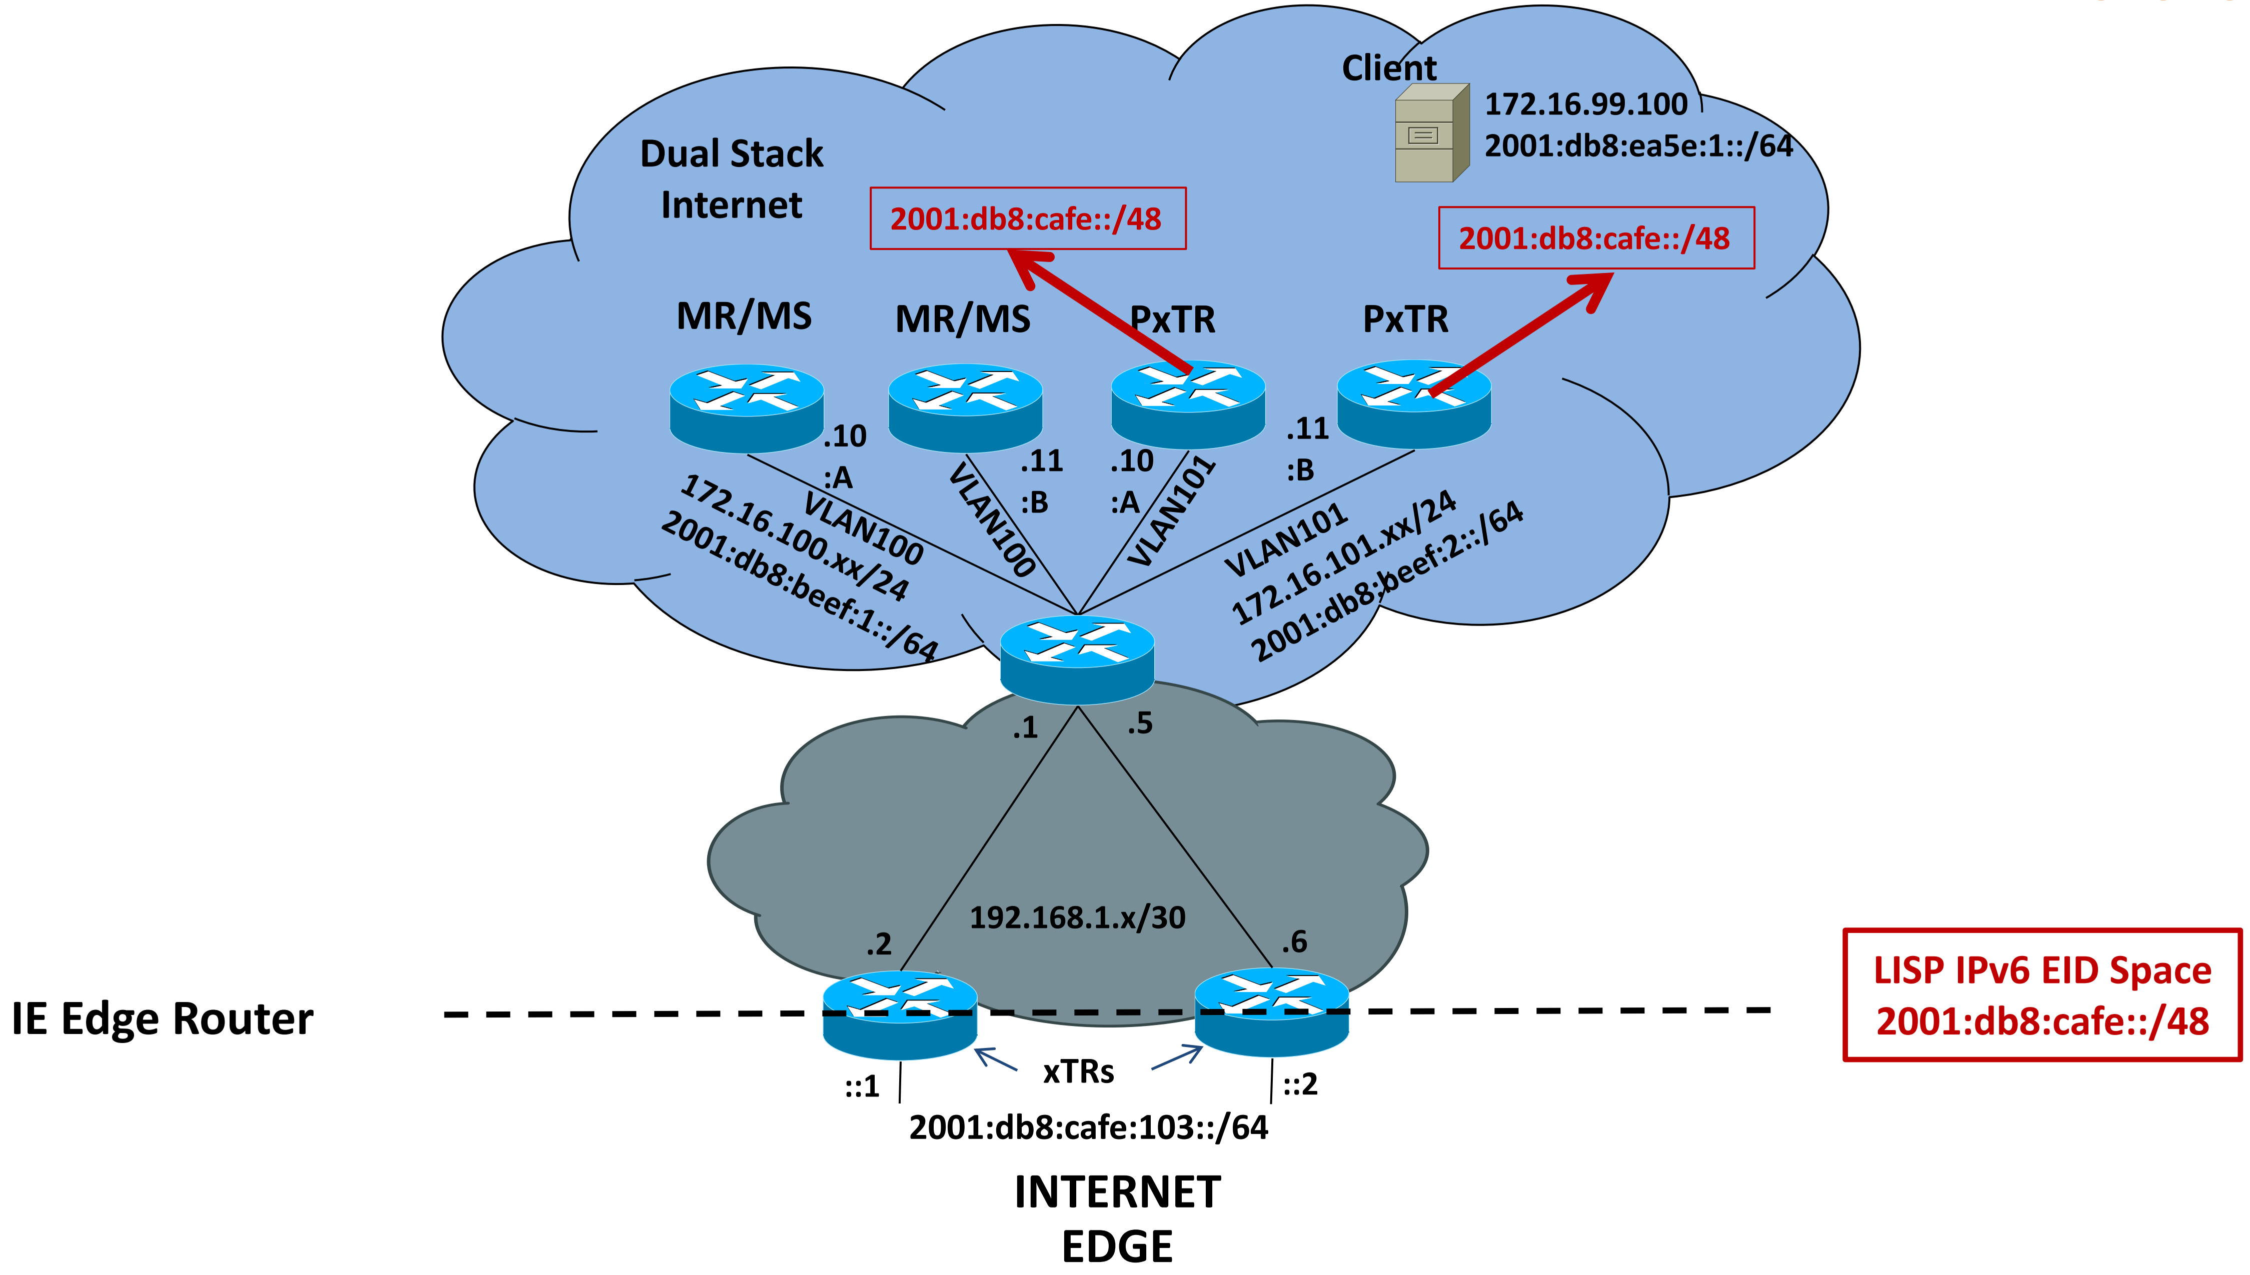

- Example addressing layout
- PxTR announces for 2001:db8:cafe::/48

xTR

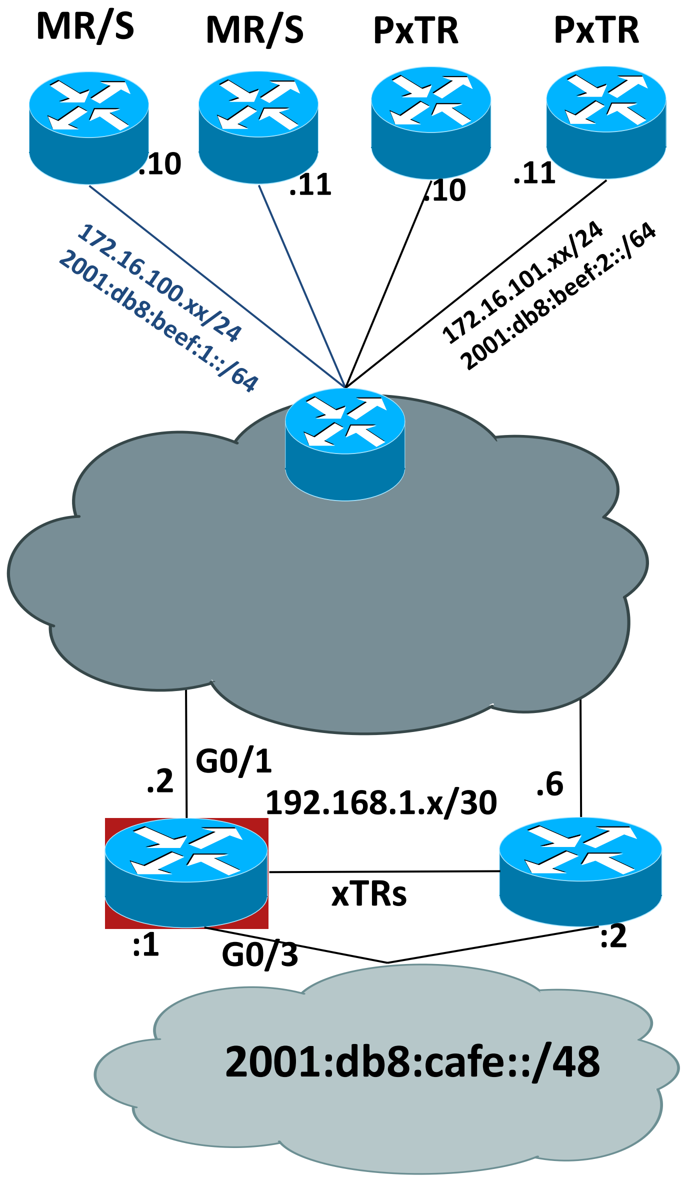


```
interface GigabitEthernet0/1
 description to ISPA (7604-1) - IPv4-ONLY
 ip address 192.168.1.2 255.255.255.252
!
interface GigabitEthernet0/3
 description to Enterprise Internet Edge IPv4/IPv6
 ip address 192.168.1.66 255.255.255.224
 ipv6 address 2001:DB8:CAFE:103::1/64
!
#BGP config excluded
!
router lisp
  eid-table default instance-id 0
   database-mapping 2001:DB8:CAFE::/48 192.168.1.2 priority 1 weight 1
   database-mapping 2001:DB8:CAFE::/48 192.168.1.6 priority 1 weight 1
   exit
  !
  ipv6 use-petr 172.16.101.10
