Advances in EIGRP

BRKIPM-3008

Iain Thomson

(ithomson@cisco.com)

Cisco Networkers 2007

HOUSEKEEPING

- We value your feedback, don't forget to complete your online session evaluations after each session and complete the Overall Conference Evaluation which will be available online from Friday.
- Visit the World of Solutions on Level -01!
- Please remember this is a 'No Smoking' venue!
- Please switch off your mobile phones!
- Please remember to wear your badge at all times including the Party!
- Do you have a question? Feel free to ask them during the Q&A section or write your question on the Question form given to you and hand it to the Room Monitor when you see them holding up the Q&A sign.

Agenda

- **Scaling Enhancements**
- **Neighbour Enhancements**
- **Routing Enhancements**

EIGRP Feature Status

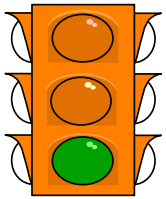

Completed and either *Pending Release* or *Released* in Cisco IOS®

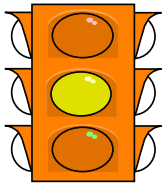

```
/*
 * Under Development
 */
```

Scaling Enhancements

Scaling Enhancements

- **EIGRP Stub**
- **Stub Leaking**
- **Summary Leaking**
- **Summary Only**
- **Stub co-existence**
- **Transport enhancements**

EIGRP Stub Operation

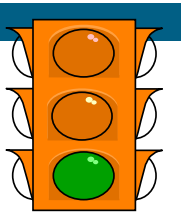

- Link down event means a local lookup for loop free paths
- If this fails we go Active for that destination and Query
- If B does not have a path it sends a Reply

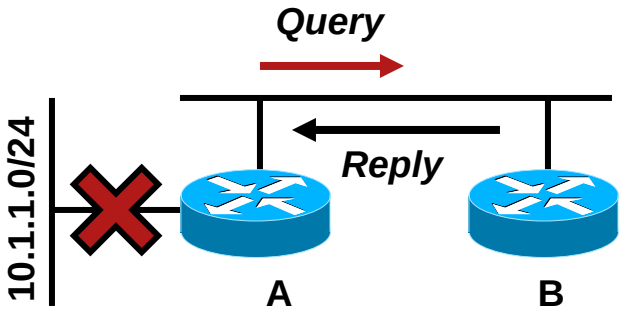


```
A#show ip eigrp topology
....
P 10.1.1.0/24, 1 successors,
       via Connected, Ethernet0/0
A#
```

```
A#show ip eigrp events
Event information for AS 65535:
1  NDB delete: 10.1.1.0/24 1
....
```

```
A#show ip eigrp events
Event information for AS 65535:
....
12 Active net/peers: 10.1.1.0/24 1
14 FC not sat Dmin/met: 4294967295..
15 Find FS: 10.1.1.0/24 128256..
18 Conn rt down: 10.1.1.0/24..
```

EIGRP Stub Operation

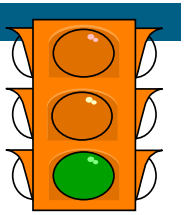

- Router A sends Queries to the spokes so they could be used as next-hops to carry traffic
- Building and processing Queries and Replies is a burden on all devices
- Configure the spokes using the EIGRP Stub feature

```
spoke#config t
spoke(config)#router eigrp 65535
spoke(config-router)#eigrp stub
spoke(config-router)#
```

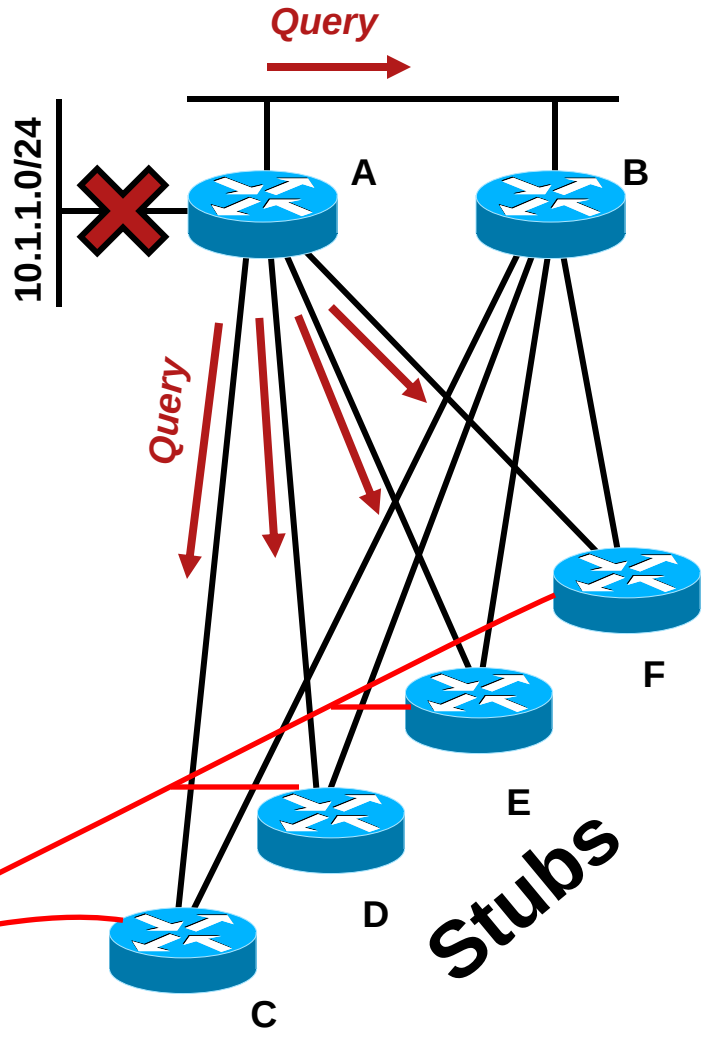

EIGRP Stub Operation

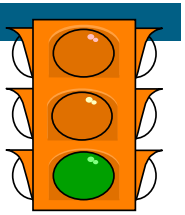

- Neither A or B will Query stubs; reducing the number of Queries to 1
- Stub has a variety of configurable options

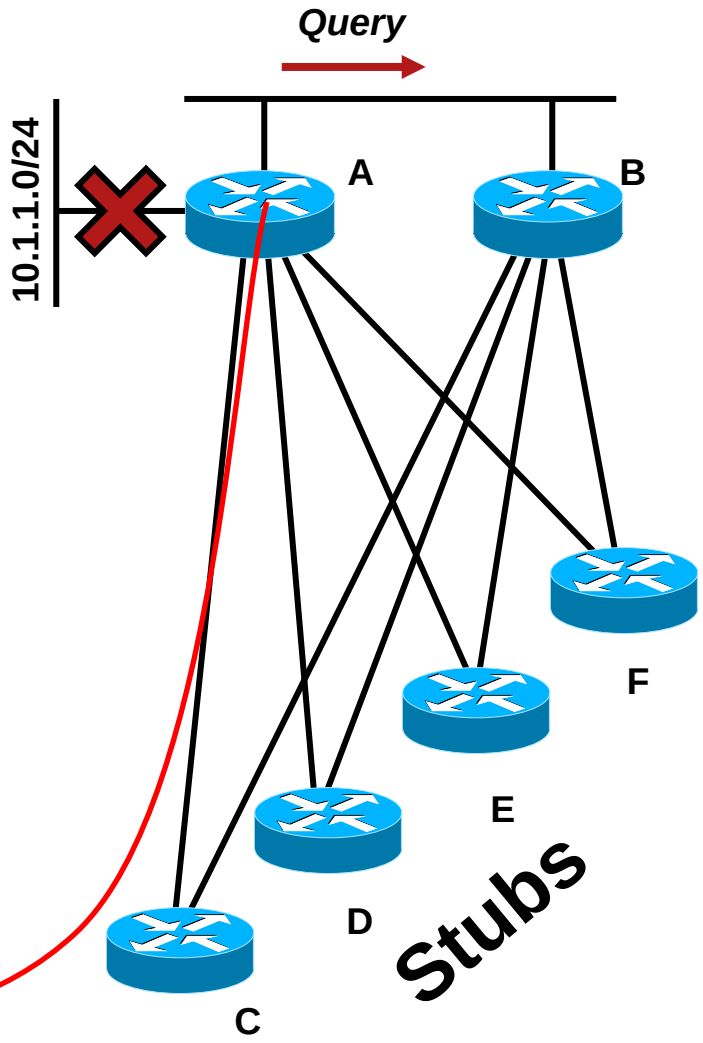


```
Spoke(config-router)#eigrp stub ?
        connected       Do advertise connected routes
        leak-map        Allow dynamic prefixes based
       receive-only    Set IP-EIGRP as receive only
        redistributed   Do advertise redistributed
        static          Do advertise static routes
        summary         Do advertise summary routes
```

```
A#show ip eigrp neighbor de
Stub Peer Advertising ( CONNECTED SUMMARY ) Routes
   Suppressing queries
```

Stub Leaking

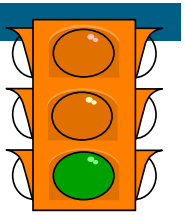

In a single remote site with two routers we want to mark the entire site as a Stub

- C and D are Stub
- A and B advertise only a default to C and D
- Because C and D are Stub they do not talk to each other and there are no advertisements

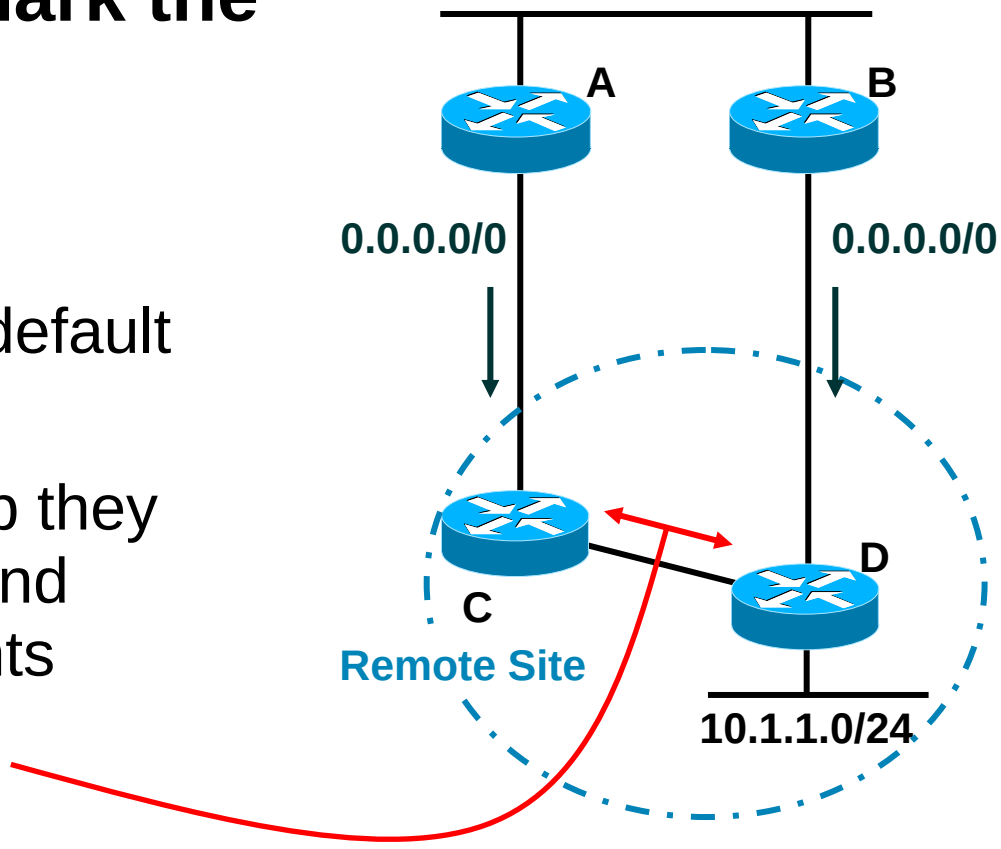

Stub Leaking

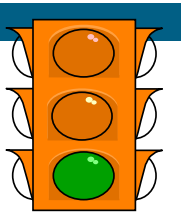

- The link from B to D fails
- Network 10.1.1.0/24 cannot be reached from A

  D isn't advertising 10.1.1.0/24 to C, since D is a Stub
- D can't reach A, or anything behind A

  C is not advertising the default route to D, since C is a Stub

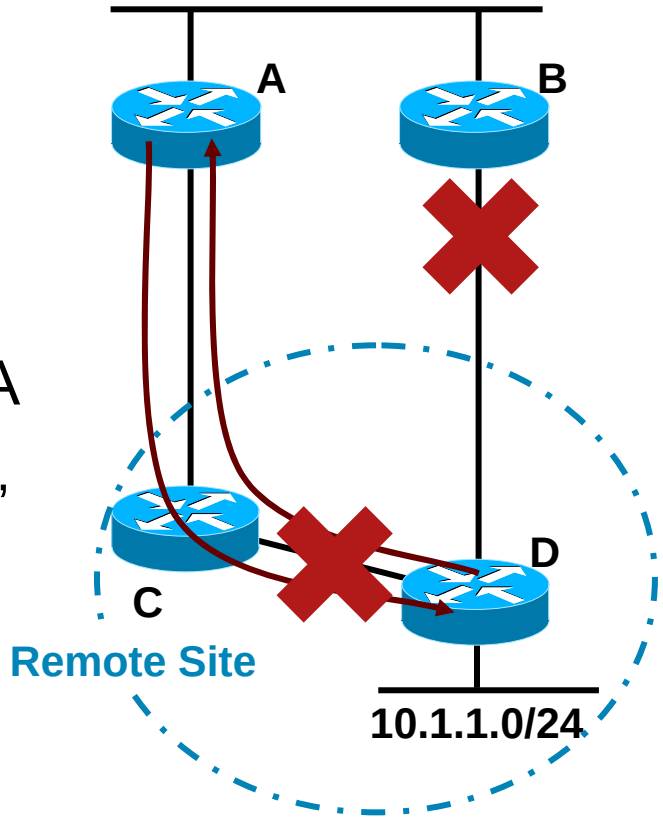

Stub Leaking

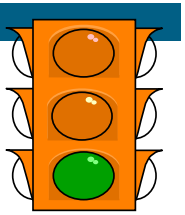

- We want that C and D advertises a subset of their learned routes, even though they are both Stub
- Stub leaking is the solution

```
router eigrp 100
 eigrp stub leak-map LeakList

route-map LeakList permit 10
   match ip address 1
   match interface e0/0

route-map LeakList permit 20
   match ip address 2
   match interface e1/0

access-list 1 permit 10.1.1.0
access-list 2 permit 0.0.0.0
```

