# Enterprise IPv6 Internet Edge Design

**Shannon McFarland**
**CCIE# 5245**
**Principal Engineer**
**Corporate Consulting Engineering**
**Research and Advanced Development**

 © 2012 Cisco Systems, Inc. All rights reserved. 

# 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/solutions/ns340/ns414/ns742/ns816/landing_br_ipv6.html
- New/Updated IPv6 Cisco Sites: http://www.cisco.com/go/ipv6 http://www.cisco.gom/go/entipv6

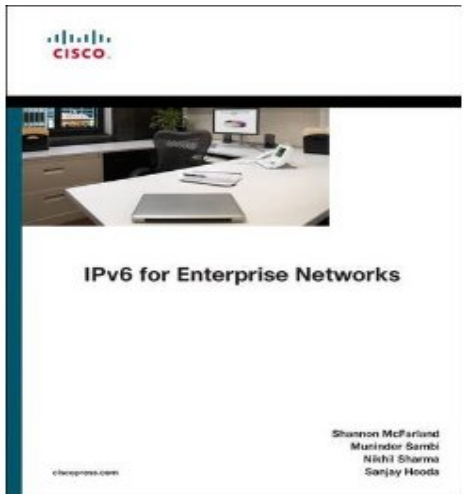


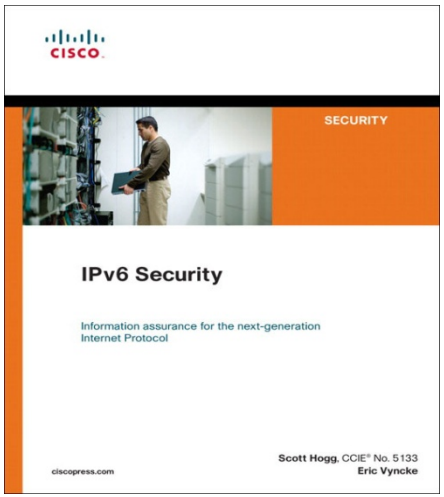



© 2012 Cisco Systems, Inc. All rights reserved.

# Enterprises Responding to Pressure

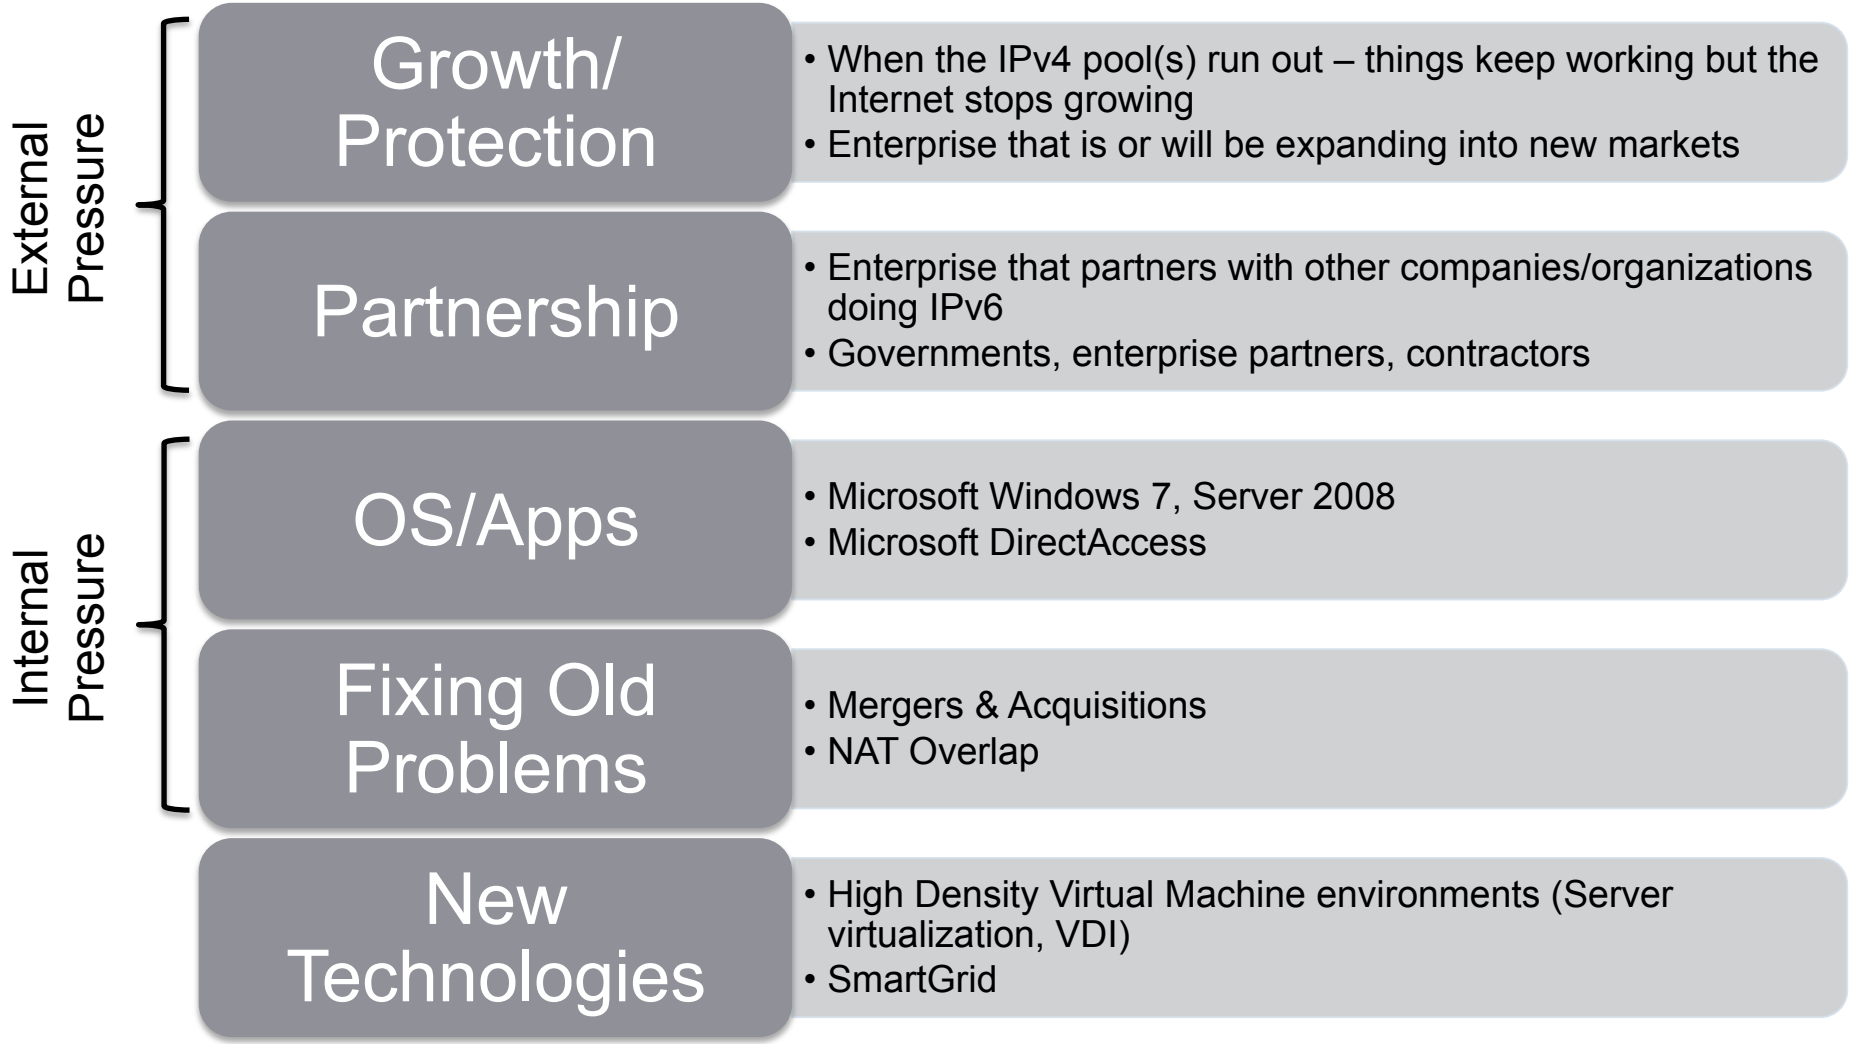



© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Requirements for any IPv6 Deployment Strategy

- Should be low-cost
- Must be low-risk
- Must co-exist with existing IPv4 infrastructure
- Must allow access to public Internet
- Must be incrementally deployable
- Must understand the cost of adding a new service
- Must not impact existing services
- End-user should not know the integration occurred (seamless)


© 2012 Cisco Systems, Inc. All rights reserved.

# Common Deployment Models for Internet Edge

## Pure Dual Stack

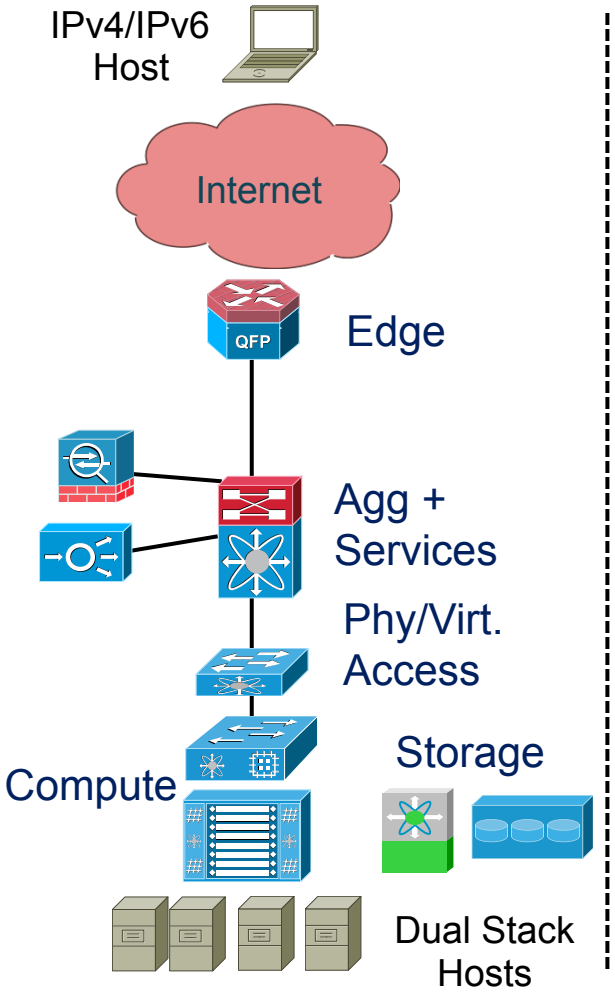


## Conditional Dual Stack

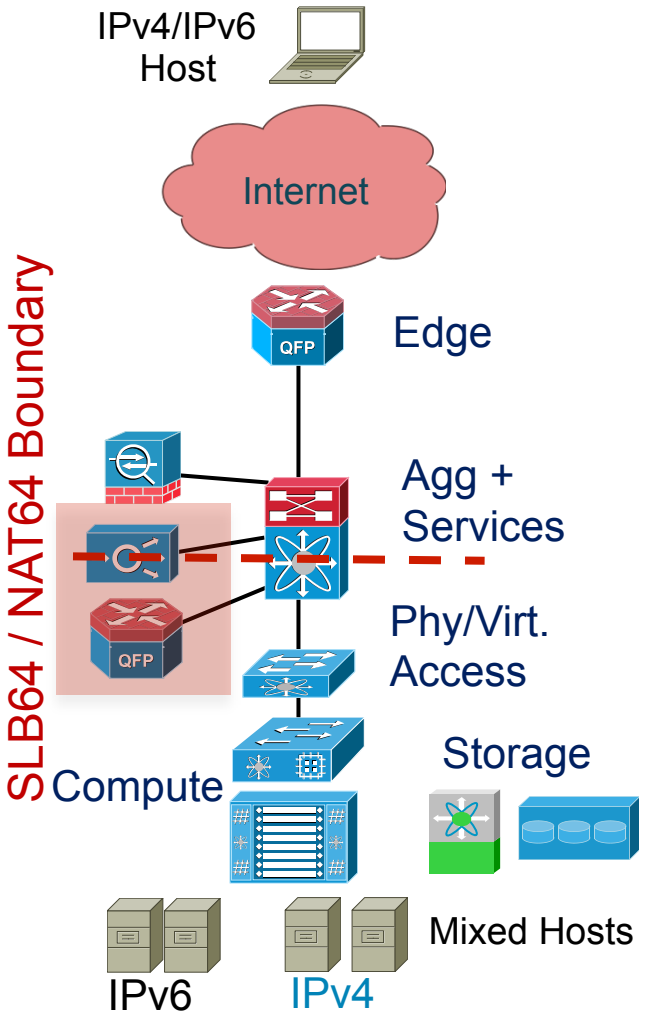


## Translation as a Service

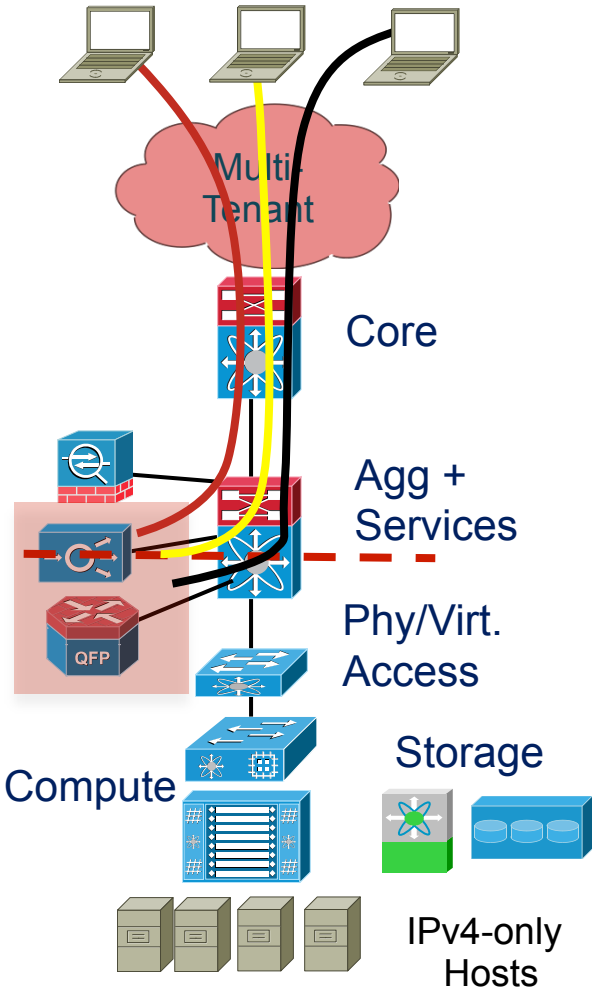



© 2012 Cisco Systems, Inc. All rights reserved.
Cisco Public

# Global Addressing Dilemma

- Today, many do NAT44 and ‘hide’ their RFC1918 space allowing for easier multi-homing scenarios
- One Provider Independent (PI) prefix for all regions or a PI per region?
- 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
- http://tools.ietf.org/html/draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat


© 2012 Cisco Systems, Inc. All rights reserved.