  ipv6 use-petr 172.16.101.11
  ipv6 itr map-resolver 172.16.100.10
  ipv6 itr map-resolver 172.16.100.11
  ipv6 itr
  ipv6 etr map-server 172.16.100.10 key CISCO
  ipv6 etr map-server 172.16.100.11 key CISCO
  ipv6 etr
  exit
!
ipv6 route ::/0 Null0
```

MR/MS

Reference

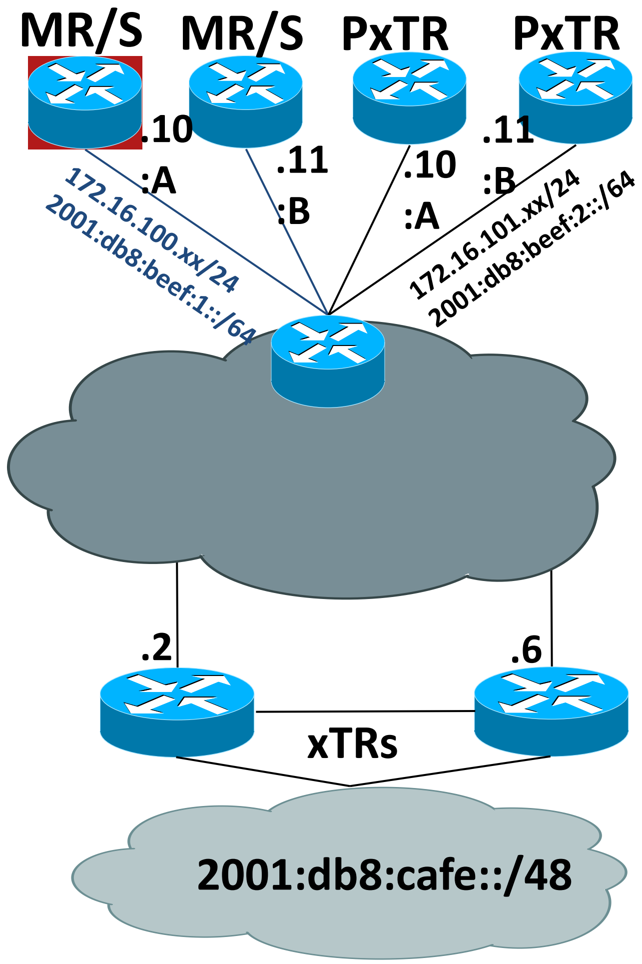


```
interface LISP0
!
interface GigabitEthernet0/0/0
 description Link to SP1 (RLOC)
 ip address 172.16.100.10 255.255.255.0
 ipv6 address 2001:DB8:BEEF:1::A/64
!
router lisp
 site CUST-1
  authentication-key CISCO
  eid-prefix 2001:DB8:CAFE::/48
  exit
!
 ipv6 map-server
 ipv6 map-resolver
 exit
!
ip route 0.0.0.0 0.0.0.0 172.16.100.1
!
ipv6 route ::/0 2001:DB8:CAFE:1::1
```

