What You Make Possible

TOMORROW
starts here.

Highly Available Wide Area Network Design

BRKRST-2042

TOMORROW starts here.

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
- Design and Deployment
- Final Wrap Up

Goals

- Design a WAN to efficiently utilise available bandwidth
- Design a WAN to dynamically respond to all types of disruptions
- Leverage most effective design techniques that meet the design requirements

Where Can Outages Occur?

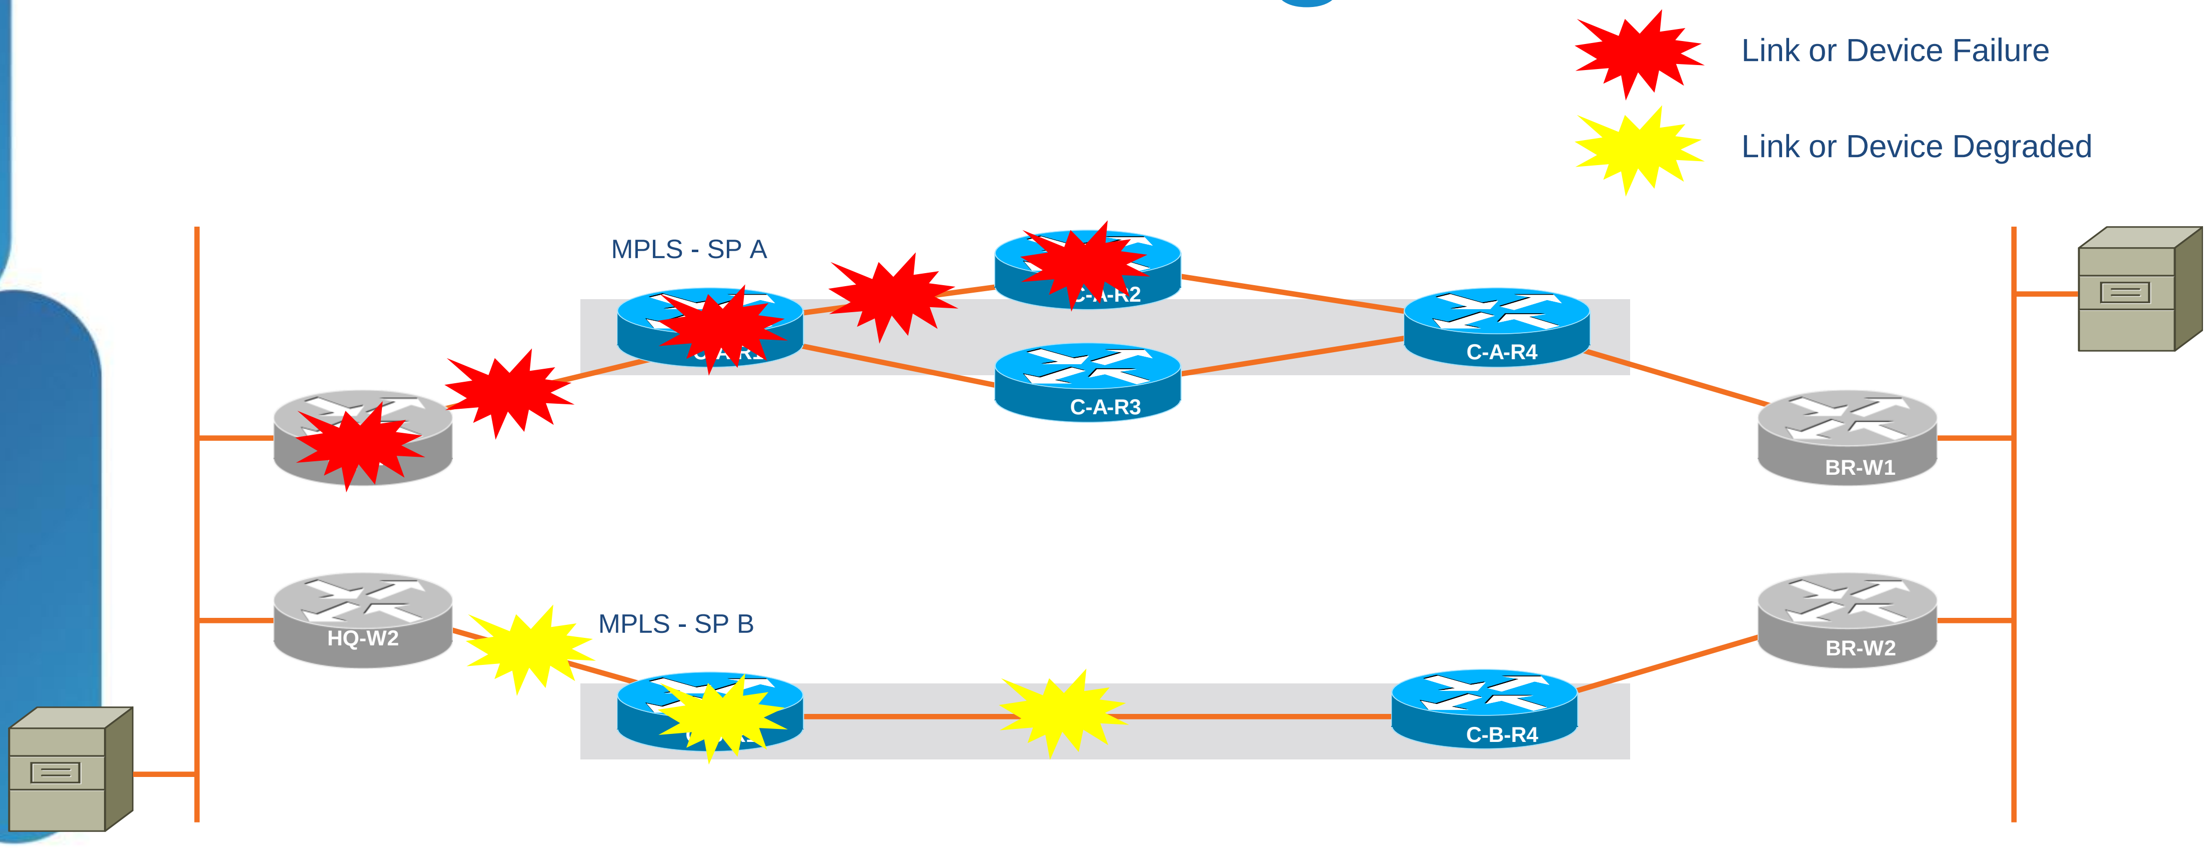

- How does outage manifest?
- How quickly can network detect?
- How long is bidirectional reconvergence?

Session Scope

- What methods are used for path selection and packet forwarding
- How does the network detect outages
- Focus on network survivability and effective utilisation rather than sub-second convergence
- Does not address “zero loss” considerations

Agenda

- Introduction
- Cisco IOS and IP Routing
  - Multiple Links/Multiple Paths
  - Load Sharing
- Convergence Techniques
- Design and Deployment
- Final Wrap Up

Routing Table Basics

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
    D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
    N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
    E1 - OSPF external type 1, E2 - OSPF external type 2
    i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
    * - candidate default, U - per-user static route, o - ODR
    P - periodic downloaded static route
```

```
D*EX  0.0.0.0/0 [170/3328] via 10.4.128.1, 21:44:37, Port-channel1
      10.0.0.0/8 is variably subnetted, 27 subnets, 6 masks
C        10.4.128.0/30 is directly connected, Port-channel1
D        10.4.128.8/30 [90/1792] via 10.4.128.1, 21:44:37, Port-channel1
D        10.4.128.128/26 [90/3072] via 10.4.128.1, 21:44:37, Port-channel1
D        10.4.128.240/32 [90/129536] via 10.4.128.1, 21:44:37, Port-channel1
C        10.4.128.241/32 is directly connected, Loopback0
D        10.4.128.244/32 [90/129792] via 10.4.128.1, 21:44:37, Port-channel1
C        10.4.142.0/29 is directly connected, GigabitEthernet0/0/4
B        10.4.142.32/30 [20/0] via 10.4.142.2, 21:44:01
B        10.4.142.144/30 [20/0] via 10.4.142.2, 21:44:01
B        10.4.143.0/29 [200/0] via 10.4.128.242, 21:44:01
```

INFORMATIONAL

Administrative Distance

- The **distance** command is used to configure a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers
- Numerically, an administrative distance is a positive integer from 1 to 255. In general, the higher the value, the lower the trust rating
- An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored

Route Source	Default Distance
Connected Interface	0
Static Route	1
EIGRP Summary Route	5
BGP external (eBGP)	20
EIGRP internal	90
OSPF	110
IS-IS	115
RIP	120
EIGRP External	170
BGP Internal (iBGP)	200
Unknown	255

Route Selection

- How is administrative distance used to
  determine which route should be installed?
- Only identical routes are compared
  - –Identical prefixes with different prefix lengths are not the same route
- The route from the protocol with the lower administrative distance is installed

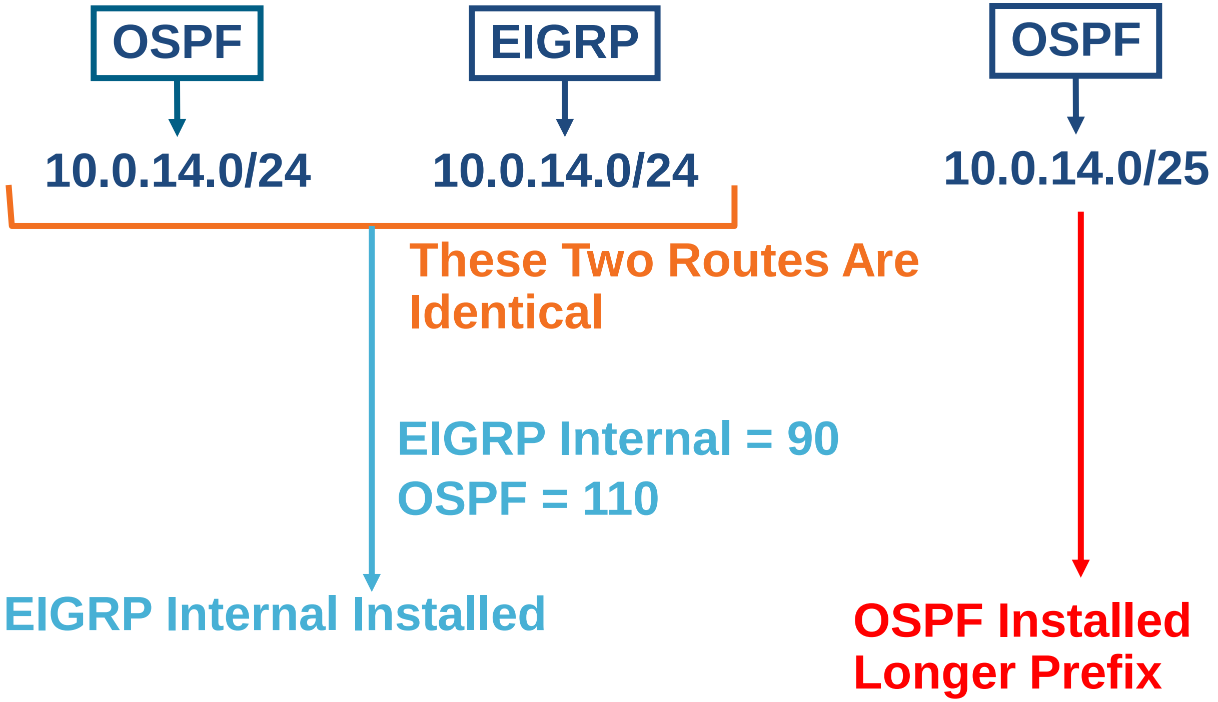

These Two Routes Are Identical

EIGRP Internal = 90
OSPF = 110

EIGRP Internal Installed

OSPF Installed
Longer Prefix

```
router#show ip route 10.0.14.0 255.255.255.0 longer-prefixes
     10.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
O IA      10.0.14.0/25 [110/40] via 10.0.67.6, 00:02:02, Ethernet0/1
D         10.0.14.0/24 [90/358400] via 10.0.37.3, 2d12h, Ethernet0/0
O IA      10.0.14.128/25 [110/40] via 10.0.67.6, 00:02:02, Ethernet0/1
```

More Specific OSPF Override EIGRP

Agenda

- Introduction
- Cisco IOS and IP Routing
  - Multiple Links/Multiple Paths
  - Load Sharing
- Convergence Techniques
- Design and Deployment
- Final Wrap Up

Load Sharing

- Assume the same routing process attempts to install two routes for the same destination in the RIB
- The routing process may allow the second route to be installed based on its own rules

	OSPF	IS-IS	EIGRP
Route Cost	Must be equal to installed route	Must be equal to installed route	Must be less than the variance times the lowest cost installed route
Maximum Paths	Must be fewer than *maximum-paths* configured under the routing process (default = 4)		

Note: BGP default value for maximum-paths = 1

CEF Load Sharing

Per-Session	Per-Packet
Default behaviour of IOS	Require “ip cef load-sharing per-packet”
Per-flow using source/destination	Per-packet using round-robin method
Packets for a given source/ destination session will take the same path	Packets for a given source/ destination session may take different paths
More effective as the number of source to destination pairs increase	Ensures traffic is more evenly distributed over multiple paths
Ensures that traffic for a given session arrives in order	Potential for packets to arrive out of sequence

Load Sharing

```
router#show ip route 192.168.239.0
Routing entry for 192.168.239.0/24
  Known via "eigrp 100", distance 170, metric 3072256, type external
  Redistributing via eigrp 100
  Last update from 192.168.245.11 on Serial3/1, 00:18:17 ago
  Routing Descriptor Blocks:
  * 192.168.246.10, from 192.168.246.10, 00:18:17 ago, via Serial3/0
      Route metric is 3072256, traffic share count is 1
      ....
    192.168.245.11, from 192.168.245.11, 00:18:17 ago, via Serial3/1
      Route metric is 3072256, traffic share count is 1
       ...
```

The *Traffic Share Count* Is Critical to Understanding the Actual Load Sharing of Packets Using These Two Routes

3072256/3072256 = 1

Load Sharing

```
router#show ip route 192.168.239.0
Routing entry for 192.168.239.0/24
  Known via "eigrp 100", distance 170, metric 3072256, type external
  Redistributing via eigrp 100
  Last update from 192.168.245.11 on Serial3/1, 00:18:17 ago
  Routing Descriptor Blocks:
  * 192.168.246.10, from 192.168.246.10, 00:18:17 ago, via Serial3/0
      Route metric is 1536128, traffic share count is 2
      ....
    192.168.245.11, from 192.168.245.11, 00:18:17 ago, via Serial3/1
      Route metric is 3072256, traffic share count is 1
       ...
```

If the Lower Metric Is Less than the Second Metric, the Traffic Share Count Will Be Something Other than 1 (Only for EIGRP and Requires Variance to Be Configured)

3072256/3072256 = 1

3072256/1536128 = 2

2x Faster Link Gets 2 Flows vs. 1 Flow

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
  - First Hop Redundancy Protocols
  - Routing Protocols
  - DDR and Static Routing
  - Performance Routing
- Design and Deployment
- Final Wrap Up

First Hop Redundancy Protocols (FHRP)

Failure Protection for the First Hop IP Router

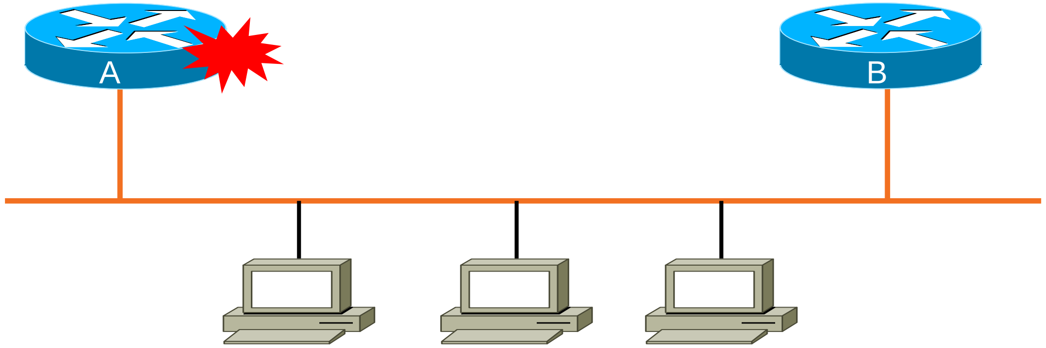

- Hot Standby Router Protocol (HSRP)
- Virtual Router Redundancy Protocol (VRRP)
- Gateway Load Balancing Protocol (GLBP)

Drivers for FHRPs

- Provide routing redundancy for access layer
  - How to handle failover when end-hosts have only a single IP default gateway and cached ARP entry
- Provide routing redundancy for devices that depend on static routing
  - Some firewalls do not support dynamic routing
- Independent of routing protocols
  - Works with any routing protocol and static routing
- Capable of providing sub-second failover
- Provides load sharing capabilities (GLBP) transparent to end host

Hot Standby Routing Protocol (HSRP)

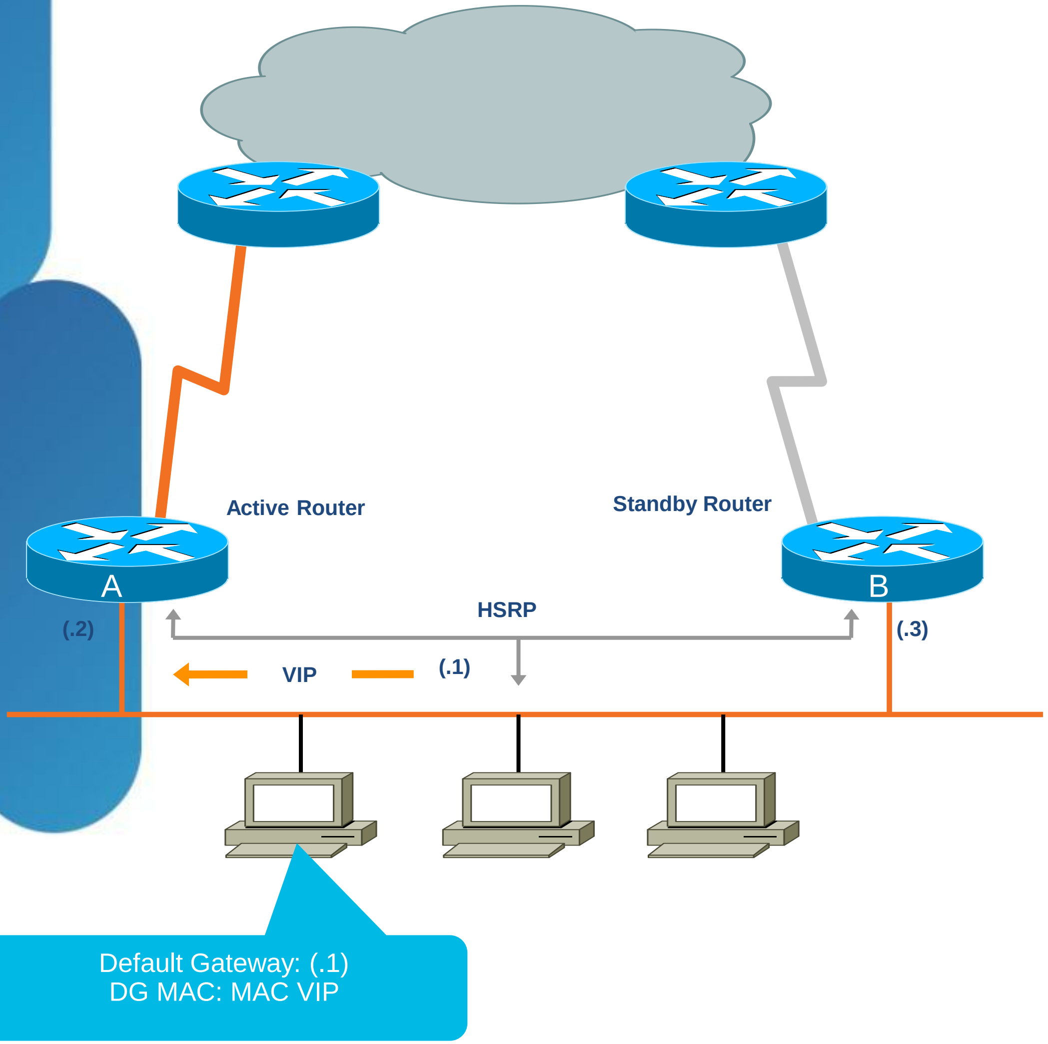

```
Router A#
interface FastEthernet0/0
 ip address 10.1.2.2 255.255.255.0
 standby 1 priority 110
 standby 1 preempt
 standby 1 ip 10.1.2.1
```

```
Router A# show standby brief


Interface Grp Prio P State   Active     Standby       Virtual IP
Fa0/0     1  110  P Active  local       10.1.2.3     10.1.2.1
```

```
Router B#
interface FastEthernet0/0
 ip address 10.1.2.3 255.255.255.0
 standby 1 priority 105
 standby 1 preempt
 standby 1 ip 10.1.2.1
```

```
Router B# show standby brief


Interface Grp Prio P State    Active      Standby      Virtual IP
Fa0/0      1   105  P Standby  10.1.2.2    local        10.1.2.1
```

Hot Standby Routing Protocol (HSRP)