# Three Speeds of Enterprise Deployment

- Based on Timeframe/Use case
- Core-to-Edge – I can wait to go end-to-end
- Edge-to-Core – I need something now inside
- Internet Edge – Business continuity

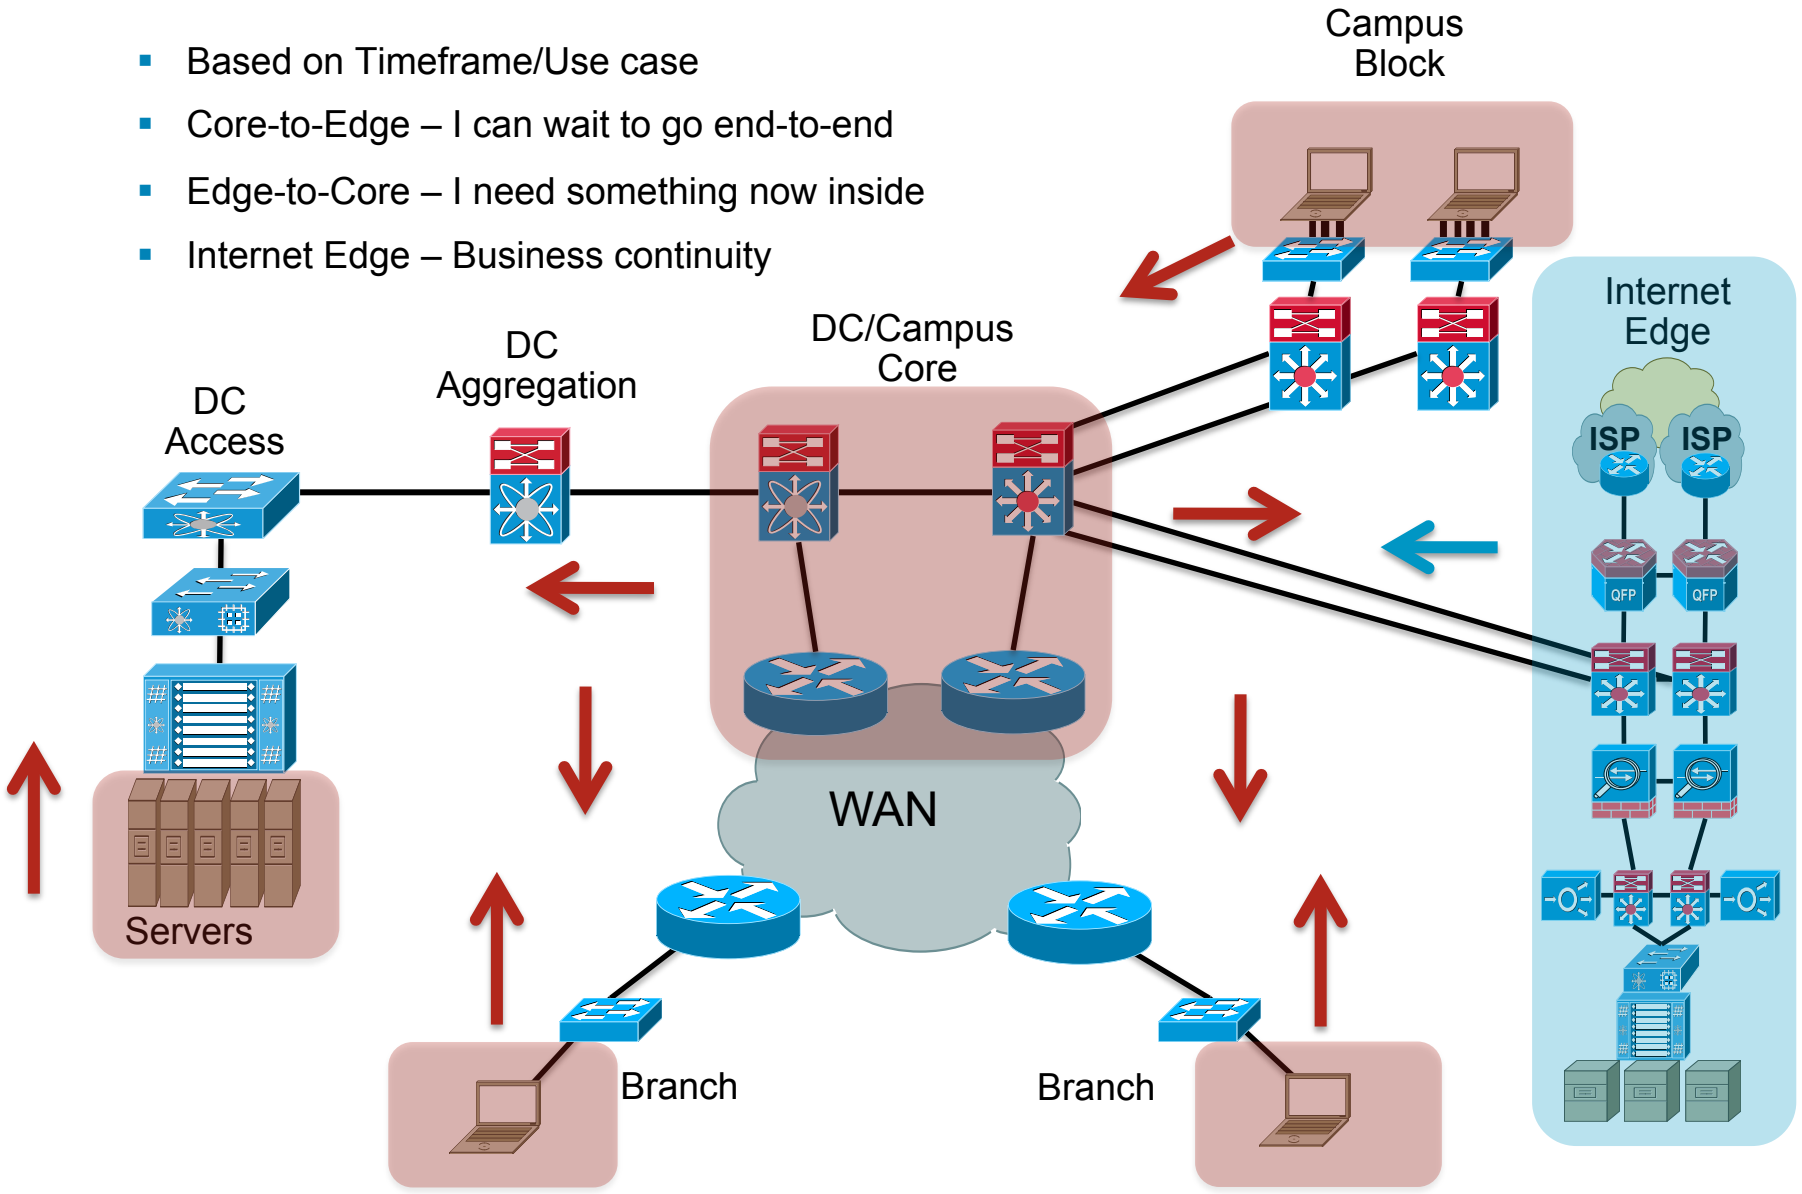



© 2012 Cisco Systems, Inc. All rights reserved.
Cisco Public

# Pick one and go

- Dual stack it all
- Dual stack as much as you can and translate
- LISP (Locator/ID Separation Protocol)
- What if your junk is in the Cloud?

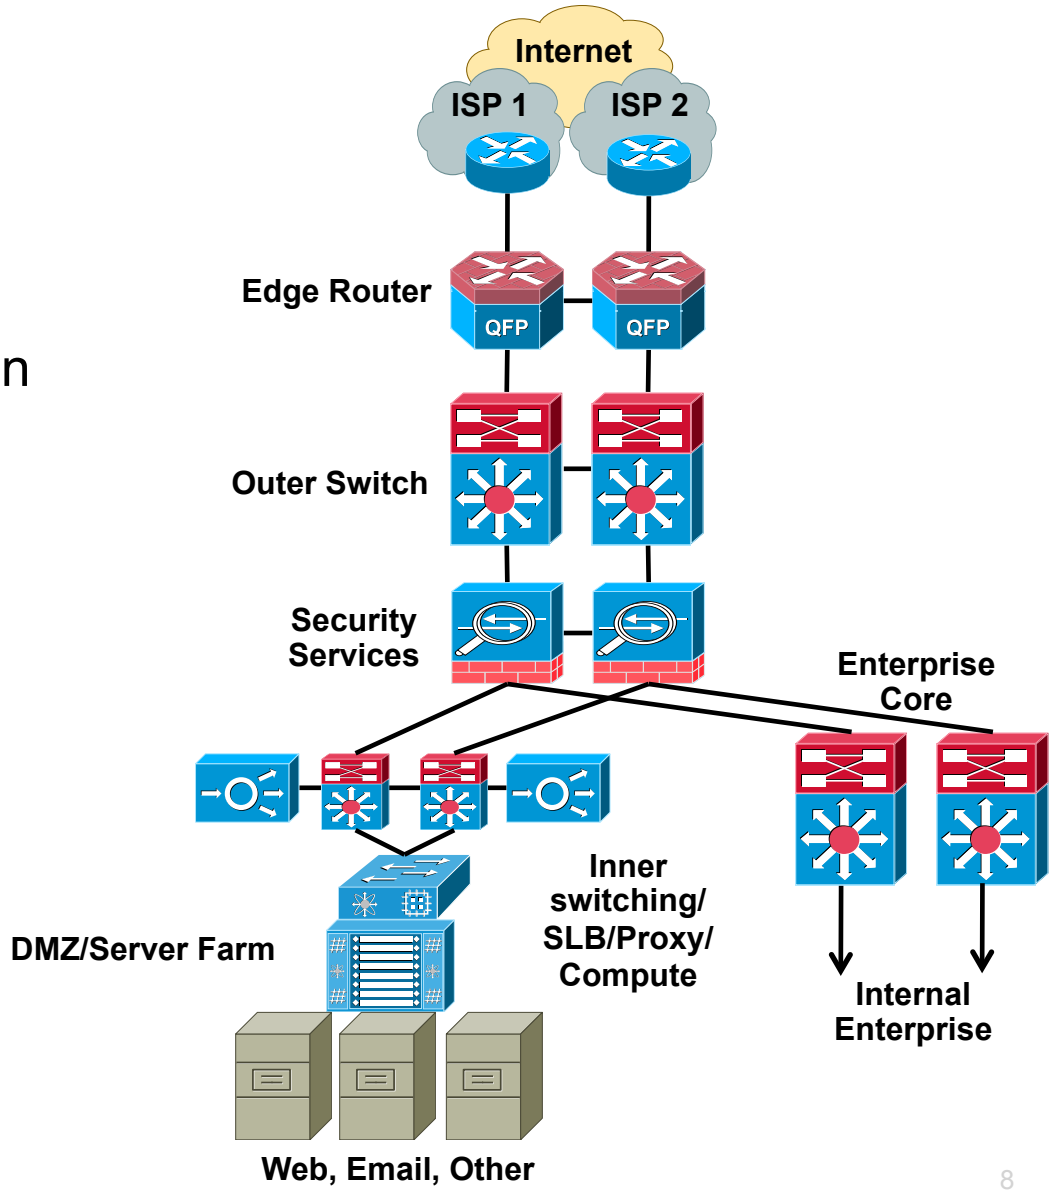



© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Multi-Homed – Dual Stack


© 2012 Cisco Systems, Inc. All rights reserved.

# Single ISP – Multi-Peer - DS

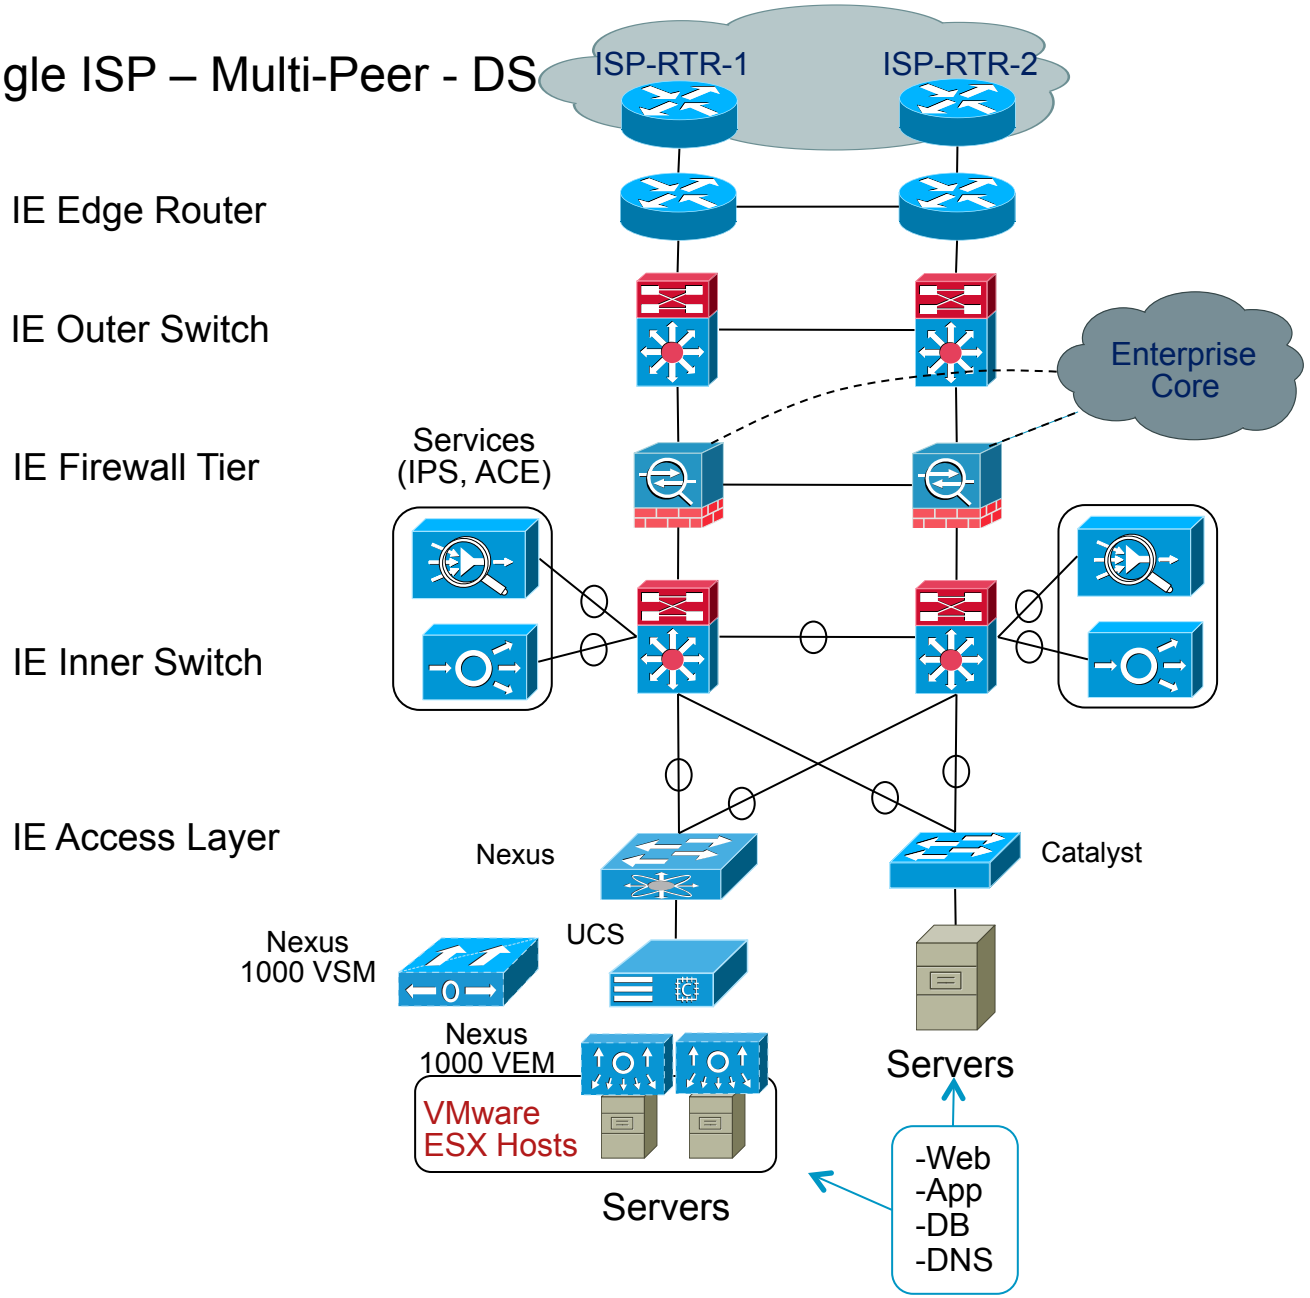


- Single ISP or multi-ISP changes BGP slightly
- PA vs. PI vs. NPTv6
- Behind the edge it all stays the same