- Redundant configurations across MR/MS routers

Reference

PxTR

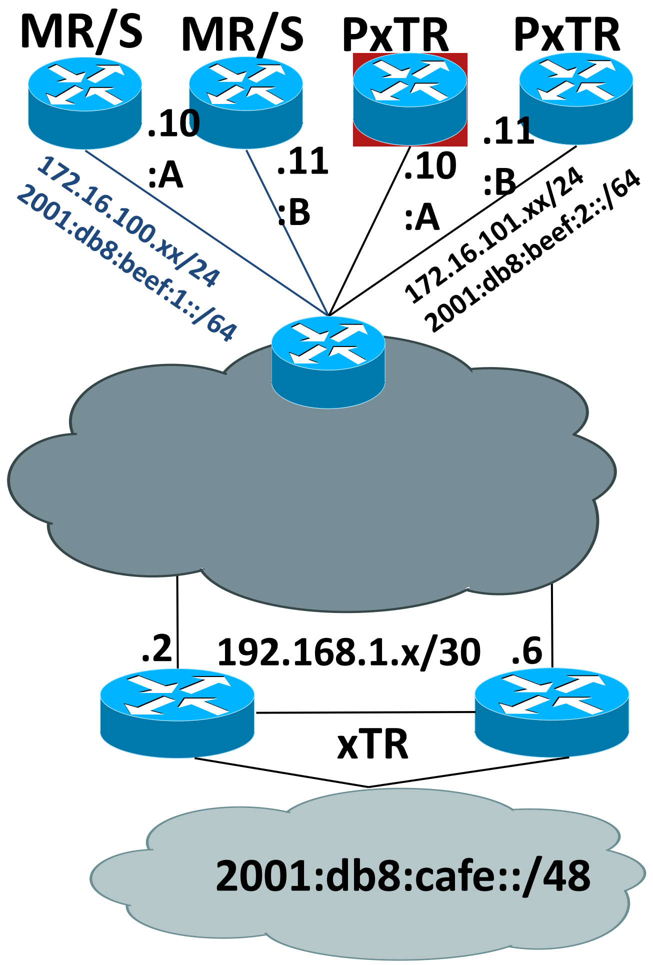


```
interface GigabitEthernet0/0/0
 description Link to Core (RLOC)
 ip address 172.16.101.10 255.255.255.0
 ipv6 address 2001:DB8:CAFE:2::A/64
!
router lisp
 eid-table default instance-id 0
  map-cache 2001:DB8:CAFE::/48 map-request
  exit
 !
 ipv6 map-request-source 2001:DB8:BEEF:2::A
 ipv6 proxy-etr
 ipv6 proxy-itr 2001:DB8:BEEF:2::A 172.16.101.10
 ipv6 itr map-resolver 172.16.100.10
 ipv6 itr map-resolver 172.16.100.11
 ipv6 itr map-resolver 2001:DB8:BEEF:1::A
 ipv6 itr map-resolver 2001:DB8:BEEF:1::B
 exit
!
ip route 0.0.0.0 0.0.0.0 172.16.101.1
ipv6 route ::/0 2001:DB8:BEEF:2::1
```

- Redundant configurations across PxTR

Putting It All Together

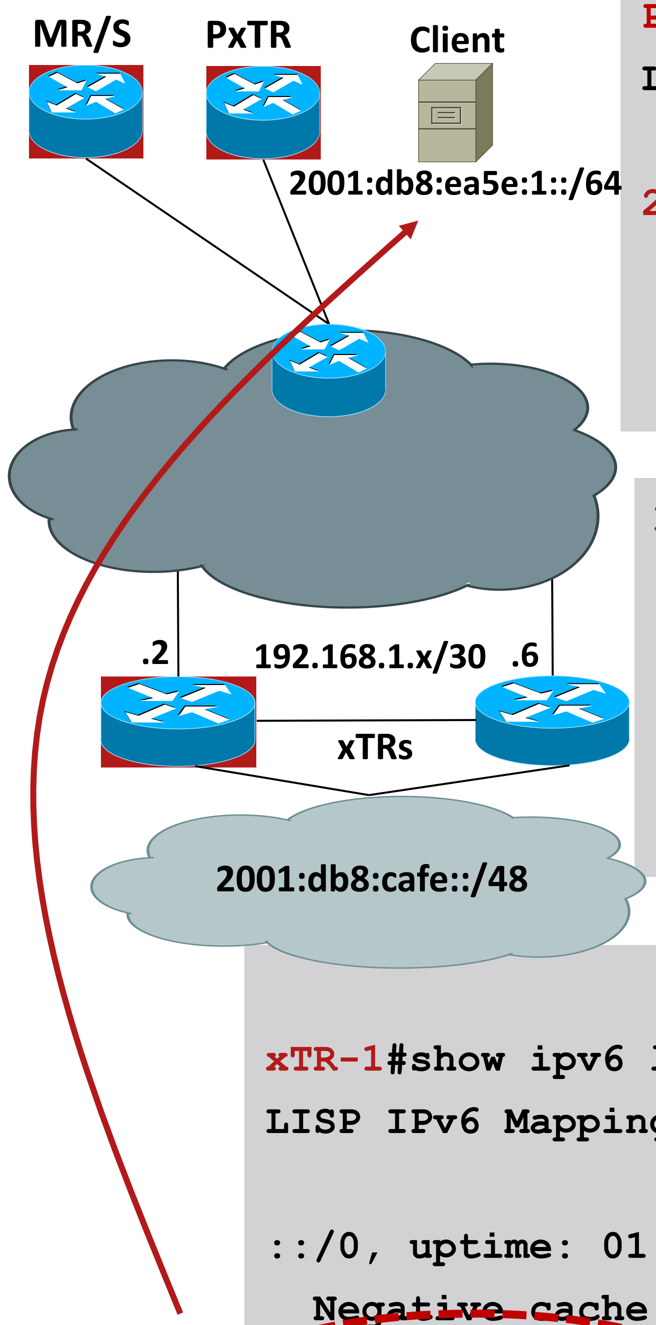


```
PxTR-1#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache for EID-table default (IID 0), 1 entries

2001:DB8:CAFE::/48, uptime: 00:55:53, expires: 23:04:52, via map-reply, complete
  Locator       Uptime    State      Pri/Wgt
  192.168.1.2   00:55:00  up           1/1
  192.168.1.6   00:55:00  up           1/1
```

```
MS-MR-1#show lisp site
LISP Site Registration Information

Site Name      Last      Up   Who Last          Inst     EID Prefix
               Register       Registered        ID
CUST-1         00:00:23  yes  192.168.1.2                2001:DB8:CAFE::/48
```

```
xTR-1#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache for EID-table default (IID 0), 2 entries

::/0, uptime: 01:01:55, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
2001:DB8:E000::/35, uptime: 00:58:48, expires: 00:00:44, via map-reply, forward-native
  Encapsulating to proxy ETR
```

Reference

Internet Edge Design Summary

Model	Benefit	Challenge
Dual Stack	-No tunnelling -No translation -No dependency on IPv4 -Best performance, scale, HA -Native visibility into IPv6 traffic	-Requires IPv6 support in all L3 aware platforms and software
SLB64	-Allows for IPv6 to IPv4-only server access -Removes immediate need to dual stack entire server farm -Higher performance and HA over sw-only reverse proxies -Leverage existing SLB platform (ACE4710/30) -Non-disruptive to IPv4 applications -Maintain IPv6 source address visibility using XFF	-Still requires IPv6 from ISP to north-facing side of ACE -Potential cost of new ACE HW -Does not support every application type or protocol today -Performance may not match dual stack design depending on traffic load
Stateful NAT64	-Allows for IPv6 to IPv4-only server access -Removes immediate need to dual stack entire server farm -Higher performance and HA over sw-only reverse proxies -Non-disruptive to IPv4 applications -No HW change needed if already using platform that supports NAT64	-Potential cost of new HW to support NAT64 -Does not support every app or protocol in the ALG function of the NAT64 feature -Performance may not match dual stack -NetFlow can be used for src IPv6 address logging but may not work with existing web analytics and logging tools.
LISP	-Provide IPv6 Internet access when ISP does not natively offer it -Quick and easy to deploy -High performance, highly available, highly scalable -Can be used with dual stack, SLB64 and NAT64 designs -Non-disruptive to existing IPv4 applications	-Requires connections to ISP provided LISP infrastructure components (MR/MS, PxTR, etc.) -Learning curve -Tunnel based