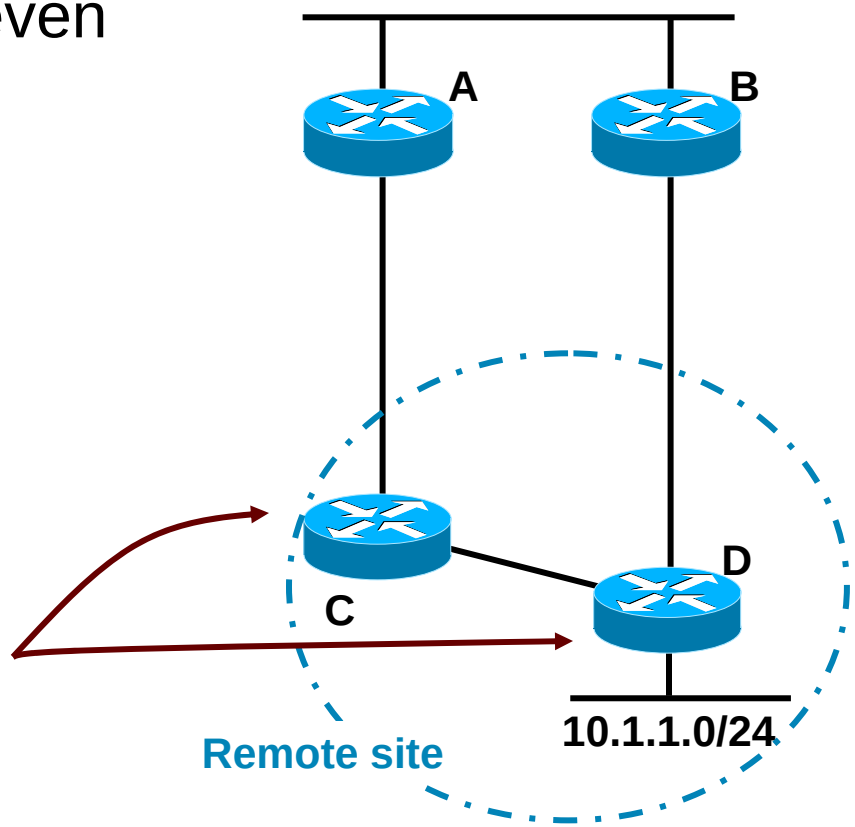

Stub Leaking

CSCec80943

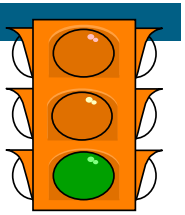

- The link from B to D fails
- D is advertising 10.1.1.0/24 to C, and from C to A, so 10.1.1.0/24 is still reachable
- C is leaking the default route to D, so D can still reach the rest of the network through D
- A and B will still not query towards the remote site as C and D are stubs

- Available 12.3(10.02)T

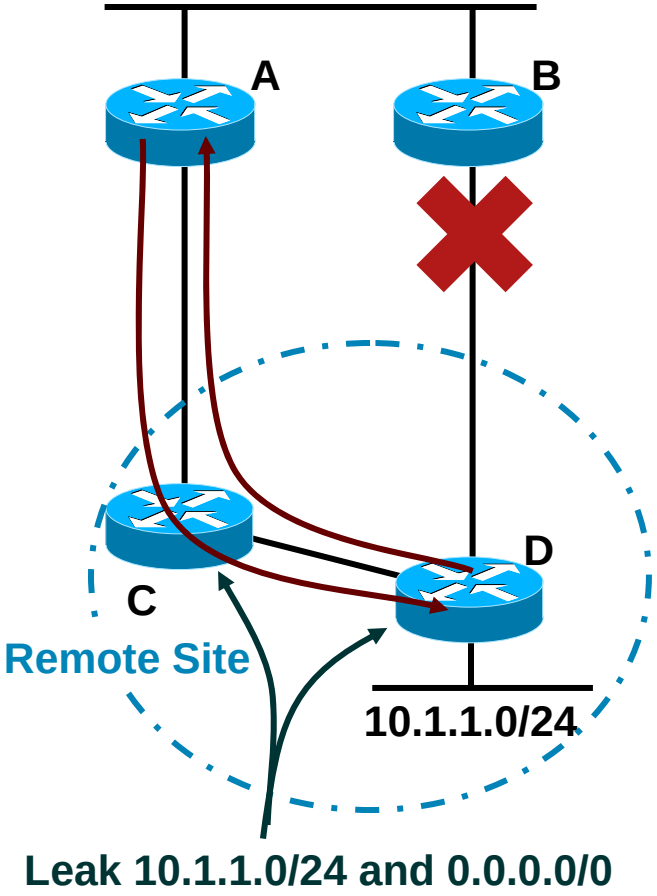

Summary Leaking

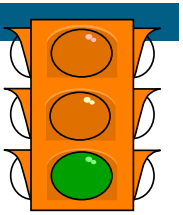

- Good design implies C should receive as few routes as possible
- We still optimally route to 10.1.1.0/24 and 10.1.2.0/24

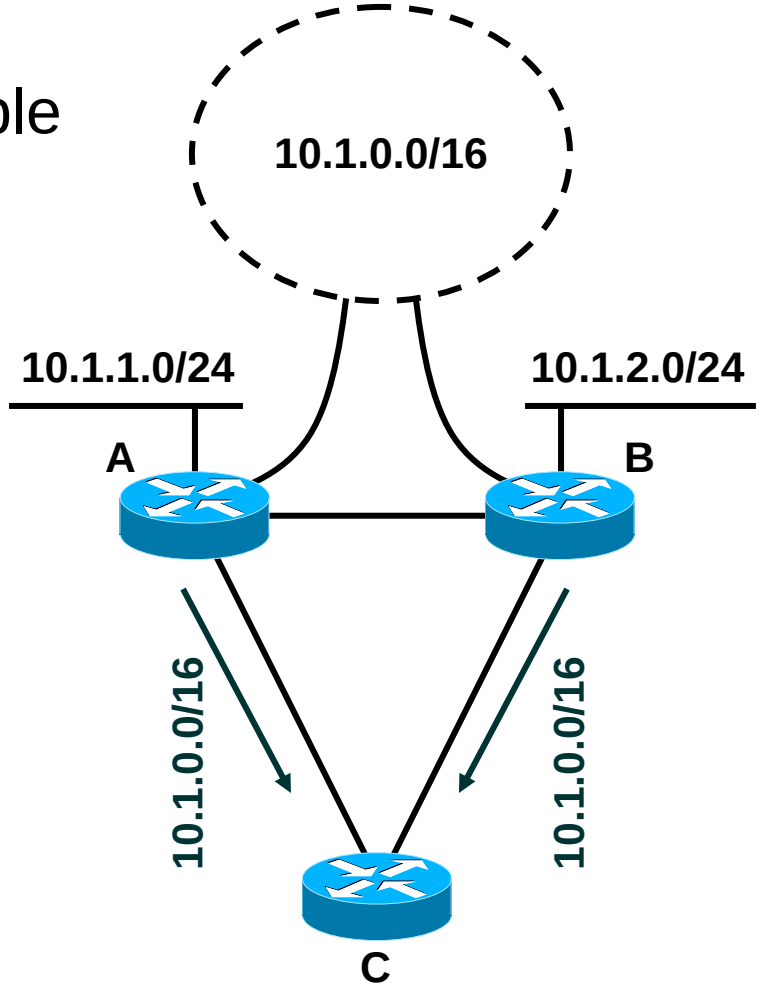

- We could use a combination of static routes and route filters to advertise both 10.1.0.0/16 and the more specific to C
- However, this is difficult for customers to maintain

Summary Leaking

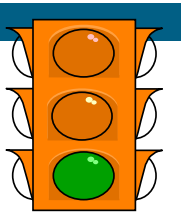

```
router eigrp 100
 redistribute static route-map aggroutes
 default-metric 1000 1 255 1 1500
 distribute-list 20 out serial0/0
!
ip route 10.1.0.0 255.255.0.0 null0
!
route-map agg-routes permit 10
 match ip address 10
 match interface serial 0/0
!
access-list 10 permit 10.1.0.0 0.0.255.255
access-list 20 permit 10.1.1.0 0.0.255.255
```

```
router eigrp 100
 redistribute static route-map aggroutes
 default-metric 1000 1 255 1 1500
 distribute-list 20 out serial0/0
!
ip route 10.1.0.0 255.255.0.0 null0
!
route-map agg-routes permit 10
 match ip address 10
 match interface serial 0/0
!
access-list 10 permit 10.1.0.0 0.0.255.255
access-list 20 permit 10.1.2.0 0.0.255.255
```

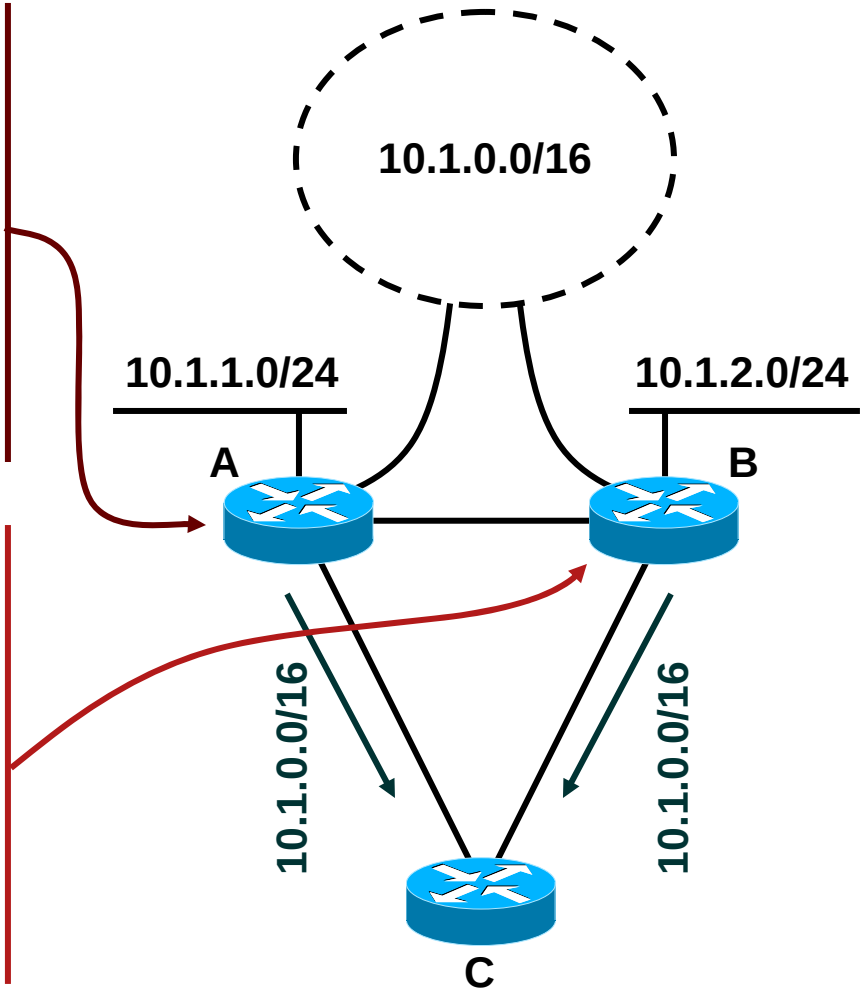

Summary Leaking

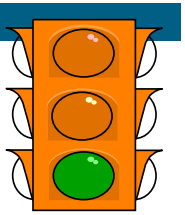

- You can also use a pair of summaries
- You need to “float” the 10.1.1.0/24 and 10.1.2.0/24 summaries
- This could remove the dynamic nature of the longer prefix optimal route advertisements

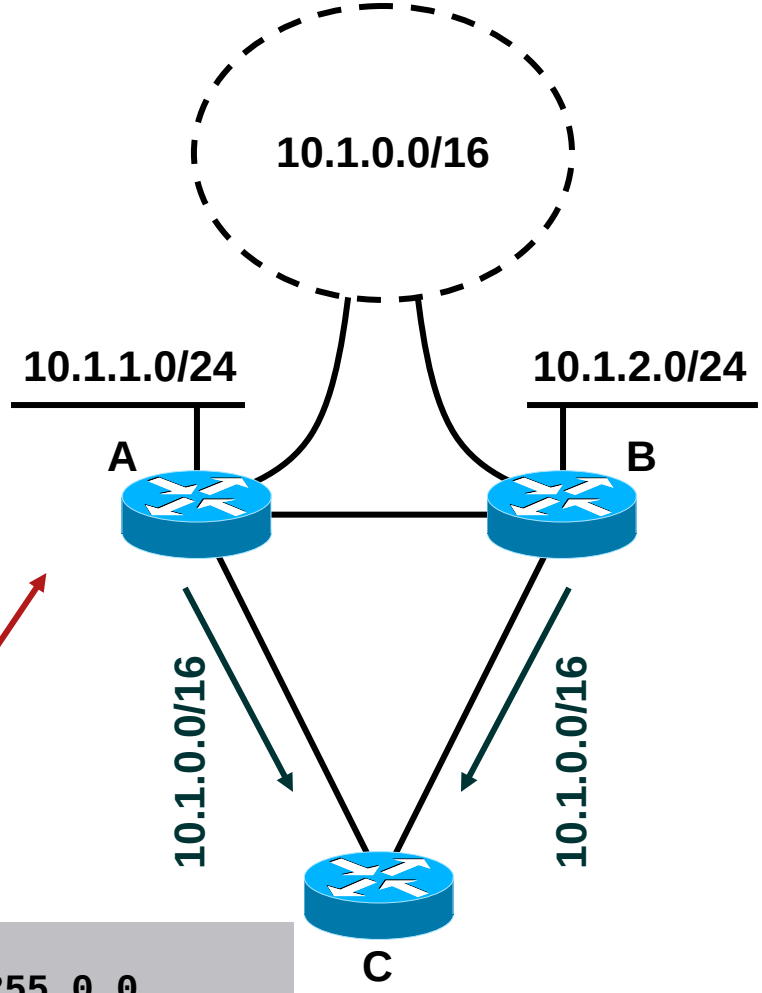


```
interface serial 0/0
 ip summary-address 10.1.0.0 255.255.0.0
 ip summary-address 10.1.1.0 255.255.255.0 200
```

Summary Leaking

CSCed01736

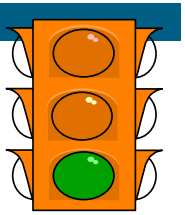

- The simplest way to handle this is to configure a leak list on the summary route
- Leak lists for summaries are available 12.3(11.01)T

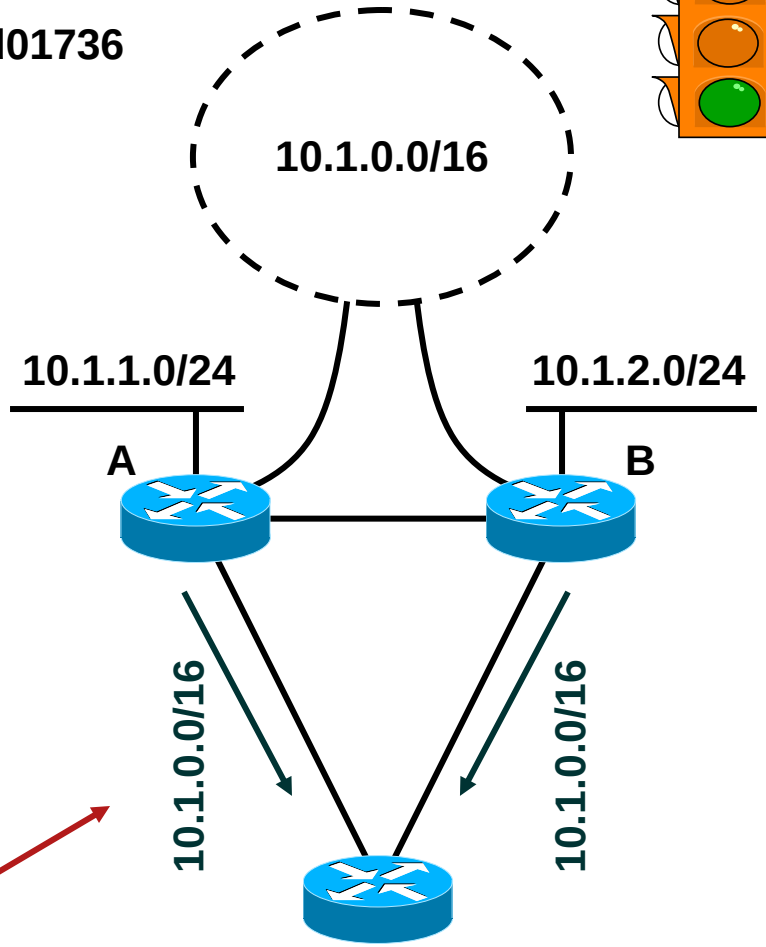


```
route-map LeakList permit 10
   match ip address 1
!
access-list 1 permit 10.1.1.0
!
interface Serial0/0
 ip summary-address eigrp 1 10.1.0.0 255.255.0.0 leak-map LeakList
```