```
Router B# show standby brief


Interface Grp Prio P State    Active       Standby       Virtual IP
Fa0/0     1    105 P Active   local        unknown       10.1.2.1
```

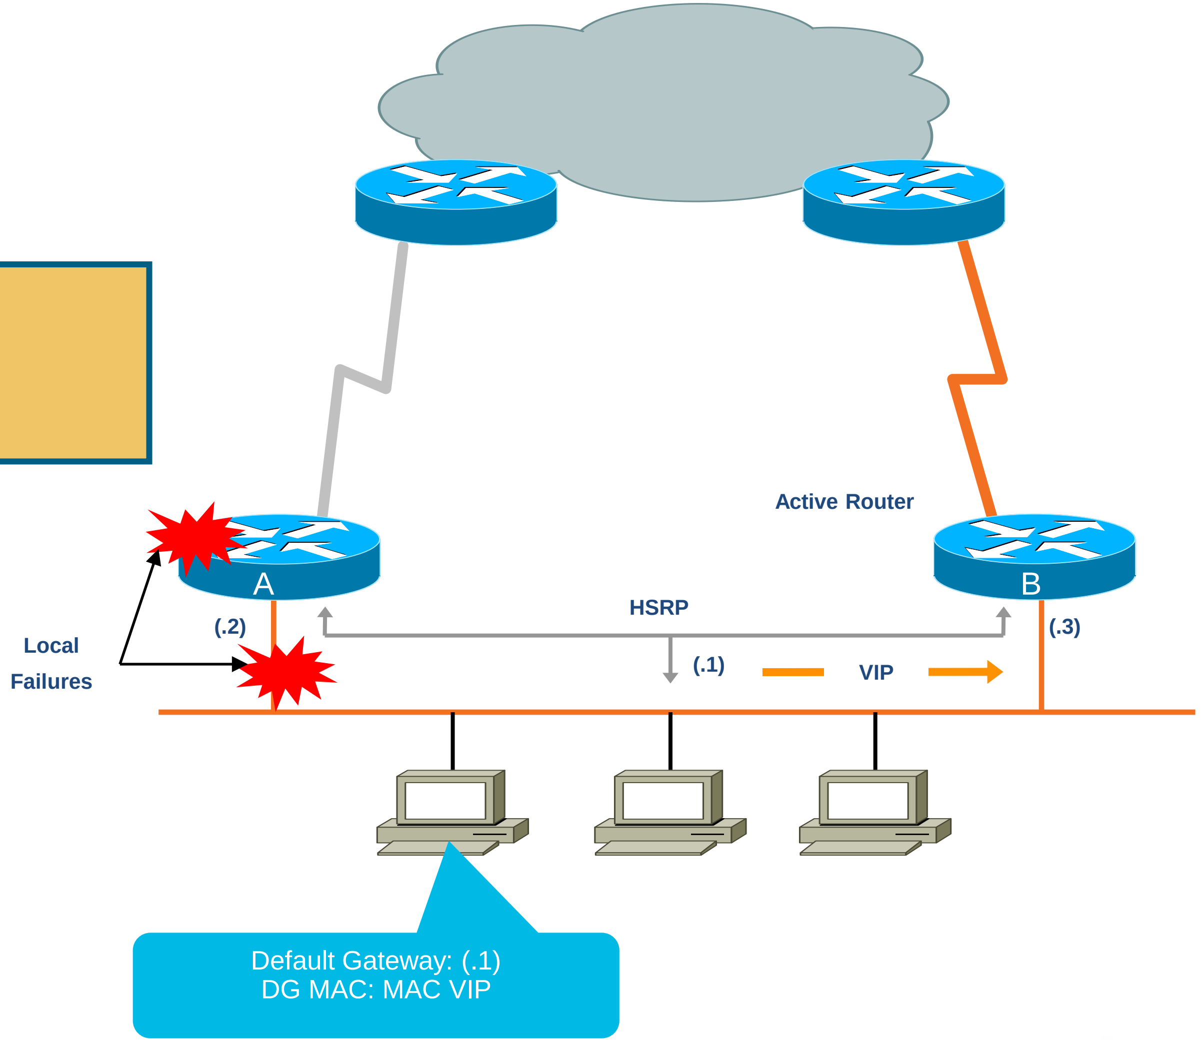

Default Gateway: (.1)
DG MAC: MAC VIP

Hot Standby Routing Protocol (HSRP)

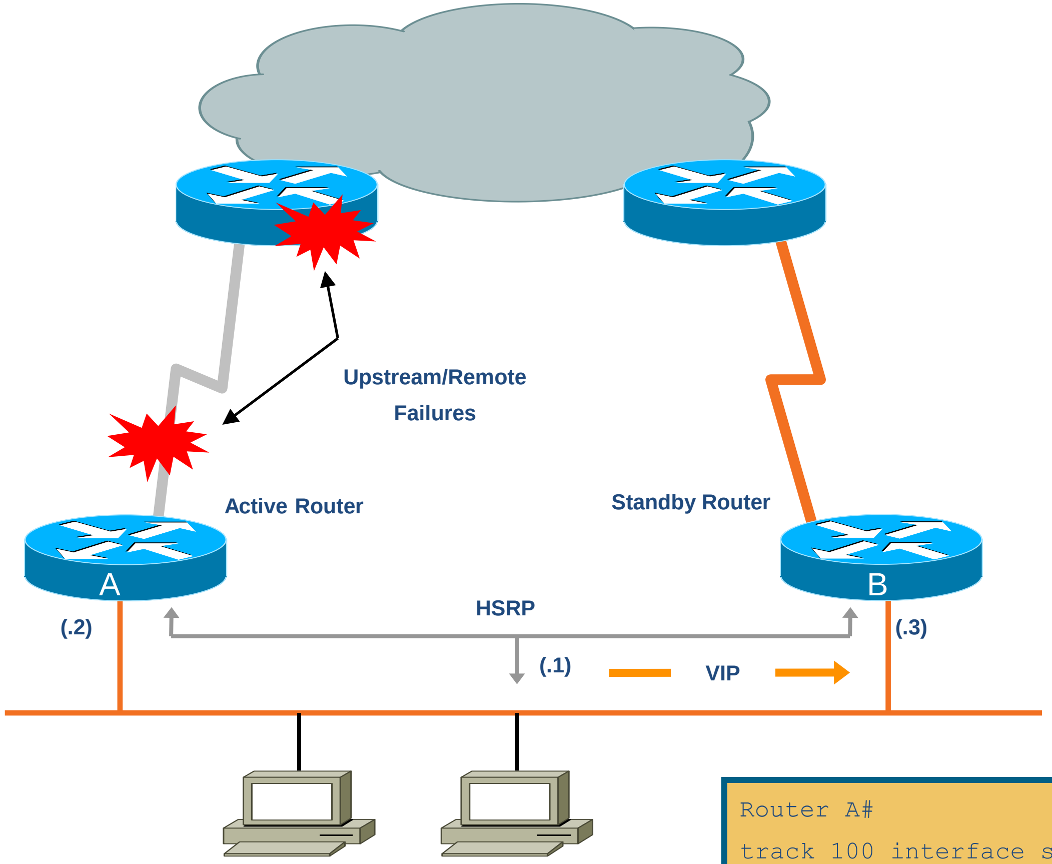

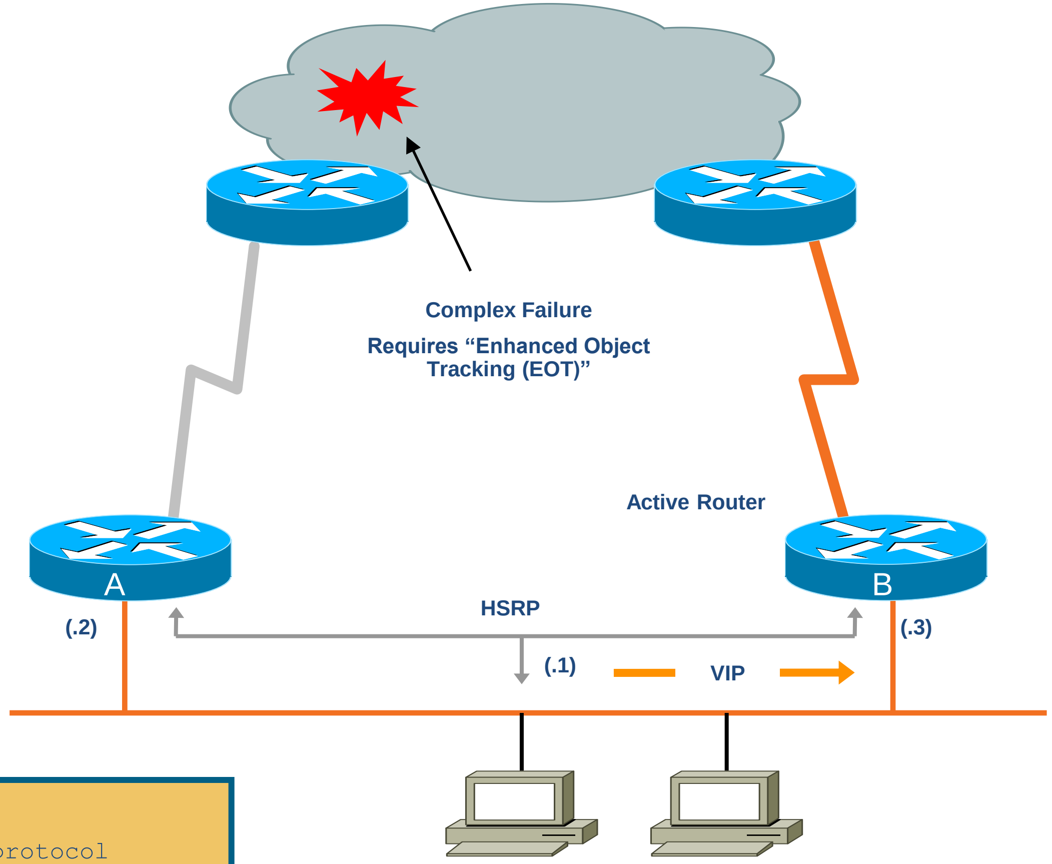

```
Router A#
track 100 interface serial1/0 line-protocol
!
interface FastEthernet0/0
 ip address 10.1.2.2 255.255.255.0
 standby 1 preempt
 standby 1 priority 110
 standby 1 track 100 decrement 10
 standby 1 ip 10.1.2.1
```

Gateway Load Balancing Protocol (GLBP)

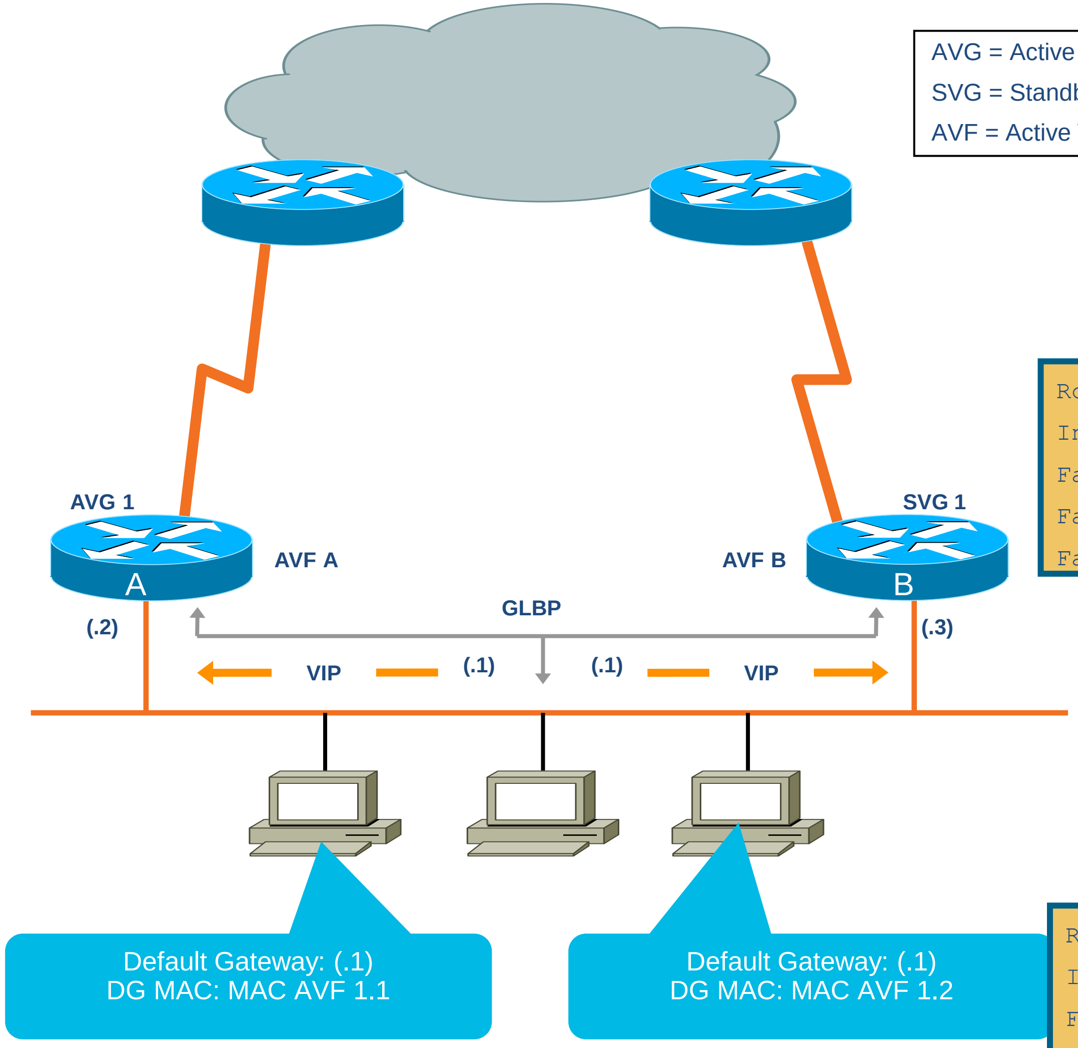

Default Gateway: (.1)
DG MAC: MAC AVF 1.1

Default Gateway: (.1)
DG MAC: MAC AVF 1.2

AVG = Active Virtual Gateway
SVG = Standby Virtual Gateway
AVF = Active Virtual Forwarder

```
Router A#
interface FastEthernet0/1
 ip address 10.1.1.2 255.255.255.0
 glbp 1 priority 110
 glbp 1 preempt
 glbp 1 ip 10.1.1.1
 glbp 1 load-balancing round-robin
```

```
Router A# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active Router  Standby Router
Fa0/1       1    -   110 Active   10.1.1.1        local          10.1.1.3
Fa0/1       1    1   -   Active   0007.b400.0101  local          -
Fa0/1       1    2   -   Listen   0007.b400.0102  10.1.2.3       -
```

```
Router B#
interface FastEthernet0/1
 ip address 10.1.1.3 255.255.255.0
 glbp 1 priority 105
 glbp 1 preempt
 glbp 1 ip 10.1.1.1
 glbp 1 load-balancing round-robin
```

```
Router B# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active Router  Standby Router
Fa0/1       1    -   105 Standby  10.1.1.1        10.1.1.2       local
Fa0/1       1    1   -   Listen   0007.b400.0101  10.1.1.2       -
Fa0/1       1    2   -   Active   0007.b400.0102  local          -
```

Gateway Load Balancing Protocol (GLBP)

AVG = Active Virtual Gateway
SVG = Standby Virtual Gateway
AVF = Active Virtual Forwarder

```
Router B#
*Mar 31 17:04:27: %GLBP-6-STATECHANGE: FastEth0/1 Grp 1 state Standby -> Active

*Mar 31 17:04:27  %GLBP-6-FWDSTATECHANGE: FastEth0/1 Grp 1 Fwd 1 state
Listen -> Active
```

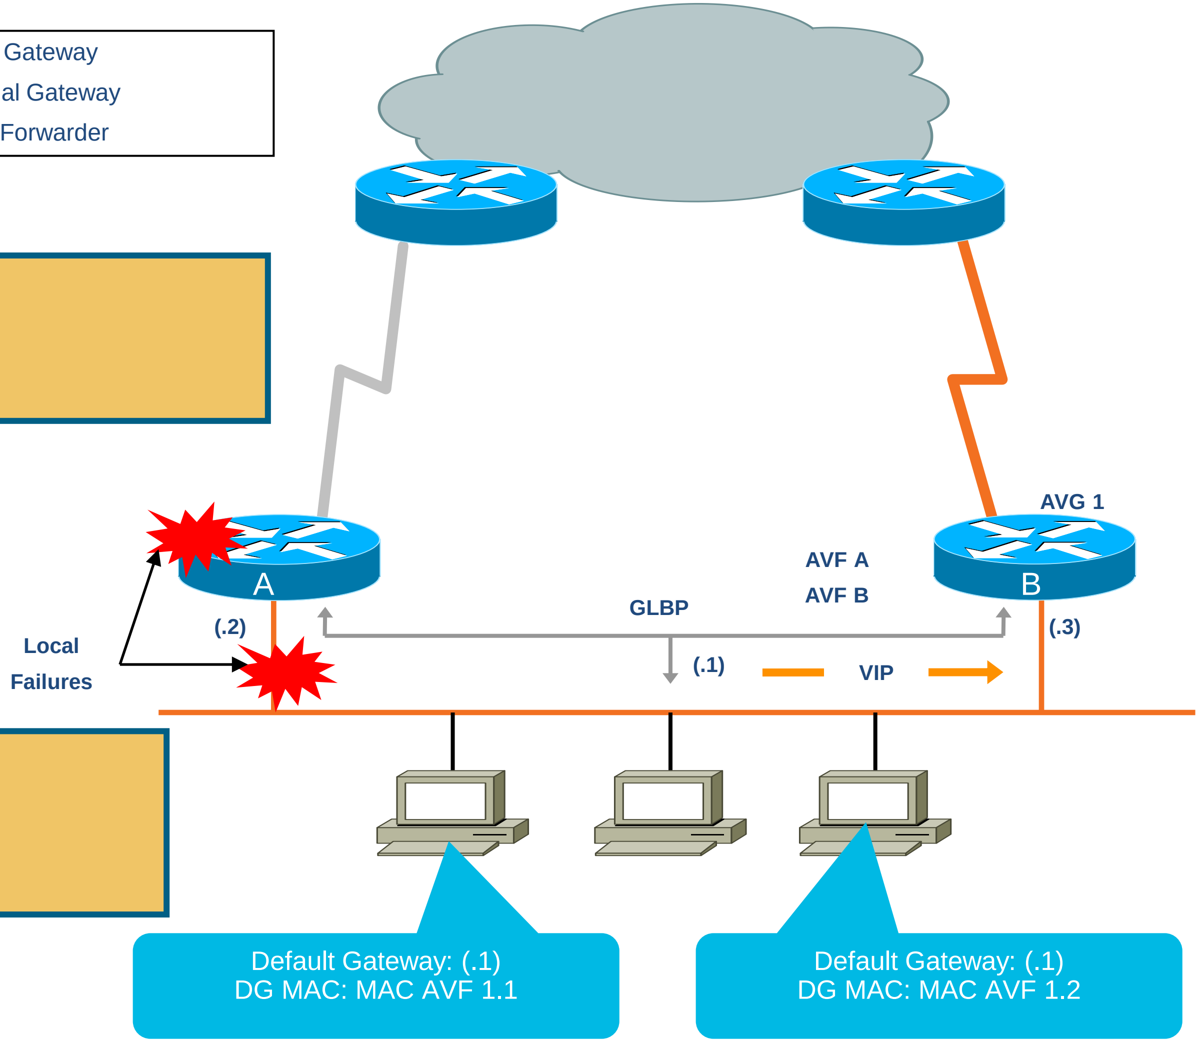

Default Gateway: (.1)
DG MAC: MAC AVF 1.1

Default Gateway: (.1)
DG MAC: MAC AVF 1.2

```
Router B# show glbp brief
Interface  Grp  Fwd Pri State    Address         Active Rtr    Standby Rtr
Fa0/1      1    -   105 Active   10.1.2.1        local            unknown
Fa0/1      1    1   -   Active   0007.b400.0101  local            -
Fa0/1      1    2   -   Active   0007.b400.0102  local            -
```

GLBP with Enhanced Object Tracking

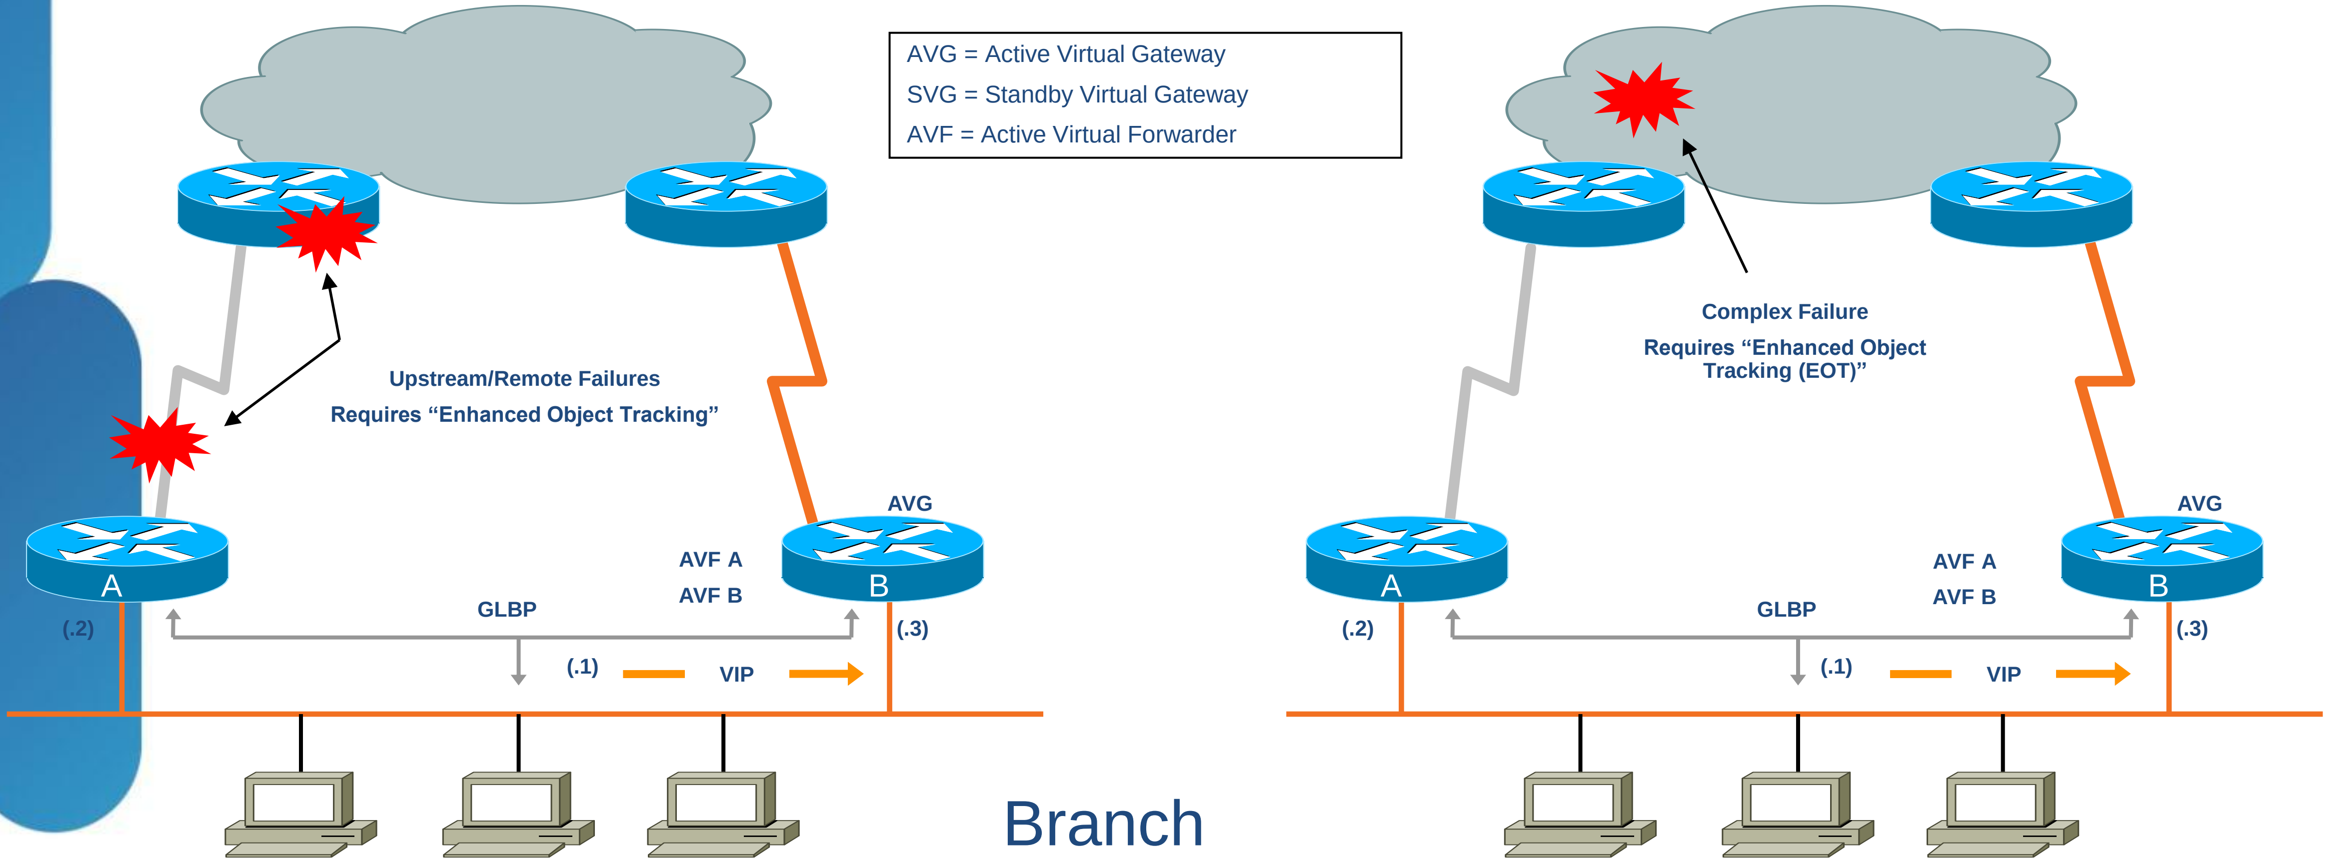

AVG = Active Virtual Gateway
SVG = Standby Virtual Gateway
AVF = Active Virtual Forwarder

Upstream/Remote Failures
Requires "Enhanced Object Tracking"

Complex Failure
Requires "Enhanced Object Tracking (EOT)"

Branch

Enhanced Object Tracking

Track Options	Syntax
Line-Protocol State of Interface	**track** *object-number* **interface** *type number* **line-protocol** track 1 interface serial 1/1 line-protocol
IP-Routing State of Interface	**track** *object-number* **interface** *type number* **ip routing** track 2 interface ethernet 1/0 ip routing
IP-Route Reachability	**track** *object-number* **ip route** *IP-Addr/Prefix-len* **reachability** track 3 ip route 10.16.0.0/16 reachability
Threshold* of IP-Route Metrics	**track** *object-number* **ip route** *IP-Addr/Prefix-len* **metric threshold** track 4 ip route 10.16.0.0/16 metric threshold

```
Router# show track 100
Track 100
Interface Serial1/1 line-protocol
Line protocol is Up
1 change, last change 00:00:05
Tracked by:
GLBP FastEthernet0/1 1
```

```
Router# show track 103
Track 103
IP route 10.16.0.0 255.255.0.0 reachability
Reachability is Up (RIP)
1 change, last change 00:02:04
First-hop interface is Ethernet0/1
Tracked by:
GLBP FastEthernet0/1 1
```

* EIGRP, OSPF, BGP, Static Thresholds Are Scaled to Range of (0 – 255)

Enhanced Object Tracking – IP SLA

Track Options	Syntax
IP SLAs Operation	**track** *object-number* **ip sla** *type number* **state** track 5 ip sla 4 state
Reachability of an IP SLAs Host	**track** *object-number* **ip sla** *type number* **reachability** track 6 ip sla 4 reachability

Types of IP SLA Probes:

dhcp
dns
ethernet
frame-relay
ftp
http
icmp-echo
icmp-jitter
mpls
path-echo
path-jitter
tcp-connect
udp-echo
udp-jitter
voip

IP SLA – UDP-Jitter Probe