WAN/Branch

- Deploying IPv6 in Branch Networks: http://www.cisco.com/univercd/cc/td/doc/solution/brchipv6.pdf

Hybrid Branch Example

- Mixture of attributes from each profile
- An example to show configuration for different tiers
- Basic HA in critical roles is the goal

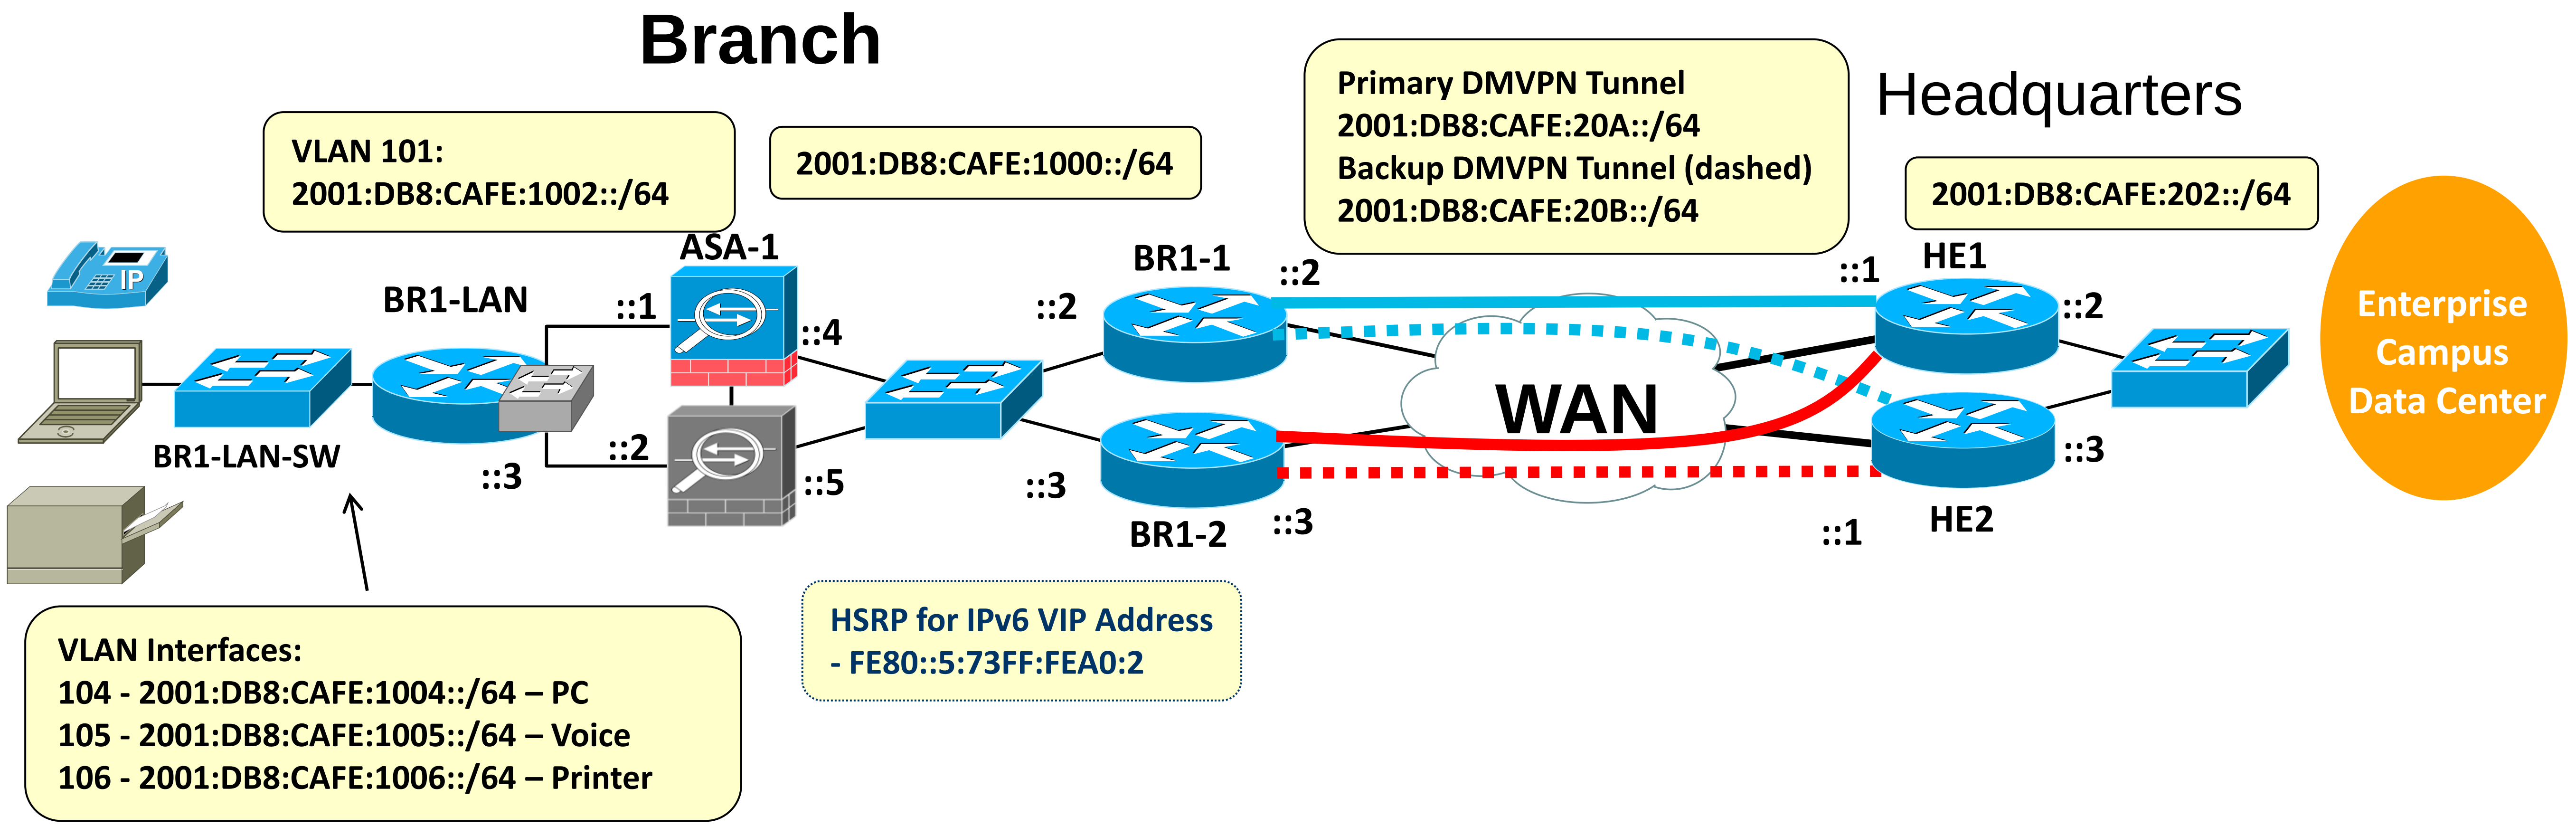

DMVPN with IPv6

Hub Configuration Example

```
crypto isakmp policy 1
 encr aes 256
 authentication pre-share
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set HUB esp-aes 256 esp-sha-hmac
!
crypto ipsec profile HUB
 set transform-set HUB
```

```
interface Tunnel0
 description DMVPN Tunnel 1
 ip address 10.126.1.1 255.255.255.0
 ipv6 address 2001:DB8:CAFE:20A::1/64
 ipv6 mtu 1416
 ipv6 eigrp 10
 ipv6 hold-time eigrp 10 35
 no ipv6 next-hop-self eigrp 10
 no ipv6 split-horizon eigrp 10
 ipv6 nhrp authentication CISCO
 ipv6 nhrp map multicast dynamic
 ipv6 nhrp network-id 10
 ipv6 nhrp holdtime 600
 ipv6 nhrp redirect
 tunnel source Serial1/0
 tunnel mode gre multipoint
 tunnel key 10
 tunnel protection ipsec profile HUB
```