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Services and Applications

- SLB66 on Cisco ACE – One arm mode
- Cisco ASA in A/A or A/S – Failover over IPv4 OR IPv6
- Cisco IPS/IDS are inline between ASA and inner switches
- Baremetal servers on Catalyst or Nexus and UCS C-Series
- Virtualized on Nexus 5000, Nexus 1000v and UCS C-Series or other combo


© 2012 Cisco Systems, Inc. All rights reserved.

# Edge Peering

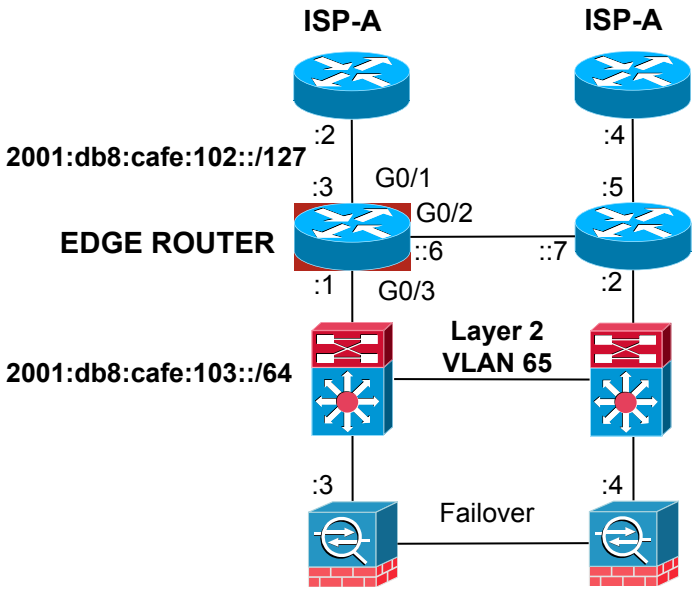


- Basic IP/Interface of left edge router
- /127s used on P2P
- /64 on shared links

```
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
 no ipv6 redirects
!
interface GigabitEthernet0/2
 description LINK to 7206-2-edge
 ipv6 address 2001:DB8:CAFE:102::6/127
 no ipv6 redirects
!
interface GigabitEthernet0/3
 description to ASA
 ipv6 address 2001:DB8:CAFE:103::1/64
 no ipv6 redirects
 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
```

HW-Dependent Support


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# BGP - Edge Router

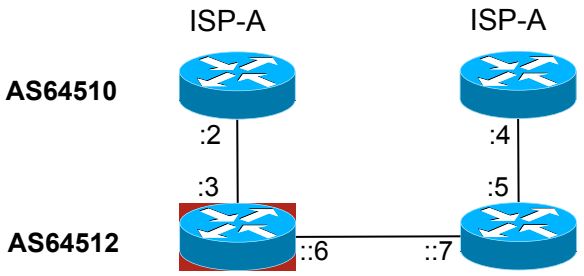


- eBGP to ISP
- iBGP to local edge router
- ‘no bgp default ipv4-unicast’ allows for multi-AF neighbor activation

```
router bgp 64512
 bgp router-id 192.168.1.33
 no bgp default ipv4-unicast
 bgp log-neighbor-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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# BGP Filters

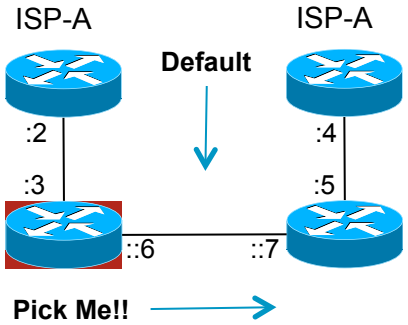


- Accepting default only
- Setting higher local pref
- ACLs for BGP

Reference

```
 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# BGP Filters - Secondary

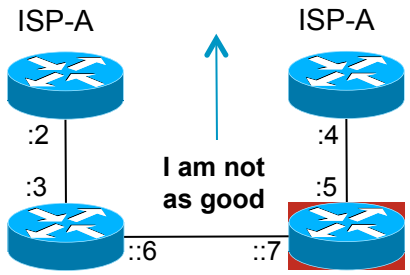


- Accepting default only
- AS PATH Prepend
- ACLs for BGP

```
 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

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

## ISP Router

```
ISPA-1#sh ip bgp ipv6 unicast
.....
 Network          Next Hop            Metric LocPrf Weight Path
*> 2001:DB8:CAFE::/48
                    2001:DB8:CAFE:102::3
                                         0             0 64512 i
*                   2001:DB8:CAFE:102::5
                                         0             0 64512 64512 i
```

AS Path Prepend


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Apply Appropriate ACLs/CoPP

- Protect infrastructure that can be hurt by control plane processing
- HbH, RH0 (http://tools.ietf.org/html/rfc5095), etc …
- Check that all networking vendors can handle /127 and/or protect against ICMP “ping pong” attacks

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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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
```

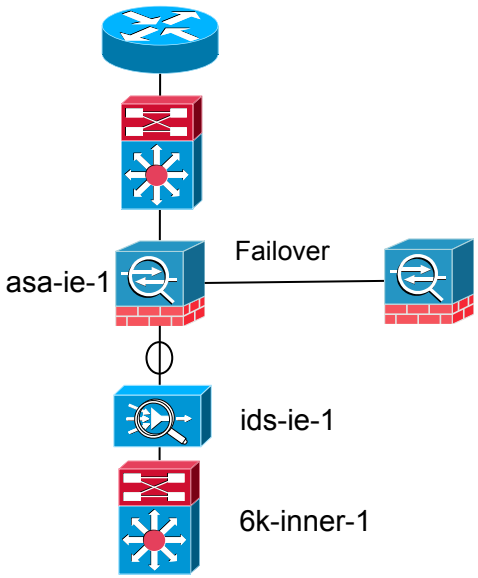


- VLANs on ASA or on inside switch
- L2 or L3 sandwich does not impact much


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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 (IPv4 or IPv6)

```
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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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?

- Object for ACE VIP
- Object for DNS
- ACL for L3/L4 stuff

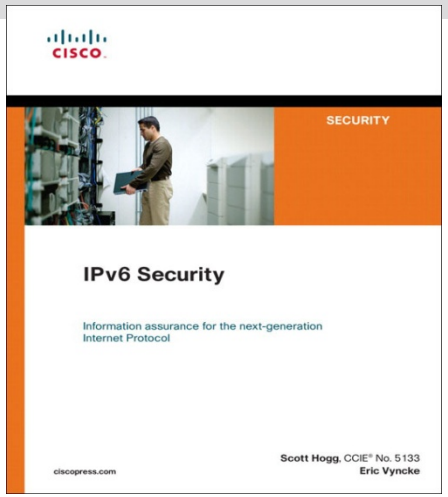



© 2012 Cisco Systems, Inc. All rights reserved.

# ASA Device Manager

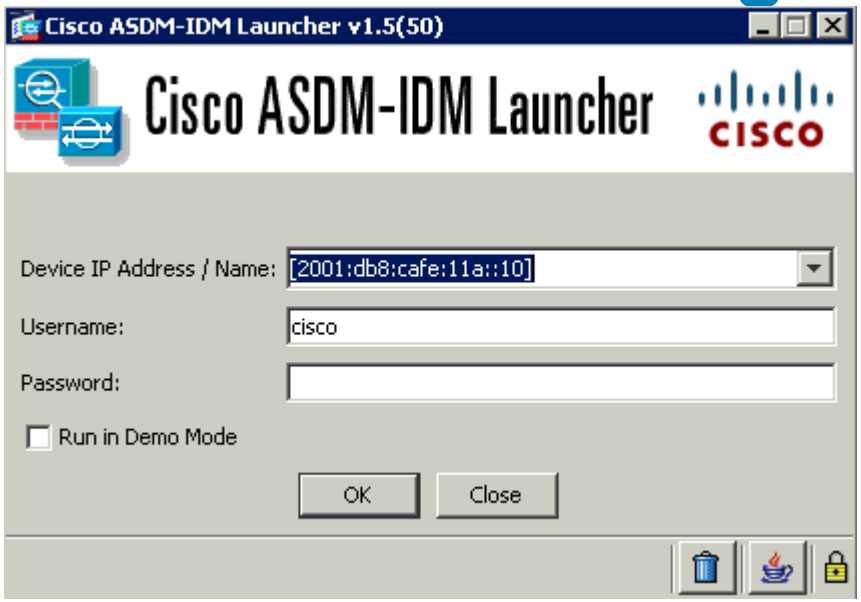


```
http server enable
http 2001:db8:cafe::/48 management
```

| # | 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 |


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# IDS/IPS

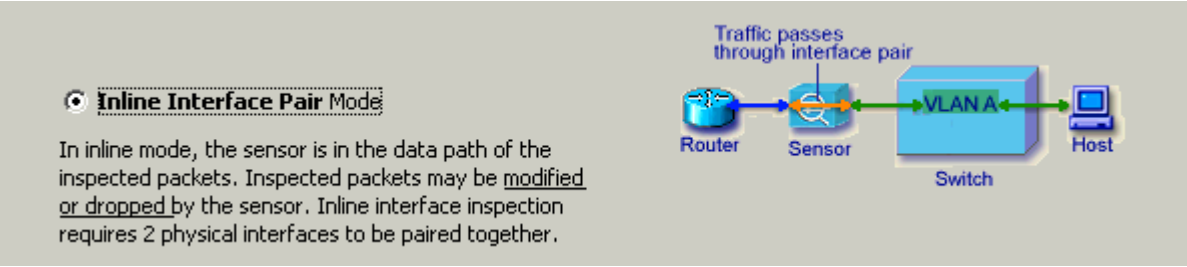


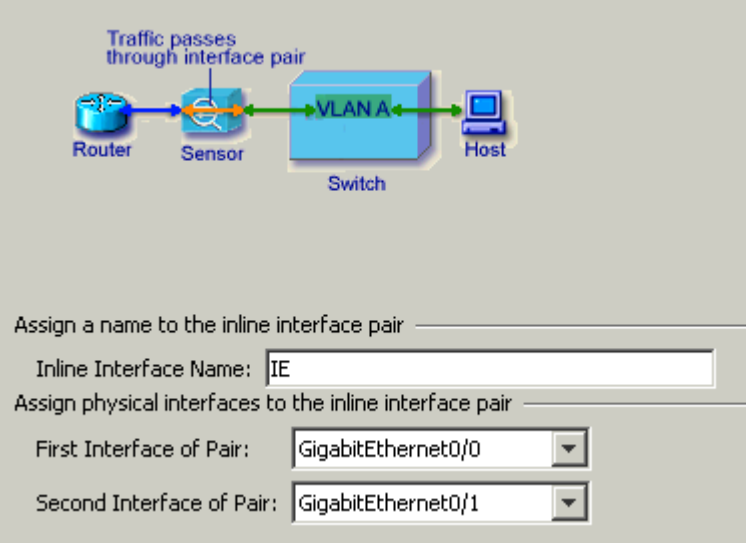


| Severity | Date | Time | Device | Sig. Name | Sig. ID | Attacker IP | Victim IP | Vicitm ... | Threa... |
|---|---|---|---|---|---|---|---|---|---|
| high | 09/27/2011 | 12:24:02 | ids-ie-1 | WWW WinNT cmd.exe Access | 5081/0 | 172.16.99.100 | 10.140.19.10 | 80 | 90 |
| high | 09/27/2011 | 12:24:42 | 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 |


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# Connecting the Inside

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

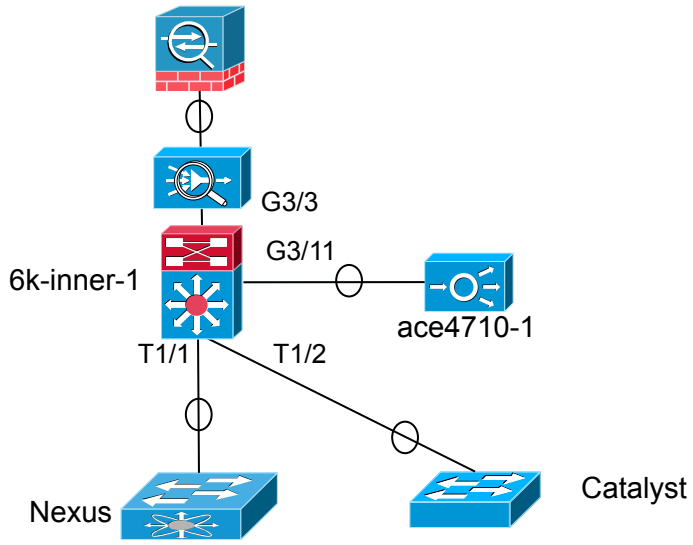


```
interface TenGigabitEthernet1/1
 description to Nexus Access Layer
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 18-25
 switchport mode trunk
 switchport nonegotiate
 spanning-tree guard root
!
interface TenGigabitEthernet1/2
 description to Catalyst Access Layer
 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# Cisco ACE – Context Definition

## Trunked Interface – One-arm Mode

```
interface port-channel 1
  description to IE-Trunk
  switchport trunk allowed vlan 19-22,24,132
  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
```


© 2012 Cisco Systems, Inc. All rights reserved.