Summary Only

CSCec10166

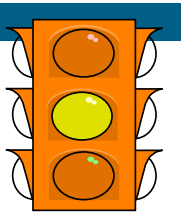

- We would like to advertise a single summary for networks with servers attached to them from A and B to C (in the range 10.1.0.0/16)
- We don't want to advertise infrastructure links to C (in the range 192.168.0.0/16)
- We could do this with a summary and filter, but this is cumbersome

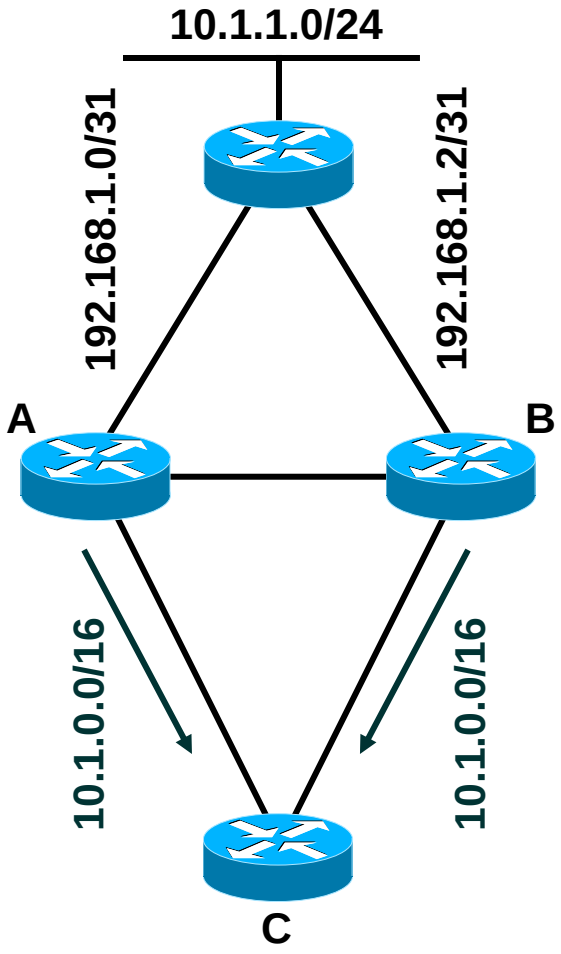

Summary Only

CSCec10166

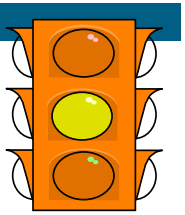

```
router eigrp 100
 distribute-list 10 out serial 0/0
 ....
!
access-list 10 permit 10.1.0.0 0.0.255.255
!
interface serial 0/0
 ip summary-address eigrp 100 10.1.0.0 255.255.0.0
```

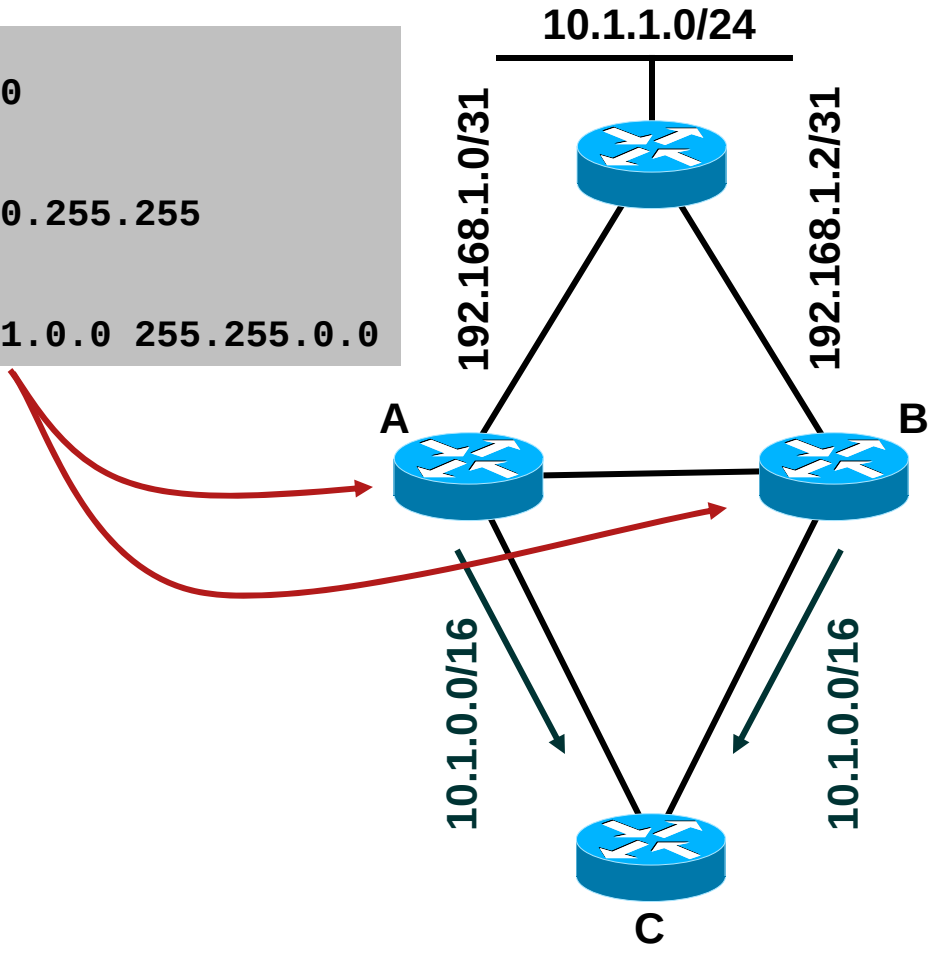

Summary Only

CSCec10166

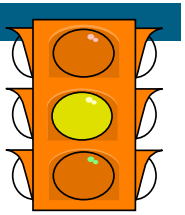

An easier way would be to define a summary that automatically blocks all routes which are not a component of the summary

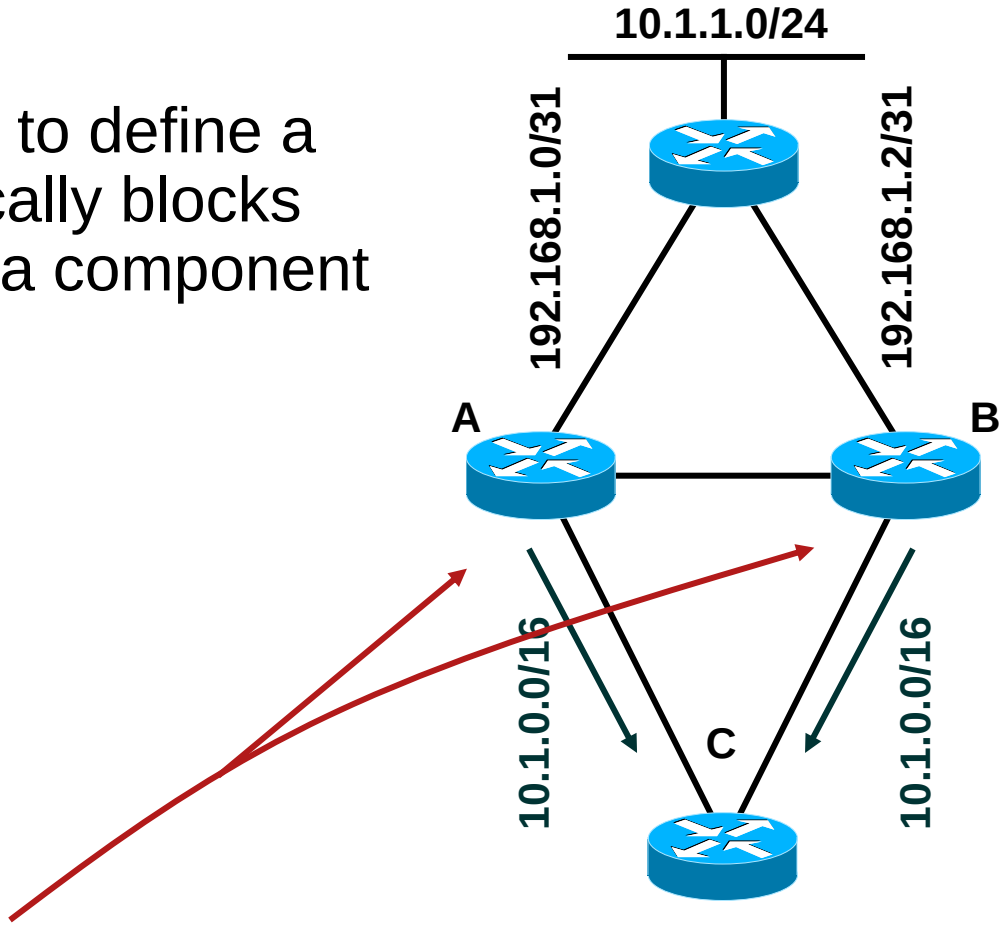


```
A(config)#interface serial 0/0
A(config-if)#ip summary-address eigrp 100 10.1.0.0 255.255.0.0
A(config-if)#ip summary-only eigrp 100
```

Stub co-existence

CSCdx74716

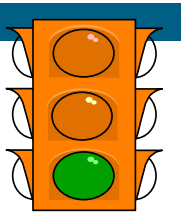

- If B and C were not Stubs A would send Queries as normal
- Once Stub is configured on B and C the restriction is all peers off router A's interface must be Stub to work in a Multi-access scenario
- If a new peer comes up that is not stub, router A's interface starts to send queries
- The restriction prevents deployment of multiple hubs

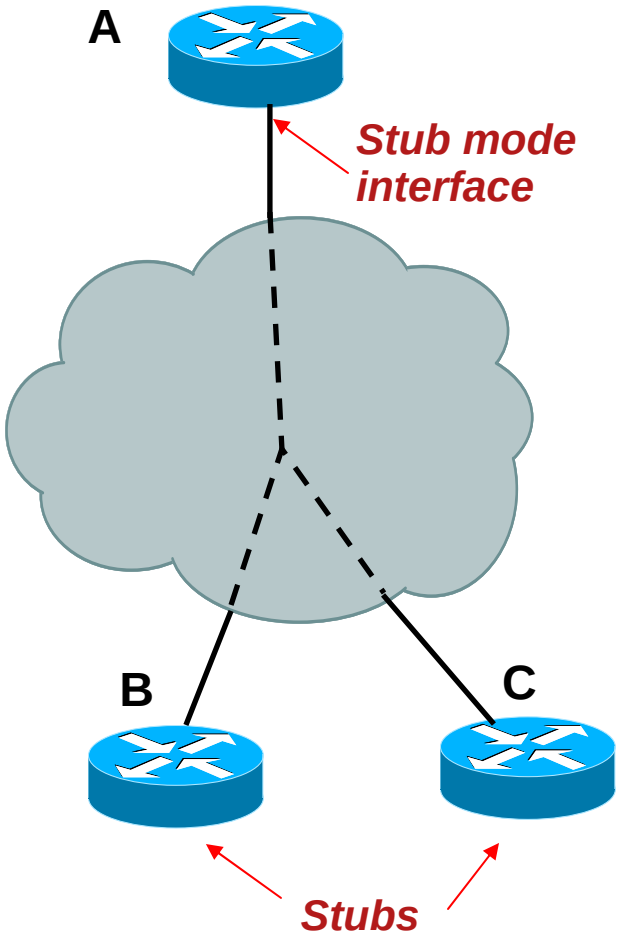

Stub co-existence

CSCdx74716

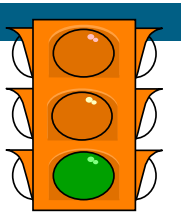

```
nw2007-A#sh ip eigrp int de ser2/0
…
  Interface has all stub peers
nw2007-A#
```

```
nw2007-A#sh ip eigrp neighbor de
…
Stub Peer Advertising ( CONNECTED SUMMARY ) Routes
   Suppressing queries
```

```
nw2007-B(config)#router eigrp 65535
nw2007-B(config-router)#eigrp stub
nw2007-B(config-router)#
```

```
nw2007-C(config)#router eigrp 65535
nw2007-C(config-router)#no eigrp stub
nw2007-C(config-router)#
```

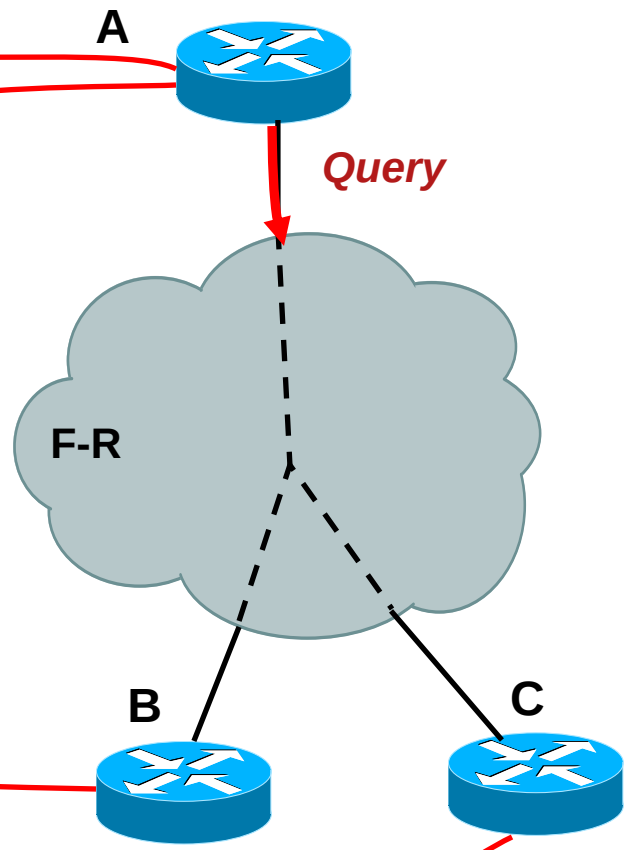

Stub co-existence Ethernet Example

CSCdx74716

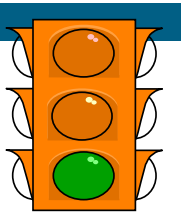

- You could use static neighbours but a router with a static neighbor rejects multicast packets from non-static neighbours
- The hub does not have to use the static neighbor command and still form adjacencies with the stub static neighbours
- With the enhancement a static neighbor now accepts and transmits multicasts

```
router eigrp 100
neighbor 10.2.25.1 ethernet0/0
neighbor 10.2.25.2 ethernet0/0
```