Primary DMVPN Tunnel
2001:DB8:CAFE:20A::/64
Backup DMVPN Tunnel (dashed)
2001:DB8:CAFE:20B::/64

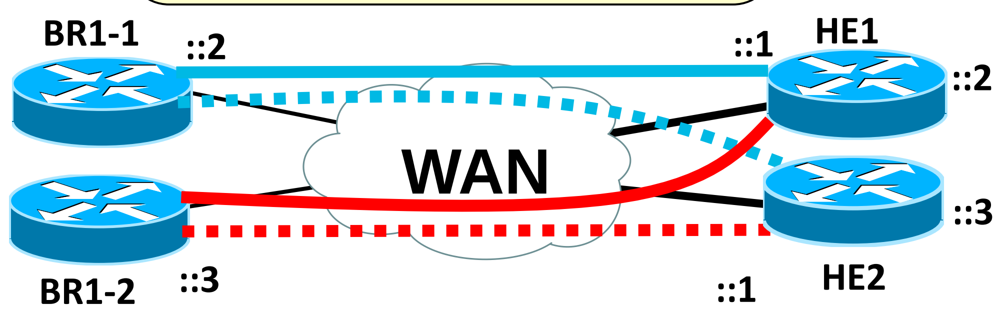

DMVPN with IPv6

Spoke Configuration Example

```
crypto isakmp policy 1
 encr aes 256
 authentication pre-share
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set SPOKE esp-aes 256 esp-sha-hmac
!
crypto ipsec profile SPOKE
 set transform-set SPOKE
```

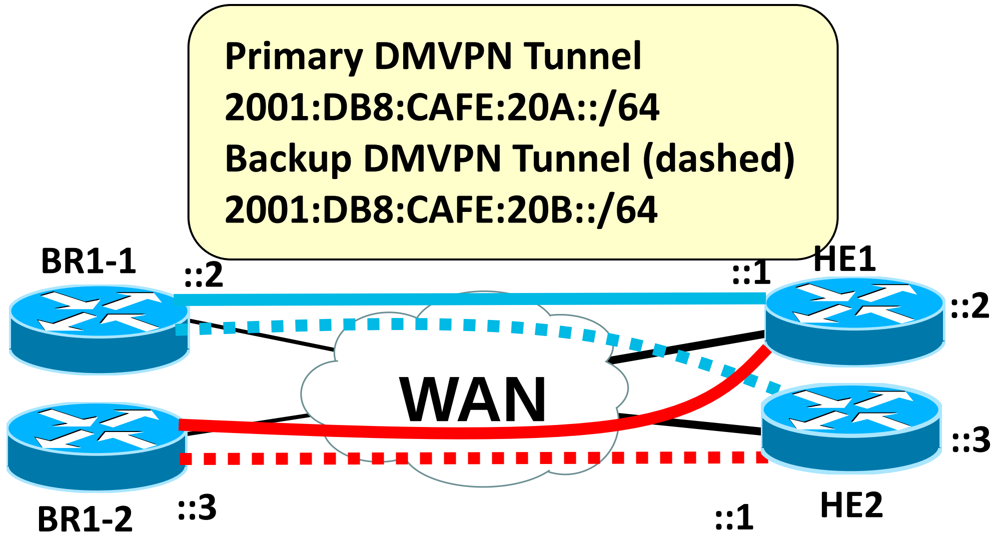


```
interface Tunnel0
 description to HUB
 ip address 10.126.1.2 255.255.255.0
 ipv6 address 2001:DB8:CAFE:20A::2/64
 ipv6 mtu 1416
 ipv6 eigrp 10
 ipv6 hold-time eigrp 10 35
 no ipv6 next-hop-self eigrp 10
 no ipv6 split-horizon eigrp 10
 ipv6 nhrp authentication CISCO
 ipv6 nhrp map 2001:DB8:CAFE:20A::1/64 172.16.1.1
 ipv6 nhrp map multicast 172.16.1.1
 ipv6 nhrp network-id 10
 ipv6 nhrp holdtime 600
 ipv6 nhrp nhs 2001:DB8:CAFE:20A::1
 ipv6 nhrp shortcut
 tunnel source Serial1/0
 tunnel mode gre multipoint
 tunnel key 10
 tunnel protection ipsec profile SPOKE
```

ASA with IPv6

Snippet of Full Config – Examples of IPv6 Usage

```
name 2001:db8:cafe:1003:: BR1-LAN description VLAN on EtherSwitch
name 2001:db8:cafe:1004:9db8:3df1:814c:d3bc Br1-v6-Server
!
interface GigabitEthernet0/0
 description TO WAN
 nameif outside
 security-level 0
 ip address 10.124.1.4 255.255.255.0 standby 10.124.1.5
 ipv6 address 2001:db8:cafe:1000::4/64 standby 2001:db8:cafe:1000::5
!
interface GigabitEthernet0/1
 description TO BRANCH LAN
 nameif inside
 security-level 100
 ip address 10.124.3.1 255.255.255.0 standby 10.124.3.2
 ipv6 address 2001:db8:cafe:1002::1/64 standby 2001:db8:cafe:1002::2
!
ipv6 route inside BR1-LAN/64 2001:db8:cafe:1002::3
ipv6 route outside ::/0 fe80::5:73ff:fea0:2
!
ipv6 access-list v6-ALLOW permit icmp6 any any
ipv6 access-list v6-ALLOW permit tcp 2001:db8:cafe::/48 host Br1-v6-Server object-group RDP
!
failover
failover lan unit primary
failover lan interface FO-LINK GigabitEthernet0/3
failover interface ip FO-LINK 2001:db8:cafe:1001::1/64 standby 2001:db8:cafe:1001::2
access-group v6-ALLOW in interface outside
```

Branch LAN

Connecting Hosts

```
ipv6 dhcp pool DATA_W7
 dns-server 2001:DB8:CAFE:102::8
 domain-name cisco.com
!
interface GigabitEthernet0/0
 description to BR1-LAN-SW
 no ip address
 duplex auto
 speed auto
!
interface GigabitEthernet0/0.104
 description VLAN-PC
 encapsulation dot1Q 104
 ip address 10.124.104.1 255.255.255.0
 ipv6 address 2001:DB8:CAFE:1004::1/64
 ipv6 nd other-config-flag
 ipv6 dhcp server DATA_W7
 ipv6 eigrp 10
!
interface GigabitEthernet0/0.105
 description VLAN-PHONE
 encapsulation dot1Q 105
 ip address 10.124.105.1 255.255.255.0
 ipv6 address 2001:DB8:CAFE:1005::1/64
 ipv6 nd prefix 2001:DB8:CAFE:1005::/64 0 0 no-autoconfig
 ipv6 nd managed-config-flag
 ipv6 dhcp relay destination 2001:DB8:CAFE:102::9
 ipv6 eigrp 10
```