Reference

# Cisco ACE – Fault Tolerance (over IPv4)

FT Interface over IPv4 on A5(1.0)

```
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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# IE-WEB Context - MGMT

Reference

```
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
```


© 2012 Cisco Systems, Inc. All rights reserved.

# 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

Don’t screw this up


© 2012 Cisco Systems, Inc. All rights reserved.

# 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Application Networking Manager 5.1

- Full Monitoring
- Configure all elements of policies
- Configure by context, filter by multiple conditions, etc..

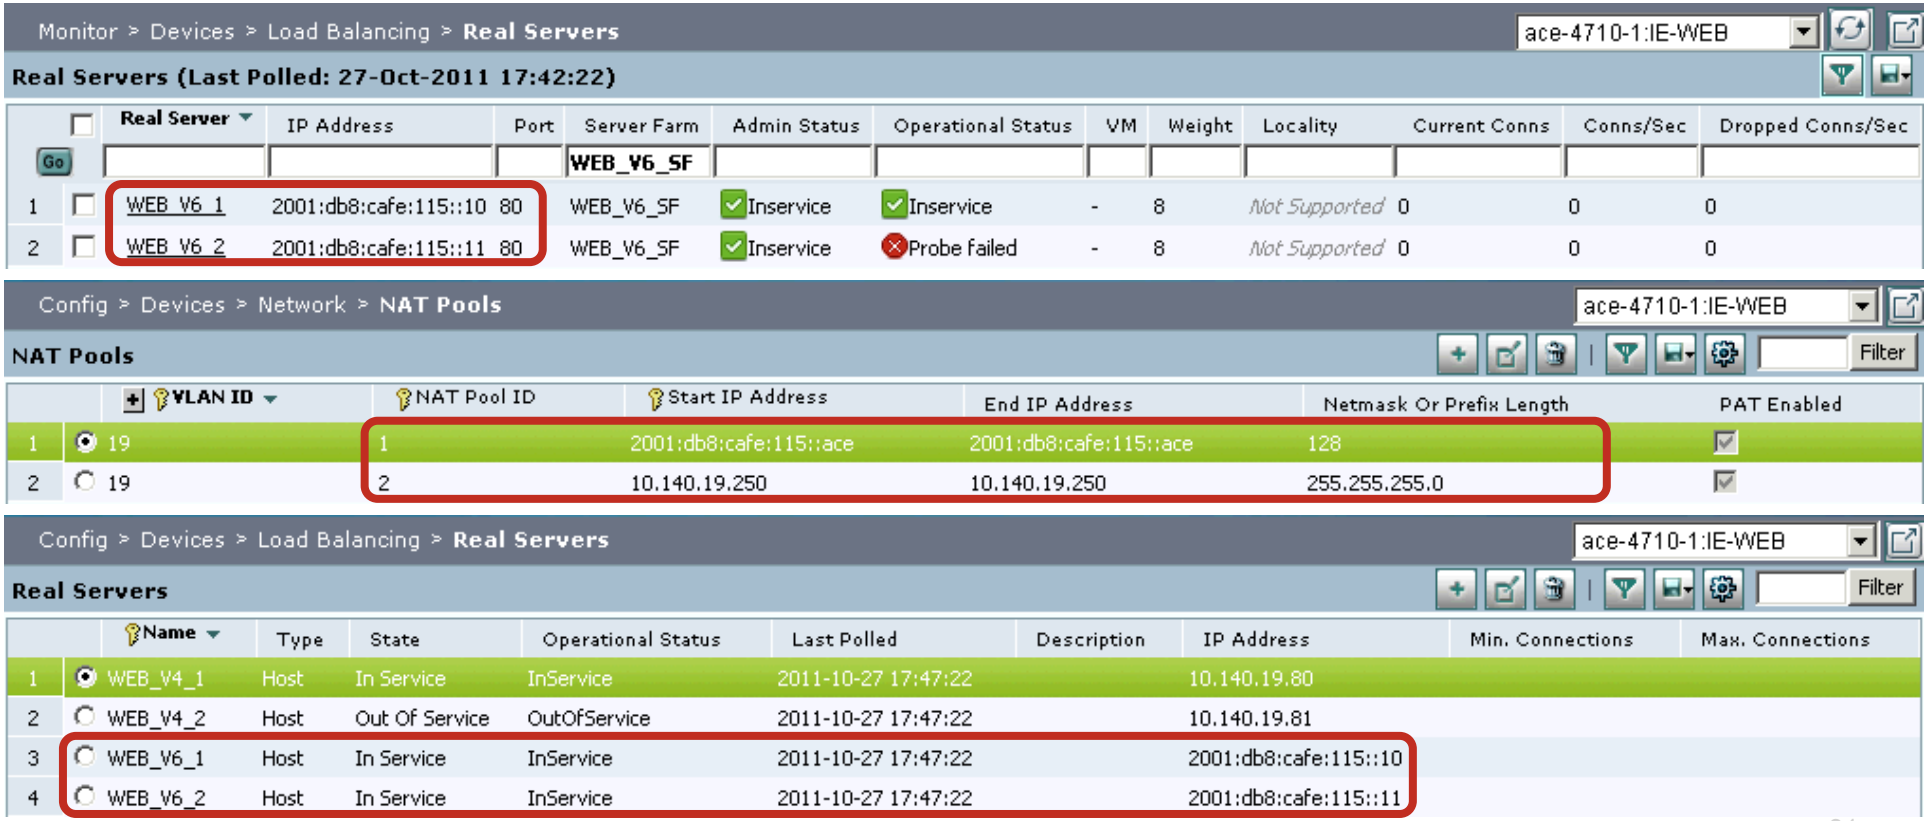



© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Access Layer Examples
## Your platform may vary

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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# VMware ESXi – IPv6 (1)


- vSphere IPv6 support since 4.1
- Static or dynamically assigned addresses
- Can restart mgmt, but should reboot host


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# VMware ESXi – IPv6 (2)

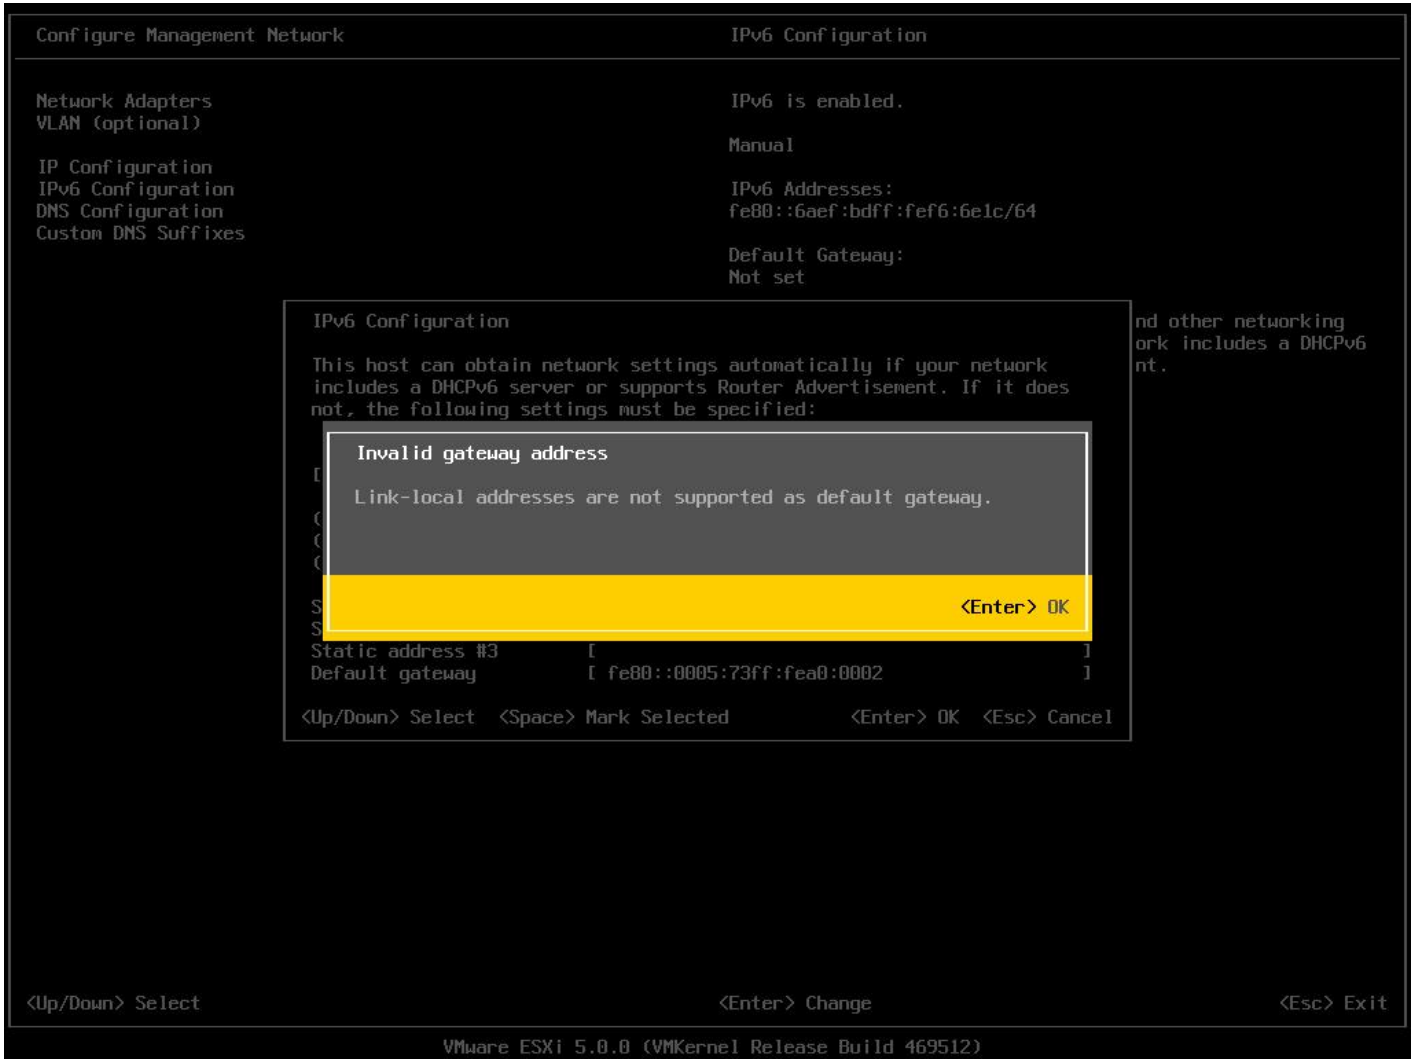


- As of ESX 5 you cannot set a LL address as a gateway
- VERY BAD
- Global or let it learn via RA


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# VMware ESXi – IPv6 (3)

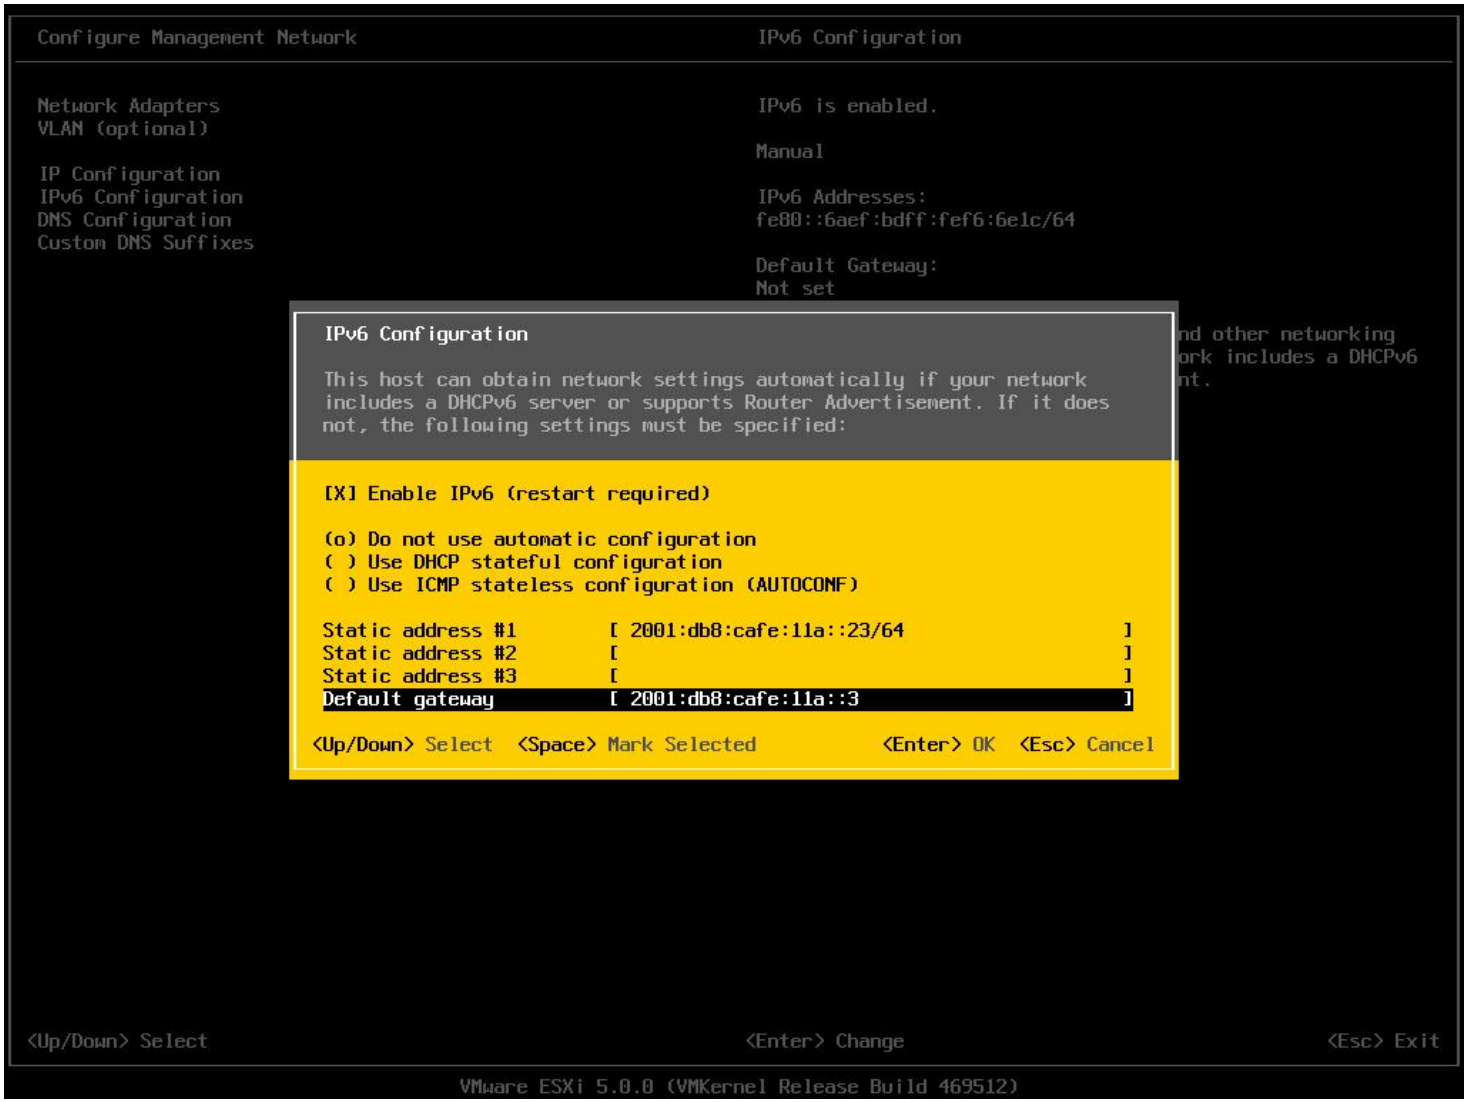


- Single GW or if GW can support FHRP on Global = OK
- If not, let host learn GW via RA (Test this!!)


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Multi-Homed – SLB64


© 2012 Cisco Systems, Inc. All rights reserved.

# Multihomed – SLB64

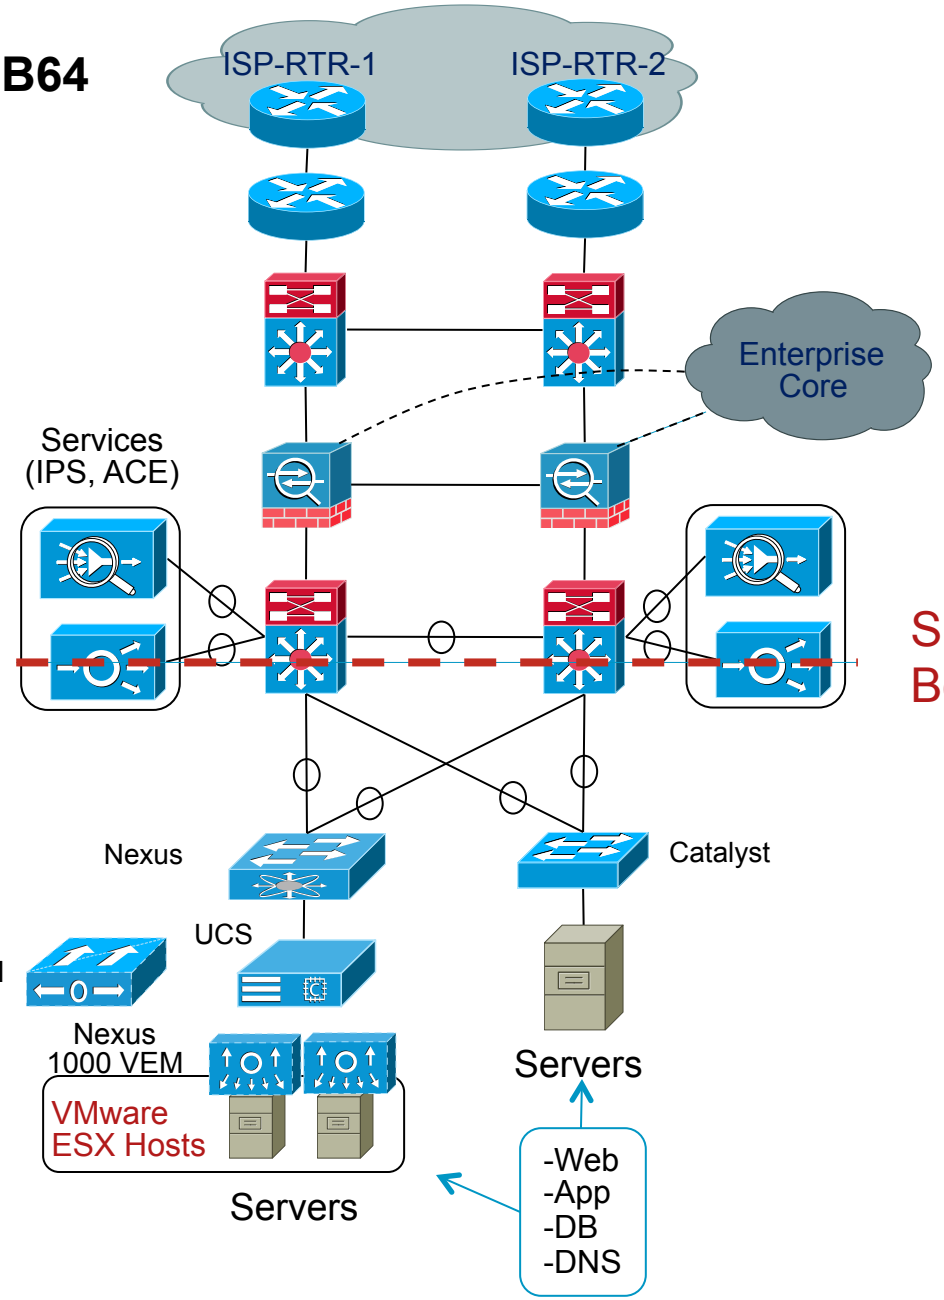


- Dual stack to the Cisco ACE
- IPv4-only South of Cisco ACE


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Services and Applications Tested

- SLB64 on Cisco ACE – One arm mode
- Cisco ASA in A/A or A/S – Failover over IPv4 OR IPv6
- Cisco IPS/IDS
- In my setup everything south of ACE is IPv4-only


© 2012 Cisco Systems, Inc. All rights reserved.

Reference

# Cisco ACE – Context Definition

Trunked Interface – One-arm Mode