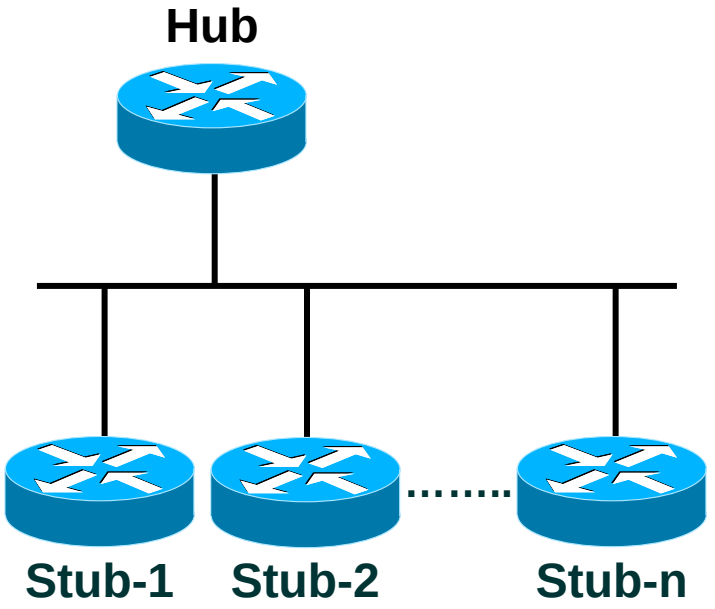

Stub co-existence Summary

CSCdx74716

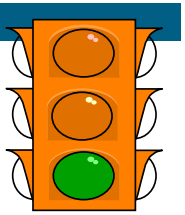

- Capability to send a query to a neighbour subset was added allowing for coexistence of both stub and non-stub neighbours on the same interface
- A hub router can now send queries to the non-stub neighbors and suppress queries to the stub neighbors
- A query can be sent unicast or multicast using conditional receive mode
- For packet transmission efficiency, queries are sent unicast if the number of non-stub peers are less than five, or less than 10% of the total number of peers on the interface; otherwise, multicast is used

Transport Enhancements

CSCei03733

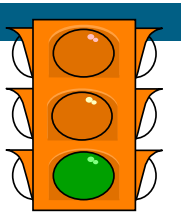

- Formerly, when bringing up an interface that has hundreds of neighbors, EIGRP may take a long time to converge
- Symptoms being neighbour resets, packet retransmission timeout, Stuck-in-Active, Hello hold time expired
- The problem is often seen with DMVPN/GRE tunnel interfaces

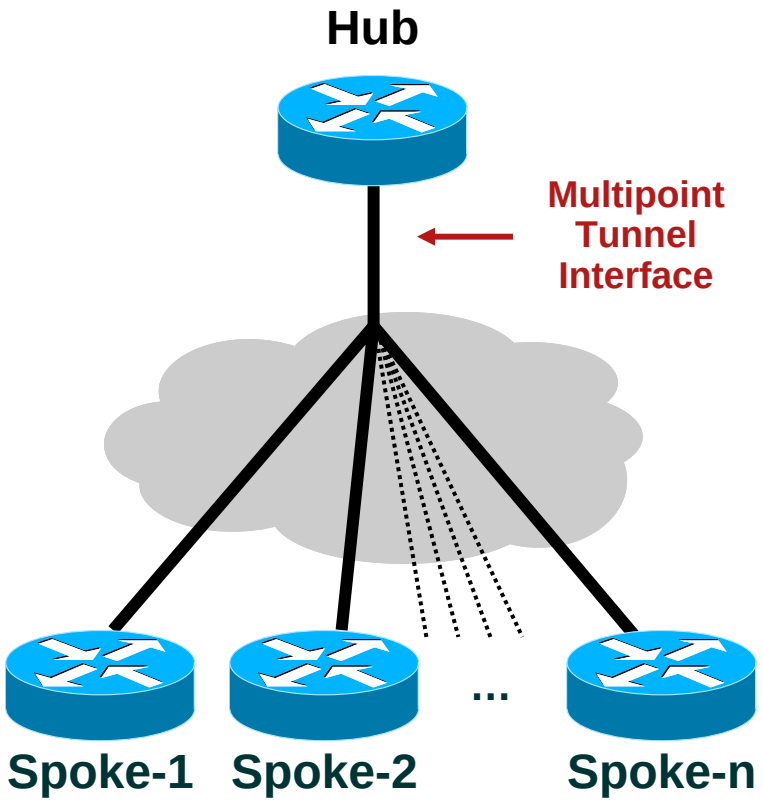

Transport Enhancements

- The new EIGRP transport module has a number of enhancements to speedup convergence and increase neighbour scaling
- The minimum packet pacing interval can be lowered to a minimum value of 1 ms by using the bandwidth or bandwidth percentage commands

```
nw2007(config-if)#ip bandwidth-percent eigrp 65535 …
```

- On a fast interface or a tunnel interface which has unreliable pacing value, EIGRP packet transmissions can also be driven using the neighbor acknowledgements (ACK-driven)
- Startup Update Packets exchanged at neighbor startup can now be sent using multicast

Neighbour Enhancements

Neighbour Enhancements

- **Graceful Restart**
- **Neighbour Startup Enhancements**
- **BFD**
- **Graceful Shutdown**

Graceful Restart

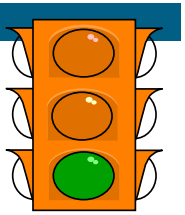

- Configuration changes reset neighbours:
  - Split-horizon, Summaries,
  - Filter changes (Distribute-lists)
  - Others
- Reset neighbours means:
  - CPU cycles
  - DUAL events
  - Neighbour re-establishment processing

Graceful Restart
Before

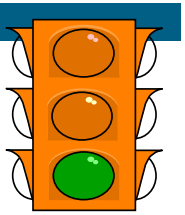

- Configuring a distribute-list

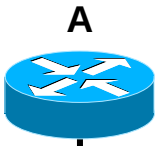


```
*Aug 27 13:06:26.758: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65535:
Neighbor 10.1.1.1 (Serial0/0) is down: peer restarted
*Aug 27 13:06:27.976: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65535:
Neighbor 10.1.1.1 (Serial0/0) is up: new adjacency
```

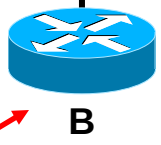


```
B#config t
B(config)#router eigrp 65535
B(config-rtr)#distribute-list 100 in serial 0/0
```

Graceful Restart
After

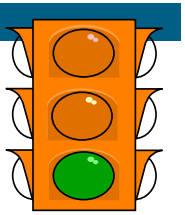

- **B does not support GR or peer resynchronization**

```
router-a#clear ip eigrp neighbor <neighbor_B> soft
%DUAL-5-NBRCHANGE: Neighbor <b> is resync: manually cleared
%DUAL-5-NBRCHANGE: Neighbor <b> is down: peer restarted
%DUAL-5-NBRCHANGE: Neighbor <b> is up: new adjacency

router-b#
%DUAL-5-NBRCHANGE: Neighbor <a> is down: peer restarted
%DUAL-5-NBRCHANGE: Neighbor <a> is up: new adjacency
```

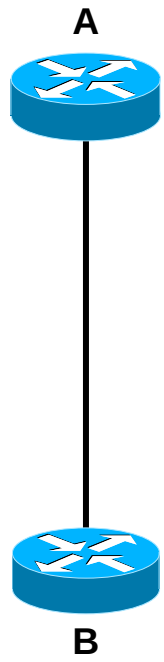

- **B supports Graceful Restart**

```
router-a#clear ip eigrp neighbor <b> soft
%DUAL-5-NBRCHANGE: Neighbor <b> is resync: manually cleared

router-b#
%DUAL-5-NBRCHANGE: Neighbor <a> is resync: peer nsf-restarted
```

- **B supports GR and peer Resynchronization**

```
router-a#clear ip eigrp neighbor <b> soft
%DUAL-5-NBRCHANGE: Neighbor <b> is resync: manually cleared

router-b#
%DUAL-5-NBRCHANGE: Neighbor <a> is resync: peer graceful restart
```

- Available in 12.3(12.06)T

Neighbour Startup Enhancements

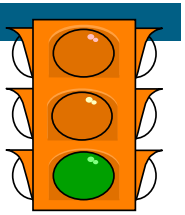

- EIGRP requires both multicast and unicast to establish and maintain neighbour relationship
- Unidirectional links or congestion could cause unstable neighbor relationship, SIA, and slow convergence
- Neighbor initialization enhancements are used to protect an EIGRP network when bringing up an unstable neighbour
- BFD and Graceful shutdown support are used to quickly terminate an unstable neighbour relationship

Neighbour Startup Enhancements Before

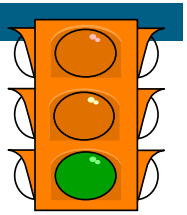

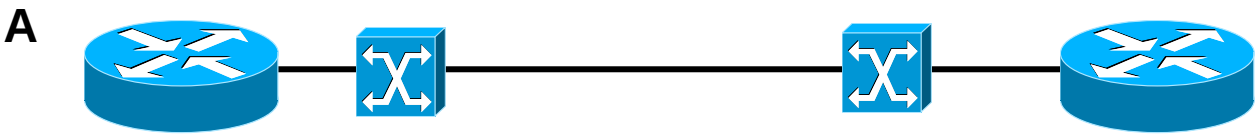

Hello	→	New peer, pending init+ack
No peer, packet dropped	←	Topo update+init
No peer, packet dropped	← - -	Retransmit
New peer, pending init+ack	←	Hello
Topo update+init	→	Accept routes, pending ack
Full peer, accept routes	←	Topo update+init+ack
Ack	→	Full peer

Neighbour Startup Enhancements
Retransmission example

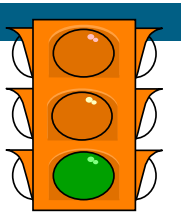

A receives the multicast HELLO from B and retransmits updates with the INIT bit set after failing to receive the ACK

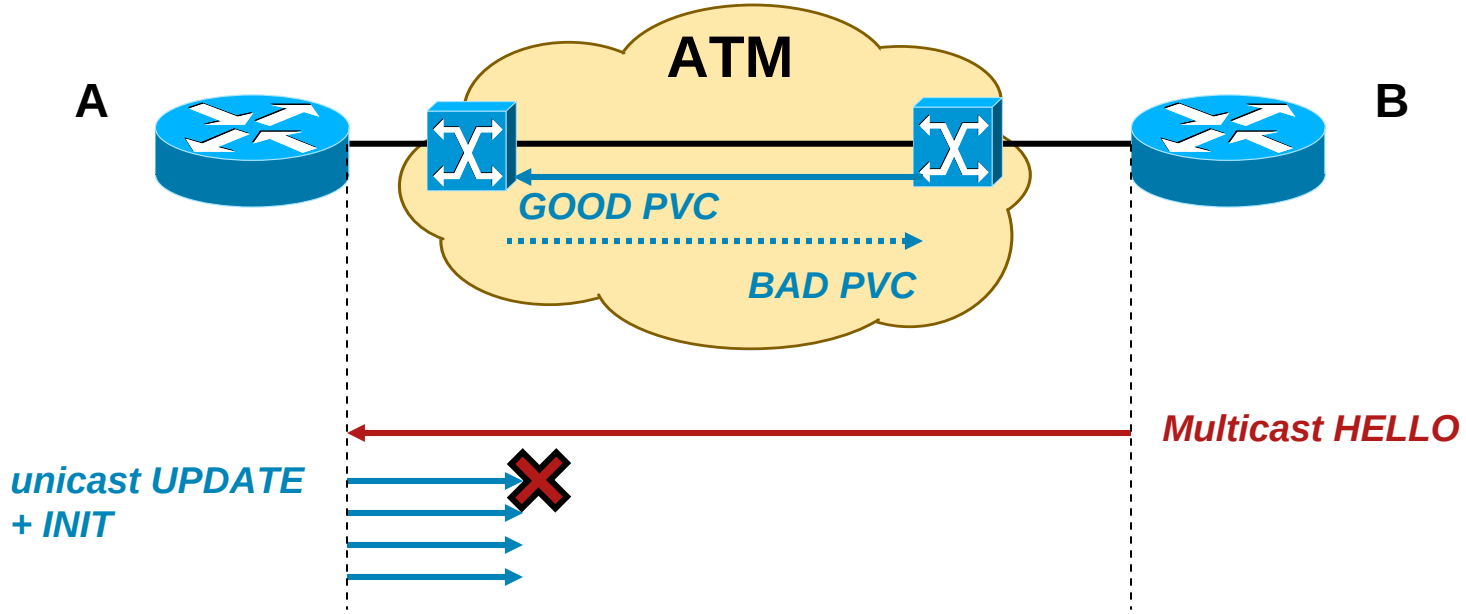

Neighbour Startup Enhancements
Retransmission Example

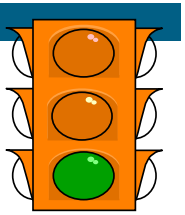

- Router B has not seen A but Router A has seen B, and is busy retransmitting reliable update packets to router B

```
A#show ip eigrp neighbor
IP-EIGRP neighbors for process 65535
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.2.2                Et0/0            13 00:00:10    1    5000  1  0
```

- The SRTT is calculated on how long it takes to get an ACK for a reliable packet. We use a weighting on using the prior SRTT and last packet
- The RTO typically uses a multiplier giving SRTT * 6. The minimum is 200 ms and maximum is 5000 ms. We take an initial RTO value of 2000 ms

Neighbour Startup Enhancements

CSCdy45118

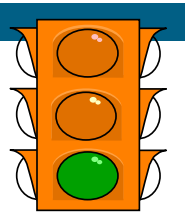

Resolves three issues with the Process

- EIGRP sends queries and expects replies from pending peer potentially causing SIA
- EIGRP exchanges topology information in the first update packet allowing routes to be propagated
- The unicast update is sometimes sent before the peer receives the first Hello causing unicast retransmission

```
A#show ip eigrp neighbor detail
IP-EIGRP neighbors for process 65535
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.2.1                Et0/0             1 00:00:13    1  5000  1  0
   Version 12.2/1.2, Retrans: 3, Retries: 4, Waiting for Init, Waiting for Init Ack
   Expecting no reply for queries
    UPDATE seq 7 ser 0-0 Sent 13184 Init Sequenced
```

- Available in 12.2(13.7)T2, 12.2(15.1)S, 12.2(16.1)B

Neighbour Startup Enhancements
New behavior

CSCdy45118

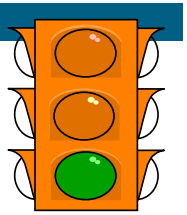

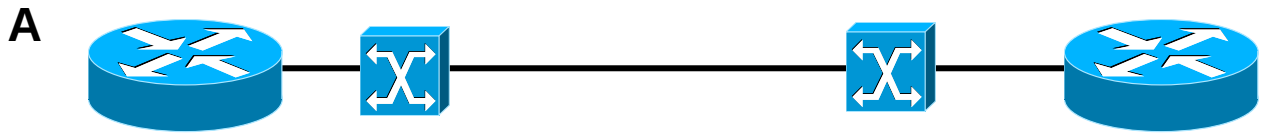

Hello	→	New peer, pending init+ack
New peer, pending init+ack	←	Expedited Hello
Pending ack	←	Null update+init
Null update+init+ack	→	Full peer
Full peer - accept routes	←	Topo update+ack
Topo update	→	Accept routes

EIGRP Convergence
Convergence Terms

- Failure detection

  How quickly a device on the network can detect and react to a failure

- Information dissemination

  How quickly the failure in the previous stage is communicated to other devices

- Repair

  How quickly a devices notified of a failure can calculate an alternate path

  Improvements any of these stages provides an improvement in overall convergence

EIGRP Convergence
EIGRP Feasible Successor

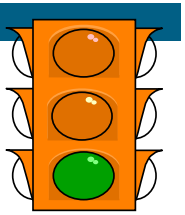

- EIGRP is the fastest converging of all IGP protocols using this technique

- EIGRP provides nearly instantaneous convergence through these pre-computed backup routes

- This prevents us going Active for a destination

- Avoid Query

- Satisfied when a neighbor advertises reachability to the destination with a metric lower than the Feasible Distance