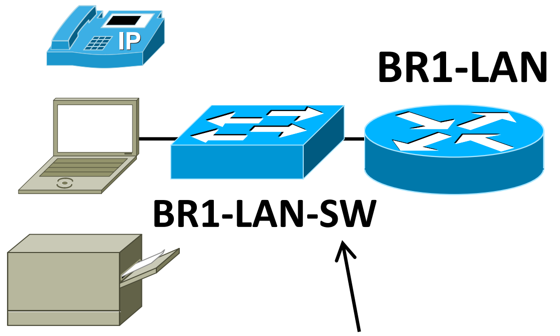

VLAN Interfaces:
104 - 2001:DB8:CAFE:1004::/64 – PC
105 - 2001:DB8:CAFE:1005::/64 – Voice
106 - 2001:DB8:CAFE:1006::/64 – Printer

IOS 15.2(1)T

DMVPN over IPv6 Transport

Spoke Configuration Example

```
interface Tunnel2
 description to HUB
 no ip address
 ipv6 address 2001:DB8:CAFE:C5C0::B/127
 ipv6 mtu 1400
 no ipv6 redirects
 ipv6 nhrp authentication CISCO
 ipv6 nhrp network-id 100
 ipv6 nhrp holdtime 300
 ipv6 nhrp nhs 2001:DB8:CAFE:C5C0::A nbma 2001:DB8:CAFE:37::B multicast
 ipv6 nhrp shortcut
 ipv6 eigrp 10
 tunnel source GigabitEthernet0/0
 tunnel mode gre multipoint ipv6
 tunnel key 100
 tunnel protection ipsec profile SPOKE
```

Remote Access

AnyConnect — SSL VPN

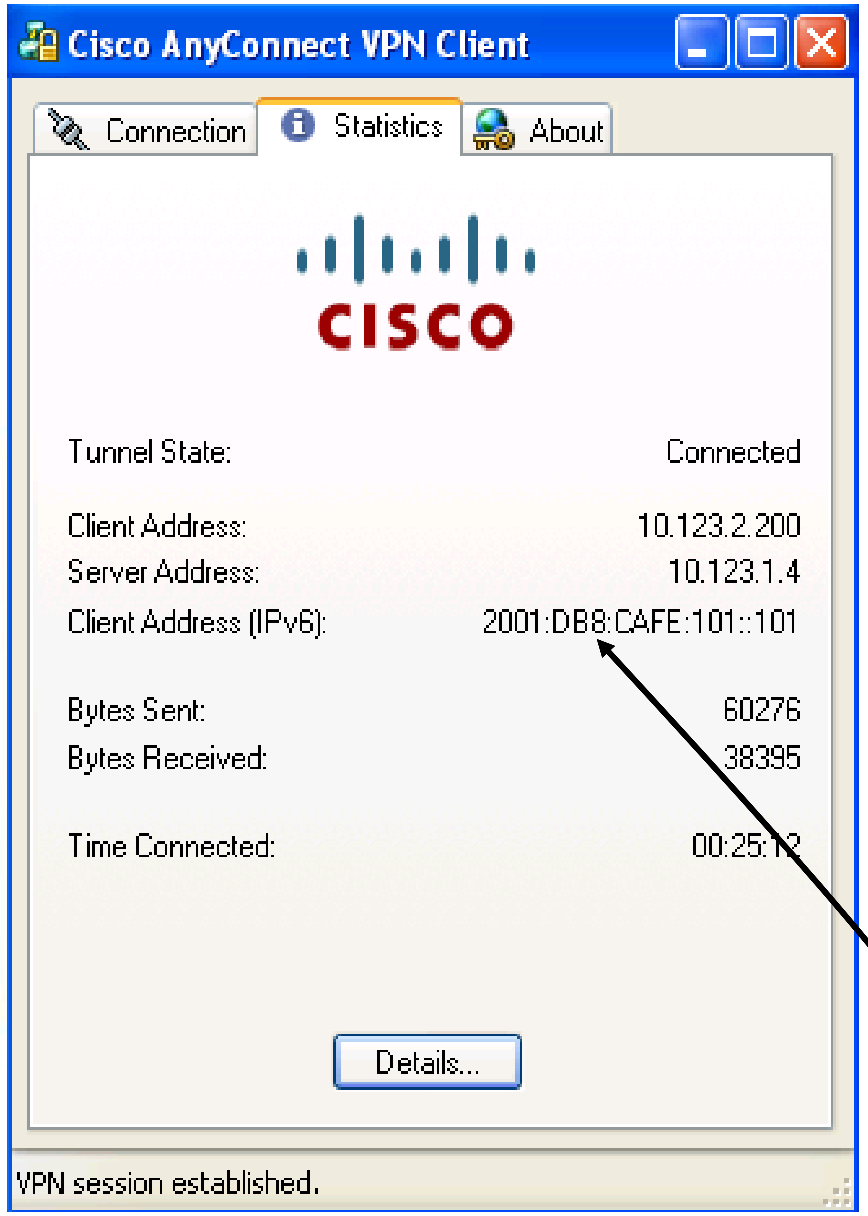


```
asa-edge-1#show vpn-sessiondb svc
Session Type: SVC
Username      : ciscoese                    Index        : 14
Assigned IP   : 10.123.2.200                Public IP    : 10.124.2.18
Assigned IPv6: 2001:db8:cafe:101::101
Protocol      : Clientless SSL-Tunnel DTLS-Tunnel
License       : SSL VPN
Encryption    : RC4 AES128                  Hashing      : SHA1
Bytes Tx      : 79763                       Bytes Rx     : 176080
Group Policy  : AnyGrpPolicy                Tunnel Group: ANYCONNECT
Login Time    : 14:09:25 MST Mon Dec 17 2007
Duration      : 0h:47m:48s
NAC Result    : Unknown
VLAN Mapping  : N/A                         VLAN         : none
```

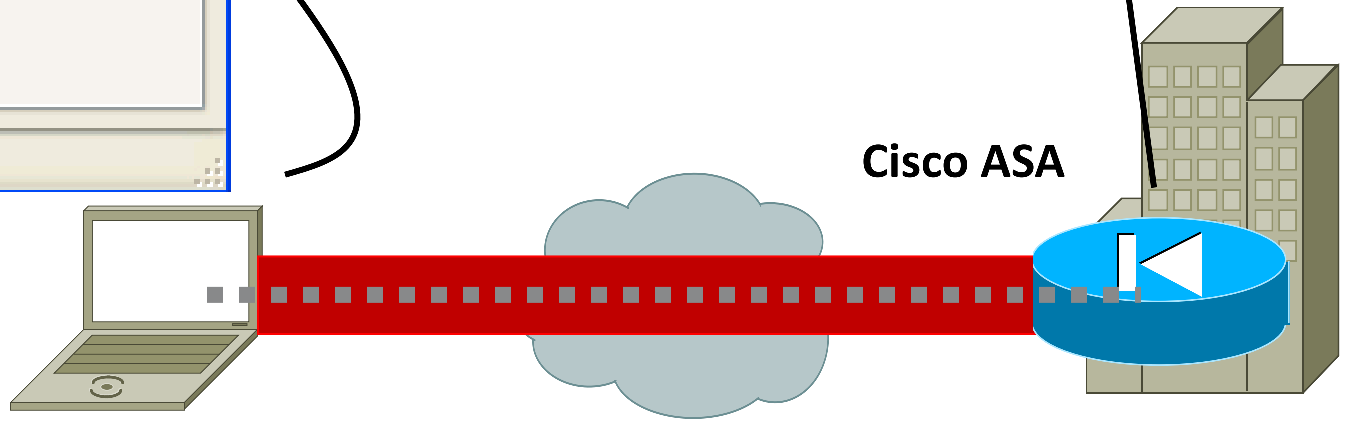

AnyConnect —Example Configuration

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.123.1.4 255.255.255.0
 ipv6 enable
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.123.2.4 255.255.255.0
 ipv6 address 2001:db8:cafe:101::ffff/64
!
ipv6 local pool ANYv6POOL 2001:db8:cafe:101::101/64 200
```