```
interface port-channel 1
  description to IE-Trunk
  switchport trunk allowed vlan 19-22,24,132
  no shutdown
```

- Nothing changes from previous SLB66 example

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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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
```

```
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
```


© 2012 Cisco Systems, Inc. All rights reserved.

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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Application Networking Manager 5.1

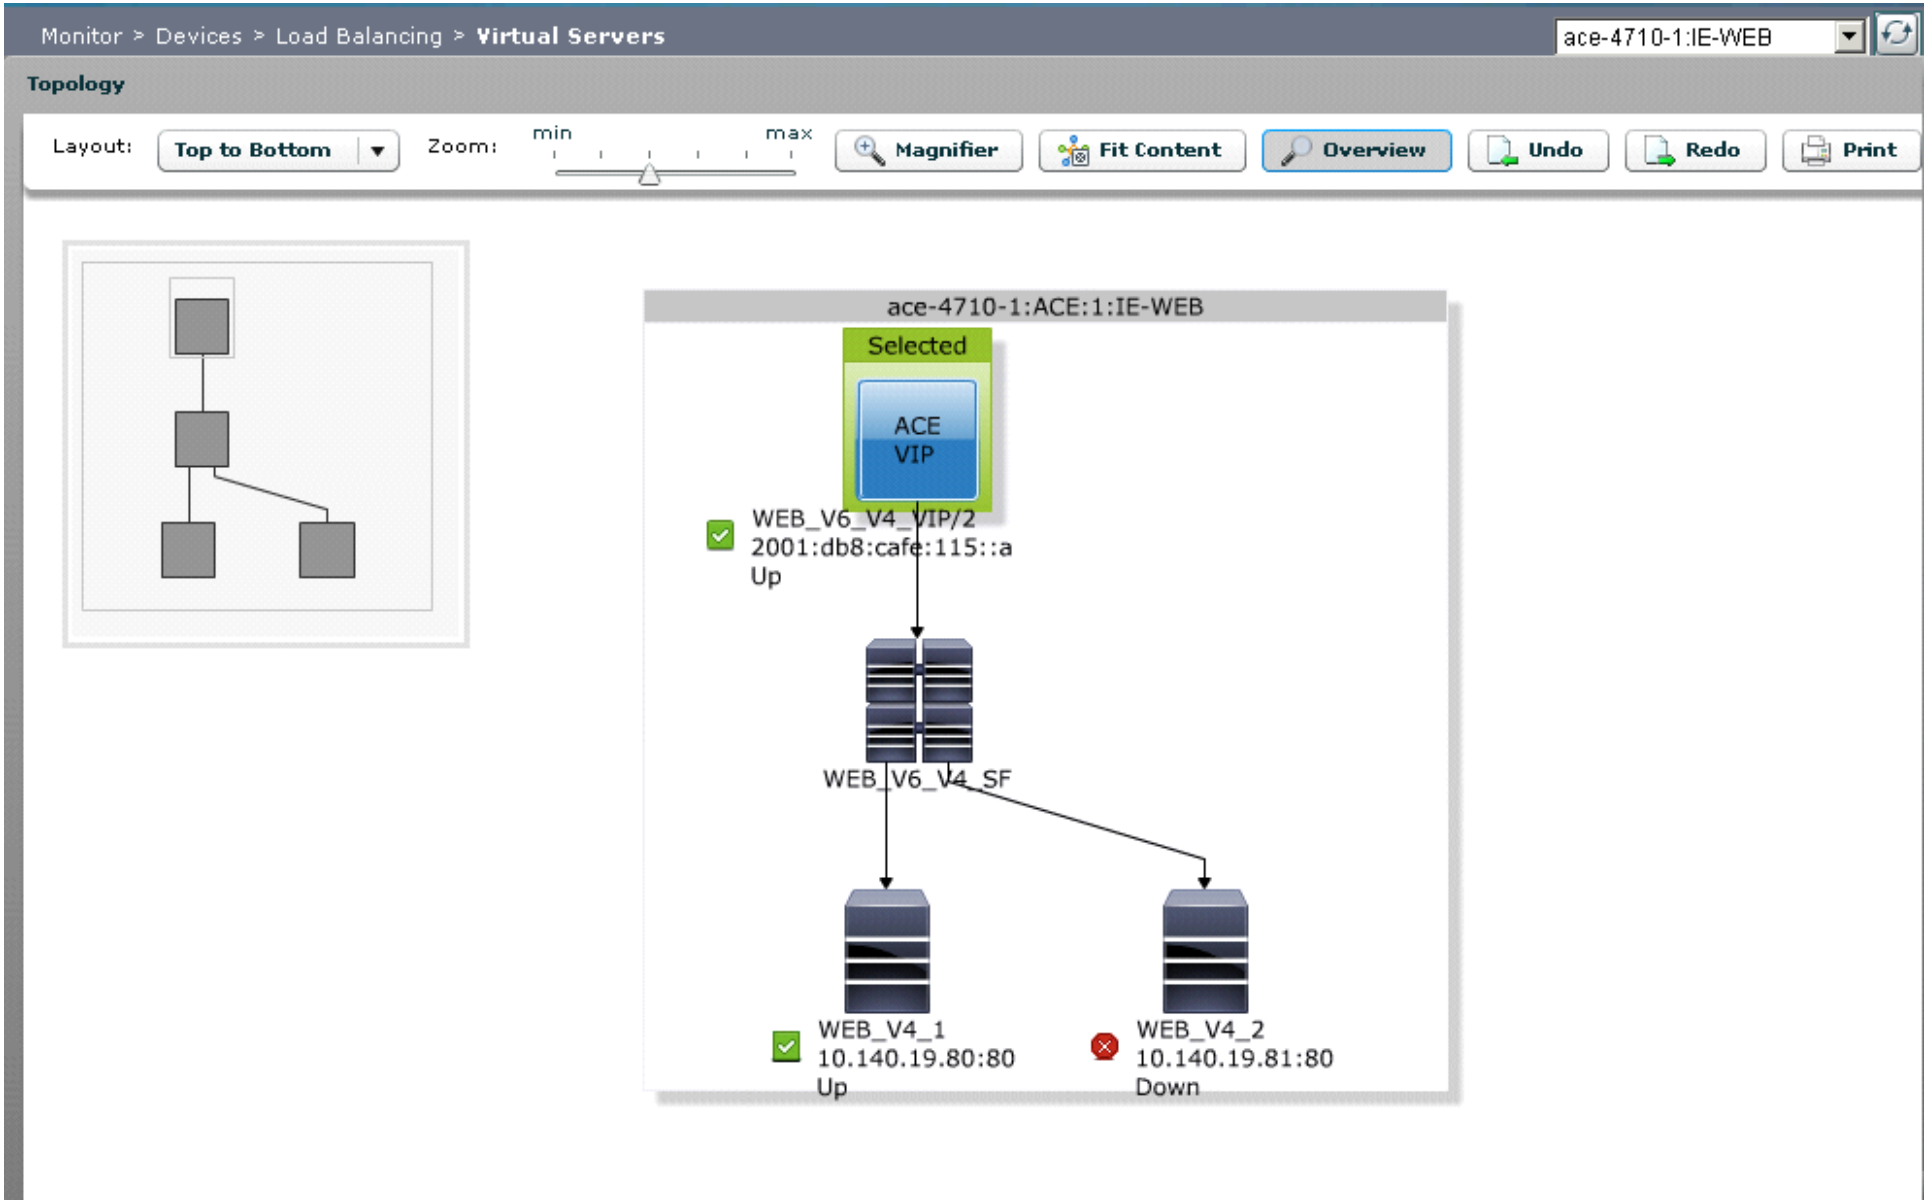

# Validation of Connection

```
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

- First connection pair are IPv6 and between client and VIP
- Second connection pair are IPv4 and between SNAT address (we are in one arm mode) and real server


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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
- Make changes to Apache LogFormat/CustomLog to get full use 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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Multi-Homed – Stateful NAT64


© 2012 Cisco Systems, Inc. All rights reserved.

# Multihomed – NAT64

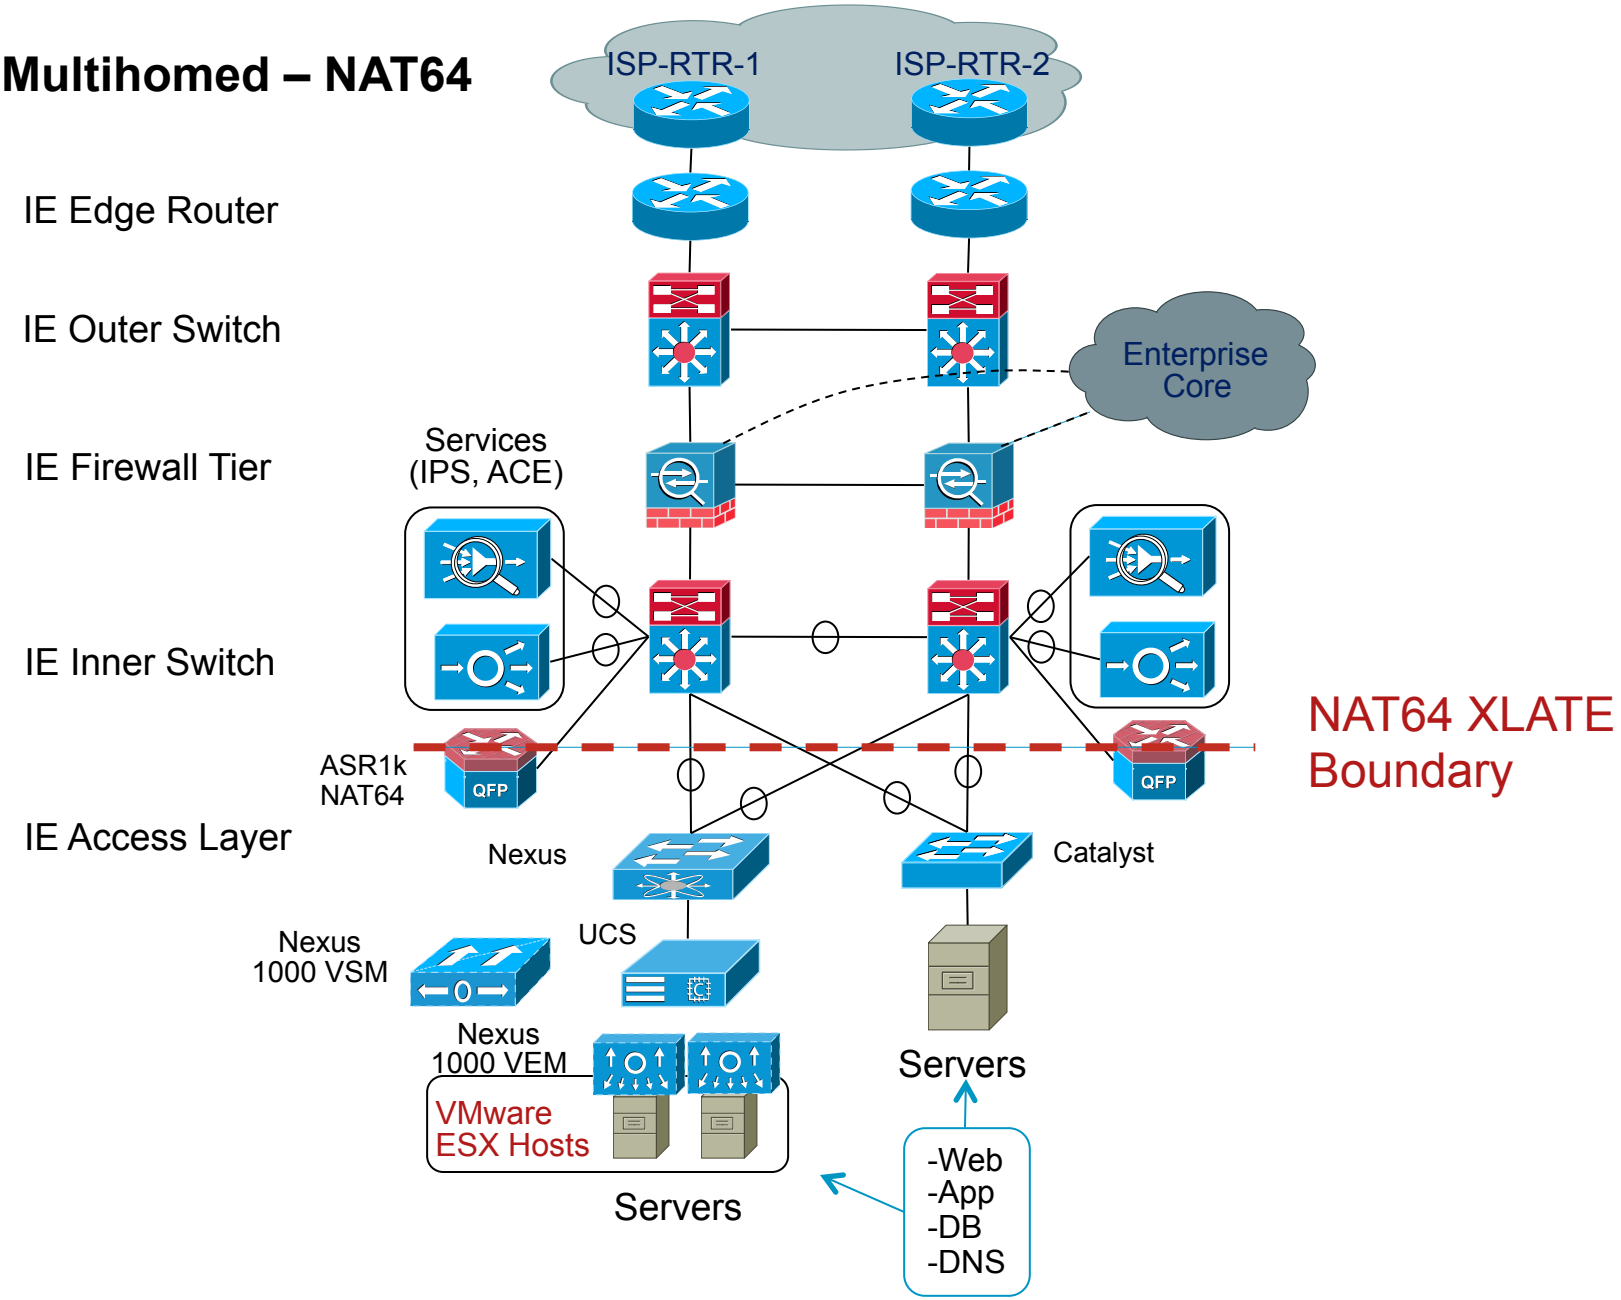



© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Services and Applications

- Reasons for Stateful NAT64 vs. SLB64:
  - Applications don't need SLB
  - SLB can't do SLB64
  - You want to do translation closer to edge router (or on it)
- Cisco ASA in A/A or A/S – Failover over IPv4 OR IPv6
- Cisco IPS/IDS
- Cisco ASR 1k is doing Stateful NAT64
- Everything South of ASR is IPv4-only
- You don't need DNS64 unless you are coming from IPv6-only to IPv4-only – Dual stacked clients can get DNS from IPv4 or IPv6-enabled DNS servers


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

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)
- New Cisco WP: http://bit.ly/poyOey


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Stateful NAT64 – Example Topology
## Static Example

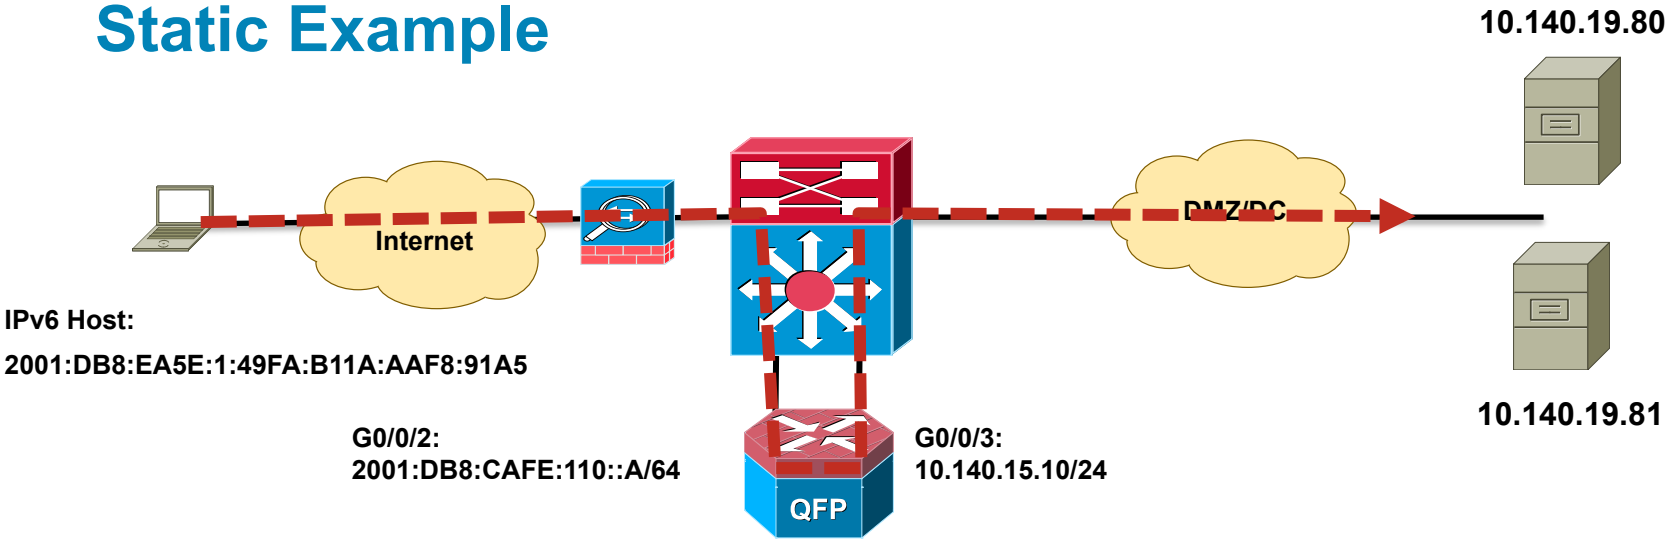


```
interface GigabitEthernet0/0/2
 description to 6k-inner-1 Outside
 no ip address
 ipv6 address 2001:DB8:CAFE:110::A/64
 nat64 enable
!
interface GigabitEthernet0/0/3
 description to 6k-inner-1 Inside
 ip address 10.140.15.10 255.255.255.0
 nat64 enable
```

Lots of nerd knobs (i.e. tune MTU)

```
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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# ASA Interfaces

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ipv6 address 2001:db8:cafe:103::3/120 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
```

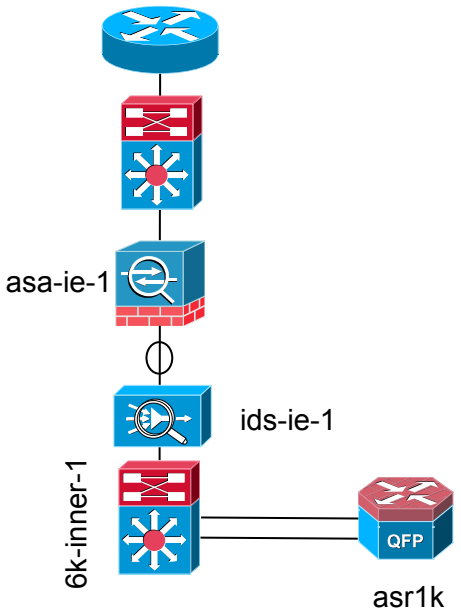


- Many connectivity types – Here, ASR is in VLAN14 that is trunked via 6k pair to the ASA pair
- If doing pure L3 P2P links to 6k then use IPv6 EIGRP to announce NAT64 prefix – here we have to do static route until ASA supports EIGRPv6 or OSPFv3


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

# ASA Object/ACL Configuration

```
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
```

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


© 2012 Cisco Systems, Inc. All rights reserved.

# 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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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
Limit Statistics
```

*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 suitable replacement for x-forwarded-for as most of your existing back-end tools are not setup for NetFlow analysis


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# NetFlow Record IPv6 Original-Input Reference

```
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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# NetFlow Export Example