```
ip sla operation-number
  type udp-jitter [hostname | ip-address] port-number [num-packets number-of-packets] [interval
inter-packet-interval]
  frequency seconds
  request-data-size bytes
```

UDP Jitter Operation Parameter	Default Value
Number of Packets (N)	10 Packets
Payload Size per Packet (S)	`32 Bytes`
Time Between Packets, in Milliseconds (T)	20 ms
Elapsed Time Before the Operation Repeats, in Seconds (F)	60 Seconds

Enhanced Object Tracking (EOT)

Tracking IP SLA

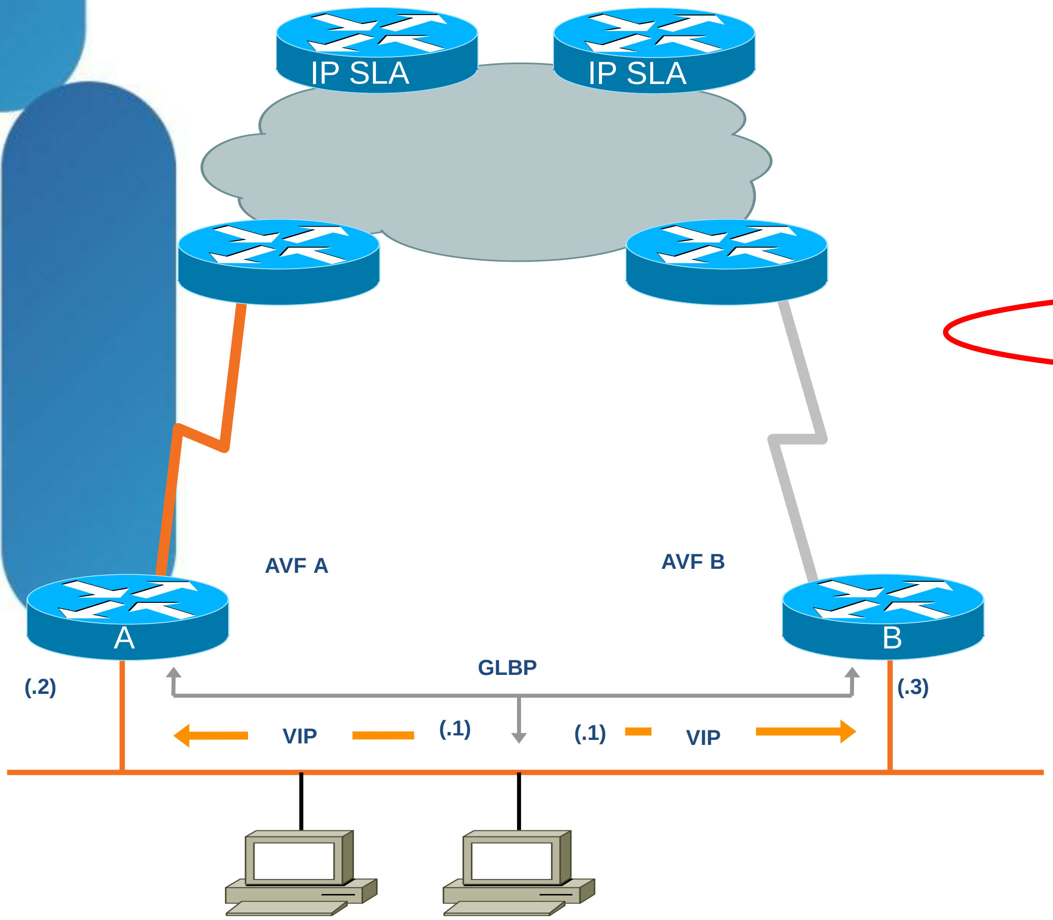

```
RouterA#

ip sla 100
 icmp-echo 10.100.100.100 source-ip 10.1.2.2
 timeout 100
 frequency 10
ip sla schedule 100 life forever start-time now
!
ip sla 200
 icmp-echo 10.100.200.100 source-ip 10.1.2.2
 timeout 100
 frequency 10
ip sla schedule 200 life forever start-time now

ip route 10.100.100.100 255.255.255.255 192.168.101.9
ip route 10.100.200.100 255.255.255.255 192.168.101.9
```

```
RouterA# show ip sla statistics

IPSLA operation id: 100
    Latest RTT: 1 milliseconds
Latest operation start time: *04:42:11.444 UTC Tue Feb 17 2009
Latest operation return code: OK
Number of successes: 46
Number of failures: 0
Operation time to live: Forever

IPSLA operation id: 200
    Latest RTT: 1 milliseconds
Latest operation start time: *04:42:11.356 UTC Tue Feb 17 2009
Latest operation return code: OK
Number of successes: 24
Number of failures: 0
Operation time to live: Forever
```

Enhanced Object Tracking

Tracking IP SLA

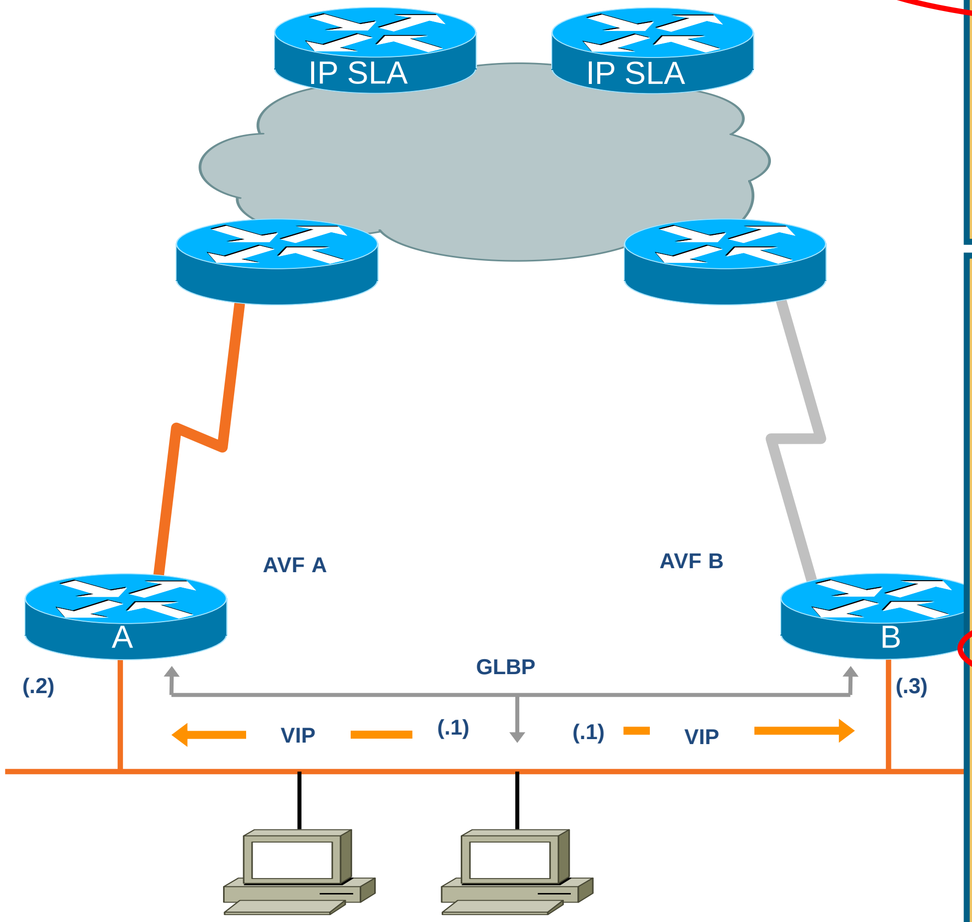

```
RouterA#
track 100 ip sla 100 reachability
track 200 ip sla 200 reachability
track 1 list boolean or
 object 100
 object 200
interface FastEthernet0/1
 ip address 10.1.2.2 255.255.255.0
 glbp 1 ip 10.1.2.1
 glbp 1 priority 110
 glbp 1 preempt
 glbp 1 weighting 120 lower 100
 glbp 1 load-balancing weighted
 glbp 1 weighting track 1 decrement 10
```

```
RouterA# show glbp
FastEthernet0/1 - Group 1
  State is Active
    1 state change, last state change 00:09:59
  Virtual IP address is 10.1.2.1
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.336 secs
  Redirect time 600 sec, forwarder timeout 14400 sec
  Preemption enabled, min delay 0 sec
  Active is local
  Standby is 10.1.2.3, priority 105 (expires in 7.808 sec)
  Priority 110 (configured)
  Weighting 120 (configured 120), thresholds: lower 100, upper 120
    Track object 1 state Up decrement 30
  Load balancing: weighted
  Group members:
    aabb.cc00.0110 (10.1.2.2) local
    aabb.cc00.0410 (10.1.2.3)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
    <SNIP>
  Forwarder 2
    State is Listen
    <SNIP>
```

Enhanced Object Tracking

Composite Failure

```
RouterA#
*Feb 17 05:17:25: %TRACKING-5-STATE: 100 ip sla 100 state Up->Down
*Feb 17 05:17:25: %TRACKING-5-STATE: 200 ip sla 200 state Up->Down
*Feb 17 05:17:26: %TRACKING-5-STATE: 1 list boolean or Up->Down
*Feb 17 05:17:26: %GLBP-6-STATECHANGE: FastEth0/1 Grp 1 state Standby -> Active
*Feb 17 05:17:38: %GLBP-6-FWDSTATECHANGE: FastEth0/1 Grp 1 Fwd 1 state Listen -> Active
```

```
RouterB#show glbp
FastEthernet0/1 - Group 1
  State is Standby
    1 state change, last state change 00:28:16
  Virtual IP address is 10.1.2.1
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.856 secs
  Redirect time 600 sec, forwarder timeout 14400 sec
  Preemption enabled, min delay 0 sec
  Active is 10.1.2.2, priority 110 (expires in 10.400 sec)
  Standby is local
  Priority 105 (configured)
  Weighting 120 (configured 120), thresholds: lower 100, upper 120
    Track object 1 state Up decrement 30
  Load balancing: weighted
  Group members:
    aabb.cc00.0110 (10.1.2.2)
    aabb.cc00.0410 (10.1.2.3) local
  There are 2 forwarders (2 active)
  Forwarder 1
    State is Active
    <SNIP>
  Forwarder 2
    State is Active
    <SNIP>
```

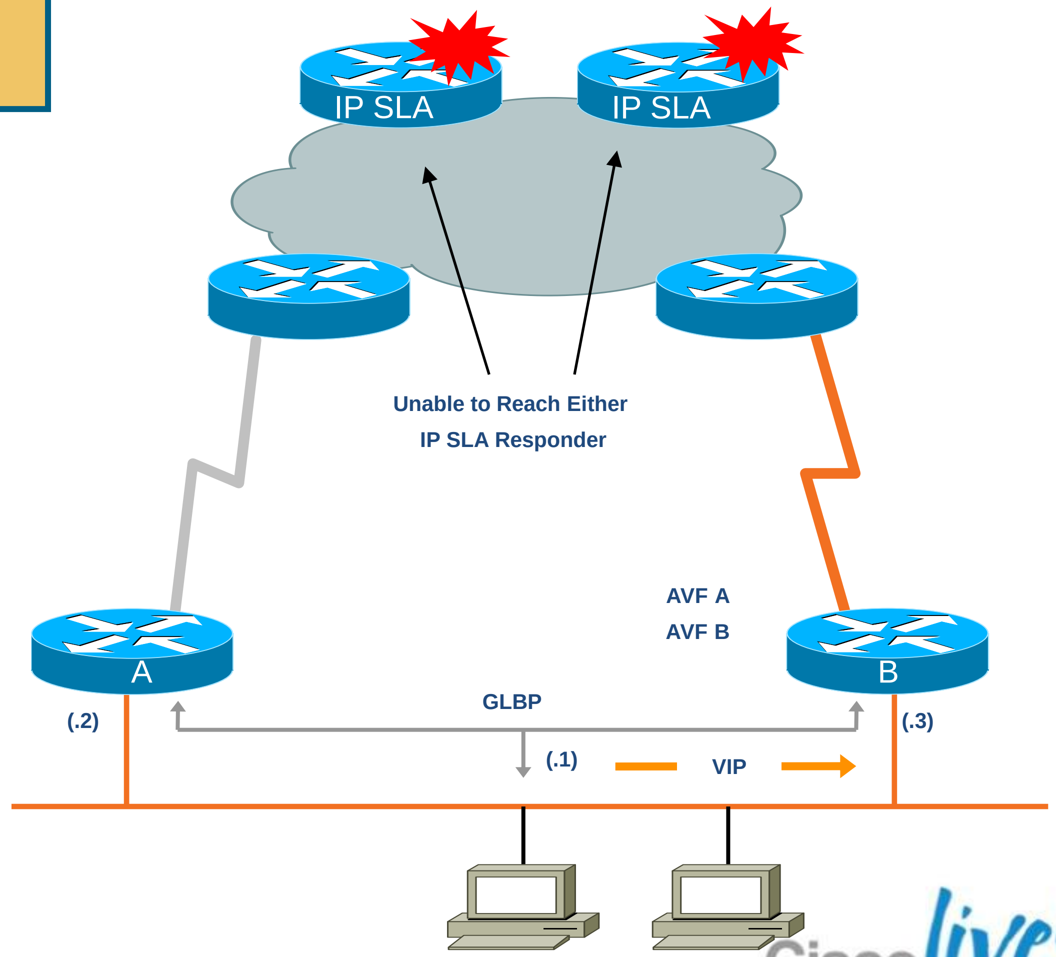

Unable to Reach Either IP SLA Responder

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
  - First Hop Redundancy Protocols
  - Routing Protocols
  - DDR and Static Routing
  - Performance Routing
- Design and Deployment
- Final Wrap Up

Routing Protocol Timers

INFORMATIONAL

	Keepalive (B) Hello (E,I,O) Update (R)	Invalid (R)	Holdtime (B,E,I) Dead (O) Holddown (R)	Flush (R)
BGP	60		180	
EIGRP (< T1)	5 (60)		15 (180)	
IS-IS (DIS)	10 (3.333)		30 (10)	
OSPF (NBMA)	10 (30)		40 (120)	
RIP/RIPv2	30	180	180	240

Note: Cisco Default Values

Routing Protocol Neighbour Behaviour

INFORMATIONAL

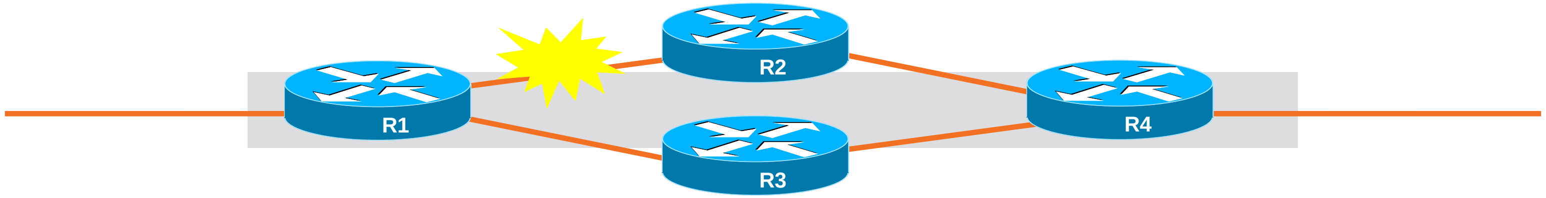

Recovery Times by Protocol

Note: Using Cisco Default Values

	Link Down Line Protocol Down	Link Up Loss 100%	Link Up Neighbour Down	Link Up Loss ~5%
BGP	~ 1 s	180	180	Never
EIGRP (< T1)	~ 1s	15 (180)	15 (180)	Never
IS-IS (DIS)	~ 1s	30 (10)	30 (10)	Never
OSPF (NBMA)	~ 1s	40 (120)	40 (120)	Never
RIP/RIPv2	~ 1s	240	240	Never

Routing Protocol Neighbour Behaviour

Adjust Hello Timers

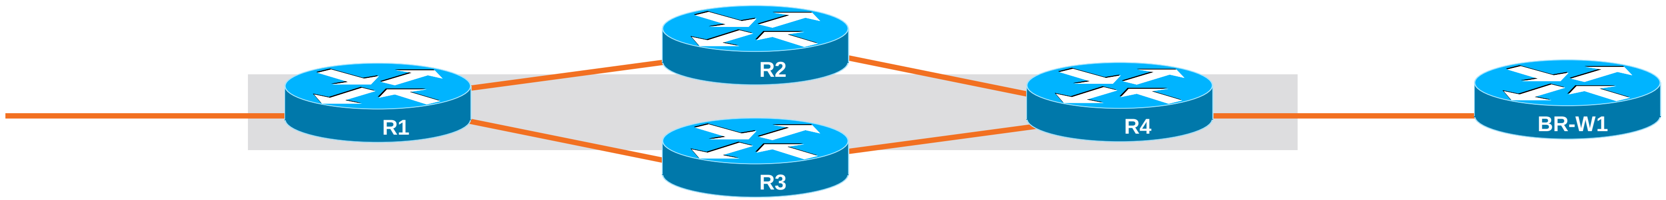

```
R4# show ip bgp vpnv4 vrf cisco neighbor
BGP neighbor is 192.168.101.10, vrf cisco, remote AS 65110, external link
 BGP version 4, remote router ID 192.168.201.10
 BGP state = Established, up for 1d10h
 Last read 00:00:19, hold time is 180, keepalive interval is 60 seconds
```

```
BR-W1#
router bgp 65110
 neighbor 192.168.101.9 timers 7 21
```

```
R4# show ip bgp vpnv4 vrf cisco neighbor
BGP neighbor is 192.168.101.10, vrf cisco, remote AS 65110, external link
 BGP version 4, remote router ID 192.168.201.10
 BGP state = Established, up for 00:01:23
 Last read 00:00:03, hold time is 21, keepalive interval is 7 seconds
```

When Configuring the *Holdtime* Argument for a Value of Less than Twenty Seconds, the Following Warning Is Displayed:

% Warning: A Hold Time of Less than 20 Seconds Increases the Chances of Peer Flapping

Introducing BFD

Bi-Directional Forwarding Detection:

- Extremely lightweight hello protocol
  - IPv4, IPv6, MPLS, P2MP
- 10s of milliseconds (technically, microsecond resolution) forwarding plane failure detection mechanism.
- Single mechanism, common and standardised
  - Multiple modes: Async (echo/non-echo), Demand
- Independent of Routing Protocols
- Levels of security, to match conditions and needs
- Facilitates close alignment with hardware

Drivers for BFD

- Link-layer detection misses some types of outages
  - e.g. Control Plane failure
- Control Plane failure detection is very conservative
  - 15-40 seconds in default configurations
- Link-layer failure detection is not consistent across media types
  - Less than 50ms on APS- protected SONET
  - A few seconds on Ethernet
  - Several seconds or more on WAN links
- Provides a measure of consistency across routing protocols
- Most current failure detection mechanisms are an order of magnitude too long for time-sensitive applications

Routing Protocol Neighbour Behaviour

Bi-Directional Forwarding Detection

```
R1#
router eigrp 65110
 network 172.16.1.0 0.0.0.255
 bfd all-interfaces
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 3
```

```
R1# show bfd neighbors detail

OurAddr    NeighAddr   LD/RD RH/RS Holddown(mult) State   Int
172.16.1.1  172.16.1.2  1/1  Up    0  (3 )   Up   Fa0/1

Session state is UP and using echo function with 50 ms interval.
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 1000000, MinRxInt: 1000000, Multiplier: 3
Received MinRxInt: 1000000, Received Multiplier: 3
Holddown (hits): 0(0), Hello (hits): 1000(311)
Rx Count: 290, Rx Interval (ms) min/max/avg: 1/1900/883 last: 328 ms ago
Tx Count: 312, Tx Interval (ms) min/max/avg: 1/1000/875 last: 244 ms ago
Elapsed time watermarks: -1 0 (last: 0)
Registered protocols: EIGRP
Uptime: 00:04:15
Last packet: Version: 1         - Diagnostic: 0
       State bit: Up        - Demand bit: 0
       Poll bit: 0         - Final bit: 0
       Multiplier: 3        - Length: 24
       My Discr.: 1        - Your Discr.: 1
       Min tx interval: 1000000  - Min rx interval: 1000000
       Min Echo interval: 50000
```

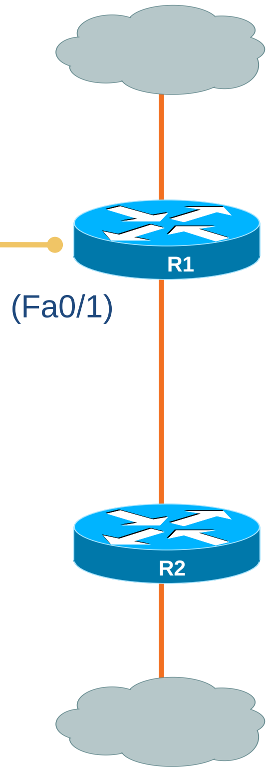

Routing Protocol Neighbour Behaviour

Detecting Unreachable Neighbour (Hello Timers vs. BFD)

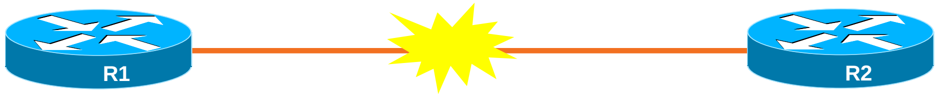

100% Packet Loss (Link Up)

EIGRP Default: Elapsed Time Between 10 – 15 Sec

```
R1# show clock
*19:43:37.646 UTC Mon Feb 16 2009

*Feb 16 19:43:48.974: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65110: Neighbor 10.1.2.220 (FastEthernet0/1) is down: holding time expired
```

BFD: Elapsed Time Between 100 - 150 ms