```
webvpn
 enable outside
 svc enable
 tunnel-group-list enable
group-policy AnyGrpPolicy internal
group-policy AnyGrpPolicy attributes
 vpn-tunnel-protocol svc
 default-domain value cisco.com
 address-pools value AnyPool
tunnel-group ANYCONNECT type remote-access
tunnel-group ANYCONNECT general-attributes
 address-pool AnyPool
 ipv6-address-pool ANYv6POOL
 default-group-policy AnyGrpPolicy
tunnel-group ANYCONNECT webvpn-attributes
 group-alias ANYCONNECT enable
```

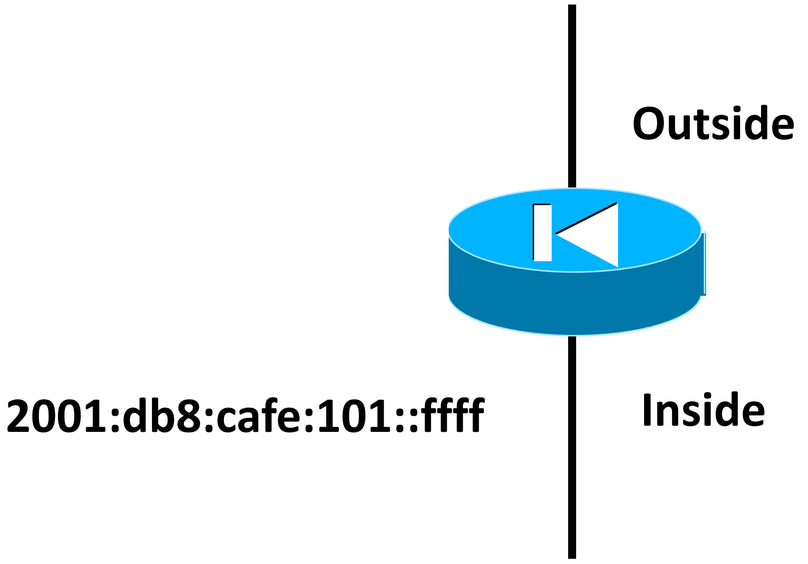

http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect20/administrative/guide/admin6.html#wp1002258

Communicating with the Service Provider

Top SP Concerns for Enterprise Accounts

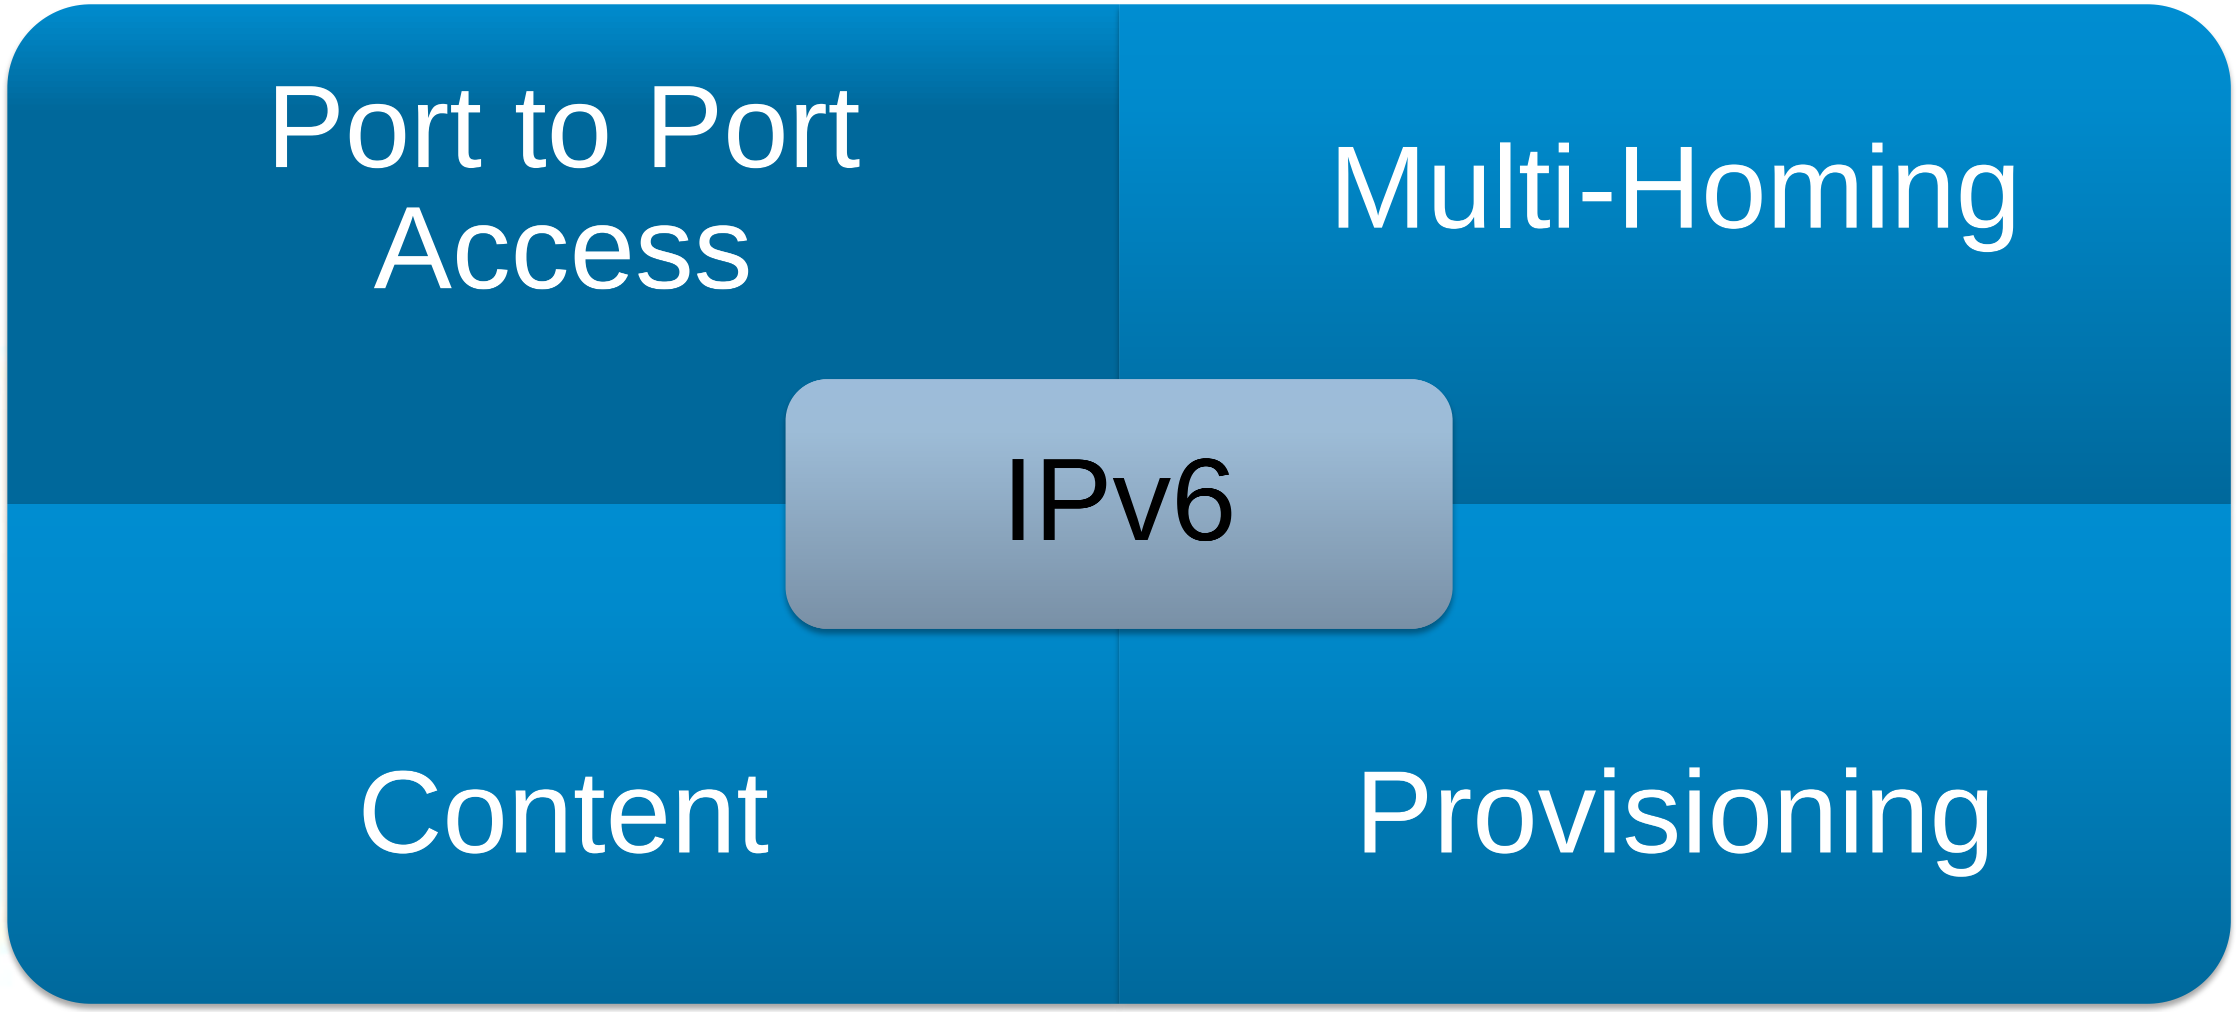

Port-to-Port Access

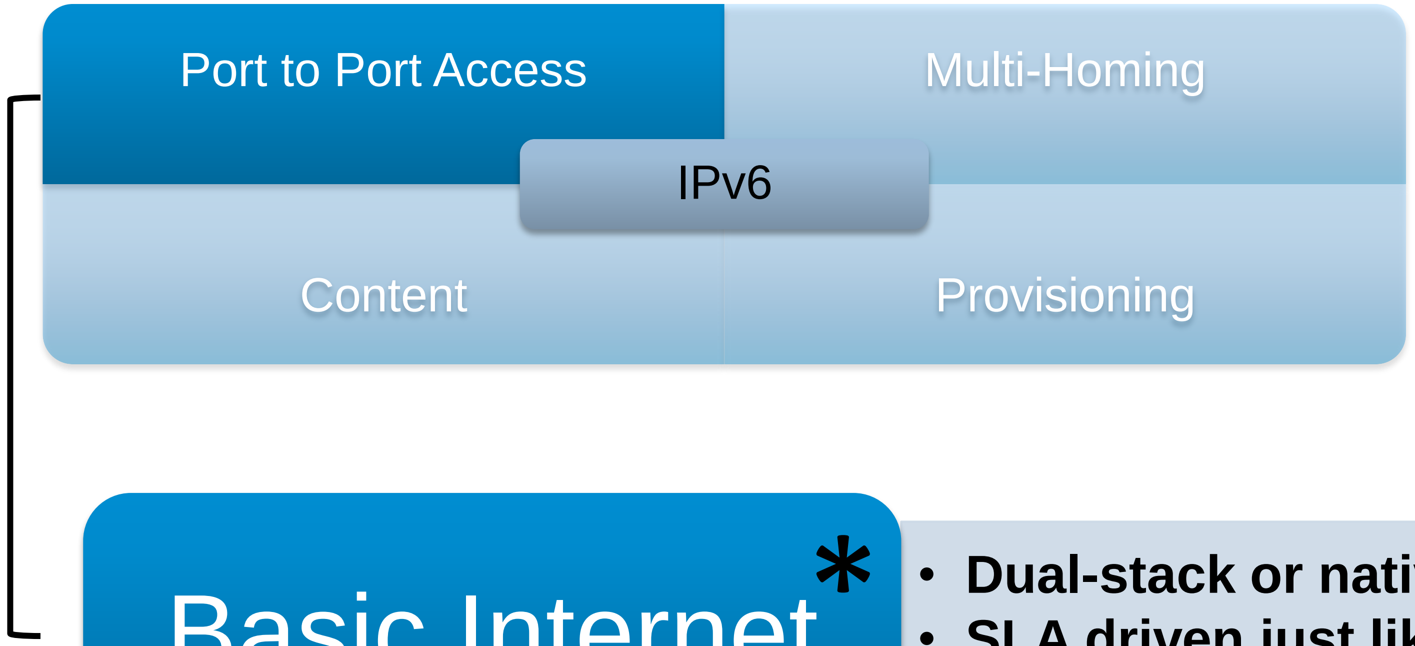

Basic Internet *

- **Dual-stack or native IPv6 at each POP**
- **SLA driven just like IPv4 to support VPN, content access**

MPLS

- **6VPE**
- **IPv6 Multicast**
- **End-to-End traceability**

Hosted (see content)

- **IPv6 access to hosted content**
- **Cloud migration (move data from Ent DC to Hosted DC)**

* = most common issue

Multi-Homing

Port to Port Access	Multi-Homing
IPv6	
Content	Provisioning

PI/PA Policy* Concerns

- PA is no good for customers with multiple providers or change them at any pace
- PI is new, constantly changing expectations and no “guarantee” an SP won’t do something stupid like not route PI space
- Customers fear that RIR will review existing IPv4 space and want it back if they get IPv6 PI

NAT

- Religious debate about the security exposure – not a multi-homing issue
- If customer uses NAT like they do today to prevent address/policy exposure, where do they get the technology from – no scalable IPv6 NAT exists today

Routing

- Is it really different from what we do today with IPv4? Is this policy stuff?
- Guidance on prefixes per peering point, per theater, per ISP, ingress/egress rules, etc.. – this is largely missing today

Content