```
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
```

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


© 2012 Cisco Systems, Inc. All rights reserved.

# 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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# LISP


© 2012 Cisco Systems, Inc. All rights reserved.

# Using LISP to Service IPv6 Access

- Everything we just talked about still applies except:

  You are leveraging LISP as a means to deal with having non-IPv6 capable providers, gear, features or all of the above

- EXTENSIVE amount of information available
- Real customer deployments are wildly successful using LISP for IPv6 (you are probably using it and not know it)
- Sites you need to bookmark

  http://lisp.cisco.com

  http://www.lisp4.net  http://www.lisp6.net

- The source of all goodness:

  http://lisp.cisco.com/lisp_tech.html


© 2012 Cisco Systems, Inc. All rights reserved.  Cisco Public

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


© 2012 Cisco Systems, Inc. All rights reserved.

# LISP Operations
## LISP Components – Ingress/Egress Tunnel Router (xTR)

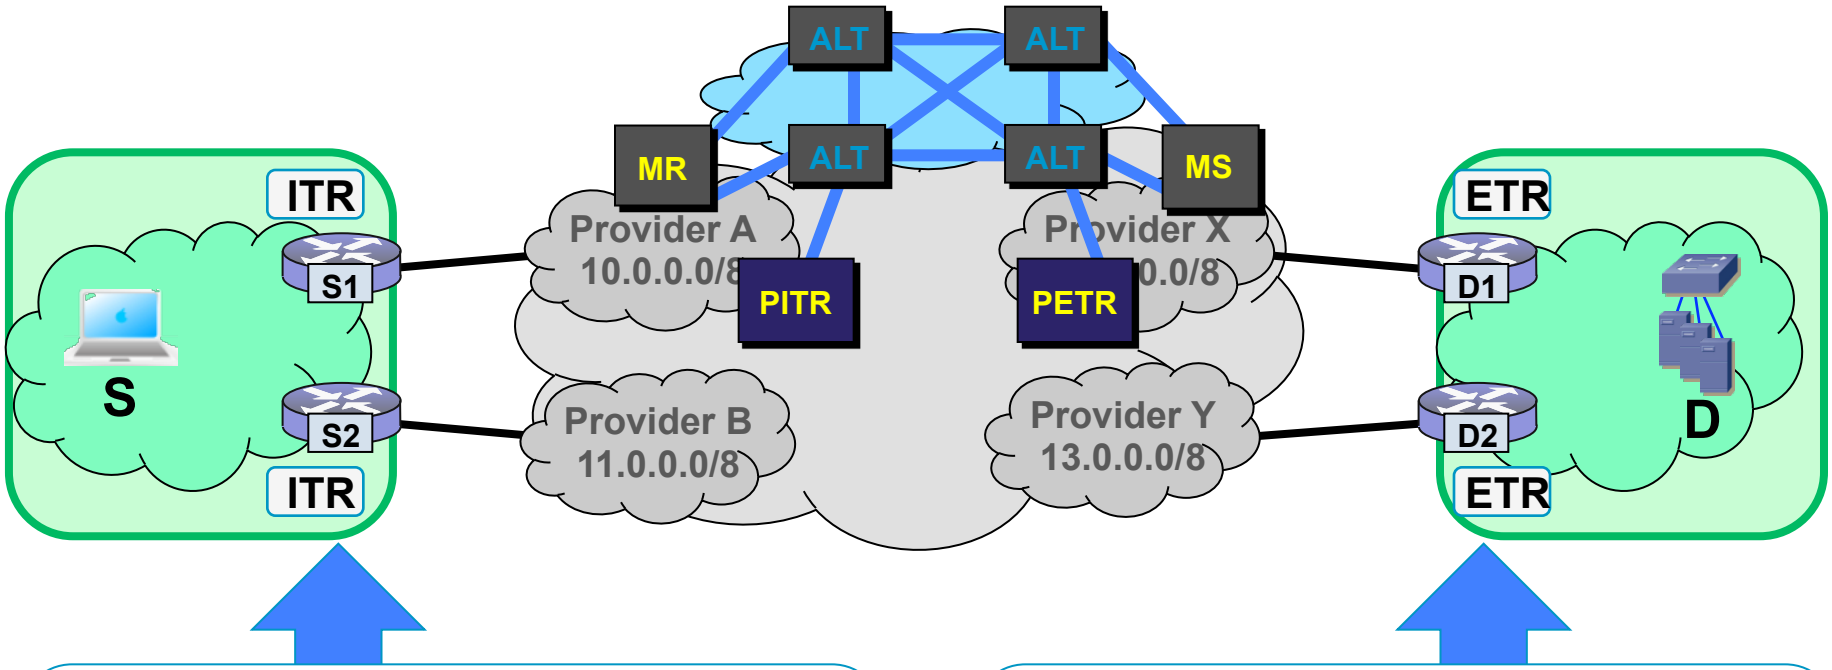


**ITR – Ingress Tunnel Router**
- Receives packets from site-facing interfaces