```
*Feb 16 19:15:41.730: bfdV1FSM e:5, s:3bfdnfy-client a:10.1.2.220, e: 1
*Feb 16 19:15:41.730: Session [10.1.2.120,10.1.2.220,Fa0/1,1], event ECHO FAILURE, state UP -> DOWN
*Feb 16 19:15:41.730: BFD: bfd_neighbor - action:DESTROY, proc/sub:2048/65110, idb:FastEthernet0/1, neighbor:10.1.2.220
*Feb 16 19:15:41.730: bfdV1FSM e:6, s:1
*Feb 16 19:15:41.730: Session [10.1.2.120,10.1.2.220,Fa0/1,1], event Session delete, state DOWN -> ADMIN DOWN
*Feb 16 19:15:41.734: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65110: Neighbor 10.1.2.220 (FastEthernet0/1) is down: BFD DOWN notification
*Feb 16 19:15:41.734: BFD: bfd_neighbor - action:DESTROY, proc/sub:2048/65110, idb:FastEthernet0/1, neighbor:10.1.2.220
```

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
  - First Hop Redundancy Protocols
  - Routing Protocols
  - DDR and Static Routing
  - Performance Routing
- Design and Deployment
- Final Wrap Up

Other Convergence Techniques

- Options using Static Routing
  - Floating Static Routes
  - Reliable Static Routing (RSR) using Enhanced Object Tracking (EOT)
- Dial on Demand Routing (DDR)
  - Backup Interface
  - Dialer Watch
  - EEM Script

- For more information:
  - http://www.cisco.com/en/US/tech/tk801/tk133/technologies_tech_note09186a008009457d.shtml

Floating Static Routes

- Advantages
  - –Independent of line protocol status
  - –Independent of encapsulation type
  - –Can backup multiple interfaces/networks on a router
- Disadvantages
  - –Requires a routing protocol and is dependent upon the routing protocol convergence times
  - –Typically only provides backup for a single router
  - –Requires “interesting” traffic to trigger DDR and to reset idle timers

Static Routes

- The concepts of administrative distance and backup routes are used to create floating static routes
- Configuring a static route with a very high administrative distance ensures it won't be installed as long as there's a dynamically learned route installed in the RIB
- Static routes can also track an SLA object to enable automatic failover

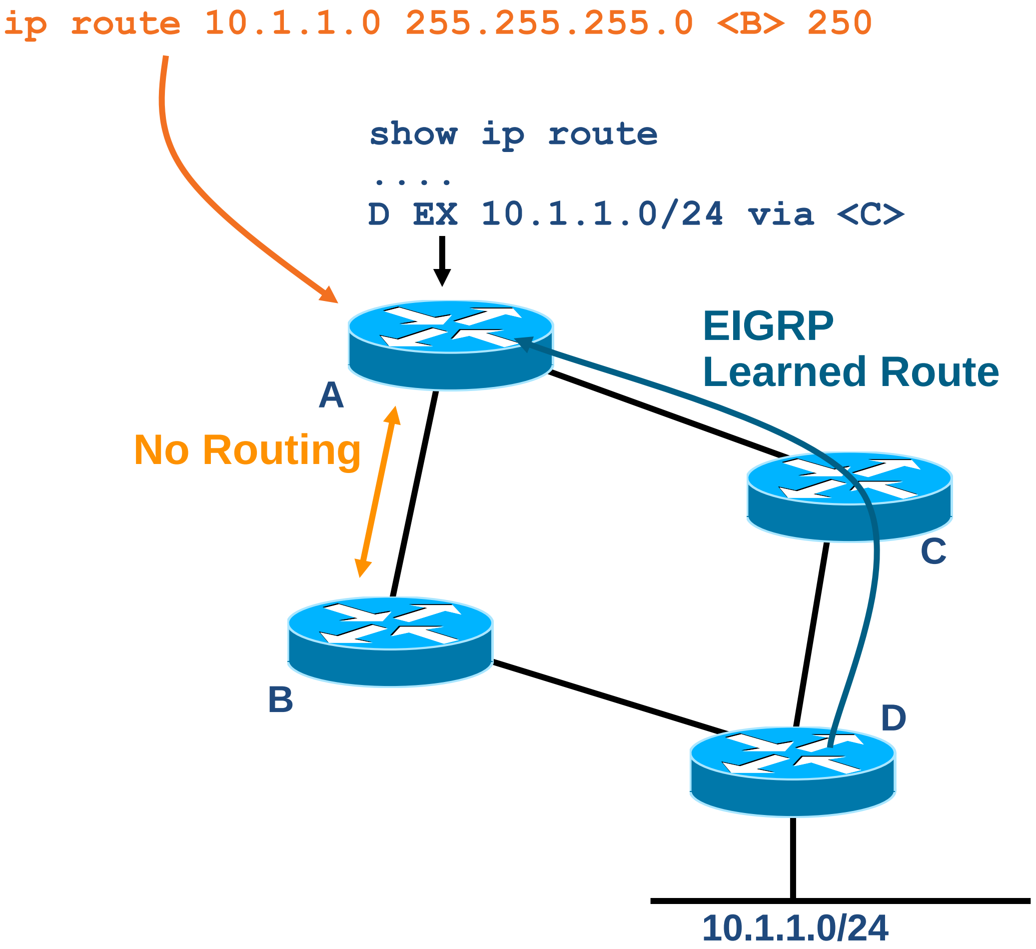

```
ip route 10.1.1.0 255.255.255.0 <B> 250
```

```
show ip route
....
D EX 10.1.1.0/24 via <C>
```

Static Routes

- When the dynamically learned route fails, the RIB calls the processes, looking for a backup route
- Since no other processes have routes to install, the static route with an administrative distance of 250 wins

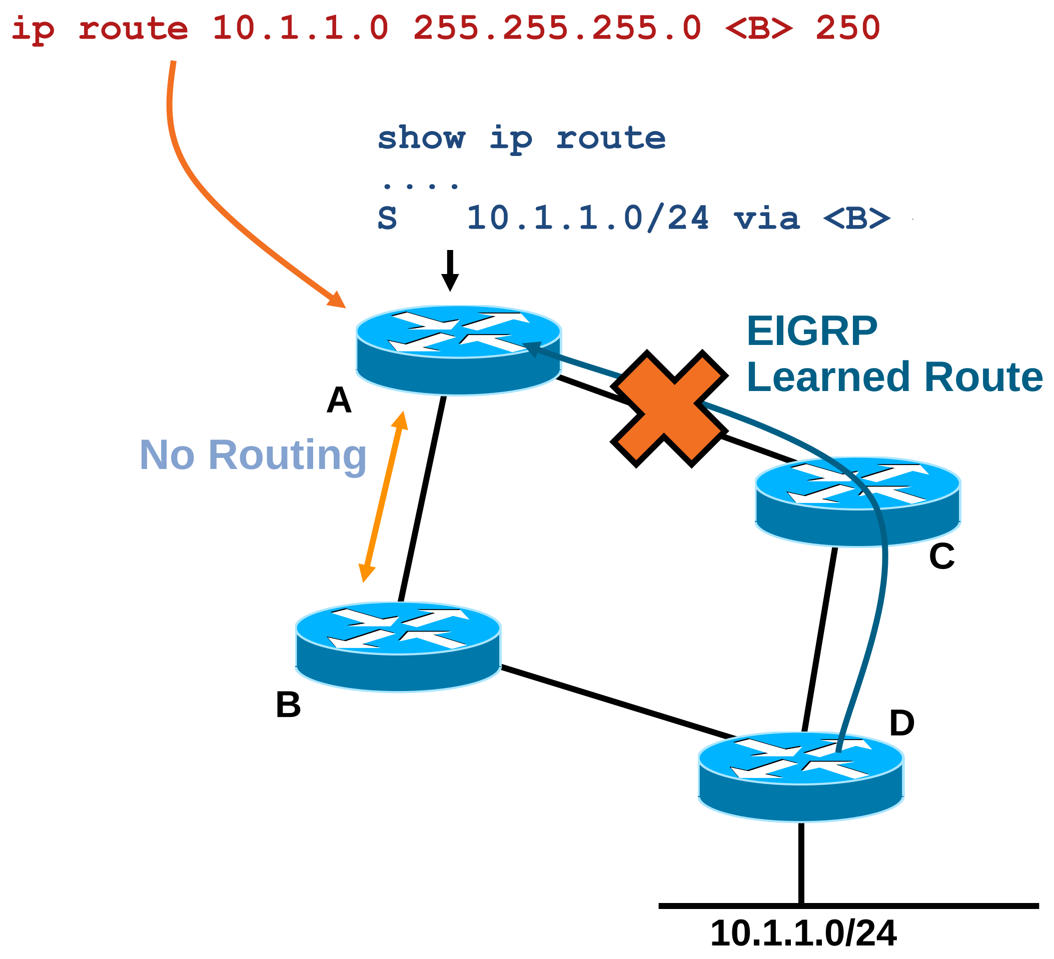

```
ip route 10.1.1.0 255.255.255.0 <B> 250
```

```
show ip route
....
S   10.1.1.0/24 via <B>
```

Reliable Static Routing

Tracking IP SLA

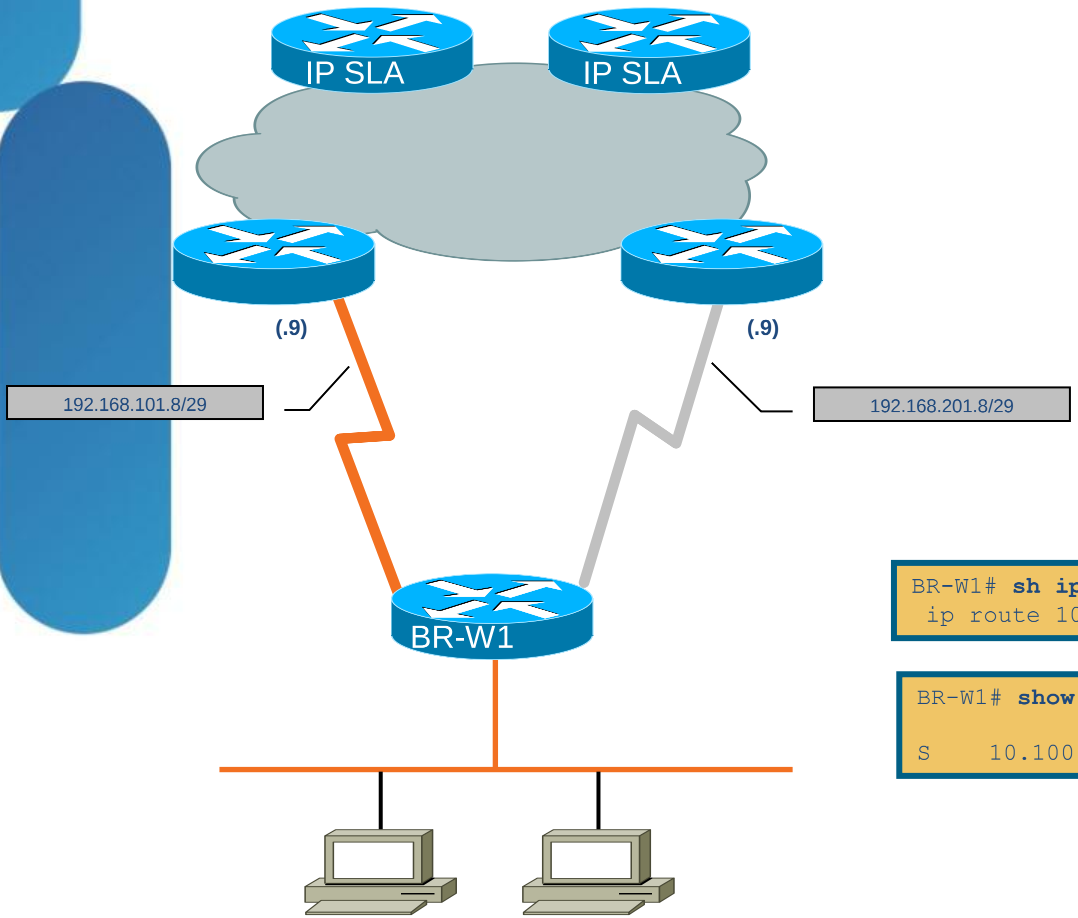

```
BR-W1#
track 100 ip sla 100 reachability
!
ip sla 100
 icmp-echo 10.100.100.100 source-ip 10.1.2.120
 timeout 100
 frequency 10
ip sla schedule 100 life forever start-time now
!
ip route 10.100.100.100 255.255.255.255 192.168.101.9
!
ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 100
ip route 10.100.0.0 255.255.0.0 192.168.201.9 200
```

```
BR-W1# sh ip route track-table
 ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 100 state is [up]
```

```
BR-W1# show ip route

S    10.100.0.0/16 [1/0] via 192.168.101.9
```

Reliable Static Routing

Tracking IP SLA

```
BR-W1#

*Mar 12 03:57:37.119: %TRACKING-5-STATE: 100 rtr 100 reachability Up->Down
```

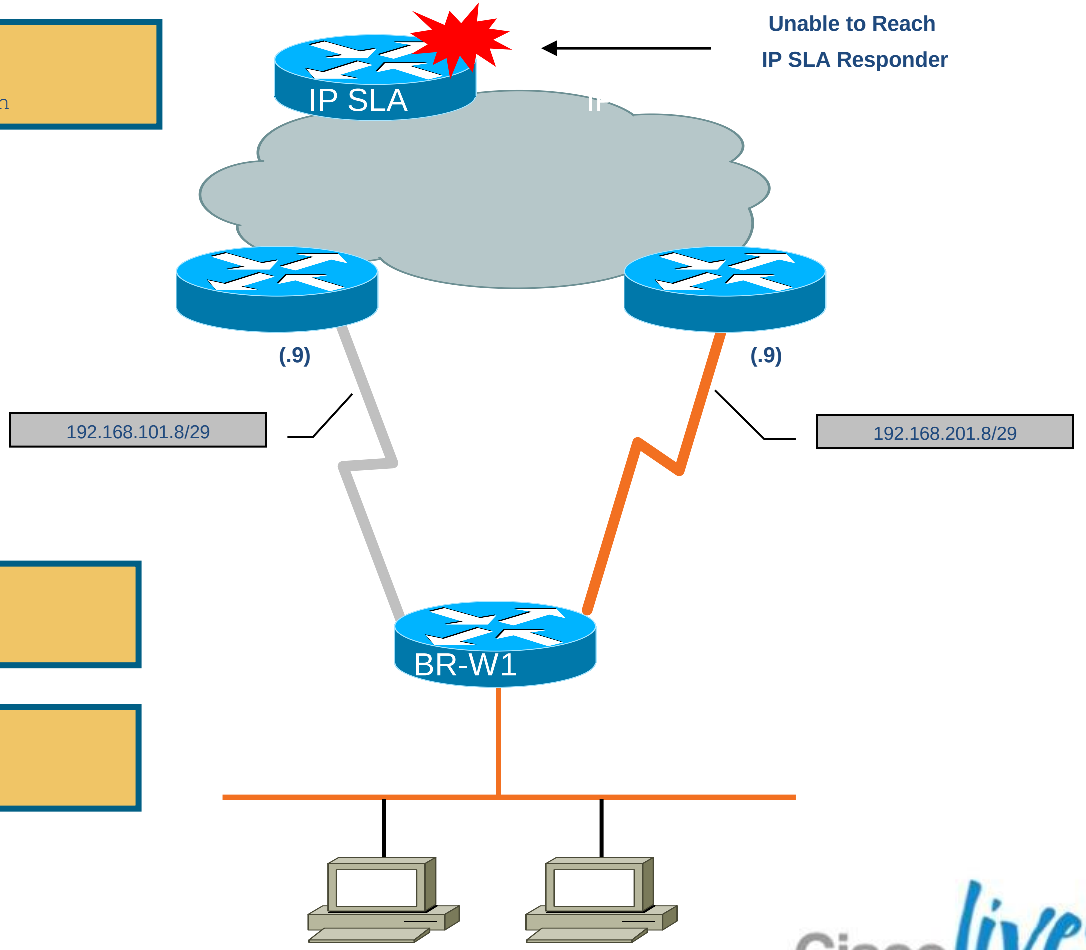