Port to Port Access | Multi-Homing

IPv6

Content | Provisioning

Hosted/Cloud Apps today *

- IPv6 provisioning and access to hosted or cloud-based services today (existing agreements)
- Salesforce.com, Microsoft BPOS (Business Productivity Online Services), Amazon, Google Apps

Move to Hosted/Cloud

- Movement from internal-only DC services to hosted/cloud-based DC
- Provisioning, data/network migration services, DR/HA

Contract/Managed Marketing/Portals

- Third-party marketing, business development, outsourcing
- Existing contracts – connect over IPv6

Provisioning

Port to Port Access	Multi-Homing
IPv6	
Content	Provisioning

SP Self-Service Portals

- Not a lot of information from accounts on this but it does concern them
- How can they provision their own services (i.e. cloud) to include IPv6 services and do it over IPv6

SLA *

- More of a management topic but the point here is that customers want the ability to alter their services based on violations, expiration or restrictions on the SLA
- Again, how can they do this over IPv6 AND for IPv6 services

Conclusion

- “Dual stack where you can – Tunnel where you must – Translate when you have a gun to your head” – It’s fun to say, but just not as practical as it used to be
- Don’t shortcut your Internet-facing deployment or it will hurt (latency, availability, security, user experience)
- There are so many options that it can be overwhelming – test and then test again
- It is all about the application and user experience
- Create a virtual team of IT representatives from every area of IT to ensure coverage for OS, Apps, Network and Operations/Management
- Now is your time to build a network your way – don’t carry the IPv4 mindset forward with IPv6 unless it makes sense

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm

Cisco live! 365

Don't forget to activate your Cisco Live 365 account for access to all session material, communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.
www.ciscoliveaustralia.com/portal/login.ww