- Encaps to remote LISP site or natively forwards to non-LISP site

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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

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

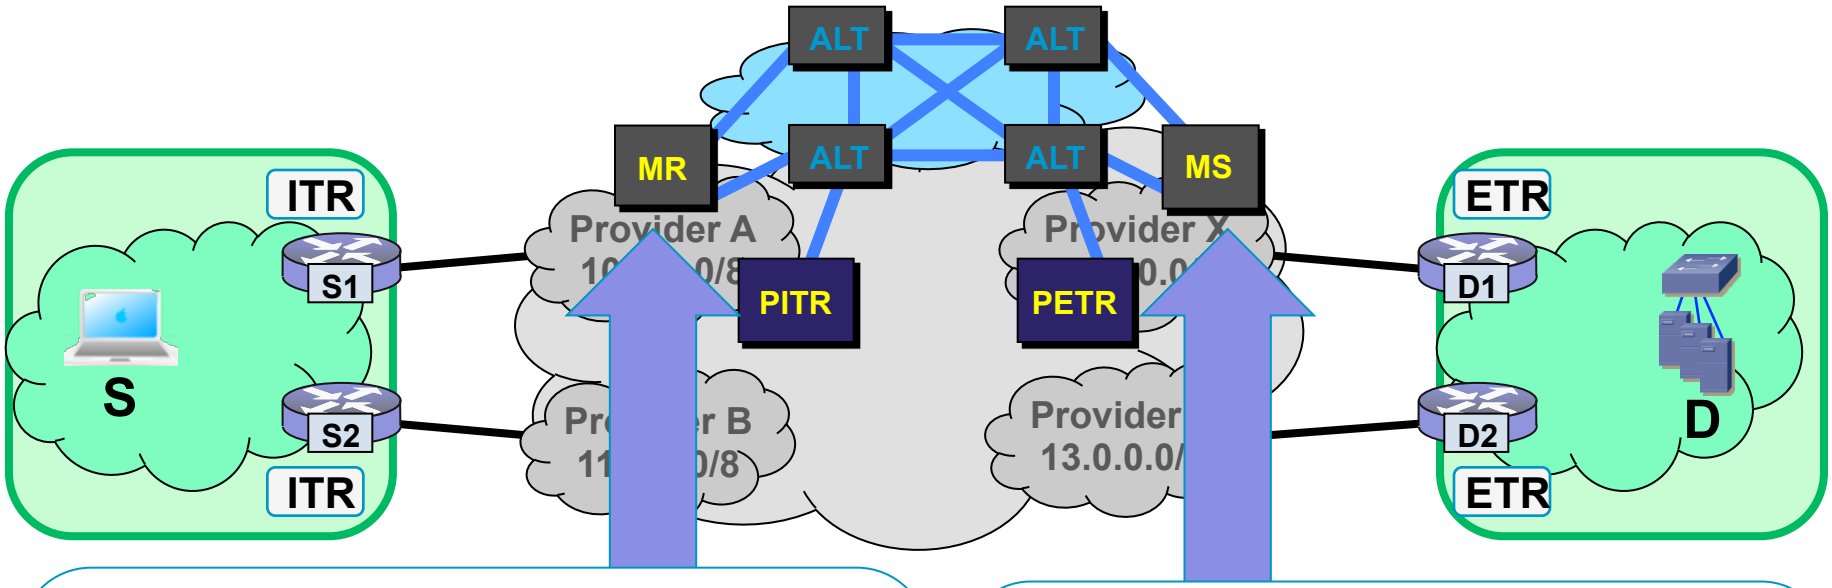


**MR – Map-Resolver**
- Receives Map-Request encapsulated from ITR
- De-caps Map-Request, forwards thru service interface onto the ALT topology
- Sends Negative Map-Replies in response to Map-Requests for non-LISP sites

**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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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 monetized service opportunity for infrastructure players


© 2012 Cisco Systems, Inc. All rights reserved.

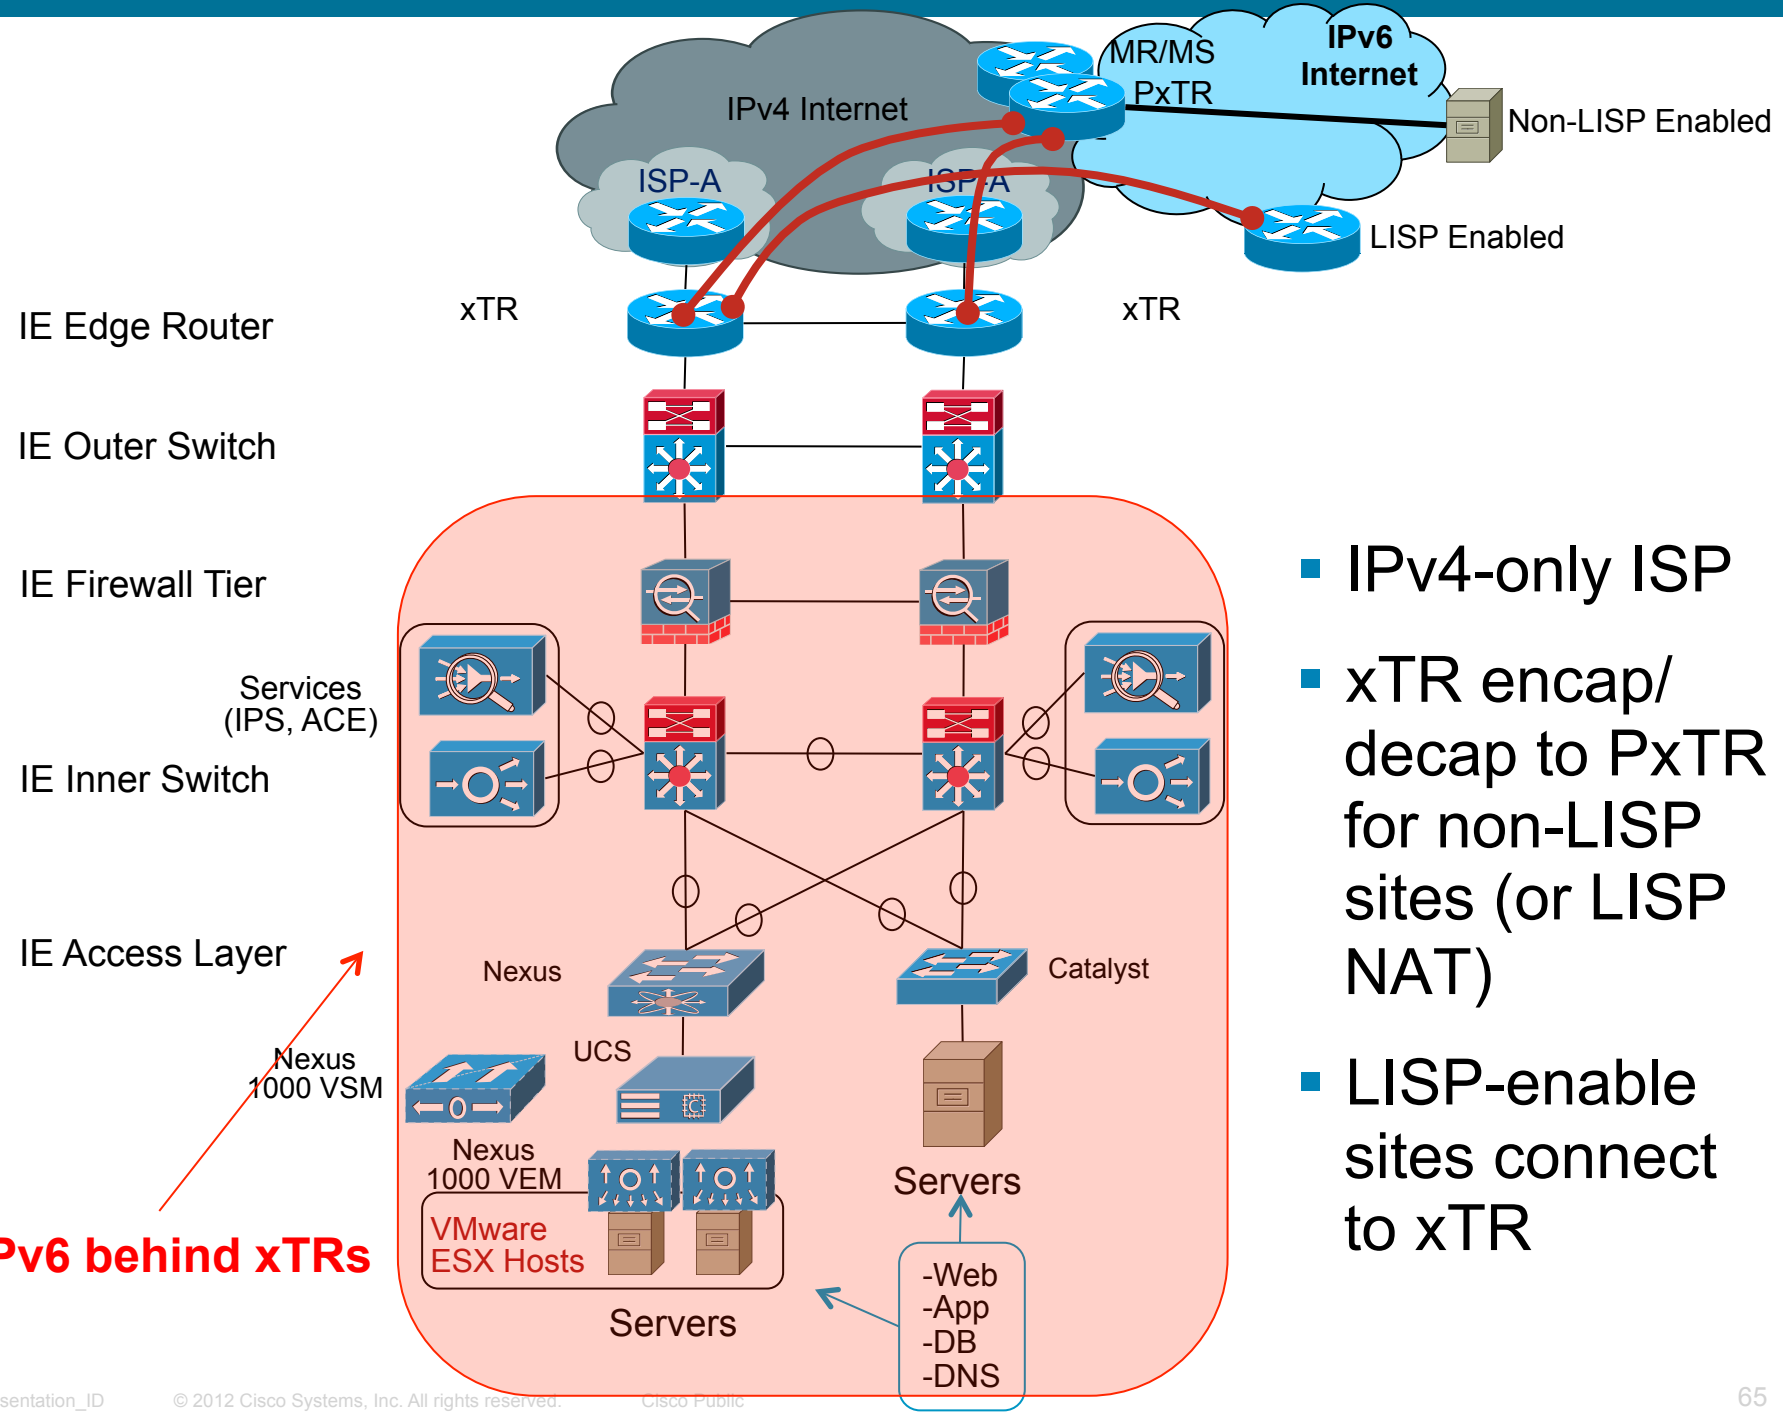


- IPv4-only ISP
- xTR encap/ decap to PxTR for non-LISP sites (or LISP NAT)
- LISP-enable sites connect to xTR


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

Reference

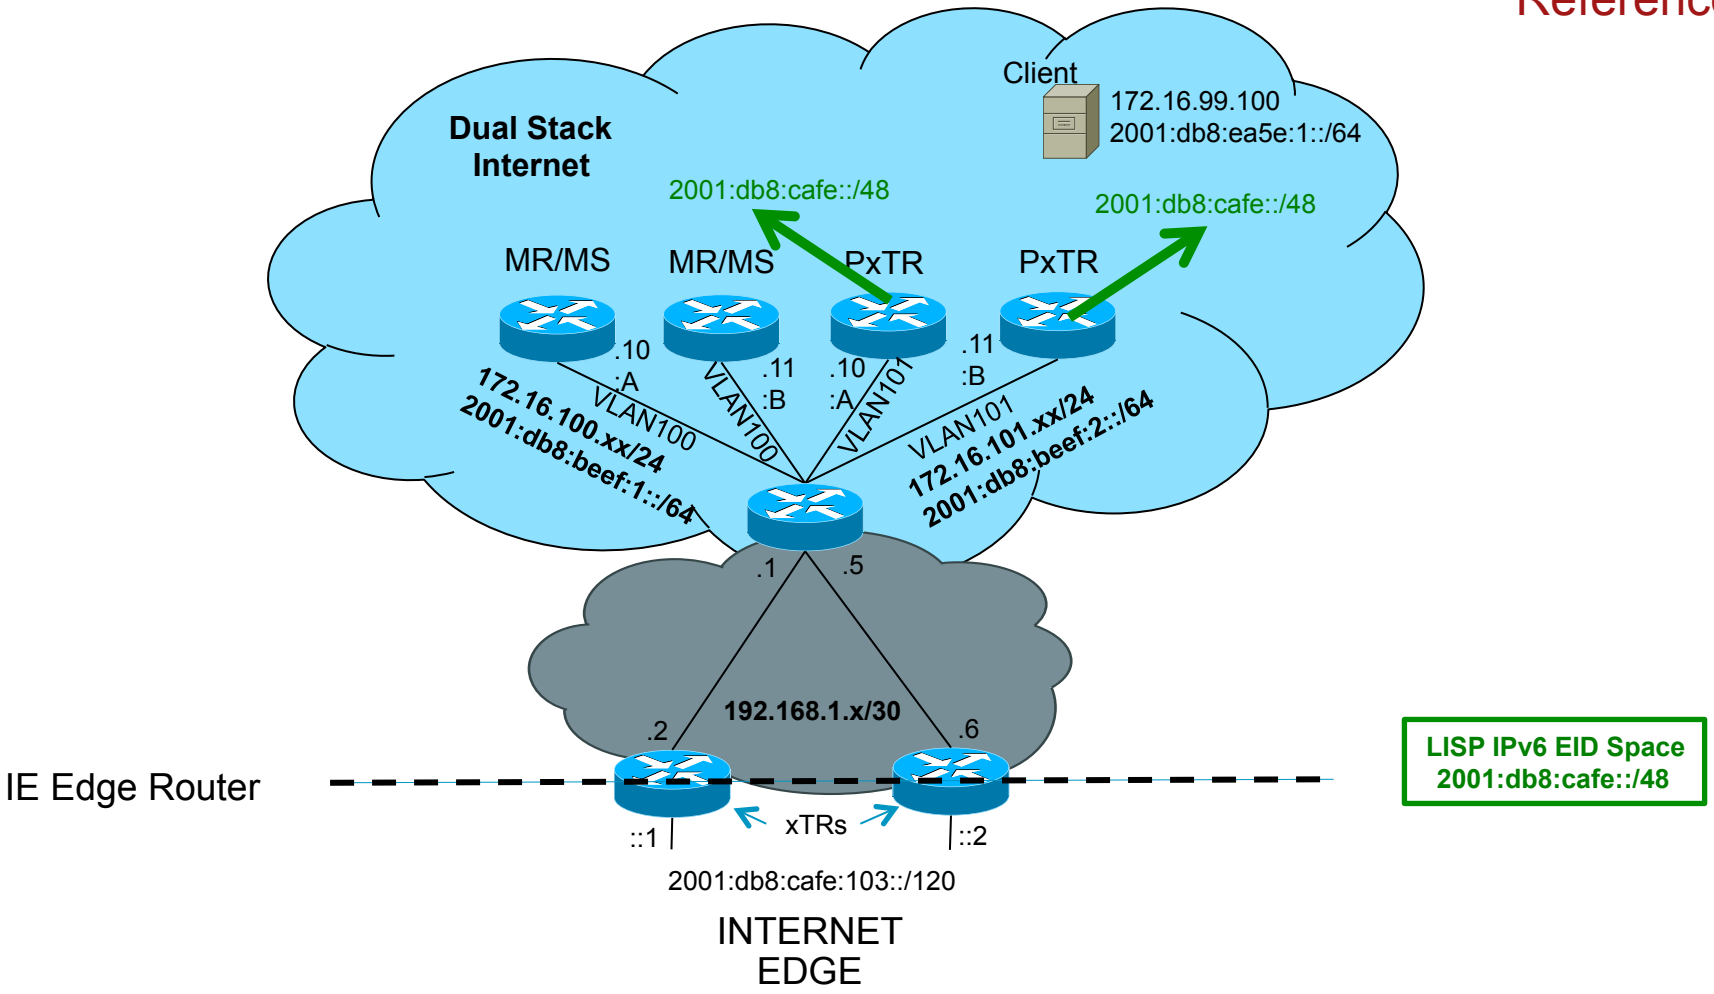


- Example addressing layout
- PxTR announces for 2001:db8:cafe::/48


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# xTR

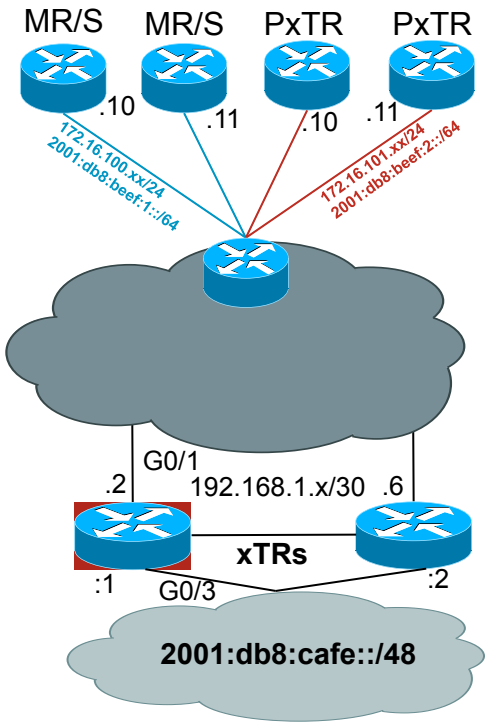