```
BR-W1# show ip route track-table

ip route 10.100.0.0 255.255.0.0 192.168.101.9 track 100 state is [down]
```

```
BR-W1# show ip route

S     10.100.0.0/16 [200/0] via 192.168.201.9
```

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
  - – Routing Protocols
  - – First Hop Redundancy Protocols
  - – DDR and Static Routing
  - – Performance Routing
- Design and Deployment
- Final Wrap Up

Enterprise WAN Challenge

Two Paths, Two Providers → WAN Availability

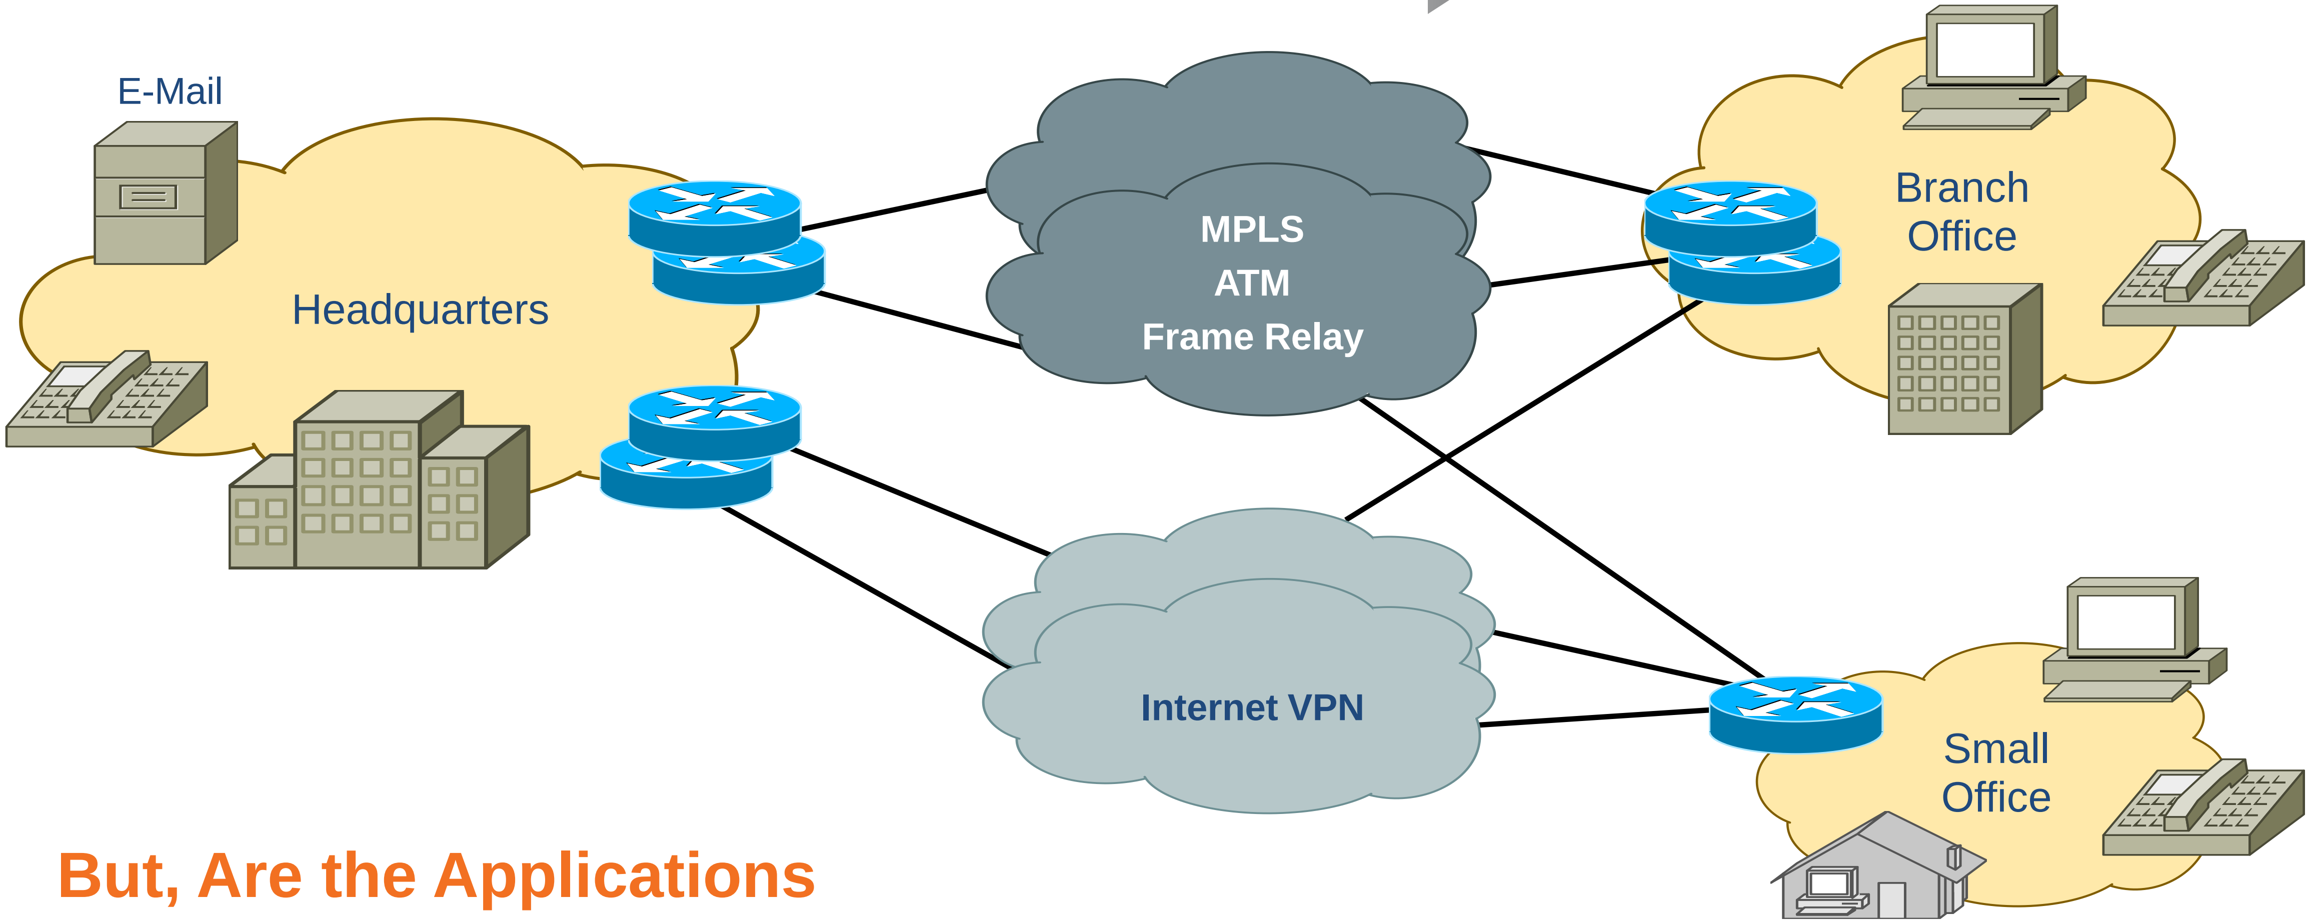

But, Are the Applications Performing Adequately?

Performance Routing (PfR)

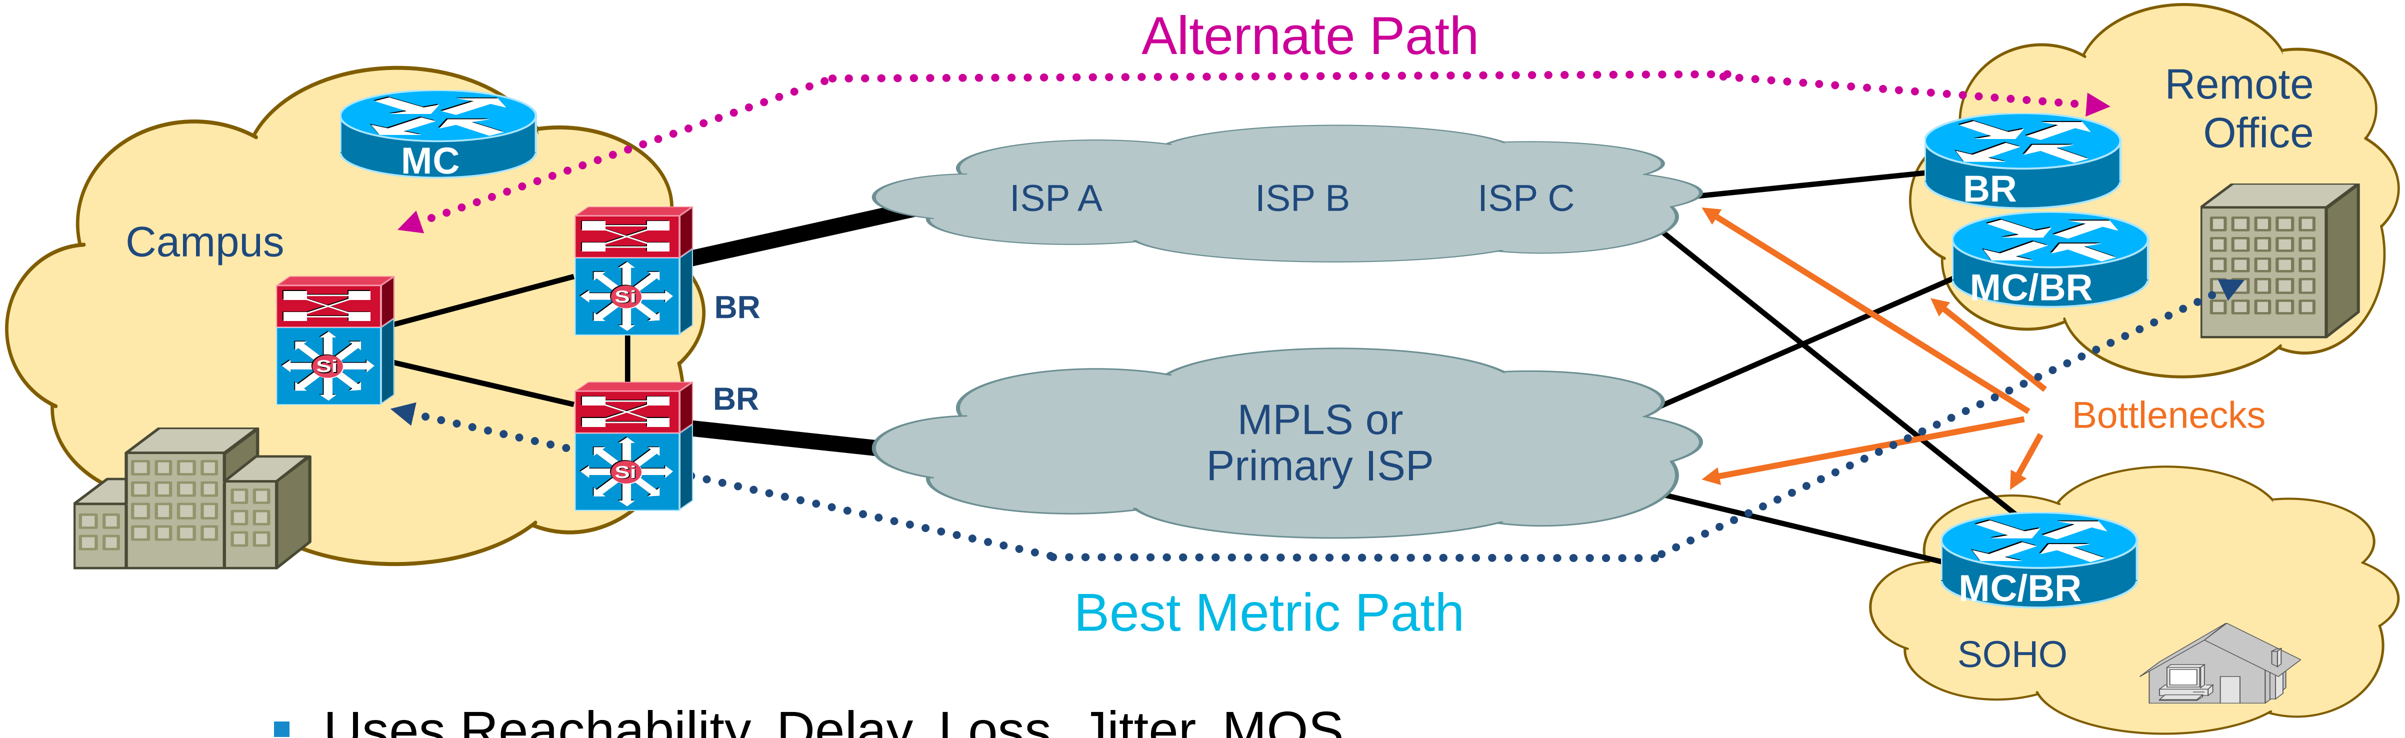

- Uses Reachability, Delay, Loss, Jitter, MOS, Load and $Cost to determine the best path
- PfR Components
  - BR—Border Router (Forwarding Path)
  - MC—Master Controller (Decision Maker)

Traditional Topology

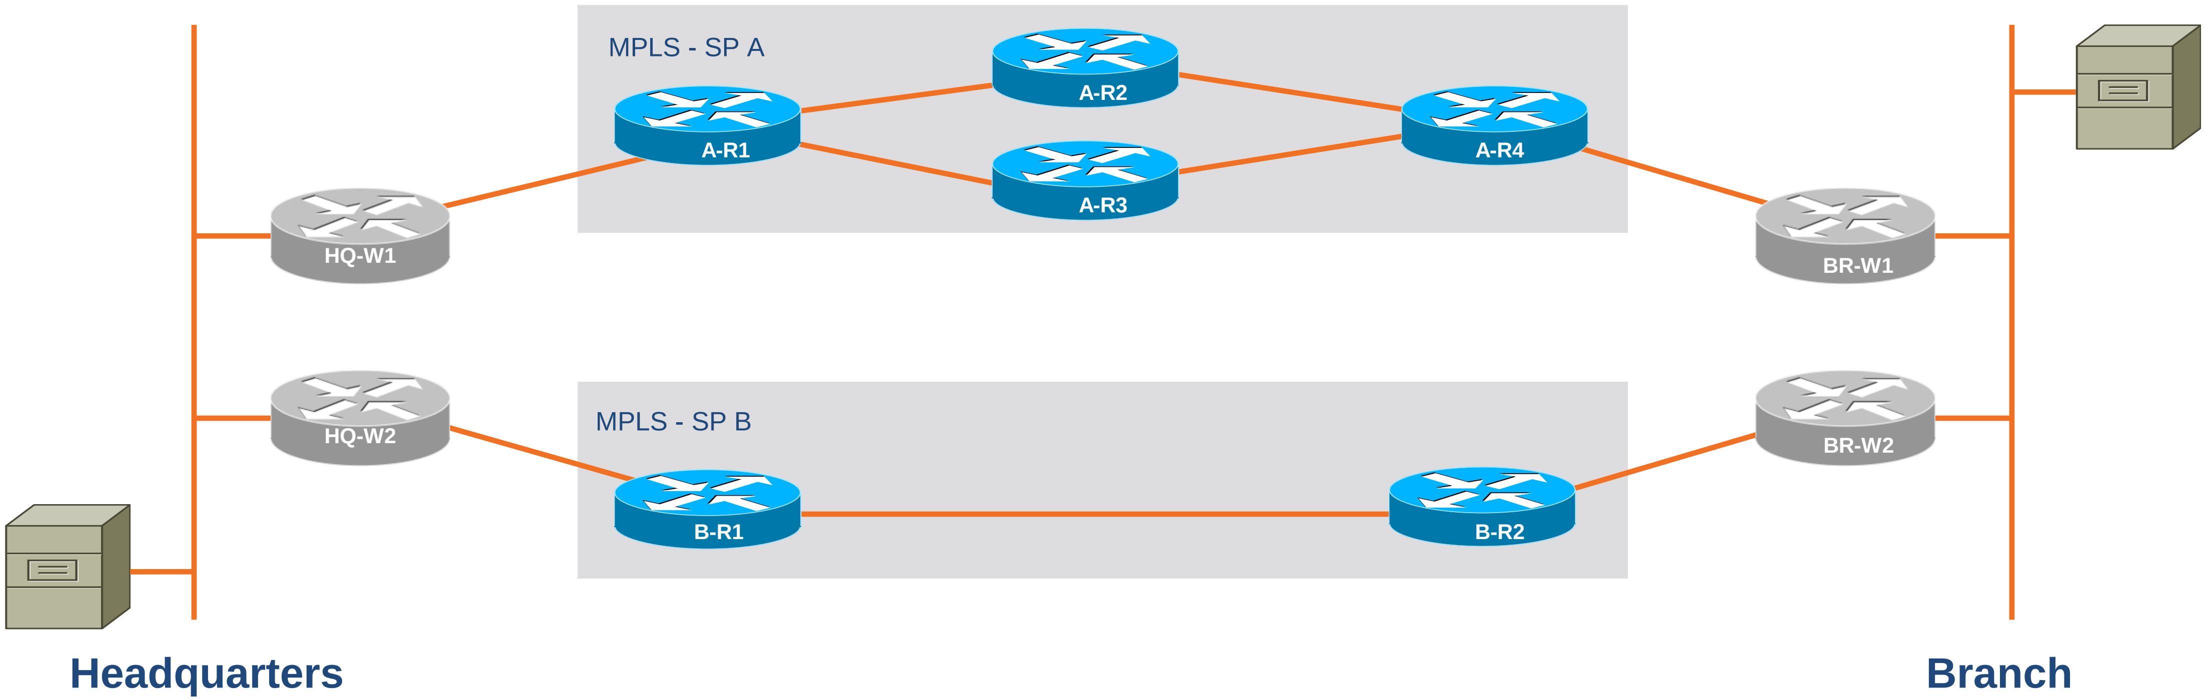

- Routing protocol selects path
- Blackhole reconvergence can take minutes
- Will not recover from brownouts

PfR Enabled Topology

L3—L7 Aware

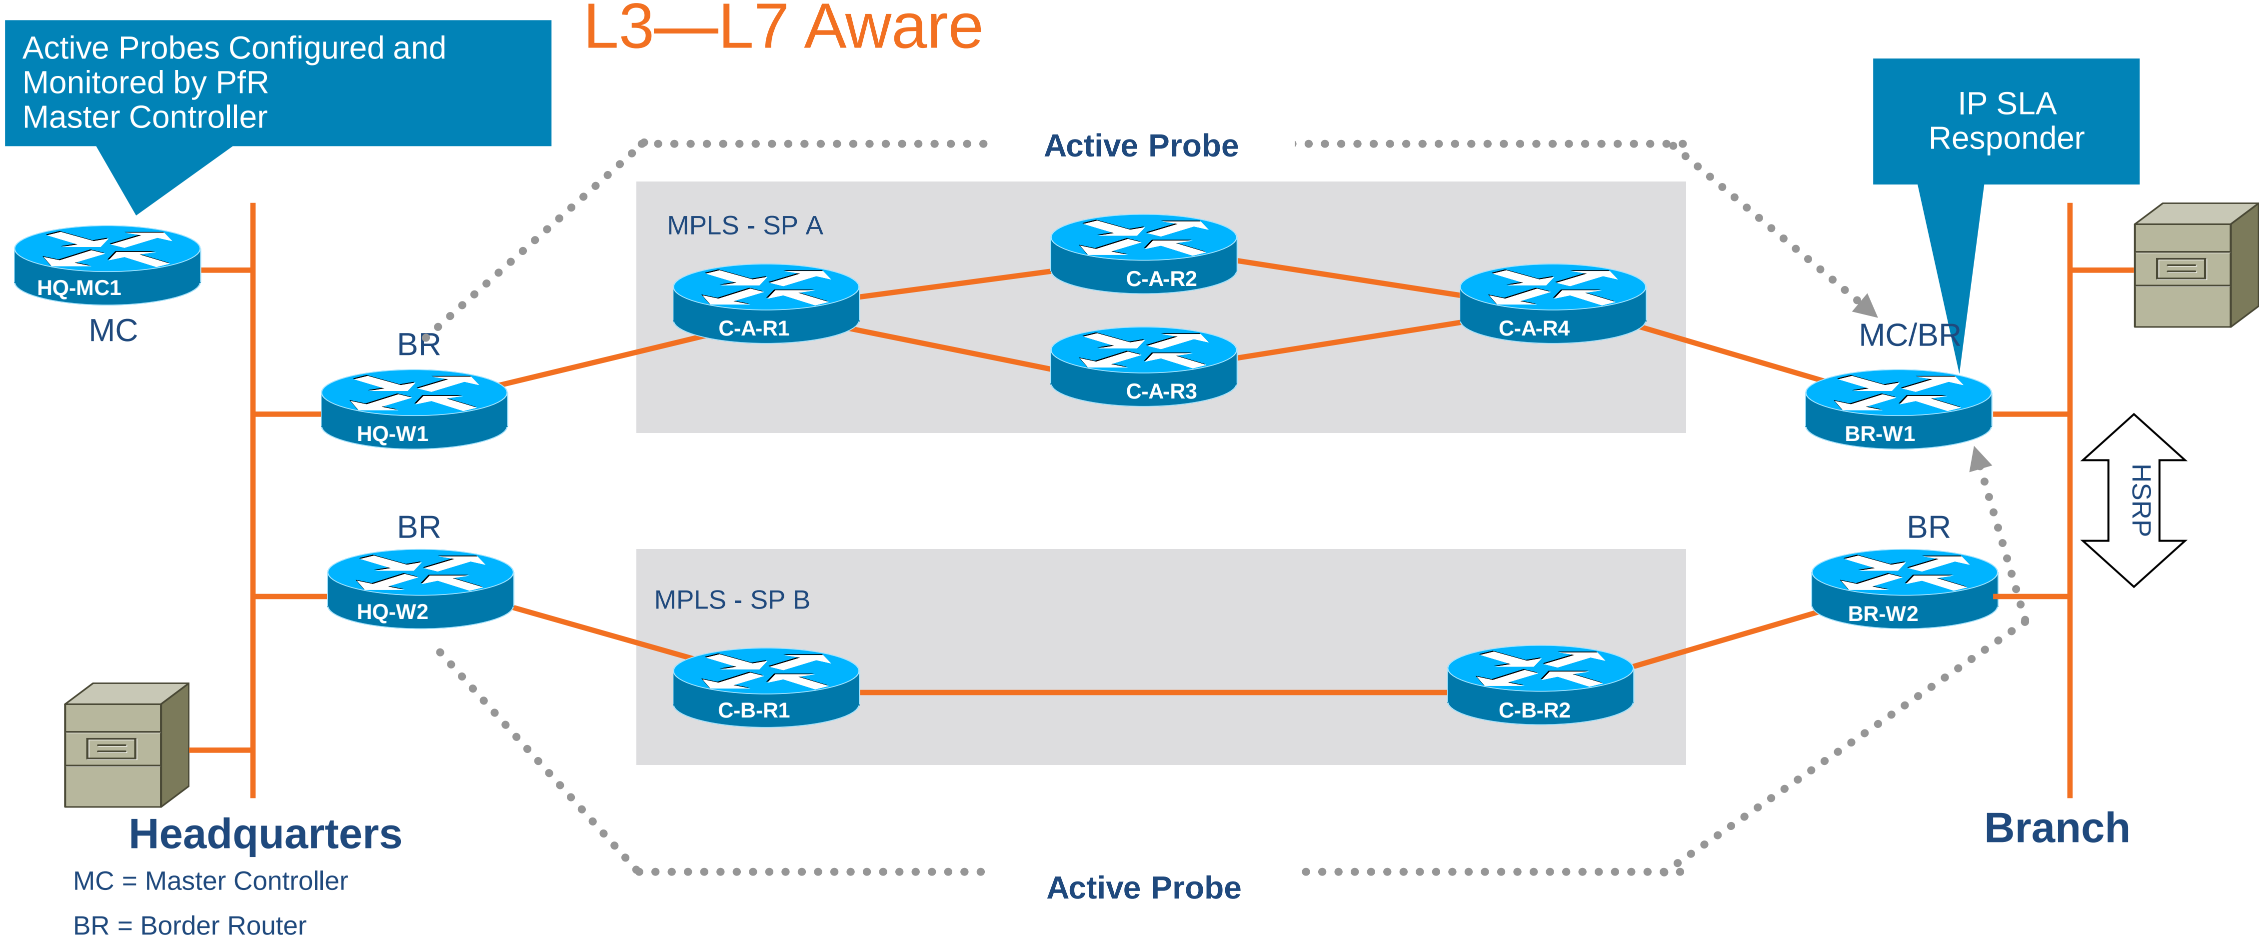

Active Probes Configured and Monitored by PfR Master Controller

IP SLA Responder

Headquarters

MC = Master Controller

BR = Border Router

Branch

- PfR can override routing protocol to select path
- Active probes significantly improve reconvergence due to blackholes and brownouts

Summary of Convergence Techniques

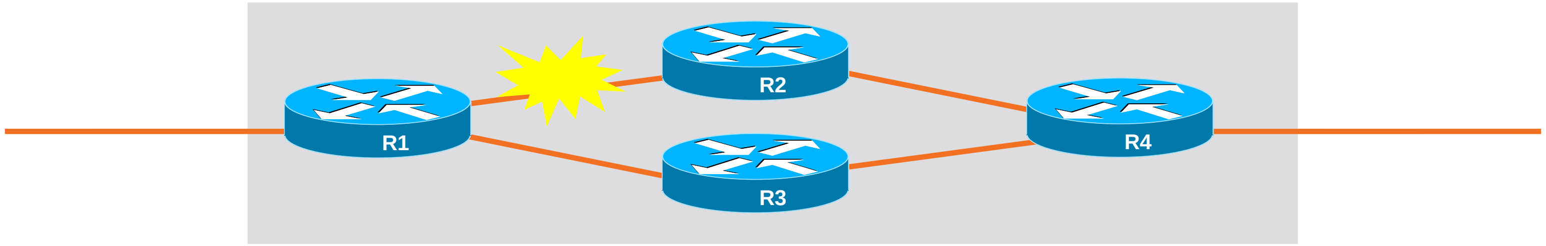

Effectiveness of Various Techniques for Different Outage Types

	Link Down	Link Up Neighbour Down	Link Up Loss ~5%	Upstream Blackhole	Upstream Brownout
Routing Protocols	🟢	🟡	🔴	🟡	🔴
BFD	🟢	🟢	🟡	N/A	N/A
EOT	🟢	🟡	🔴	🟡	🔴
EOT & RSR (w/IP SLA)	🟢	🟢	🟢	🟢	🟢
PfR	🟢	🟢	🟢	🟢	🟢

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
- Design and Deployment
  - MPLS Dual Carrier
  - MPLS + Internet
- Final Wrap Up

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- Default behaviour: 1-way load sharing
- Load is shared from HQ to Branch
- Only one link used Branch to HQ

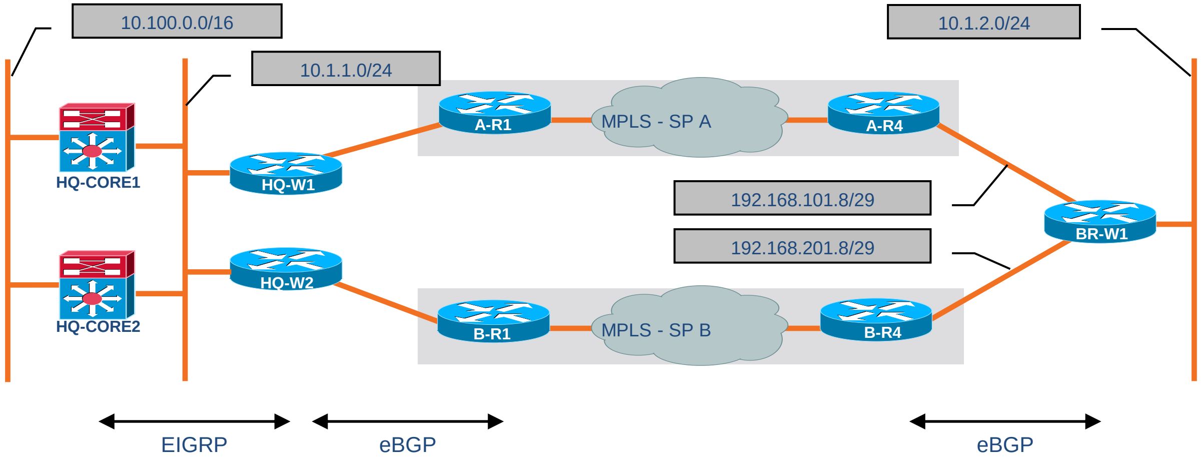

```
HQ-CORE1# show ip route
D EX  10.1.2.0/24 [170/258816] via 10.1.1.110, 02:24:22, Vlan10
                  [170/258816] via 10.1.1.210, 02:24:22, Vlan10
```

```
BR-W1# show ip route
B    10.100.0.0/16 [20/0] via 192.168.101.9, 00:34:00
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- EIGRP
  - –Routes redistributed from BGP into EIGRP (match & tag)
  - –BGP routes are treated as EIGRP external
- BGP
  - –No iBGP required between HQ-W1 & HQ-W2 (CE routers)
  - –Routes redistributed from EIGRP into BGP except those tagged as originally sourced from BGP

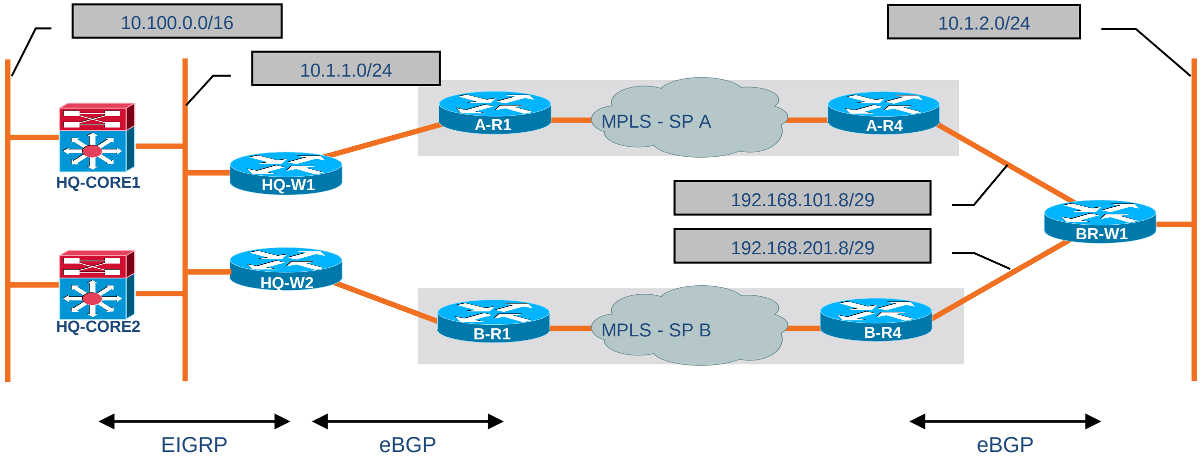

Dual WAN (MPLS—Dual Carrier)

Mutual Route Redistribution Detail

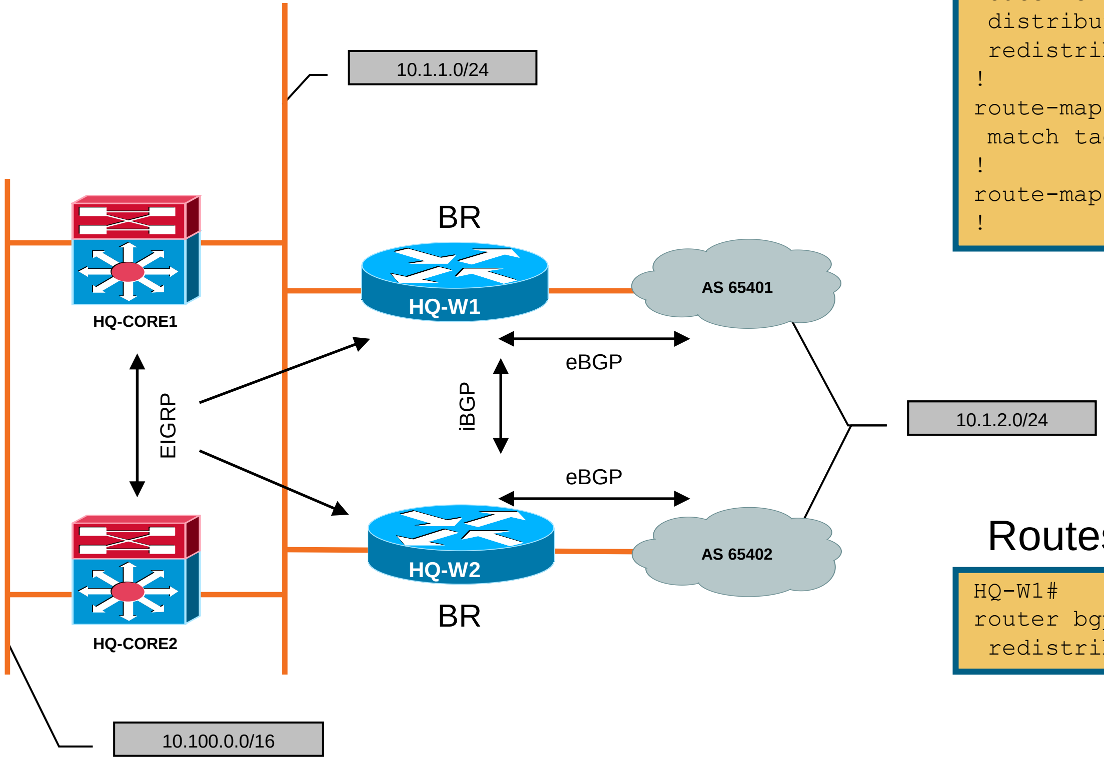

Routes into EIGRP

```
HQ-W1#
router eigrp 65110
 distribute-list route-map BLOCK-TAGGED-ROUTES in
 redistribute bgp 65110
!
route-map BLOCK-TAGGED-ROUTES deny 10
 match tag 65401 65402
!
route-map BLOCK-TAGGED-ROUTES permit 20
!
```

Routes into BGP

```
HQ-W1#
router bgp 65110
 redistribute eigrp 65110
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- EIGRP
  - –No EIGRP required on BR-W1 (collapsed routing)
- BGP
  - –Protect Branch from becoming transit AS

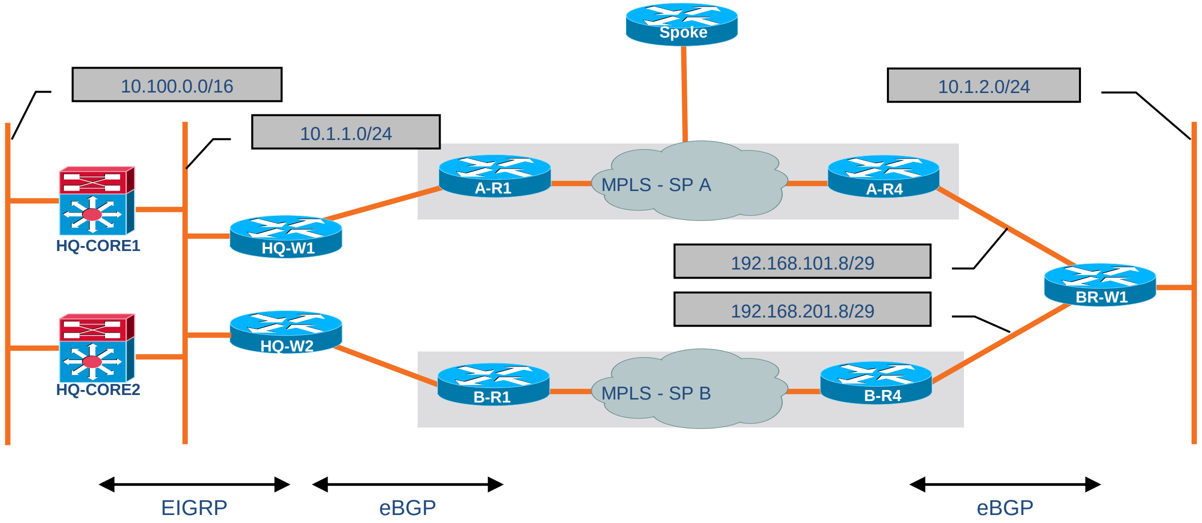