EIGRP Convergence
EIGRP Feasible Successor

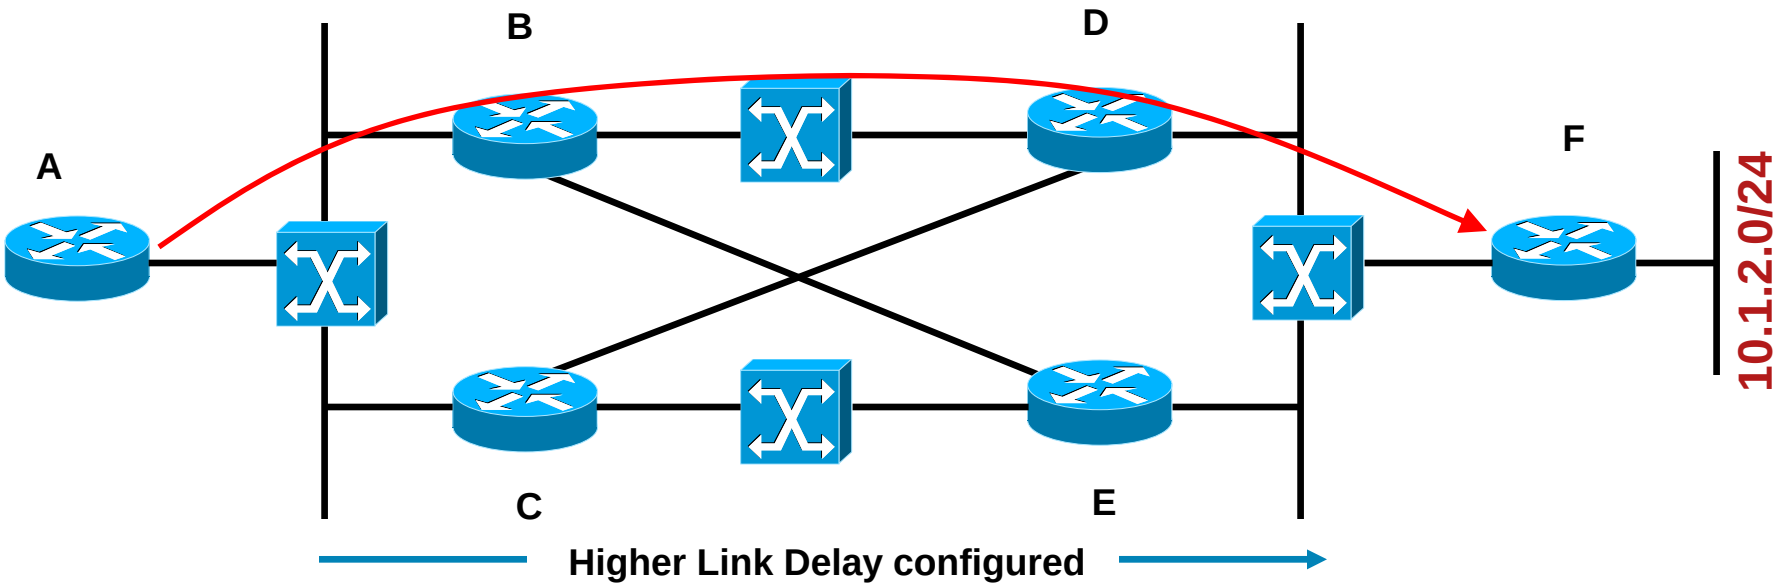


```
A#traceroute 10.1.2.1

Type escape sequence to abort.
Tracing the route to 10.1.2.1

  1 B (16.1.1.2) 20 msec 20 msec 20 msec
  2 D (14.1.1.2) 40 msec 20 msec 20 msec
  3 F (13.1.1.3) 52 msec *  28 msec
A#
```

EIGRP Convergence
EIGRP Feasible Successor

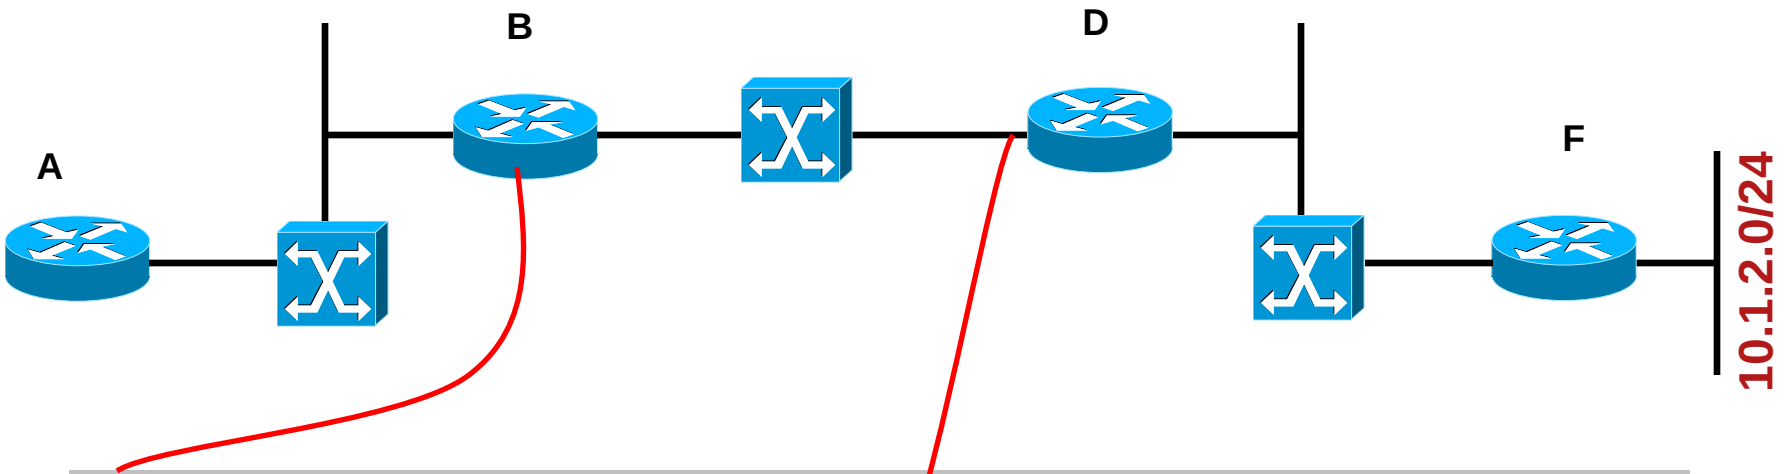


```
B#show ip route 10.1.2.0
Routing entry for 10.1.2.0/24
  Known via "eigrp 65535", distance 90, metric 332800, type internal
  Redistributing via eigrp 65535
  Last update from 14.1.1.2 on Ethernet1/0, 00:00:22 ago
  Routing Descriptor Blocks:
  * 14.1.1.2, from 14.1.1.2, 00:00:22 ago, via Ethernet1/0
      Route metric is 332800, traffic share count is 1
      Total delay is 3000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 252/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

EIGRP Convergence
EIGRP Feasible Successor

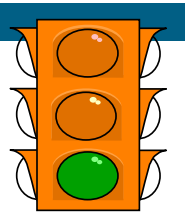

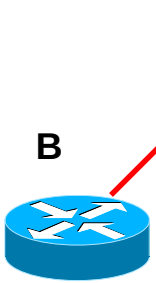


```
B#show ip eigrp topology
IP-EIGRP Topology Table for AS(65535)/ID(17.1.1.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.1.2.0/24, 1 successors, FD is 332800
        via 14.1.1.2 (332800/307200), Ethernet1/0
        via 17.1.1.2 (345600/307200), Ethernet0/2
P 13.1.1.0/24, 1 successors, FD is 307200
        via 14.1.1.2 (307200/281600), Ethernet1/0
        via 17.1.1.2 (320000/281600), Ethernet0/2
P 14.1.1.0/24, 1 successors, FD is 281600
        via Connected, Ethernet1/0
P 15.1.1.0/24, 2 successors, FD is 307200
        via 17.1.1.2 (320000/281600), Ethernet0/2
        via 16.1.1.3 (320000/294400), Ethernet0/0
P 16.1.1.0/24, 1 successors, FD is 281600
        via Connected, Ethernet0/0
P 17.1.1.0/24, 1 successors, FD is 294400
        via Connected, Ethernet0/2
B#
```

EIGRP Convergence
EIGRP Feasible Successor

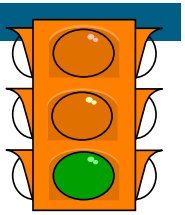

```
B#show ip eigrp events
Event information for AS 65535:
....
Route install: 10.1.2.0/24 17.1.1.2
RDB delete: 10.1.2.0/24 14.1.1.2
FC sat rdbmet/succmet: 345600 307200
FC sat nh/ndbmet: 17.1.1.2 332800
Find FS: 10.1.2.0/24 332800
Peer down: 14.1.1.2 Ethernet1/0
```

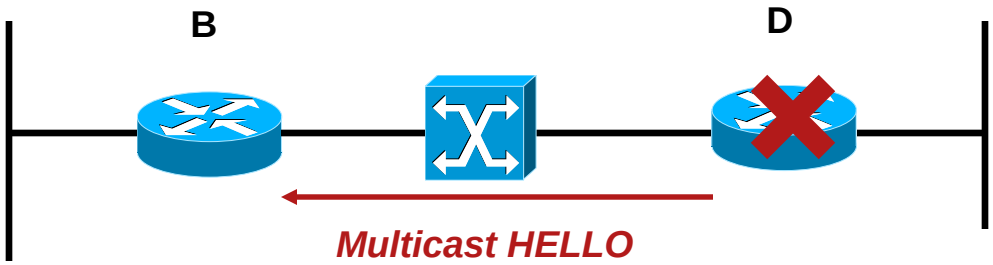


```
B#
03:37:50: IP: s=14.1.1.2 (Ethernet1/0), d=224.0.0.10, len 60, rcvd 2
03:37:55: IP: s=14.1.1.2 (Ethernet1/0), d=224.0.0.10, len 60, rcvd 2
03:38:10: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65535: Neighbor 14.1.1.2 (Ethernet1/0) is down:
holding time expired
03:38:11: RT: delete route to 10.1.2.0 via 14.1.1.2, eigrp metric [90/332800]
03:38:11: RT: no routes to 10.1.2.0
03:38:11: RT: delete subnet route to 10.1.2.0/24
03:38:11: RT: delete network route to 10.0.0.0
03:38:11: RT: add 10.1.2.0/24 via 17.1.1.2, eigrp metric [90/345600]
```

- We did not go Active; using our Feasible Successor
- The delay was the initial failure detection; waiting for EIGRP Hold-time to expire for the peer

EIGRP Convergence
EIGRP Feasible Successor

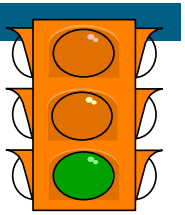

- EIGRP Hello timers can be tuned to a minimum of 1 second. This is not configurable to sub-second

```
B(config-if)#ip hello-interval eigrp 65535 ?
  <1-65535>  Seconds between hello transmissions

B(config-if)#ip hello-interval eigrp 65535 1
```

- There are reasons for not recommending this and also for us not offering such low values; a recommended alternative is discussed in the following section

EIGRP Convergence
EIGRP Bidirectional Forwarding

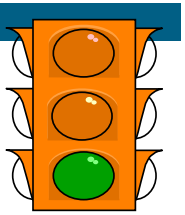

- BFD exhibits lower overhead than aggressive hellos
- BFD is a heartbeat at Layer 2.5
- BFD can provide sub-second failure detection

  http://www.ietf.org/internet-drafts/draft-ietf-bfd-generic-02.txt

  http://www.ietf.org/internet-drafts/draft-ietf-bfd-base-05.txt
- BFD works on most media
- For SONET/SDH alarm detection BFD can provide close to the same reaction time
- 12.2(18)SXE, 12.4T, 12.0(31)S

EIGRP Convergence
EIGRP Bidirectional Forwarding

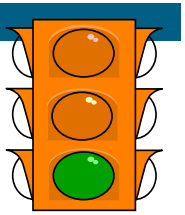

- BFD relies on EIGRP to tell it about neighbours
- BFD peers exchange control packets
- BFD peers negotiate timers
- The detect timer determines a failure
- We should receive a BFD control packet within the detect-timer

  [(Required Minimum RX Interval) * (Detect Multiplier)]

EIGRP Convergence
EIGRP Bidirectional Forwarding

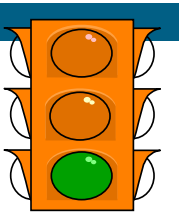

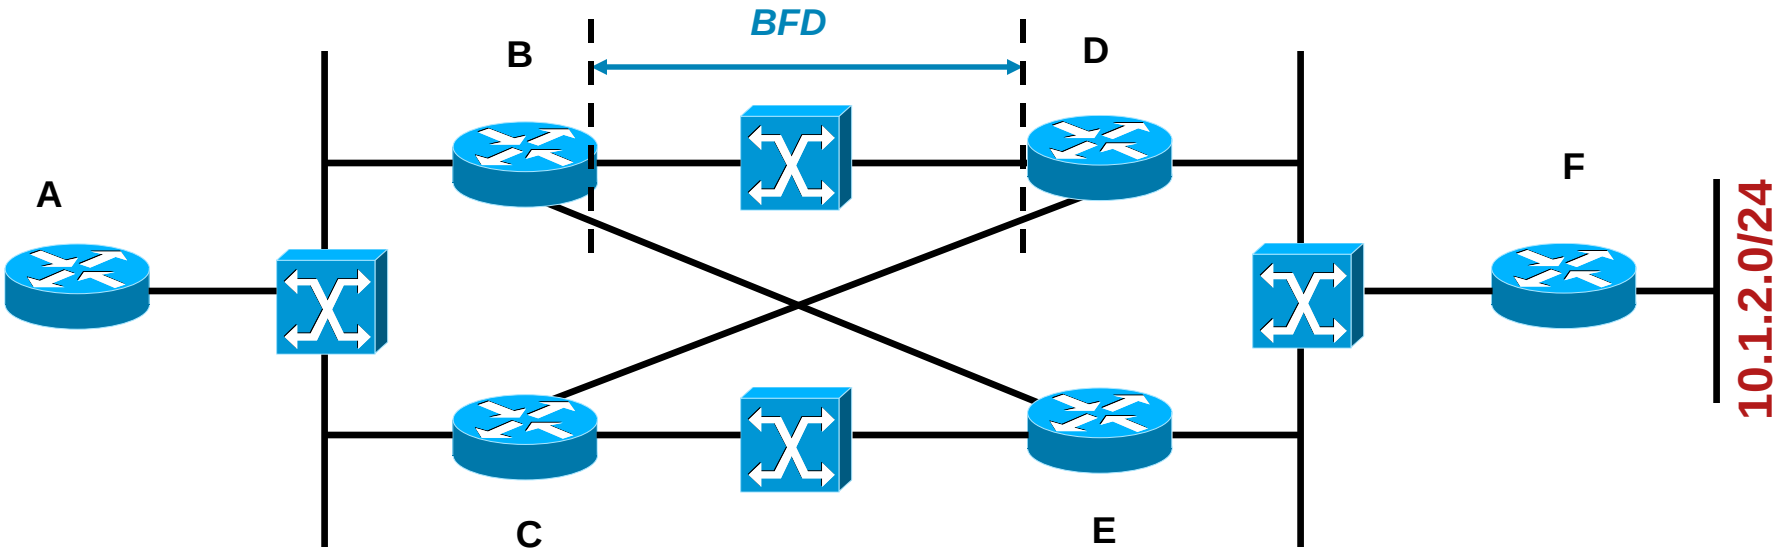

- BFD working together with EIGRP as the upper layer protocol

- To tell us quickly about changes in Layer 2 state
- Let us apply BFD now under EIGRP

EIGRP Convergence
EIGRP Bidirectional Forwarding

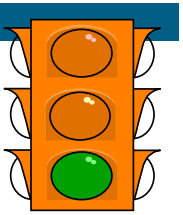

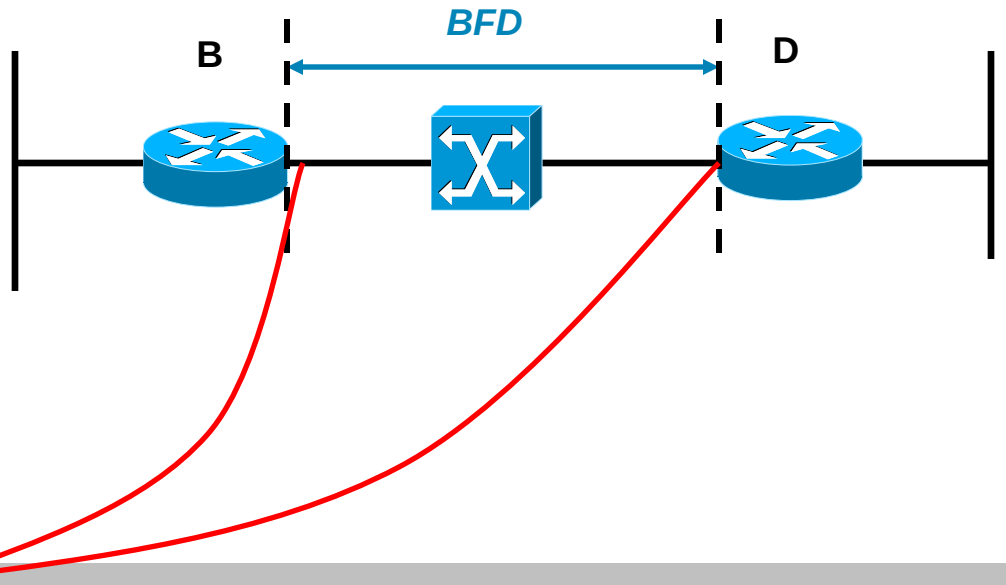