```
interface GigabitEthernet0/1
 description to ISPA (7604-1) - IPv4-ONLY
 ip address 192.168.1.2 255.255.255.252
!
interface GigabitEthernet0/3
 description to Enterprise Internet Edge IPv4/IPv6
 ip address 192.168.1.66 255.255.255.224
 ipv6 address 2001:DB8:CAFE:103::1/120
!
#BGP config excluded
!
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
```


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

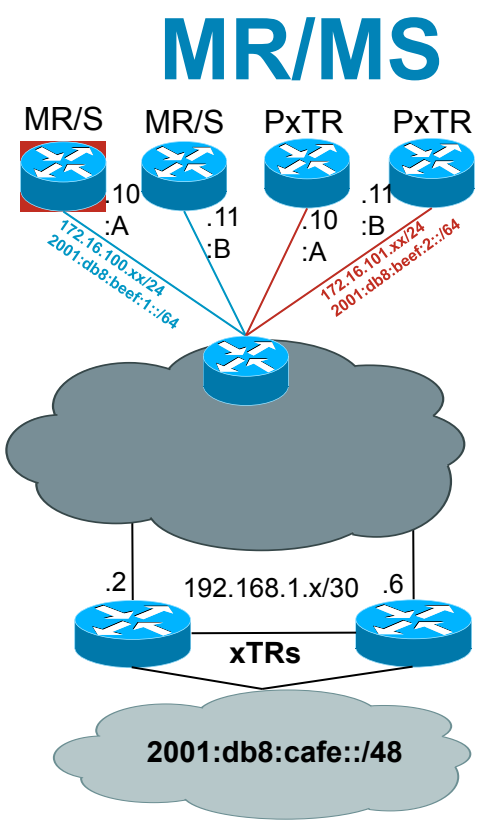


Reference

```
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


© 2012 Cisco Systems, Inc. All rights reserved.

# PxTR

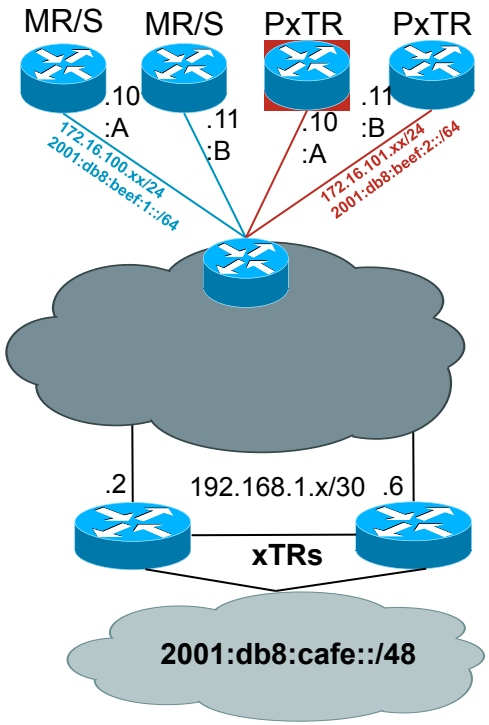


Reference

```
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


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Putting It All Together

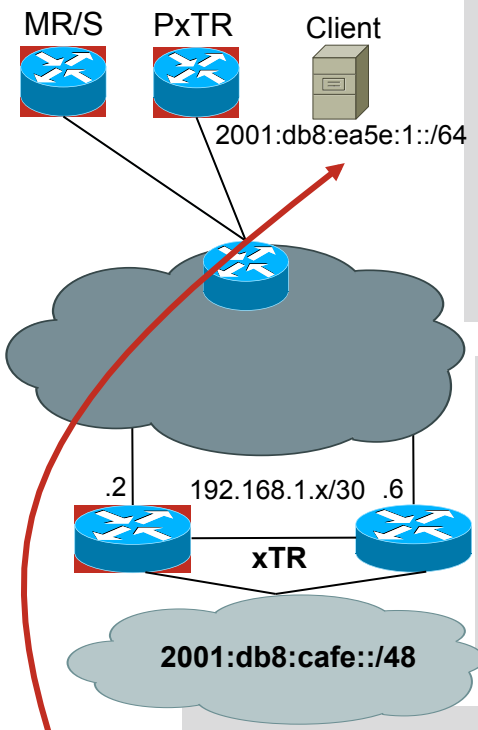


```
PxTR-1#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache for EID-table default (IID 0), 1 entries

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
```

Aggregate map-cache


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Summary


© 2012 Cisco Systems, Inc. All rights reserved.

# Areas of IPv6 Deployment in the Internet Edge – Stuff we talked about

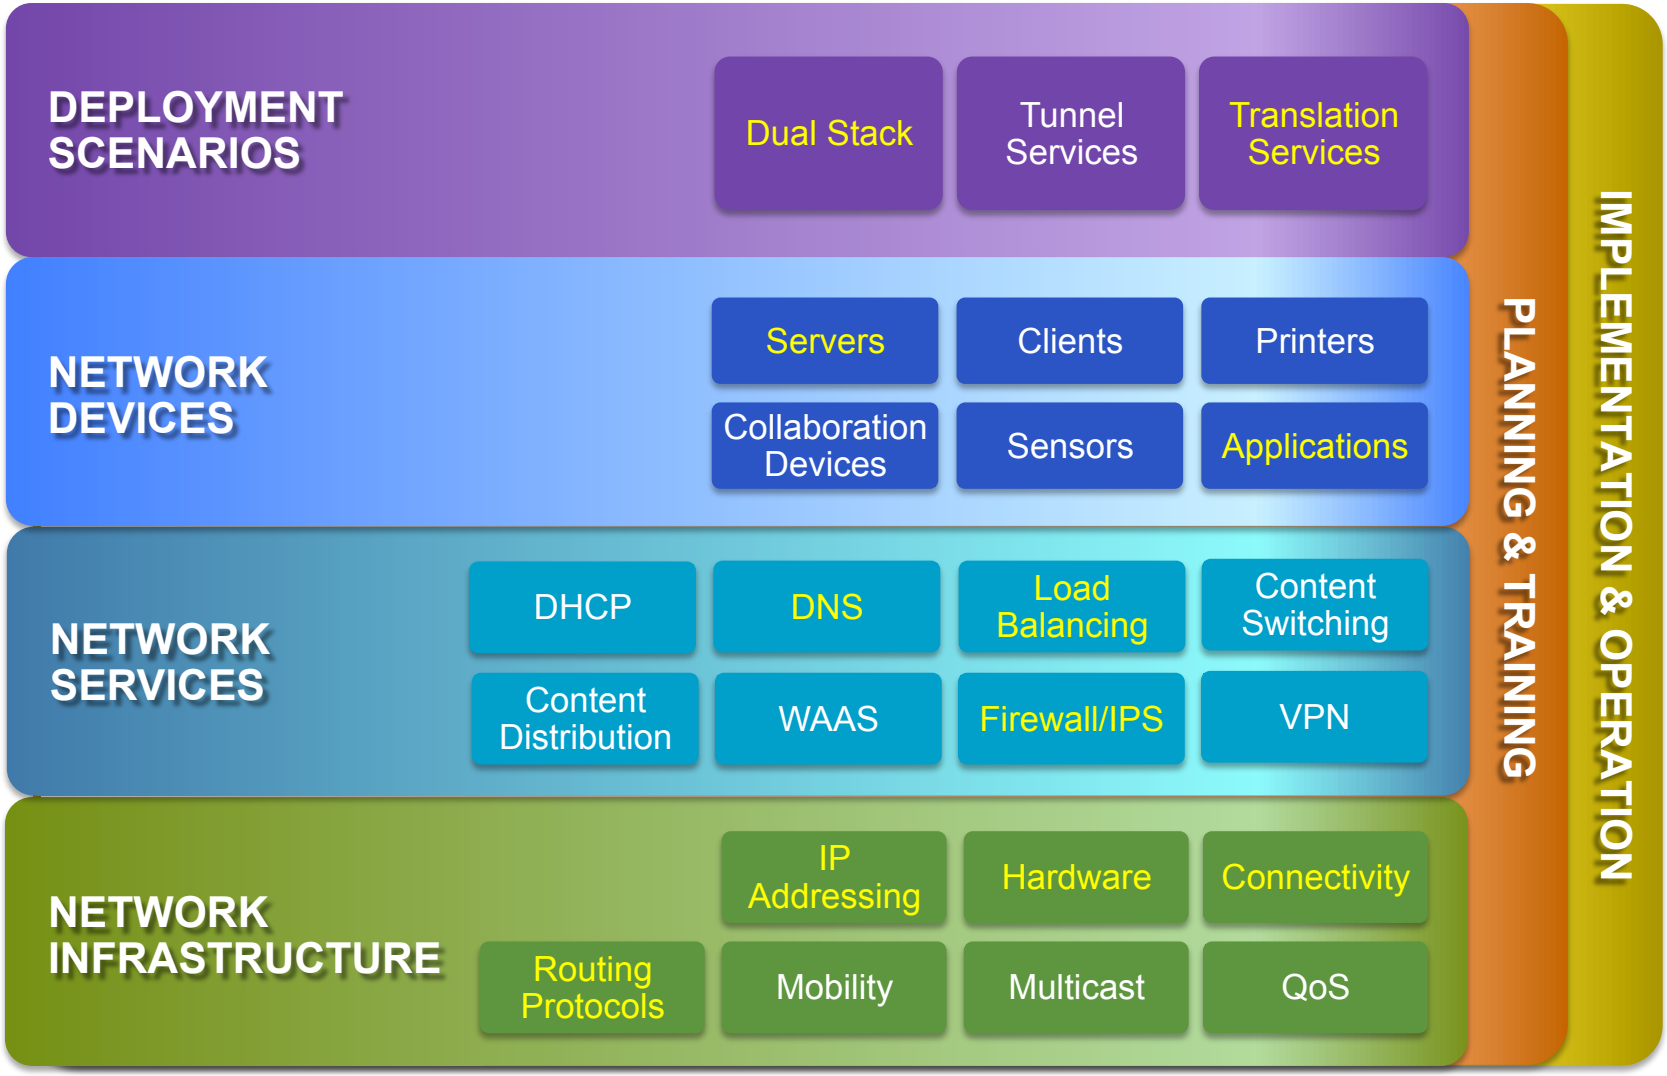



© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# Other Stuff

- Network Management – You will manage the same kind of stuff regardless of protocol
- NetFlow, Deep Packet Inspection, etc..
- Email, DNS, other apps
- More comprehensive security recommendations
  - Blocking routing type 0
  - 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
- NPTv6 for single address space multi-homing configurations
  - http://tools.ietf.org/html/draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat


© 2012 Cisco Systems, Inc. All rights reserved. Cisco Public

# 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


© 2012 Cisco Systems, Inc. All rights reserved.