```
BR-W1# show ip bgp

  Network          Next Hop         Metric LocPrf Weight Path

* 10.100.0.0/16    192.168.201.9                       0 65200 65200 ?
*>                 192.168.101.9                       0 65100 65100 ?
```

```
BR-W1#
router bgp 65110
 neighbor 192.168.101.9 route-map NO-TRANSIT-AS out
 neighbor 192.168.201.9 route-map NO-TRANSIT-AS out
....
ip as-path access-list 1 permit ^$

route-map NO-TRANSIT-AS permit 10
 match as-path 1
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- Is it possible to load share from Branch to HQ?
- BGP Multipath
  - –Allows installation of multiple BGP paths to same destination
  - –Requirements (all must be equal)
    - Neighbour AS or AS-PATH
    - Weight
    - Local Pref
    - AS-PATH length
    - Origin
    - Med

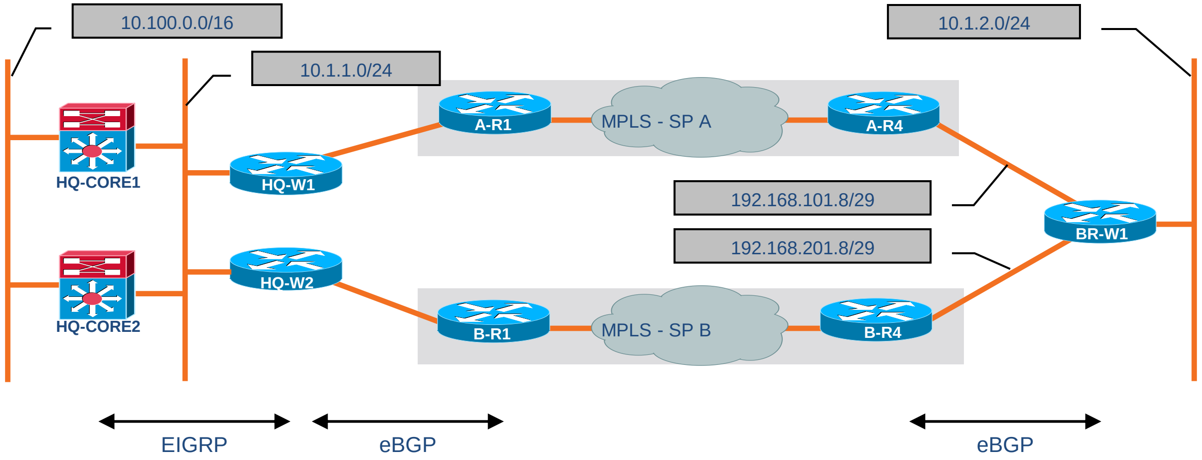

```
BR-W1# show ip bgp

      Network        Next Hop            Metric LocPrf Weight Path

*     10.100.0.0/16  192.168.201.9                          0 65200 65200 ?
*>                   192.168.101.9                          0 65100 65100 ?
```

```
BR-W1# show ip route
B    10.100.0.0/16 [20/0] via 192.168.101.9, 00:34:00
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: BGP

- Is it possible to load share from Branch to HQ?
  - maximum-paths 2
- Requires hidden command:
  - bgp bestpath as-path multipath-relax

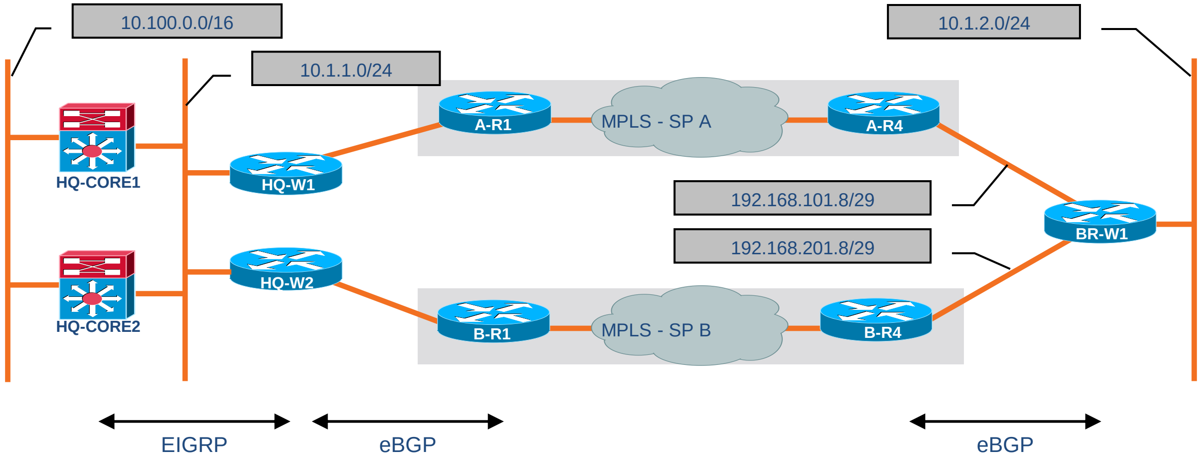

```
BR-W1#
router bgp 65110
 bgp bestpath as-path multipath-relax
 maximum-paths 2
```

```
BR-W1# show ip route
B     10.100.0.0/16 [20/0] via 192.168.201.9, 00:03:44
                    [20/0] via 192.168.101.9, 00:03:44
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: EIGRP

- Default behaviour: 2-way load sharing
  - Load is shared from HQ to Branch
  - Load is shared from Branch to HQ

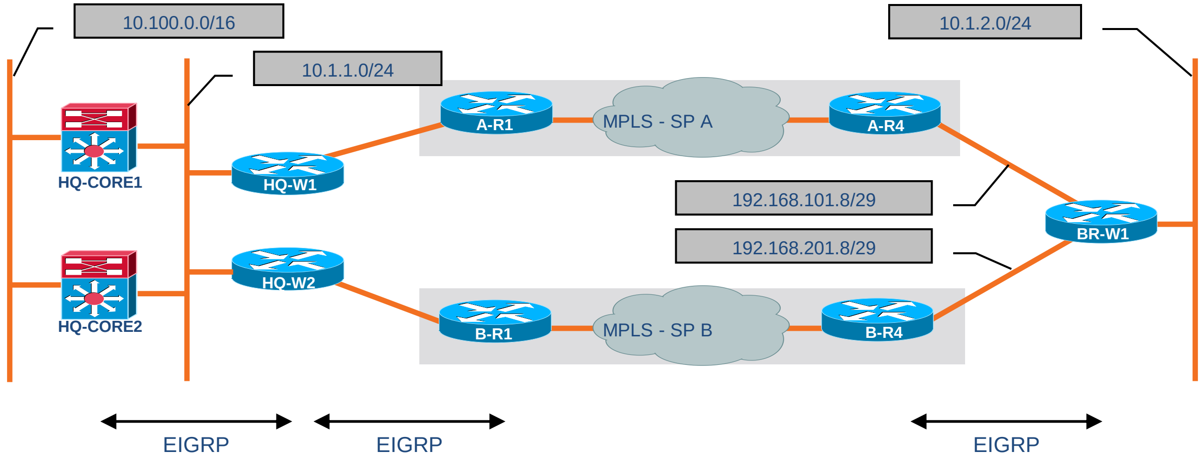

```
HQ-CORE1# show ip route
D    10.1.2.0/24 [90/33536] via 10.1.1.110, 00:15:29, Vlan10
               [90/33536] via 10.1.1.210, 00:15:29, Vlan10
```

```
BR-W1# show ip route
D    10.100.0.0/16
       [90/161280] via 192.168.201.9, 00:00:12, FastEthernet0/0.220
       [90/161280] via 192.168.101.9, 00:00:12, FastEthernet0/0.120
```

Dual WAN (MPLS—Dual Carrier)

PE-CE Protocol: EIGRP

- EIGRP
  - No route redistribution required on CE routers
  - Protect Branch from becoming transit network
- BGP
  - PE routers handle mutual route redistribution

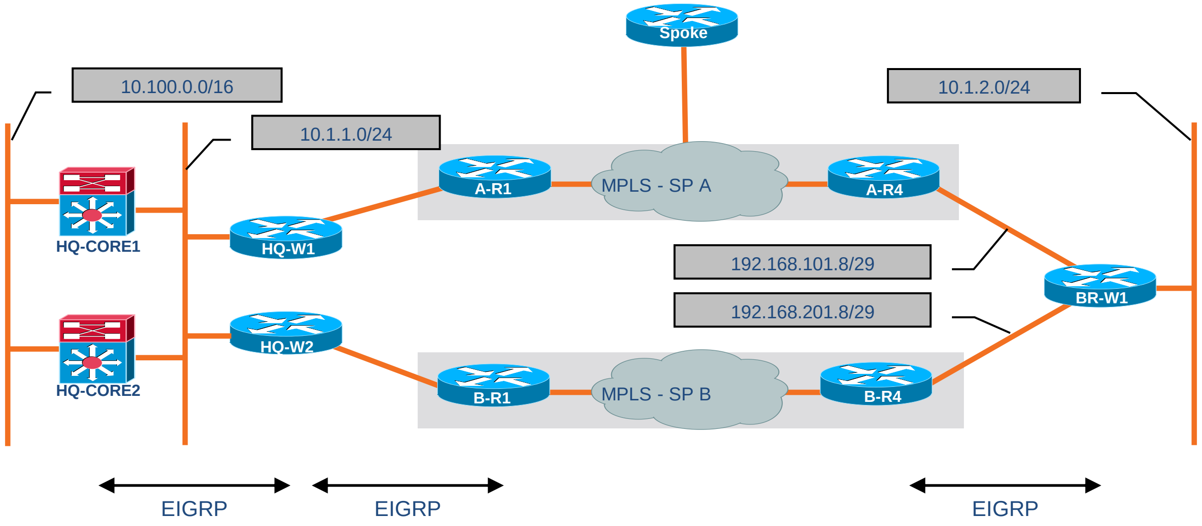

```
BR-W1#
router eigrp 65110
 eigrp stub connected
```

If Using Summaries (Optional):

```
BR-W1#
router eigrp 65110
 eigrp stub connected summary
```

Performance Routing (PfR)

Basic Configuration—Dedicated MC, BRs

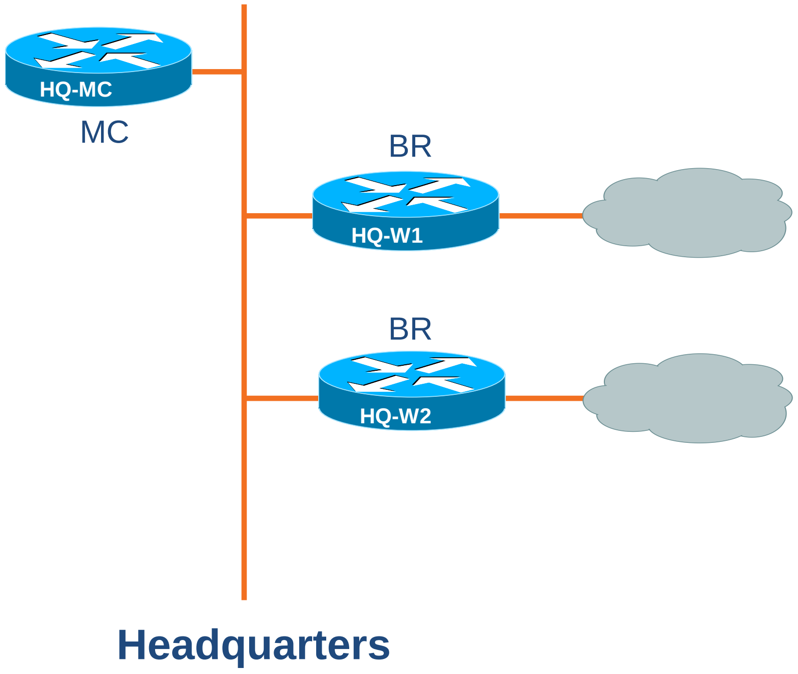

```
HQ-MC#

key chain PFR-KEYCHAIN
 key 1
  key-string cisco123
!
pfr master
 !
 border 10.1.1.110 key-chain PFR-KEYCHAIN
 interface GigabitEthernet0/0 internal
 interface GigabitEthernet0/1 external
 !
 border 10.1.1.210 key-chain PFR-KEYCHAIN
 interface GigabitEthernet0/0 internal
 interface GigabitEthernet0/1 external
```

```
HQ-W1# (*and* HQ-W2)

key chain PFR-KEYCHAIN
 key 1
  key-string cisco123
!
pfr border
 local GigabitEthernet0/0
 master 10.1.1.10 key-chain PFR-KEYCHAIN
```

MC = Master Controller
BR = Border Router
PfR = Performance Routing
OER = Optimized Edge Routing

pfr Keyword in Examples. oer Prior to IOS 15.1 Versions.

Performance Routing

iBGP Configuration—Multiple BRs

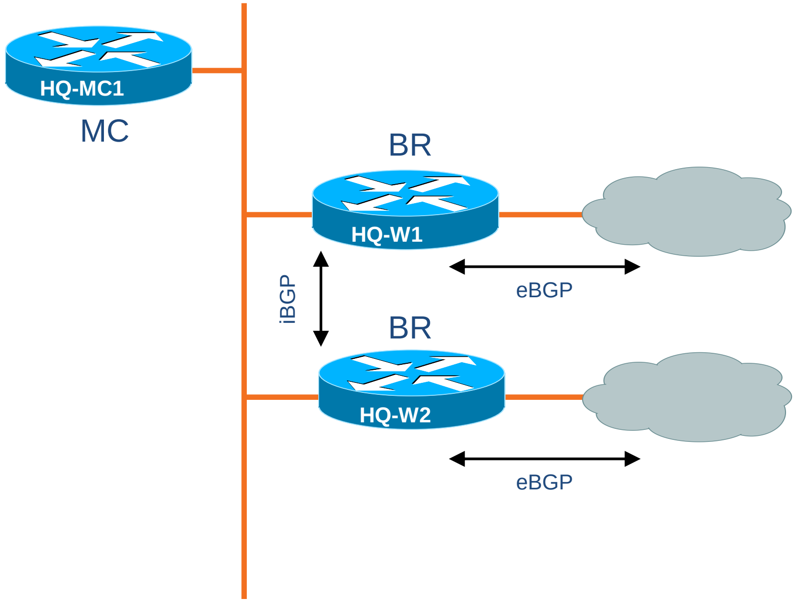

```
HQ-W1#

router bgp 65110
 neighbor 10.1.1.210 remote-as 65110
 neighbor 10.1.1.210 next-hop-self
 neighbor 10.1.1.210 send-community
```

```
HQ-W1# show ip bgp
  Network     Next Hop      Metric LocPrf Weight Path
* i10.1.2.0/24   10.1.1.210        0  100   0 65200 65200 i
*>               192.168.101.2              0 65100 65100 i
```

```
HQ-W2#

router bgp 65110
 neighbor 10.1.1.110 remote-as 65110
 neighbor 10.1.1.110 next-hop-self
 neighbor 10.1.1.110 send-community
```

```
HQ-W2# show ip bgp
  Network     Next Hop      Metric LocPrf Weight Path
* i10.1.2.0/24   10.1.1.110        0  100   0 65100 65100 i
*>               192.168.201.2              0 65200 65200 i
```

Headquarters

MC = Master Controller

BR = Border Router

Performance Routing

Basic Configuration—Combined MC and BR

```
BR-W1#
key chain PFR-KEYCHAIN
 key 1
  key-string cisco123
!
pfr master
 !
 border 10.1.2.120 key-chain PFR-KEYCHAIN
 interface FastEthernet0/1 internal
 interface FastEthernet0/0.120 external
 interface FastEthernet0/0.220 external
!
pfr border
 local FastEthernet0/1
 master 10.1.2.120 key-chain PFR-KEYCHAIN
 active-probe address source interface FastEthernet0/1
!
interface FastEthernet0/0
 load-interval 30
!
interface FastEthernet0/0.120
 bandwidth 4000
!
interface FastEthernet0/0.220
 bandwidth 4000
```

- Load-interval affects moving average calculations
- Bandwidth affects utilisation calculations

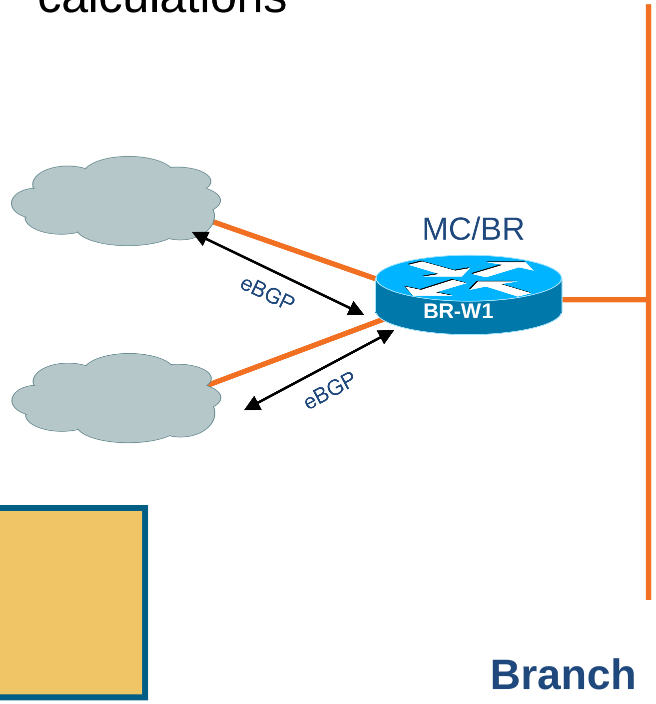

MC = Master Controller

BR = Border Router

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.100.0/24  192.168.101.9                          0 65100 65100 ?
*                   192.168.201.9                          0 65200 65200 ?
*> 10.100.200.0/24  192.168.101.9                          0 65100 65100 ?
*                   192.168.201.9                          0 65200 65200 ?
```

```
BR-W1# show ip route
B    10.100.100.0 [20/0] via 192.168.101.9, 03:32:30
B    10.100.200.0 [20/0] via 192.168.101.9, 03:32:30
```

Performance Routing

Load Sharing Configuration—Link Utilisation

If Traffic Goes Above the **max-xmit-utilisation** Threshold, PfR Tries to Move the Traffic from this Exit Link to Another Underutilised Exit Link

```
BR-W1#

pfr master
 interface FastEthernet0/0.120 external
  max-xmit-utilization percentage 50
 interface FastEthernet0/0.220 external
  max-xmit-utilization percentage 50
 !
 learn
  throughput
  periodic-interval 1
  monitor-period 1
 mode route observe
 mode select-exit best
 resolve utilization priority 1 variance 5
 no resolve delay
```

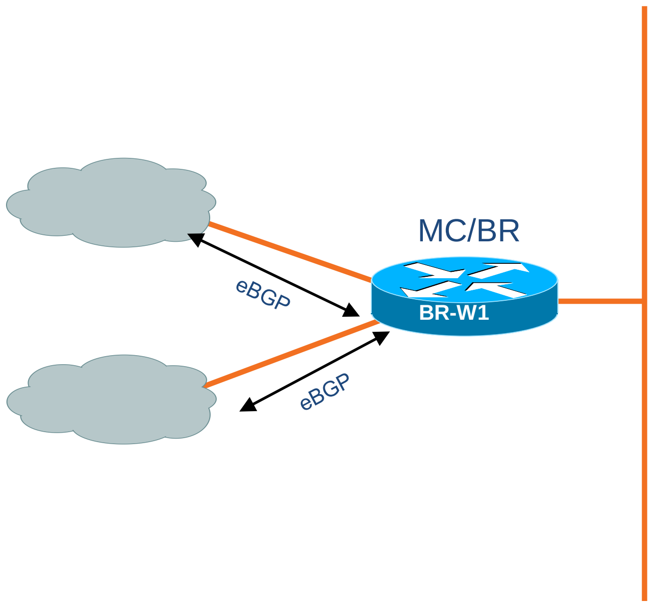

The **Variance** Keyword Configures the Allowable Percentage that an Exit Link Can Vary from the User-Defined Policy Value and Still Be Considered Equivalent

Branch

MC = Master Controller

BR = Border Router

Performance Routing

Load Sharing Example: PfR Enabled (Observe Only)

- Example Load
  - ~546 Kbps UDP Bi-Directional

    10.1.2.100 to 10.100.100.100

    10.1.2.100 to 10.100.200.201
- MPLS – SP B is not currently being utilised for Branch to HQ traffic

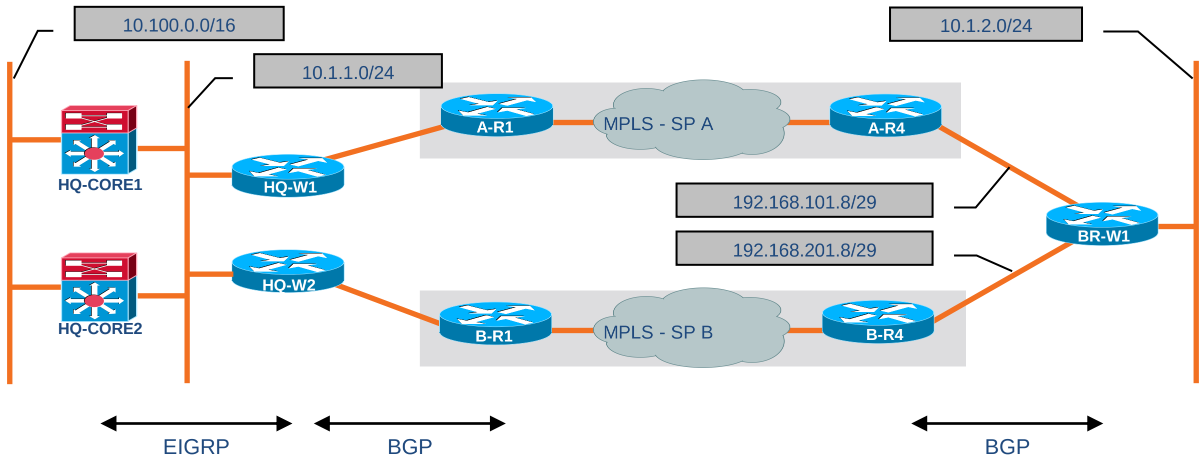

```
BR-W1# show pfr master border detail
Border           Status   UP/DOWN             AuthFail  Version
10.1.2.120       ACTIVE   UP       02:30:02          0  2.2
 Fa0/1           INTERNAL UP
 Fa0/0.120       EXTERNAL UP
 Fa0/0.220       EXTERNAL UP

 External            Capacity      Max BW   BW Used    Load  Status   ExitId
 Interface            (kbps)       (kbps)    (kbps)    (%)
 ---------           ---------     ------   -------- --------  ------   ------
 Fa0/0.120       Tx      4000        2000      1093      27    UP      2
                 Rx                  4000       547      13
 Fa0/0.220       Tx      4000        2000         0       0    UP      1
                 Rx                  4000       546      13
```

Performance Routing

Load Sharing Example: PfR Enabled (Route Control)

```
BR-W1#

pfr master
 mode route control
```

- Both MPLS carriers are now being utilised (in both directions)
- More prefixes and flows result in better load sharing

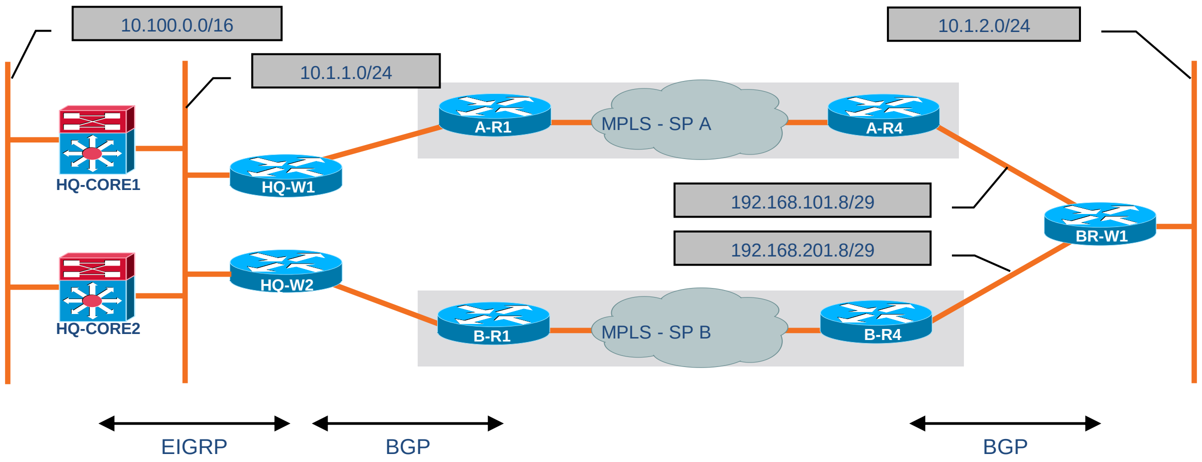