```
B#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
B(config)#router eigrp 65535
B(config-router)#bfd ?
  all-interfaces  Enable BFD on all interfaces
  interface       Enable BFD on specific interface
```

EIGRP Convergence
EIGRP Bidirectional Forwarding

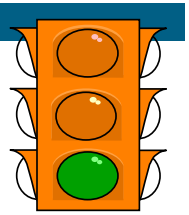

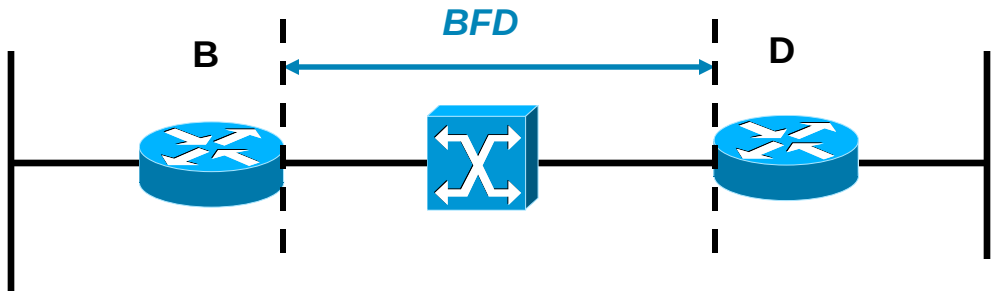


```
B#show bfd neighbors
OurAddr         NeighAddr      LD/RD RH   Holdown(mult)  State      Int
14.1.1.1        14.1.1.2       5/1   1    252   (3 )     Up         E1/0
B#
```

Verbose output

```
B#show bfd neighbor detail | begin Registered
Registered protocols: EIGRP
Uptime: 00:06:33
```

```
B#show ip eigrp int de e1/0
IP-EIGRP interfaces for process 65535
BFD is enabled
B#
```

EIGRP  BFD

EIGRP Convergence
EIGRP Bidirectional Forwarding

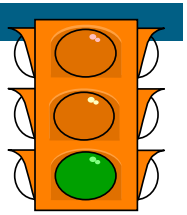

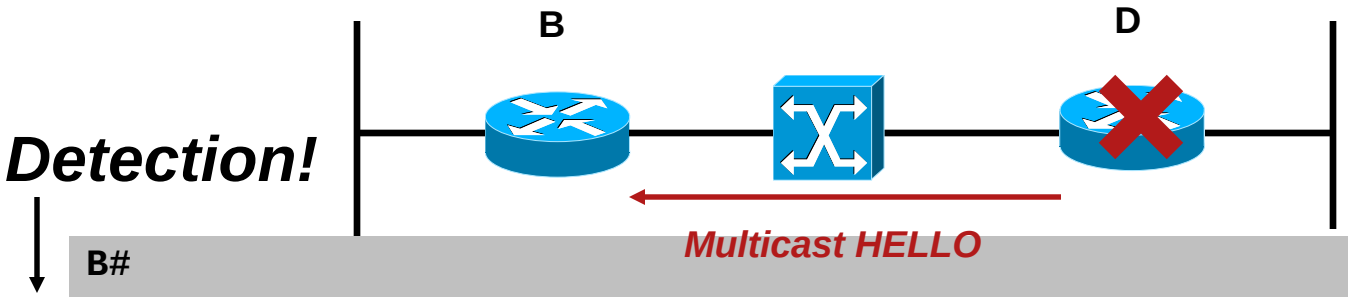


```
B#
17:14:21: IP: s=14.1.1.2 (Ethernet1/0), d=224.0.0.10, len 60, rcvd 2
17:14:21: Session [14.1.1.1,14.1.1.2,E1/0,7], event DETECT TIMER EXPIRED,
state UP -> FAILING
17:14:21: BFD: bfd_neighbor - action:DESTROY, proc:1024, idb:Ethernet1/0,
neighbor:14.1.1.2
17:14:21: Session [14.1.1.1,14.1.1.2,E0/0,7], event Session delete, state
FAILING -> ADMIN DOWN
17:14:21: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 65535: Neighbor 14.1.1.2
(Ethernet1/0) is down: BFD DOWN notification
17:14:21: RT: delete route to 10.1.2.0 via 14.1.1.2, eigrp metric [90/332800]
17:14:21: RT: no routes to 10.1.2.0
17:14:21: RT: delete subnet route to 10.1.2.0/24
17:14:21: RT: delete network route to 10.0.0.0
17:14:21: RT: add 10.1.2.0/24 via 17.1.1.2, eigrp metric [90/345600]
```

EIGRP Convergence
Graceful Shutdown

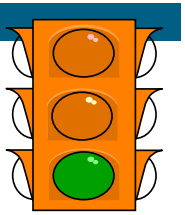

- Previously when problems occurred at router B neighbours waited on hold time before switching traffic
- During the hold time packets on the wire will be lost
- It's better to get A and D to gracefully route around B while B can still forward traffic
- This feature facilitates a `Goodbye` message to be sent

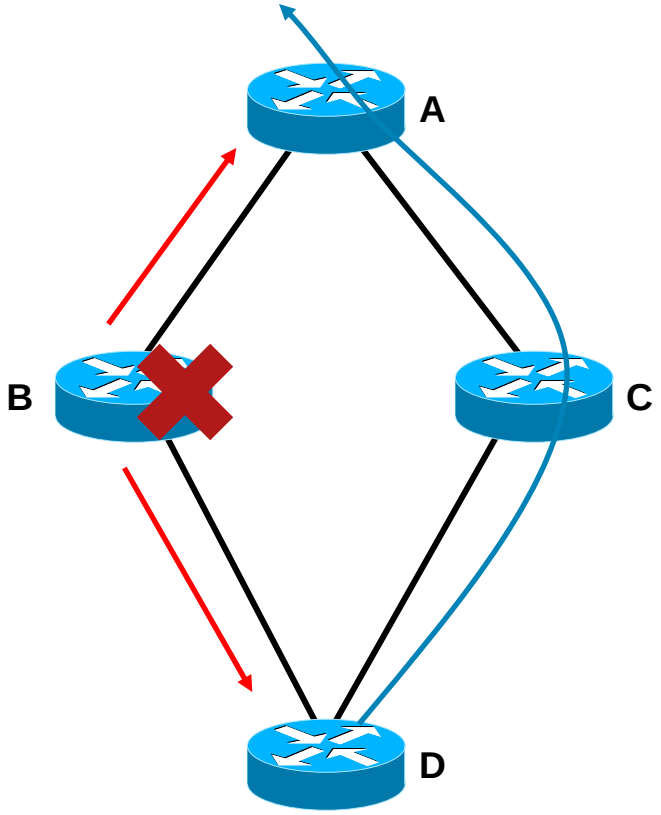

EIGRP Convergence
Graceful Shutdown

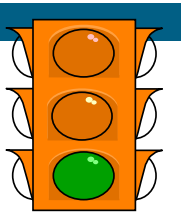

Graceful Shutdown scenarios

- Unidirectional Link Failures
- Router reload
- Removing EIGRP Process
- EIGRP IPv6 on manual protocol shutdown

EIGRP Convergence
Graceful Shutdown

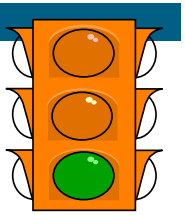

- Code prior to Graceful shutdown exhibits this

```
Neighbor 1.1.1.2 (Ethernet0/0) is down: K- value mismatch
```

- Newer code with the enhancement displays

```
Neighbor 1.1.1.2 (Ethernet0/0) is down: Interface Goodbye received
```

- Or this, depending on the event

```
Neighbor 1.1.1.1 (Ethernet0/0) is down: Peer goodbye received
```

- Available in 12.3(2.3)B 12.3(1.4)T 12.3(1.4)

Routing Enhancements

Routing Enhancements

- EIGRP IPv6
- Third Party Next Hop
- Route-map Enhancements
- SNMP
- MPLS PE/CE
- Multi Topology Routing
- PIX EIGRP Support
- OER EIGRP Support
- EIGRP MANET

IPv6 Industry Drivers

- **US Government DOD Memo DoD Memo (June 9, 2003)**
  - All systems comply with IPv6 by the year 2008

The DoD goal is to complete the transition to IPv6 for all inter and intra networking across the DoD by FY 2008. To enable this transition it is DoD policy for all Information Technology (IT) and National Security Systems (NSS) which make up the GIG (ref a) that:

- As of October 1, 2003, all GIG assets being developed, procured or acquired shall be IPv6 capable (in addition to maintaining interoperability with IPv4 systems/capabilities). This explicitly includes all acquisitions that reach Milestone C after October 1, 2003. The next version of the Joint Technical Architecture (JTA) will reflect this requirement.

Source: U.S. Assistant Secretary of Defense - John Stenbit

EIGRP IPv6 Goals

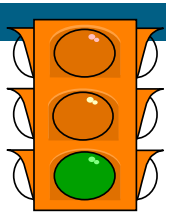

- A Protocol Dependent Module (PDM) to route IPv6
- A familiar *Look and Feel* means incumbent EIGRP operational expertise can be leveraged
- Add new TLV's (Type, Length, Value) in EIGRP packets to carry IPv6 prefixes
- Use the tried and trusted Reliable Transport Protocol (RTP) for reliable delivery of packets
- DUAL performs route computations for IPv6 without modifications

EIGRP Components

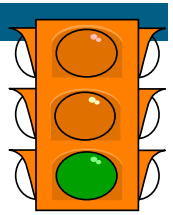

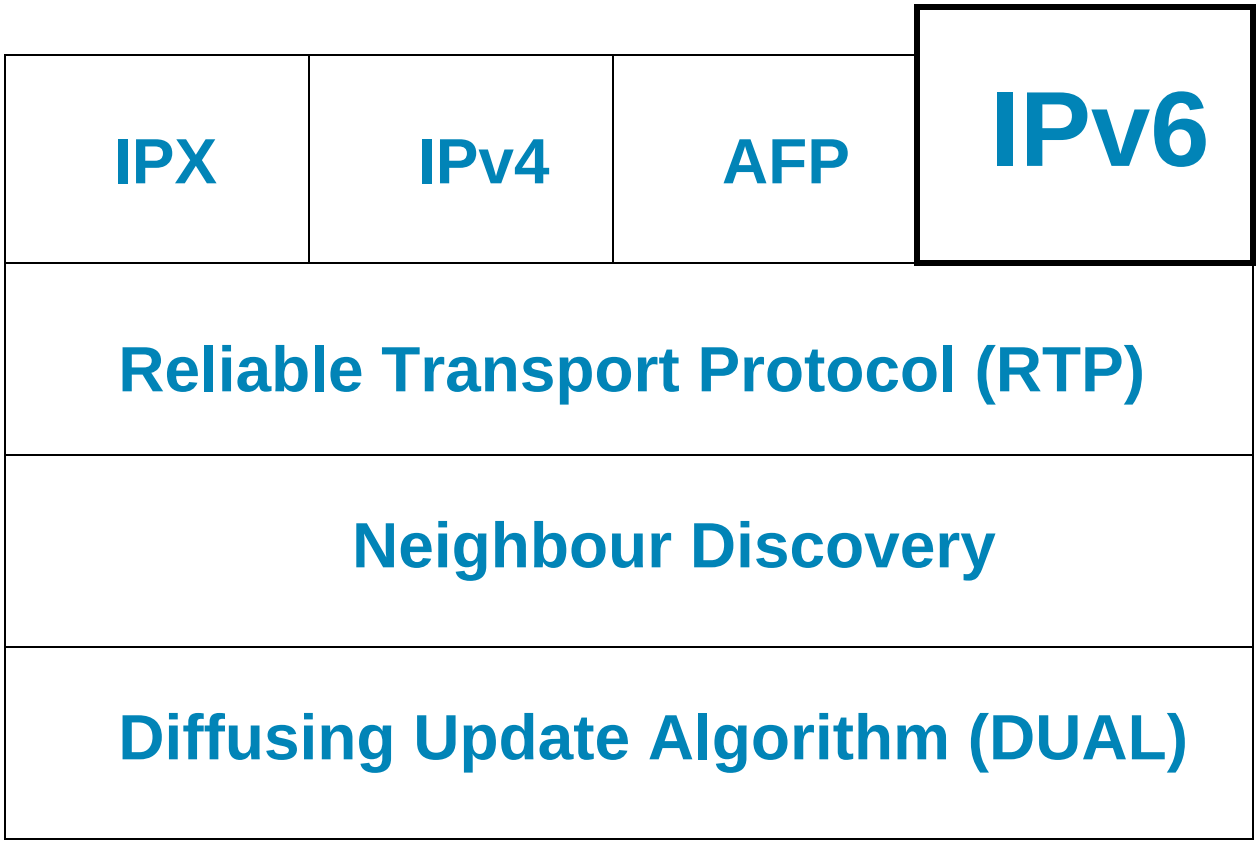

Note: AFP (appletalk eigrp) has been removed in 12.2(13T) and 12.2S images

IPv6 Addressing
Concepts

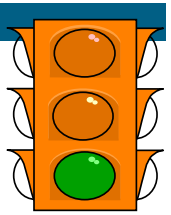

- An IPv6 address is an extended 128-bit / 16 bytes address that gives
  - 2^{128} possible addresses (3.4×10^{38})
- A typical unicast IPv6 address is
  - 64 bits for the subnet ID, 64 bits for the interface ID
- IPv6 addresses
  - Separated into 8 * 16-bit Hexadecimal numbers
  - Each block is separated by a colon :
  - :: can replaced leading, trailing or consecutive zeros
  - :: can only appear once
  - x:x:x:x:x:x:x:x
- Example:
  - 2003:0000:130F:0000:0000:087C:876B:140B
  - 2003:0:130F::87C:876B:140B

IPv6 Addressing
Link-local

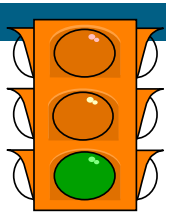

- This IPv6 Link-local address is never routed
- An IPv6 link-local is prefixed by fe80 and has a prefix length of /10
- You can configure this manually on an interface