```
BR-W1# show pfr master border detail
Border           Status   UP/DOWN             AuthFail  Version
10.1.2.120       ACTIVE   UP       02:40:38          0  2.2
 Fa0/1           INTERNAL UP
 Fa0/0.120       EXTERNAL UP
 Fa0/0.220       EXTERNAL UP

 External            Capacity      Max BW   BW Used    Load  Status  ExitId
 Interface            (kbps)       (kbps)    (kbps)    (%)
 ---------           ---------     ------   ------- -------  ------  ------
 Fa0/0.120       Tx      4000        2000       547      13  UP       2
                 Rx                  4000       546      13
 Fa0/0.220       Tx      4000        2000       546      13  UP       1
                 Rx                  4000       548      13
```

Performance Routing

Load Sharing Example: PfR Enabled (Route Control)

- BGP route selection is influenced by PfR
- BGP change is also reflected with update
  to routing table

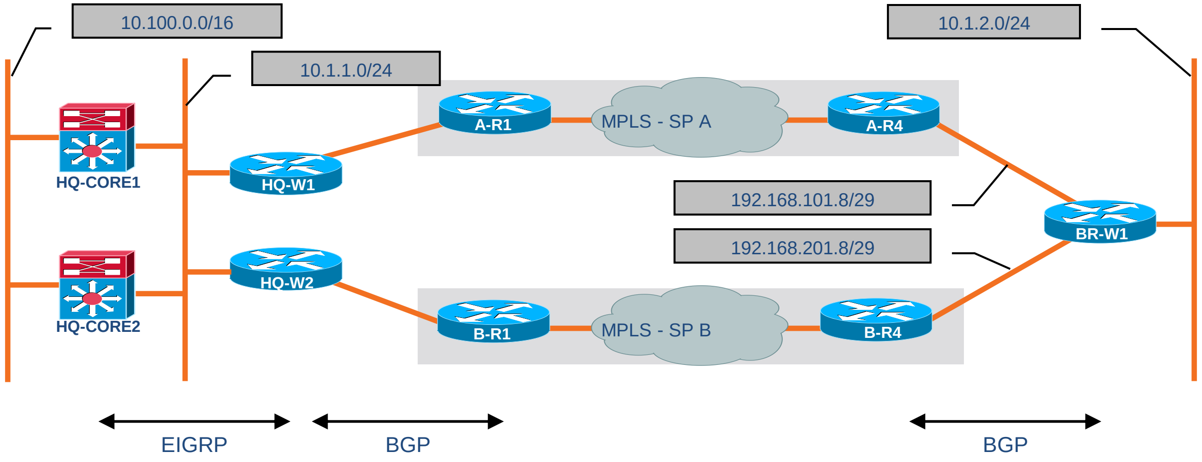

Unchanged by PfR

Controlled by PfR

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.100.0/24  192.168.101.9                          0 65100 65100 ?
*                   192.168.201.9                          0 65200 65200 ?
*  10.100.200.0/24  192.168.101.9                          0 65100 65100 ?
*>                  192.168.201.9                          0 65200 65200 ?
```

```
BR-W1# show ip route
B    10.100.100.0 [20/0] via 192.168.101.9, 03:38:43
B    10.100.200.0 [20/0] via 192.168.201.9, 03:45:13
```

Performance Routing

Load Sharing Example: PfR Enabled (Route Control)

- Dual Router WAN Edge
  - HSRP facing LAN hosts
  - Requires iBGP config (similar to HQ)
  - PfR influences outbound traffic using BGP local-preference (5000)

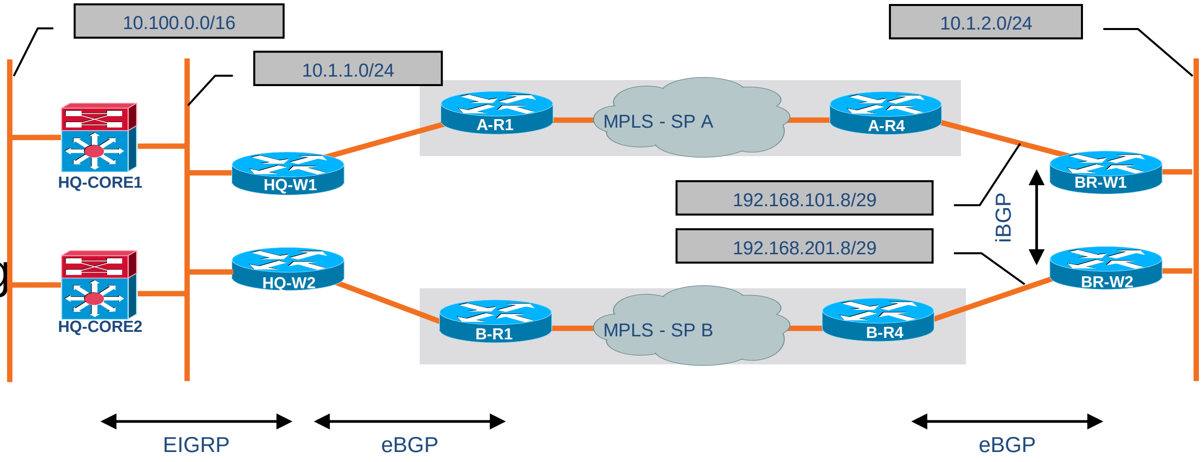

```
BR-W1#show pfr master
<snip>
Global Settings:
  max-range-utilization percent 20 recv 0
  mode route metric bgp local-pref 5000
  mode route metric static tag 5000
  trace probe delay 1000
  no logging
  exit holddown time 60 secs, time remaining 0
Default Policy Settings: <snip>
Learn Settings: <snip>
```

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
* i10.100.100.0/24  10.1.2.220               0    100      0 65200 65200 ?
*>                  192.168.101.9                          0 65100 65100 ?
*>i10.100.200.0/24  10.1.2.220               0   5000      0 65200 65200 ?
*                   192.168.101.9                          0 65100 65100 ?
```

```
BR-W1# show ip route
B       10.100.100.0 [20/0] via 192.168.101.9, 01:08:43
B       10.100.200.0 [200/0] via 10.1.2.220, 00:03:22
```

Performance Routing

Multiple Paths—Select Best Path by Destination Prefix

- Monitor relevant path characteristics (round trip delay, loss, jitter, …)
  - Example:
  - path A: <5 ms delay, 0% loss
  - path B: < 50 ms delay, 0% loss
- Accurate measurement of most parameters requires active probes (which leverage IP SLA)
- Each path must be evaluated in each direction independently
- Craft a policy to take advantage of unique link characteristics
  - Example:
  - If both paths are lossless, then prefer the path with lower delay.
  - However, if loss begins to exceed .01% then prefer the lossless path even if it has increased delay.

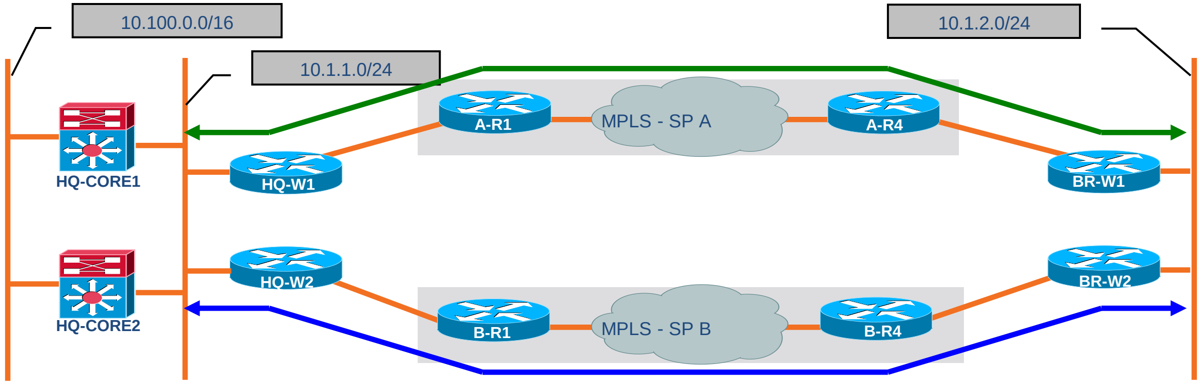

Performance Routing

Active Probe Configuration

```
RTR#
int FastEthernet0
 ip address 10.100.100.100 255.255.255.0
!
ip sla responder
```

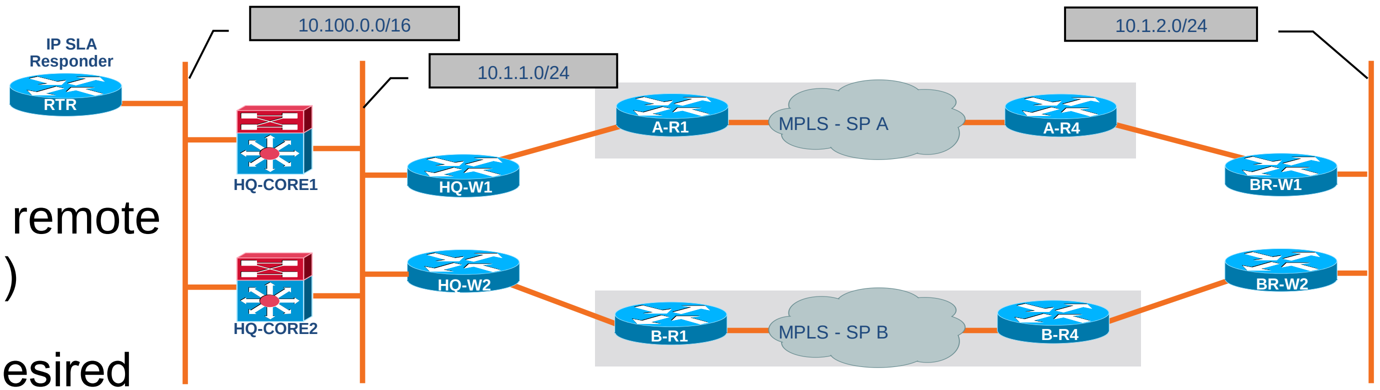

- Configure IP SLA Responder on remote router (consider “shadow router”)
- Configure “pfr-map” to monitor desired remote IP prefix(es)
- For performance sensitive traffic, use “mode monitor fast”
- Probes are sourced from Border Routers and routed via external interfaces. Probe return traffic returns via traditional routed path (likely asymmetric)

```
BR-W1#
ip prefix-list HQ-CRITICAL-100 seq 10 permit 10.100.100.0/24
!
pfr-map PFR-HQ-FAST-FAILOVER 10
 match traffic-class prefix-list HQ-CRITICAL-100
 set mode route observe
 set mode monitor fast
 set active-probe jitter 10.100.100.100 target-port 5555
 set probe frequency 2
!
pfr master
 policy-rules PFR-HQ-FAST-FAILOVER
```

Performance Routing

Active Prefix Monitoring

```
BR-W1# show pfr master policy 10
pfr-map PFR-HQ-FAST-FAILOVER 10
 match ip prefix-lists: HQ-CRITICAL-100
 backoff 300 3000 300
 delay relative 50
 holddown 300
 periodic 0
 *probe frequency 2
 *mode route observe
 *mode monitor fast
 mode select-exit best
 loss relative 10
 jitter threshold 20
 mos threshold 3.60 percent 30
 unreachable relative 50
 next-hop not set
 forwarding interface not set
 resolve utilization priority 1 variance 5
* Overrides Default Policy Setting
```

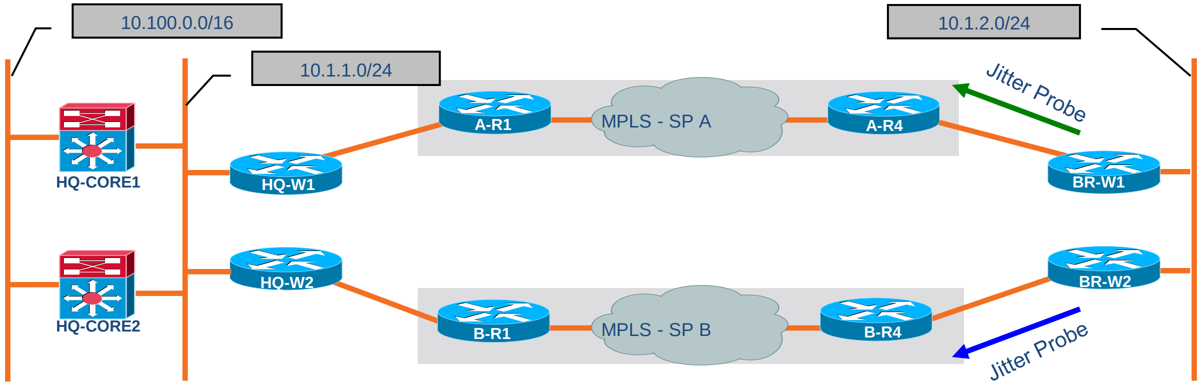

```
BR-W1# show pfr master prefix detail
Prefix: 10.100.100.0/24
   State: INPOLICY*   Time Remaining: @0
   Policy: 10

   Most recent data per exit
   Border          Interface          PasSDly  PasLDly  ActSDly  ActLDly
  *10.1.2.120      Fa0/0                    0        0        3        3
   10.1.2.220      Fa0/0                    0        0       28       28

   Most recent voice data per exit
   Border          Interface          ActSJit  ActPMOS  ActSLos  ActLLos
  *10.1.2.120      Fa0/0                    0        0        0        0
   10.1.2.220      Fa0/0                    0        0        0        0
```

Short/Long Term Delay

Short/Long Term Loss

Performance Routing

Controlling a Configured Prefix

- Compare loss characteristics of each exit (within 5% considered “same”)
- Next, compare delay characteristics of each exit (within 5% considered “same”)
- Maximum packet loss for an exit is limited to 100 (packets per million) or 0.01%
- Maximum delay for an exit is 100 ms
- Unreachable policy is always considered “highest priority”

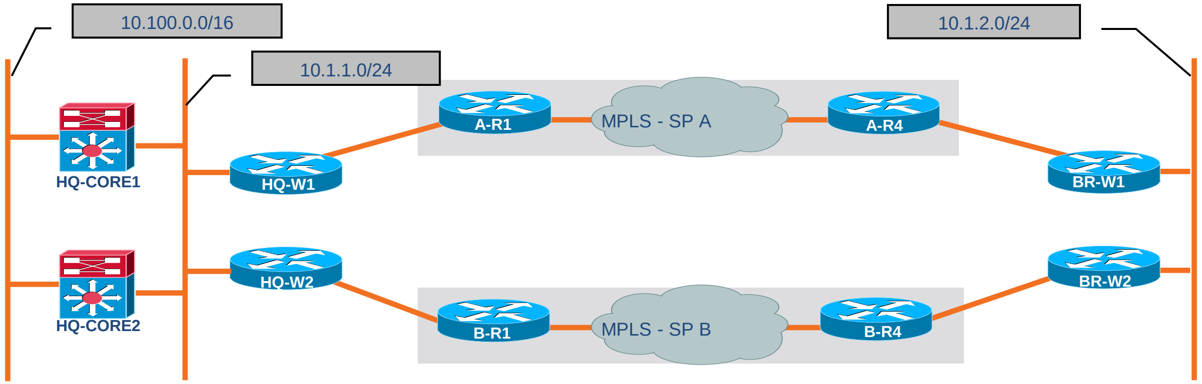

```
BR-W1#

pfr-map PFR-HQ-FAST-FAILOVER 10
 no set resolve utilization
 set resolve loss priority 1 variance 5
 set resolve delay priority 2 variance 5
 set loss threshold 100
 set delay threshold 100
 !
 set holddown 90
 set periodic 90
 set mode select-exit best
 !
 set mode route control
!
pfr master
 policy-rules PFR-HQ-FAST-FAILOVER
```

Performance Routing

Controlling a Configured Prefix

Branch Route Tables – Monitor Only

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.0.0/16    192.168.101.9                          0 65100 65100 ?
* i                 10.1.2.220               0    100      0 65200 65200 ?
```

```
BR-W2# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
* i10.100.0.0/16    10.1.2.120               0    100      0 65100 65100 ?
*>                  192.168.201.9                          0 65200 65200 ?
```

PfR Moves the (More Specific) Prefix 10.100.100.0/24 to the Path with Lower Delay

Branch Route Tables – PfR Route Control

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.0.0/16    192.168.101.9                          0 65100 65100 ?
* i                 10.1.2.220                             0 65200 65200 ?
*> 10.100.100.0/24  192.168.101.9                          0 65100 65100 ?
```

Prefix Inserted by PfR

```
BR-W2# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
* i10.100.0.0/16    10.1.2.120               0    100      0 65100 65100 ?
*>                  192.168.201.9                          0 65200 65200 ?
*>i10.100.100.0/24  10.1.2.120               0   5000      0 65100 65100 ?
```

Prefix Inserted by PfR

Performance Routing

Path Disruption: Loss

```
HQ-MC#
*Mar 3 21:18:53.247: %OER_MC-5-NOTICE: Active ABS Loss OOP Prefix 10.1.2.0/24, loss 5025,
BR 10.1.1.110, i/f Gi0/1
*Mar 3 21:18:55.263: %OER_MC-5-NOTICE: Active ABS Loss OOP Prefix 10.1.2.0/24, loss 15228,
BR 10.1.1.110, i/f Gi0/1
*Mar 3 21:18:55.267: %OER_MC-5-NOTICE: Route changed Prefix 10.1.2.0/24, BR 10.1.1.210,
i/f Gi0/1, Reason Loss, OOP Reason Loss
```

2% Packet Loss

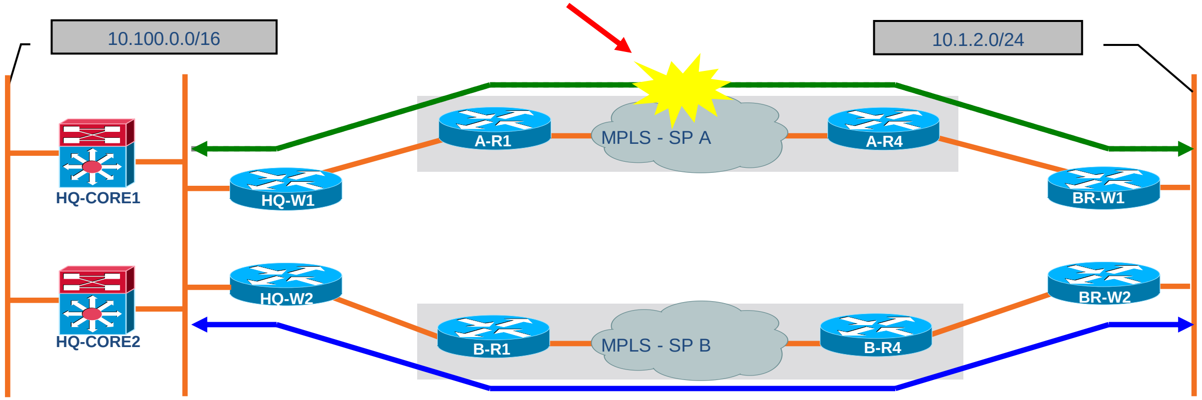

```
BR-W1#
*Mar 3 21:18:53.847: %OER_MC-5-NOTICE: Active ABS Loss OOP Prefix 10.100.100.0/24, loss 4016, BR 10.1.2.120, i/f Fa0/0
*Mar 3 21:18:55.863: %OER_MC-5-NOTICE: Active ABS Loss OOP Prefix 10.100.100.0/24, loss 5025, BR 10.1.2.120, i/f Fa0/0
*Mar 3 21:18:55.867: %OER_MC-5-NOTICE: Route changed Prefix 10.100.100.0/24,
BR 10.1.2.220, i/f Fa0/0, Reason Loss, OOP Reason Loss
```

Performance Routing

Controlling a Configured Prefix

Branch Route Tables – PfR Route Control

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.0.0/16    192.168.101.9                          0 65100 65100 ?
* i                 10.1.2.220               0    100      0 65200 65200 ?
*> 10.100.100.0/24  192.168.101.9                          0 65100 65100 ?
```

```
BR-W2# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
* i10.100.0.0/16    10.1.2.120               0    100      0 65100 65100 ?
*>                  192.168.201.9                          0 65200 65200 ?
*>i10.100.100.0/24  10.1.2.120               0   5000      0 65100 65100 ?
```

PfR Moves the Prefix 10.100.100.0/24 to the Loss Free Path

Branch Route Tables – PfR Route Control

Prefix Controlled by PfR

```
BR-W1# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.100.0.0/16    192.168.101.9                          0 65100 65100 ?
* i                 192.168.201.9            0    100      0 65200 65200 ?
*>i10.100.100.0/24  10.1.2.220               0   5000      0 65200 65200 ?
```

Prefix Controlled by PfR

```
BR-W2# show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
* i10.100.0.0/16    10.1.2.120               0    100      0 65100 65100 ?
*>                  192.168.201.9                          0 65200 65200 ?
*> 10.100.100.0/24  192.168.201.9                          0 65200 65200 ?
```

Performance Routing

Multiple Paths—Select Best Path by Application

- Monitor relevant path characteristics (round trip delay, loss, jitter, …)
  - Example:
  - path A: <5 ms delay, 0% loss, 0% jitter
  - path B: < 50 ms delay, 0% loss, 0% jitter
- Craft a policy to take advantage of unique link characteristics
  - Example:
  - If both paths are free of loss and jitter, then prefer the path with lower delay.
  - However, if jitter begins to exceed 20ms, then prefer jitter free path even if it has increased delay
  - If loss begins to exceed .01% then prefer the lossless path even if it has increased delay or jitter.