```
nw2007(config-if)#ipv6 address ?
  X:X:X:X::X          IPv6 link-local address
  X:X:X:X::X/<0-128>  IPv6 prefix
  autoconfig          Obtain address using autoconfiguration
```

- A IPv6 Link-local address is used by EIGRP to source Hello packets and establish an adjacency

IPv6 Addressing
Link-local

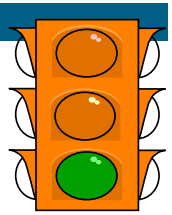

Don't worry about addresses!

- There is an easy way to get started

```
nw2007#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
nw2007(config)#int e1/0
nw2007(config-if)#ipv6 enable
```

```
nw2007#show ipv6 int e1/0
Ethernet1/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:201
```

- More in the configuration section

EIGRP IPv6

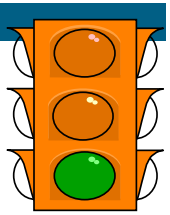

- EIGRP IPv6 uses the same familiar packet types
  - HELLO
  - UPDATE
  - QUERY
  - REPLY
  - ACK (Essentially a HELLO)
- Periodic Hello packets are unreliably multicast and sourced from the link-local IPv6 address
- EIGRP IPv6 Multicast transport

  FF02:0:0:0:0:0:0:A or abbreviated to FF02::A

Consecutive zeros replaced by ::

EIGRP IPv6
New TLV's

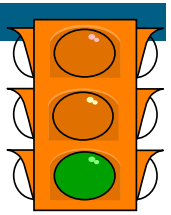

- New TLV (Type, Length, Value) in EIGRP packets
- Existing IPv4 TLV's
  - Internal routes TLV (Type 0x0102)
  - External routes TLV (Type 0x0103)
- New IPv6 TLV's
  - Internal routes TLV (Type 0x0402)
  - External routes TLV (Type 0x0403)
- For IPv6 TLV formats are similar except for
  - A 128-bit next-hop
  - A variable Prefix (up to 128-bits)

EIGRP IPv6 Configuration
Global and Interface Mode

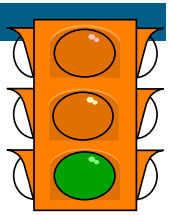

Minimum IPv6 Configuration

- Enables IPv6 packet forwarding and must be configured first under global configuration

```
nw2007(config)#ipv6 unicast
```

- Specifies that EIGRP IPv6 process 65535 runs over this interface

```
nw2007(config)#ipv6 unicast
nw2007(config)#int Ethernet 0/0
nw2007(config-if)#ipv6 enable
nw2007(config-if)#ipv6 eigrp 65535
```

EIGRP IPv6 Configuration Router Mode

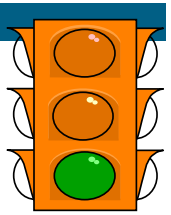

Minimum IPv6 Configuration

- IPv6 router mode configuration

```
nw2007(config)#ipv6 router eigrp 65535
```

- A 32-bit router-id can be configured explicitly or derived implicitly from an existing IP address

```
nw2007(config-rtr)#router-id 1.1.1.1
nw2007(config-rtr)#
```

- The Implicit approach will select the highest loopback IP address

EIGRP IPv6 Configuration
Missing Router-id

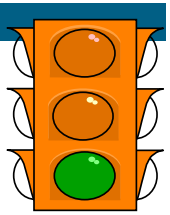

- No Router-id means Hello packets will not be transmitted. We issue a warning for this

```
nw2007#show ipv6 eigrp topology
EIGRP Topology Table for AS(65535)/ID(0.0.0.0)
% No router ID for EIGRP 65535
nw2007#
```

- The originating router-id is required and is part of the IPv6 external routes TLV; used to detect routing loops

```
nw2007#show ipv6 eigrp events
07:06:12.323 Ignored route, metric: 2040:3333::31:113:0/112 281600
07:06:12.323 Ignored route, neighbor info: FE80::A8BB:CCFF:FE00:200 Ether..
07:06:12.323 Ignored route, dup router: FE80::A8BB:CCFF:FE00:200 s
```

EIGRP IPv6 Configuration
Protocol Shutdown

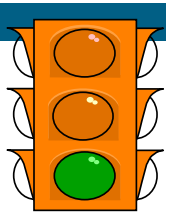

Minimum IPv6 Configuration

- The process starts in shutdown and we need to bring it up to complete the configuration

```
nw2007(config)#ipv6 router eigrp 65535
nw2007(config-rtr)#no shutdown
nw2007(config-rtr)#
```

- As with a missing router-id, if we do not do this it means EIGRP IPv6 is not ready to go

```
nw2007#show ipv6 eigrp topology
EIGRP Topology Table for AS(65535)/ID(1.1.1.1)
% EIGRP 65535 is in SHUTDOWN
nw2007#
```

EIGRP IPv6 Configuration Router-Mode

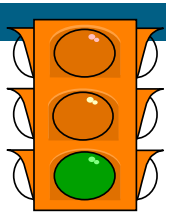

- Some commands are New

```
nw2007(config-rtr)#?
      default                 Set a command to its defaults
      default-metric          Set metric of redistributed routes
      distance                Administrative distance
      distribute-list         Filter networks in routing updates
      log-neighbor-changes    Enable/Disable EIGRP neighbor logging
      log-neighbor-warnings   Enable/Disable EIGRP neighbor warnings
      maximum-paths           Forward packets over multiple paths
      metric                  Modify EIGRP routing metrics and parameters
      neighbor                Specify a neighbor router
      passive-interface       Suppress routing updates on an interface
      redistribute            Redistribute IPv6 prefixes from another routing protocol
      router-id               router-id for this EIGRP process
      shutdown                Shutdown protocol
      stub                    Set EIGRP as stubbed router
      timers                  Adjust routing timers
      variance                Control load balancing variance
```

- Some commands are gone

EIGRP IPv6 Topology Table

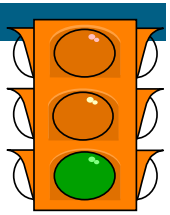

- The Topology show commands are congruent with IPv4

```
nw2007#show ipv6 eigrp topology
IPv6-EIGRP Topology Table for AS(65535)/ID(1.1.1.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status

P 2040:3333::31:113:0/112, 1 successors, FD is 281600
        via FE80::A8BB:CCFF:FE00:200 (281600/256), Ethernet0/0
P 2040:3333::31:114:0/112, 1 successors, FD is 281600
        via FE80::A8BB:CCFF:FE00:200 (281600/256), Ethernet0/0
nw2007#
```

- The next-hop is the neighbours 128-bit link-local

EIGRP IPv6 Topology Table

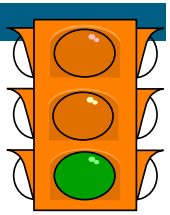

- The information source and next-hop 128-bit address

```
nw2007#show ipv6 eigrp topology 2040:3333::31:113:0/112
IPv6-EIGRP (AS 65535): Topology entry for 8:1:1::1/128
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
  Routing Descriptor Blocks:
  FE80::A8BB:CCFF:FE00:200 (Ethernet0/0), from FE80::A8BB:CCFF:FE00:200, Send flag is 0x0
      Composite metric is (281600/256), Route is External
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 1000 microseconds
        Reliability is 0/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      External data:
        Originating router is 2.2.2.2
        AS number of route is 0
        External protocol is Static, external metric is 0
        Administrator tag is 0 (0x00000000)
```

EIGRP IPv6 Event Log

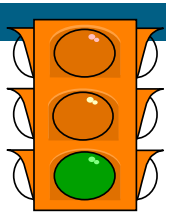

- The IPv6 EIGRP event log

```
nw2007#show ipv6 eigrp eve
1    06:27:52.115 Change queue emptied, entries: 1
2    06:27:52.115 Metric set: 2040:3333::31:113:0/112 281600
3    06:27:52.115 Update reason, delay: new if 4294967295
4    06:27:52.115 Update sent, RD: 2040:3333::31:113:0/112 4294967295
5    06:27:52.115 Update reason, delay: metric chg 4294967295
6    06:27:52.115 Update sent, RD: 2040:3333::31:113:0/112 4294967295
```

EIGRP IPv6 Debugs

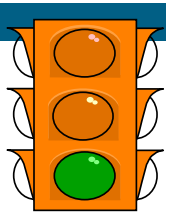

EIGRP IPv6 Specific Debugging

```
nw2007#debug ipv6 eigrp ?
  <1-65535>      Autonomous System
  neighbor       EIGRP neighbor debugging
  notifications  EIGRP event notifications
  summary        EIGRP summary route processing
  <cr>
```

```
nw2007#debug ipv6 eigrp
IP-EIGRP Route Events debugging is on
nw2007#
```

EIGRP IPv6 Debugs

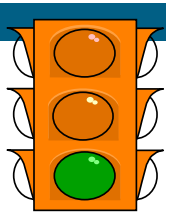

EIGRP IPv6 information in existing debugs

```
nw2007#debug eigrp ?
  fsm         EIGRP Dual Finite State Machine events/actions
  neighbors   EIGRP neighbors
  nsf         EIGRP Non-Stop Forwarding events/actions
  packets     EIGRP packets
  transmit    EIGRP transmission events
```

```
nw2007#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP,
PROBE, ACK, STUB, SIAQUERY, SIAREPLY)
nw2007#
00:52:47: EIGRP: Received HELLO on Ethernet1/0 nbr
FE80::A8BB:CCFF:FE00:401
00:52:47:   AS 65535, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ
un/rely 0/0 peerQ un/rely 0/0
```

EIGRP IPv6 Differences: Summaries

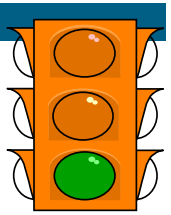

- Auto-summary is not configurable in EIGRP IPv6 because IPv6 is essentially classless
- Manual summarization is supported, as it is with EIGRP IPv4, and can therefore be configured at any point in the network

```
nw2007(config-if)#ipv6 summary-address eigrp 65535 ?
  X:X:X:X::X/<0-128>  IPv6 prefix
```

EIGRP IPv6 Differences: Protocol Shutdown

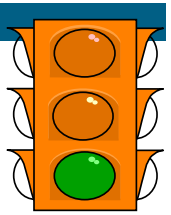

- EIGRP IPv6 supports a protocol shutdown command under router configuration mode
- The IPv6 legacy is a per-interface configuration. We want to start EIGRP IPv6 when router mode commands are executed E.g. Distribute-lists
- Initially, the process will be shutdown

```
!
ipv6 router eigrp 65535
 shutdown
!
```

- For maintenance, Protocol Shutdown initiates Graceful shutdown to inform neighbours quickly

EIGRP IPv6
Differences: Default-information

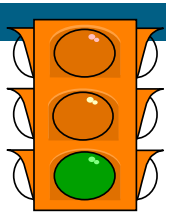

- IPv6 does not currently support configuration of default networks other than ::/0, the IPv4 0.0.0.0/0.

- Therefore EIGRP IPv6 does not support the “default-information” EIGRP legacy command

- We will support the “default-information originate” command for both IPv4 and IPv6 PDM’s (CSCdr75703)

define: PDM - protocol dependent module

EIGRP Third Party Next Hop Example 1

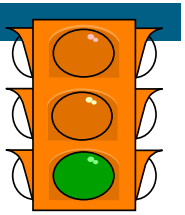

- A, B and C share the same broadcast segment
  - A redistributes RIP into EIGRP
  - B isn't running RIP
  - C isn't running EIGRP
- For redistributed RIP routes B normally shows A as next hop despite a direct connection to C
- A now sends updates to B with C as the next-hop

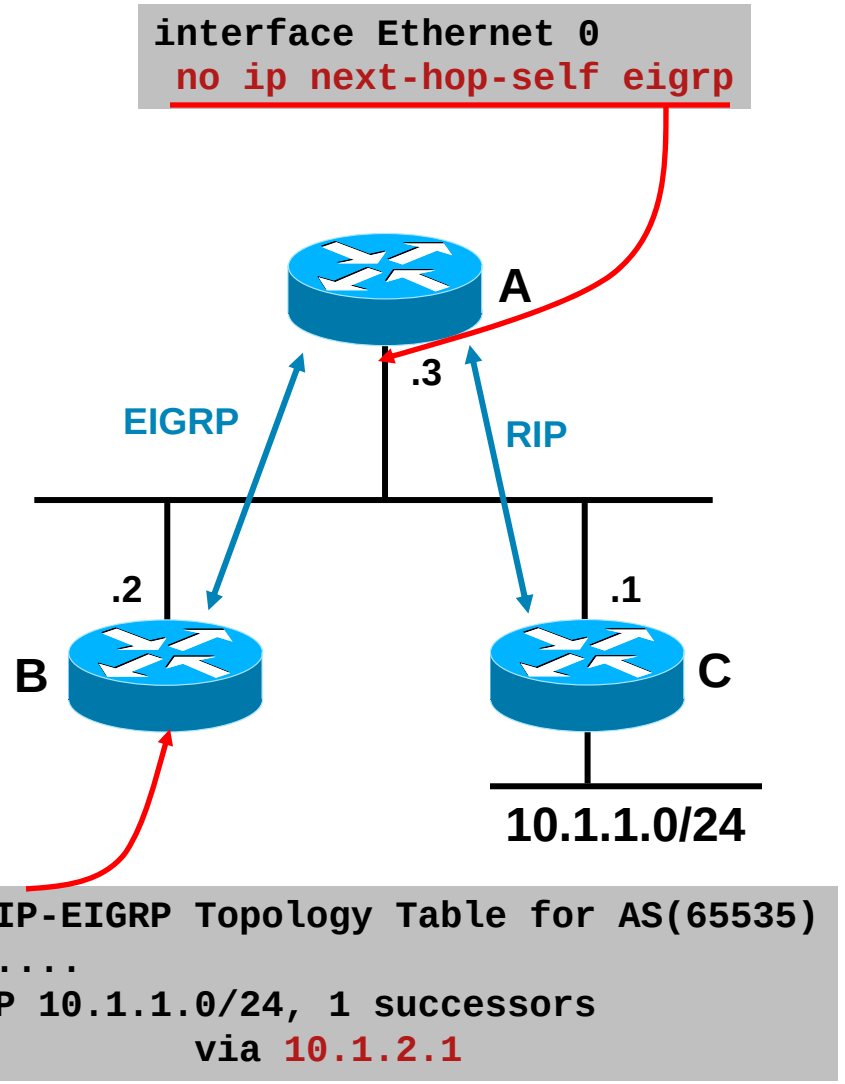

EIGRP Third Party Next Hop Example 2

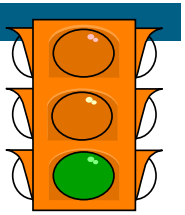

- The multipoint PVC between A, B and C means B learns the IPv6 prefix from both A and C