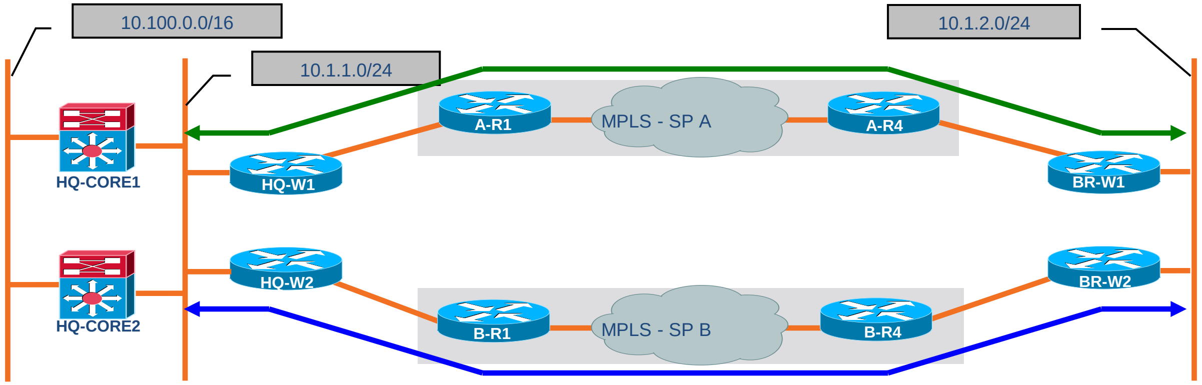

Performance Routing

Application Specific Monitoring

- Characterise the traffic of interest
- Configure “pfr-map” to monitor desired application (and src/dst)
- For performance sensitive traffic, use “mode monitor fast”

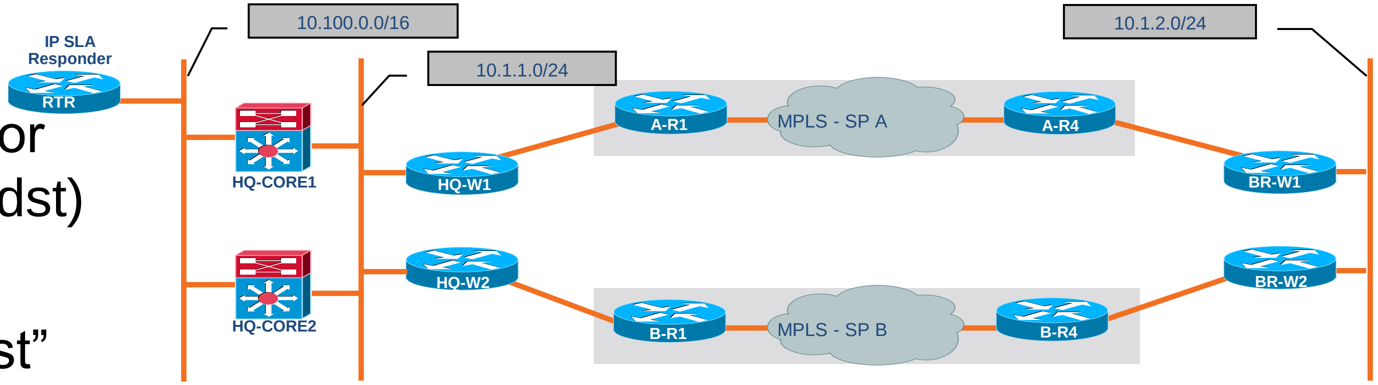

```
BR-W1#
ip access-list extended VOICE-ACL
 permit udp any range 16384 32767 10.100.100.0 0.0.0.255 range 16384 32767
!
pfr-map PFR-HQ-VOICE 10
 match traffic-class access-list VOICE-ACL
 set mode route observe
 set mode monitor fast
 set active-probe jitter 10.100.100.100 target-port 22345
 set probe frequency 2
!
pfr master
 policy-rules PFR-HQ-VOICE
```

Performance Routing

Application Specific Monitoring

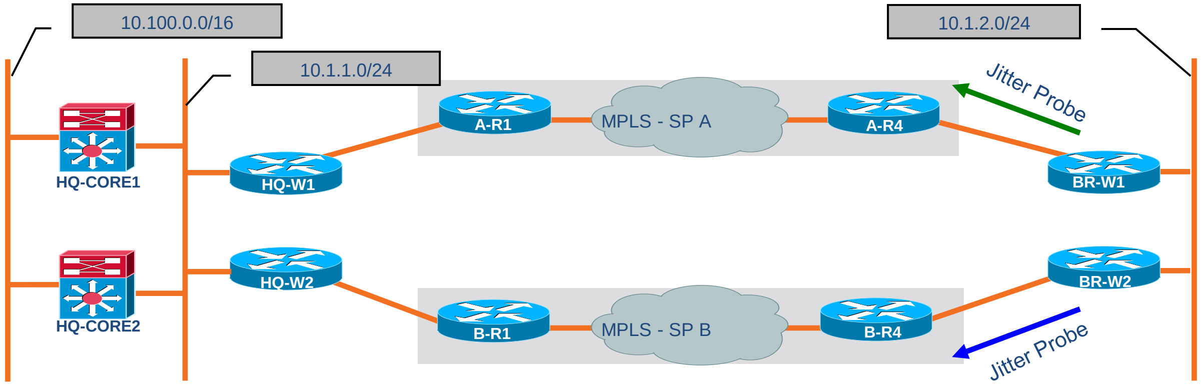

Short Term Jitter

```
BR-W1# show pfr master traffic-class detail
Prefix: 10.100.100.0/24  Protocol: 17  Port: [16384, 32767] [16384, 32767] DSCP: 0
   State: INPOLICY    Time Remaining: @22
   Policy: 10

   Most recent data per exit
   Border          Interface         PasSDly  PasLDly  ActSDly  ActLDly
  *10.1.2.120      Fa0/0                   0        0        3        3
   10.1.2.220      Fa0/0                   0        0       28       28

   Most recent voice data per exit
   Border          Interface         ActSJit  ActPMOS  ActSLos  ActLLos
  *10.1.2.120      Fa0/0                   0        0        0        0
   10.1.2.220      Fa0/0                   0        0        0        0
```

Short/Long Term Delay

Short/Long Term Loss

Performance Routing

Controlling a Configured Application

- Compare loss characteristics of each exit (within 5% considered “same”)
- Next, compare jitter characteristics of each exit (within 5% considered “same”)
- Finally, compare delay characteristics of each exit (within 5% considered “same”)
- Maximum packet loss for an exit is limited to 100 (packets per million) or 0.01%
- Maximum jitter for an exit is 20 ms
- Maximum delay for an exit is 100 ms
- Unreachable policy is always considered “highest priority”

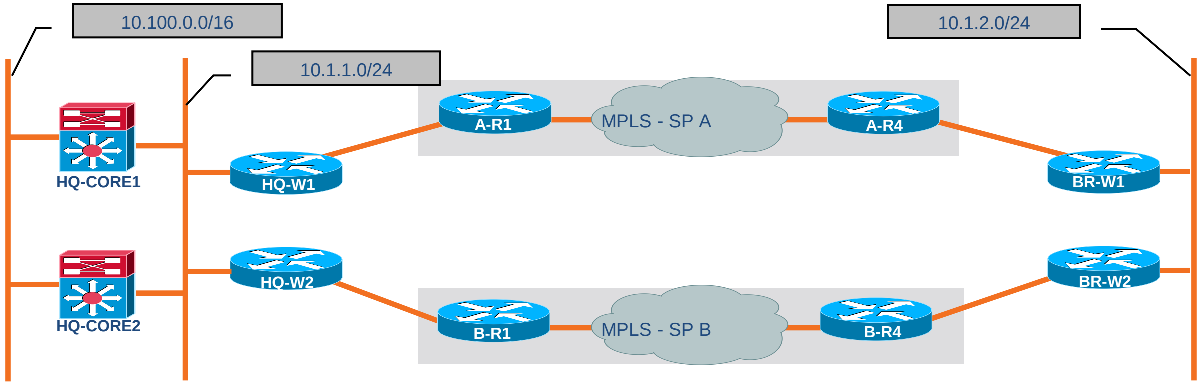

```
BR-W1#

pfr-map PFR-HQ-VOICE 10
 no set resolve utilization
 set resolve loss priority 1 variance 5
 set resolve jitter priority 2 variance 5
 set resolve delay priority 3 variance 5
 set loss threshold 100
 set jitter threshold 20
 set delay threshold 100
 !
 set holddown 90
 set periodic 90
 set mode select-exit best
 !
 set mode route control
!
pfr master
 policy-rules PFR-HQ-VOICE
```

Performance Routing

Controlling a Configured Application

Chosen Exit

Points to A-R4 (via External)

```
BR-W1# show ip policy
Interface    Route map
Fa0/1      OER-03/04/09-17:43:17.387-F-OER (Dynamic)

BR-W1# show route-map dynamic
route-map OER-03/04/09-17:43:17.387-F-OER, permit, sequence 0, identifier 1200584152
 Match clauses:
  ip address (access-lists): oer#15
 Set clauses:
  ip next-hop 192.168.101.9
  interface FastEthernet0/0
 Policy routing matches: 1040 packets, 7690 bytes
 Current active dynamic routemaps = 1

BR-W1# show ip access-lists dynamic
Extended IP access list oer#15
  1073741823 permit udp any range 16384 32767 10.100.100.0 0.0.0.255 range 16384 32767
```

Inactive Exit

Points to BR-W1 (via Internal)

```
BR-W2# show ip policy
Interface    Route map
Fa0/1      OER-03/04/09-17:43:17.979-22-OER (Dynamic)

BR-W2# show route-map dynamic
route-map OER-03/04/09-17:43:17.979-22-OER, permit, sequence 0, identifier 1194973244
 Match clauses:
  ip address (access-lists): oer#15
 Set clauses:
  ip next-hop 10.1.2.120
  interface FastEthernet0/1
 Policy routing matches: 0 packets, 0 bytes
 Current active dynamic routemaps = 1

BR-W2# show ip access-lists dynamic
Extended IP access list oer#15
  1073741823 permit udp any range 16384 32767 10.100.100.0 0.0.0.255 range 16384 32767
```

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
- Design and Deployment
  - MPLS Dual Carrier
  - MPLS + Internet
- Final Wrap Up

DUAL WAN (MPLS + Backup)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- Headquarters WAN Edge
  - W1 learns Branch route via eBGP
  - W2 learns Branch route via EIGRP
- Headquarters Core
  - W1 redistributes eBGP into EIGRP, results in EIGRP external
  - W2 does not require redistribution, results in EIGRP internal
  - Core1, Core2 install Branch route via W2

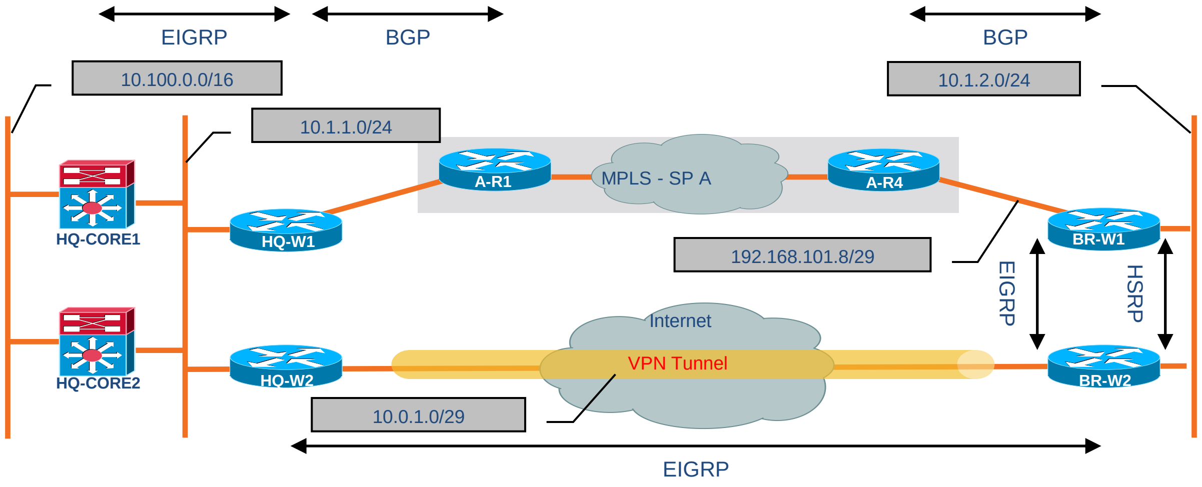

HQ to Branch Traffic Flows Across Tunnel

```
HQ-W1# show ip route
B    10.1.2.0/24 [20/0] via 192.168.101.2, 05:24:01
```

```
HQ-W2# show ip route
D    10.1.2.0/24 [90/26882560] via 10.0.1.2, 00:00:04, Tunnel1
```

```
HQ-CORE1# show ip route
D    10.1.2.0/24 [90/26882816] via 10.1.1.210, 00:02:32, Vlan10
```

DUAL WAN (MPLS + Backup)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- Branch WAN Edge
  - W1 learns HQ route via eBGP
  - W2 learns HQ route via EIGRP
  - No redistribution configured
  - HSRP Primary is on W1

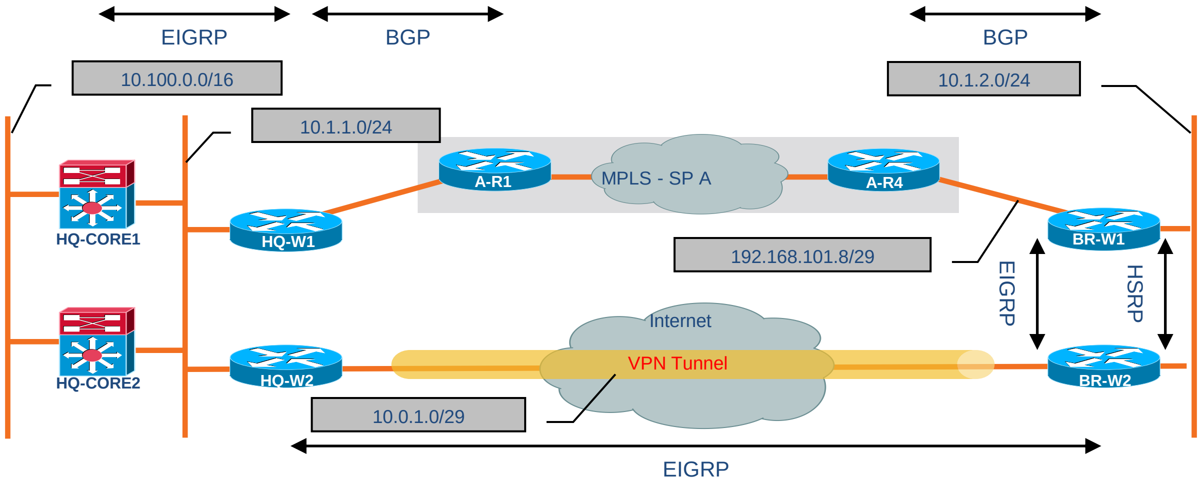

Branch to HQ Traffic Flows Across MPLS

```
BR-W1# show ip route
B    10.100.100.0/24 [20/0] via 192.168.101.9, 04:48:58
B    10.100.200.0/24 [20/0] via 192.168.101.9, 03:44:06
```

```
BR-W2# show ip route
D    10.100.100.0/24 [90/26882816] via 10.0.1.1, 00:10:56, Tunnel1
D    10.100.200.0/24 [90/26882816] via 10.0.1.1, 00:10:57, Tunnel1
```

```
BR-W1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Pri P State   Active      Standby     Virtual IP
Fa0/1     1  110 P Active local      10.1.2.220   10.1.2.1
```

DUAL WAN (MPLS + Backup)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- How to force HQ to Branch traffic across MPLS (primary)?
  - Adjust administrative distance of EIGRP routes learned via tunnel
  - Ensure new distance is higher than that of EIGRP external (170)

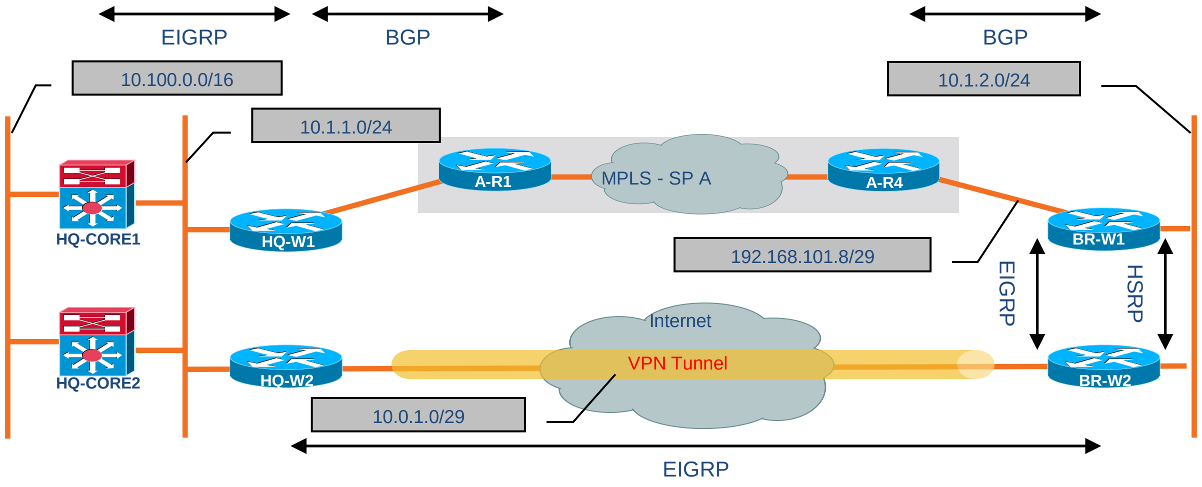

```
HQ-W2#
router eigrp 65110
 network 10.0.1.0 0.0.0.7
 distance 195 10.0.1.0 0.0.0.7
```

Now:

HQ to Branch Traffic Flows Across MPLS

```
HQ-W1# show ip route
B    10.1.2.0/24 [20/0] via 192.168.101.2, 05:24:01
```

```
HQ-W2# show ip route
D EX  10.1.2.0/24 [170/261120] via 10.1.1.110, 00:07:25, GigE0/0
```

```
HQ-CORE1# show ip route
D EX  10.1.2.0/24 [170/258816] via 10.1.1.110, 00:08:44, Vlan10
```

DUAL WAN (MPLS + Backup)

MPLS Failure

- Failure within MPLS cloud
- Worst Case
  - Primary dependency is BGP timers
  - Results in end to end convergence time as long as BGP Holdtime
  - Could be much lower with BGP tuning and use of BFD

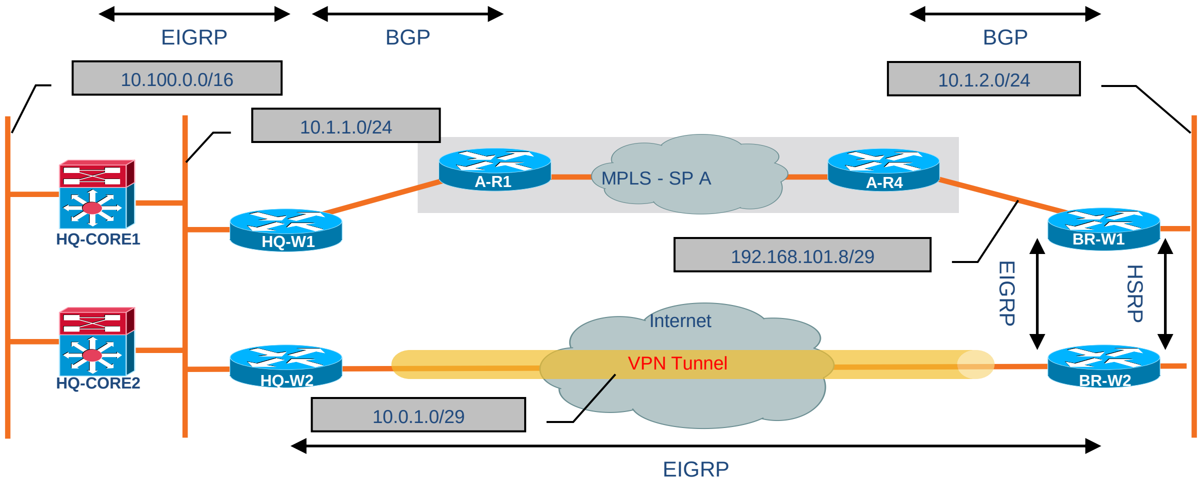

After Failure:

HQ to Branch Traffic Flows Across Tunnel

HQ Route Tables

```
HQ-W2# show ip route
D    10.1.2.0/24 [195/26882560] via 10.0.1.2, 00:06:46, Tunnel1
```

```
HQ-CORE1# show ip route
D    10.1.2.0/24 [90/26882816] via 10.1.1.210, 00:09:18, Vlan10
```

DUAL WAN (MPLS + Backup)

MPLS Failure

- Failure within MPLS cloud
- Suboptimal routing at Branch
  - HSRP primary remains unchanged at BR-W1
  - Could use EOT and move HSRP primary to BR-W2

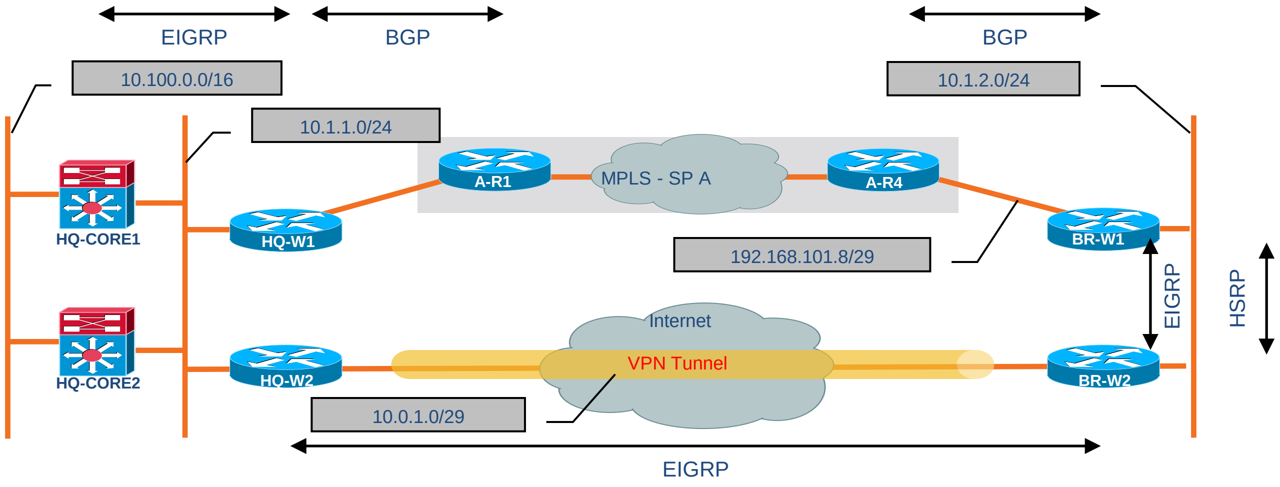

After Failure:

Branch to HQ Traffic Flows Across Tunnel

Branch Route Tables

```
BR-W1# show ip route
D    10.100.100.0/24
       [90/26885376] via 10.1.2.220, 00:22:42, FastEthernet0/1
D    10.100.200.0/24
       [90/26885376] via 10.1.2.220, 00:22:42, FastEthernet0/1
```

```
BR-W2# show ip route
D    10.100.100.0/24 [90/26882816] via 10.0.1.1, 01:08:44, Tunnel1
D    10.100.200.0/24 [90/26882816] via 10.0.1.1, 01:08:45, Tunnel1
```

DUAL WAN (MPLS + Backup)

PE-CE Protocol: BGP, Tunnel Protocol: EIGRP

- Options for PfR with Multiple Routing Protocols
  - PIRO – Protocol Independent Route Optimisation
  - EIGRP Route Control (requires EIGRP only as route source for WAN)

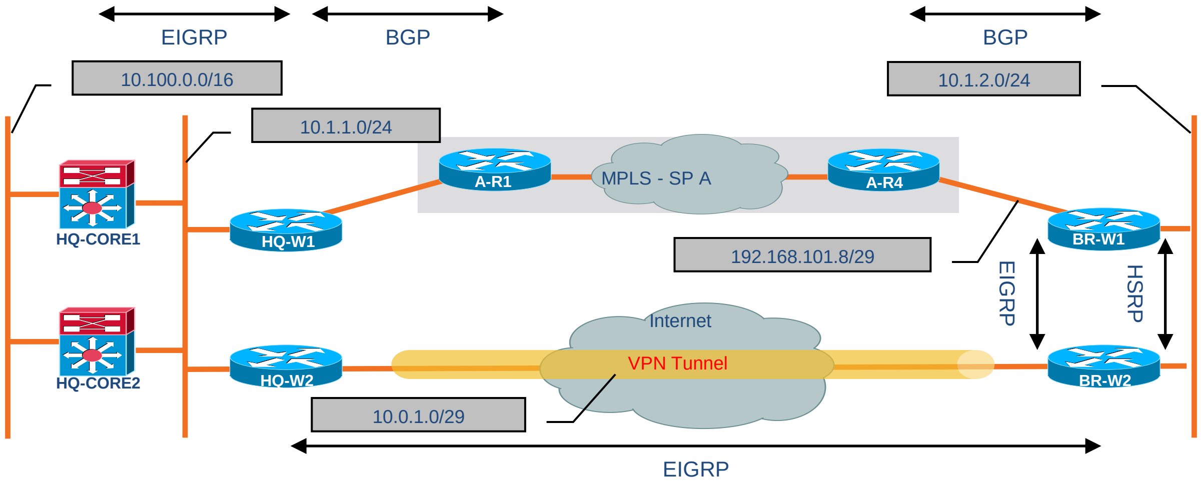

PIRO [12.4(24)T]

Supports Application Specific Monitoring (Dynamic Policy Routing)

Supports Hybrid BGP/EIGRP Topology with “mode route protocol pbr” - Requires 15.0(1)M4

Requires BR-BR Direct Neighbour Relationship

EIGRP Route Control [15.0(1)M]

Requires iBGP & EIGRP between border routers

Agenda

- Introduction
- Cisco IOS and IP Routing
- Convergence Techniques
- Design and Deployment
- Final Wrap Up
  - – Key Takeaways

Key Takeaways

- Outages can manifest in many different ways. Network design should be based on application requirements to survive various outages.
- Cisco IOS has inherent load sharing capabilities. Analyse your network topology and use these to your advantage.
- End-to-end convergence time is a critical metric. Understand how localised topology changes affect end-to-end resiliency.
- Multiple links/paths not only increase network reliability but can improve application performance.

Key Takeaways

- IP SLA based monitoring can detect outage types that are virtually undetectable by traditional “hello based” techniques.
- Performance Routing permits path selection based on current real time characteristics.
- Most effective network designs incorporate a combination of convergence techniques

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

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