```
interface serial2/0
 no ipv6 next-hop-self eigrp
```

- Next-hop and the source of that information source are visible in the topology table

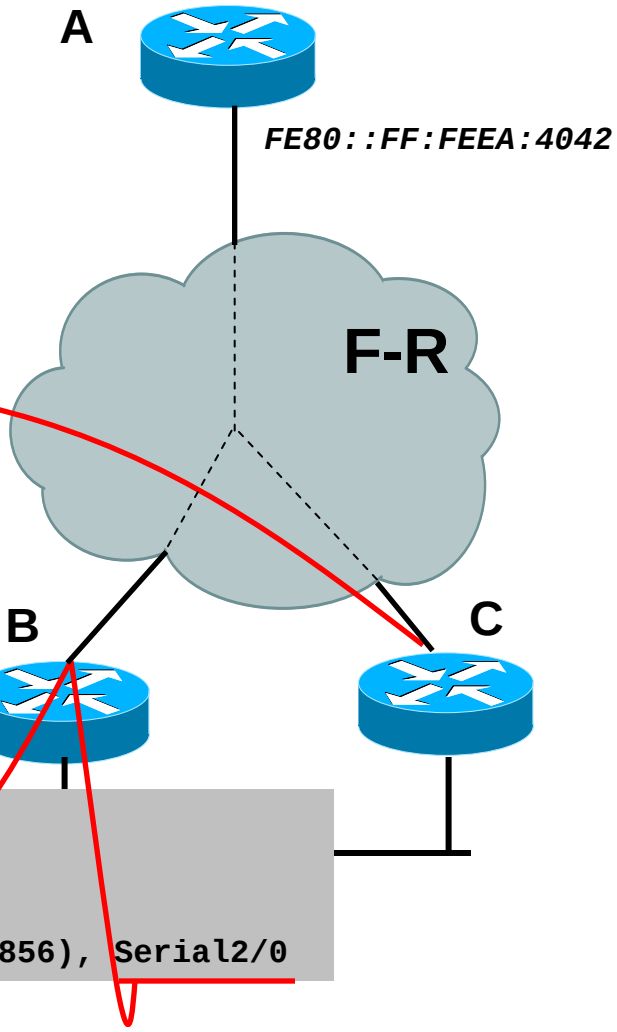


```
P 2040:6666:5555:6666::/90, 1 successors, FD is 2681856
  via FE80::FF:FEEA:4042 (2681856/2169856), Serial2/0
  via FE80::A8BB:CCFF:FE00:1601 (2707456/2681856), Ethernet1/0
  FE80::FF:FEEA:4042 via FE80::A8BB:CCFF:FE00:1601 (3193856/2681856), Serial2/0
```

EIGRP Third Party Next Hop

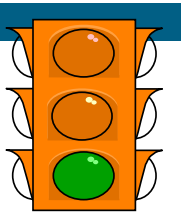

- Applications for third party next hop include Dynamic Multipoint Virtual Private Networks where we preserve the next hop in redistribution from broadcast networks

- Available in 12.3(07)XI 12.2(23.01)S 12.3(02.03)B 12.3(01.02)T 012.003(001.003)

Route Map Enhancements

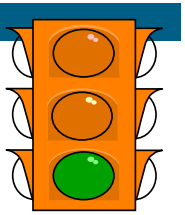

- CSCdw22585 provides enhanced support of route maps for EIGRP
- This allows using a route map to prefer one path over another, for instance

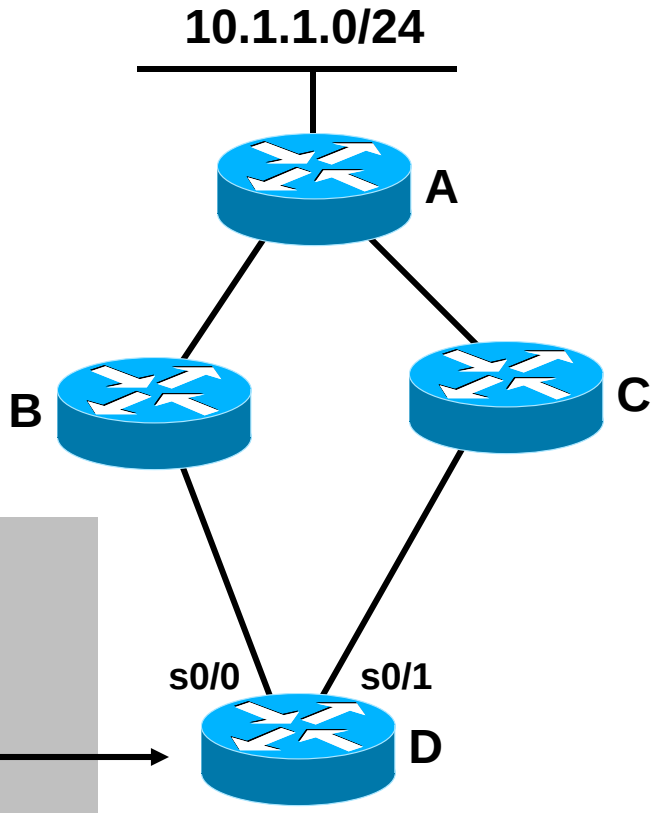


```
route-map setmetric permit 10
 match interface serial 0/0
 set metric 1000 1 255 1 1500

route-map setmetric permit 20
 match interface serial 0/1
 set metric 2000 1 255 1 1500
route-map setmetric permit 30
....
router eigrp 100
 distribute-list route-map setmetric in
```

Route Map Enhancements

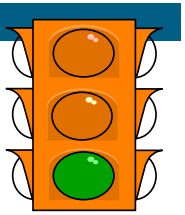

match ip address	Matches routes from prefix list or access list
match ip route-source	Matches routes based on source or neighbour list
match ip route-source redistribution-source	Matches external routes based on originating-router router-id
match interface	Matches routes based on the interface used for next-hop
match tag	Matches internal and external routes based on tag
match ip next-hop	Matches routes based on next-hop field

Route Map Enhancements

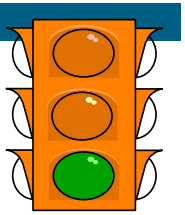

match metric [+-]	Matches routes based on metric, with deviation (+-)
match metric external [+-]	Matches routes based on external protocol metric
match source-protocol	Matches external routes based on external protocol and AS
set metric	Sets metric components
set tag	Sets the tag on internal routes

SNMP EIGRP MIB Support

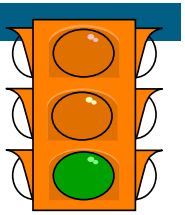

- EIGRP supports 68 MIB objects in 4 major tables
  - EIGRP traffic statistics
  - EIGRP topology data
  - EIGRP neighbour data
  - EIGRP interface data
- A fifth table, the EIGRP VPN table, is included for indexing
- eigrpRouteSIA and eigrpAuthFailure can trigger SNMP traps

SNMP EIGRP MIB Support

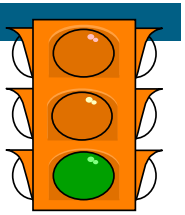

- EIGRP Traffic Statistics
  - AS Number
  - Hellos Sent/Received
  - Updates Sent/Received
  - Queries Sent/Received
  - Replies Sent/Received
- EIGRP Topology Data
  - Destination Net/Mask
  - Active State
  - Feasible Successors
  - Origin Type
  - Distance
  - Reported Distance
- EIGRP Interface Data
  - Peer Count
  - Reliable/Unreliable Queues
  - Pacing
  - Pending Routes
  - Hello Interval
- EIGRP Peer Data
  - Peer Address
  - Peer Interface
  - Hold Time
  - Up Time
  - SRTT/RTO
  - Version

And Many More…

SNMP EIGRP MIB Support
HP OpenView

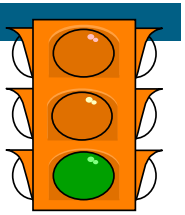

- IP Address
- OID
- Community

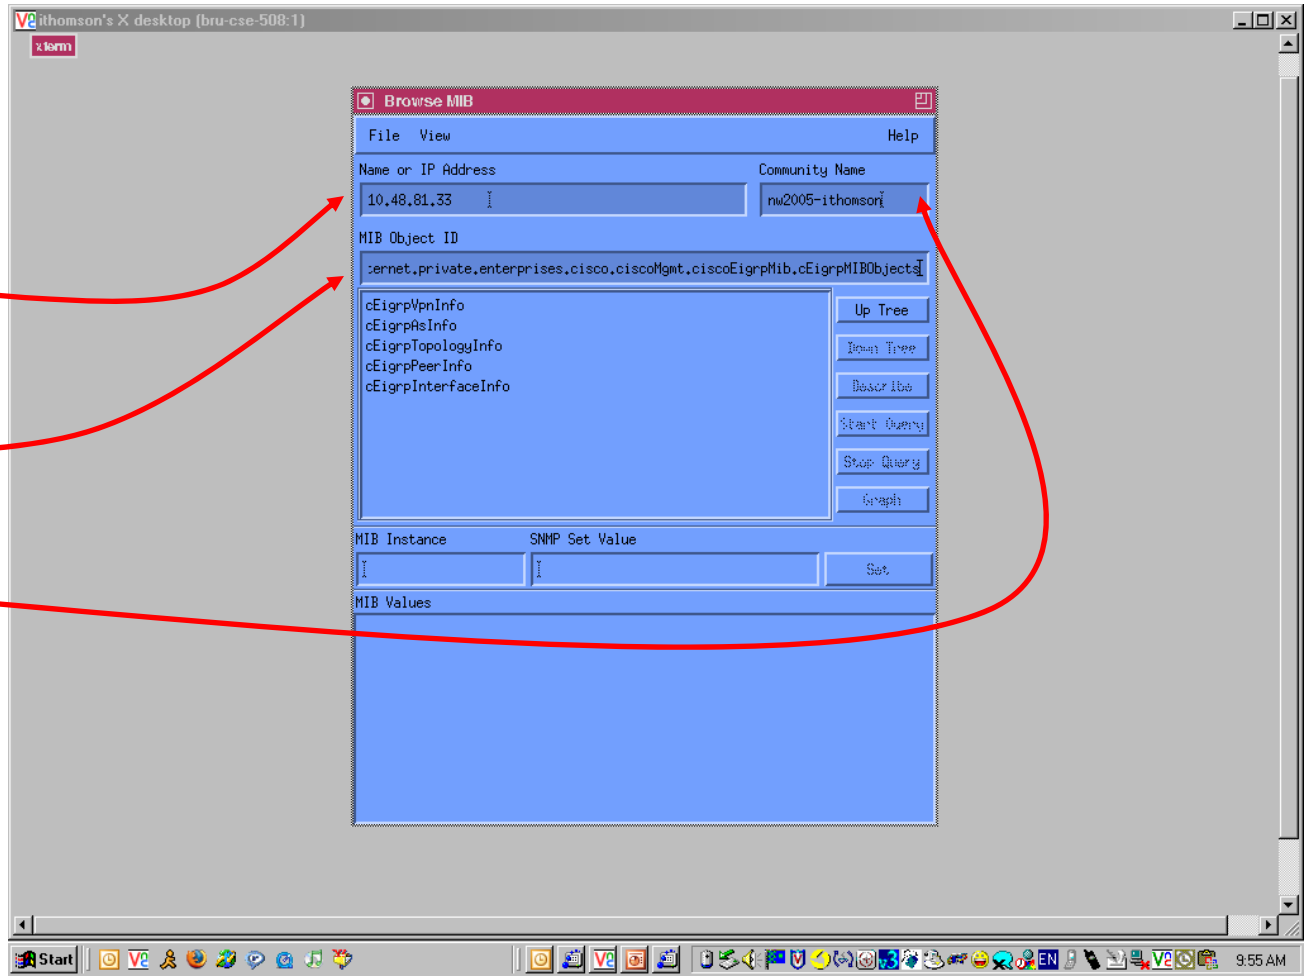

SNMP EIGRP MIB Support

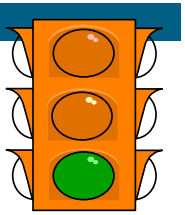

- CISCO EIGRP MIB resources

  http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml
  http://www.cisco.com/go/mibs
  ftp://ftp.cisco.com/pub/mibs/oid/

- CISCO-EIGRP-MIB.oid
- The OID determines how to query values

```
"ciscoEigrpMIB"         "1.3.6.1.4.1.9.9.449"
"cEigrpAsInfo"          "1.3.6.1.4.1.9.9.449.1.2"
"cEigrpTopologyInfo"    "1.3.6.1.4.1.9.9.449.1.3"
"cEigrpPeerInfo"        "1.3.6.1.4.1.9.9.449.1.4"
"cEigrpInterfaceInfo"   "1.3.6.1.4.1.9.9.449.1.5"
```

MPLS VPN PE / CE Concept

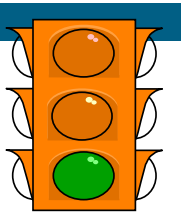

- MP-BGP: Multi-protocol BGP
- PE: Provider Edge
- CE: Customer Edge
- VPN: Virtual Private Network
- VRF: Virtual Routing and Forwarding Instance
- Backdoor link: Link between sites not a VPN

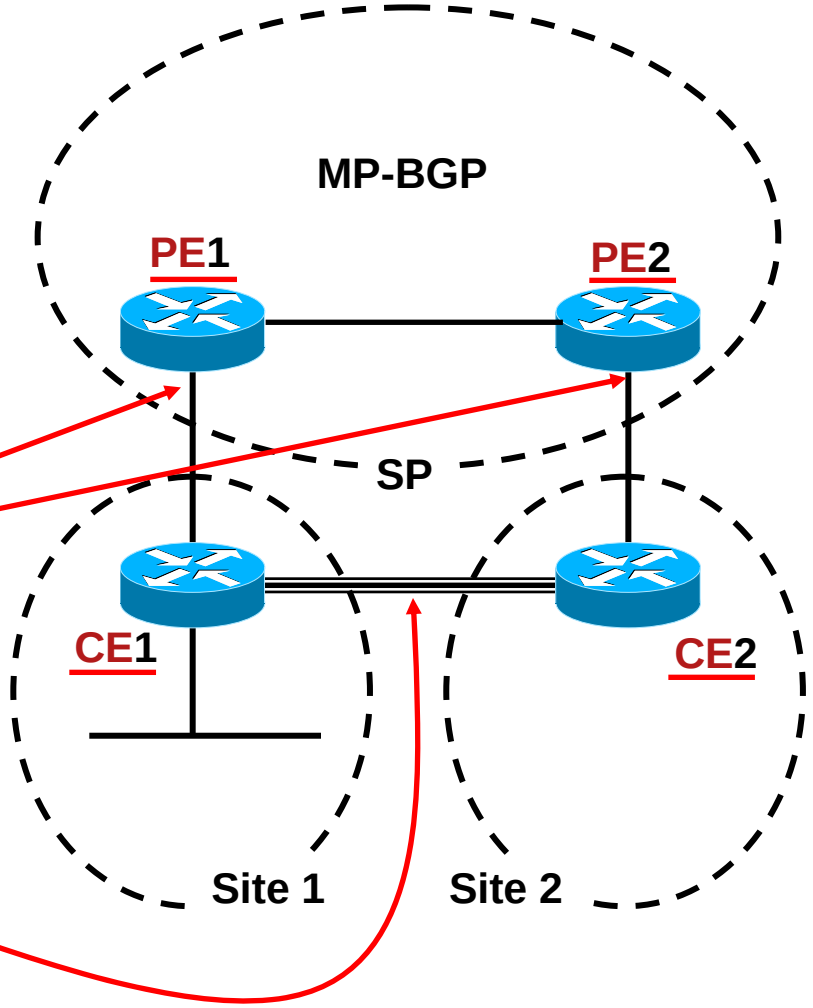

MPLS VPN PE / CE Concept

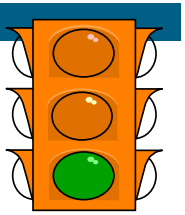

- Routes are redistributed into BGP on PE1

  Extended communities containing the EIGRP attributes are attached

- Routes are redistributed back into EIGRP at PE2

  Extended communities are used to reconstruct the routes as internals preserving metrics

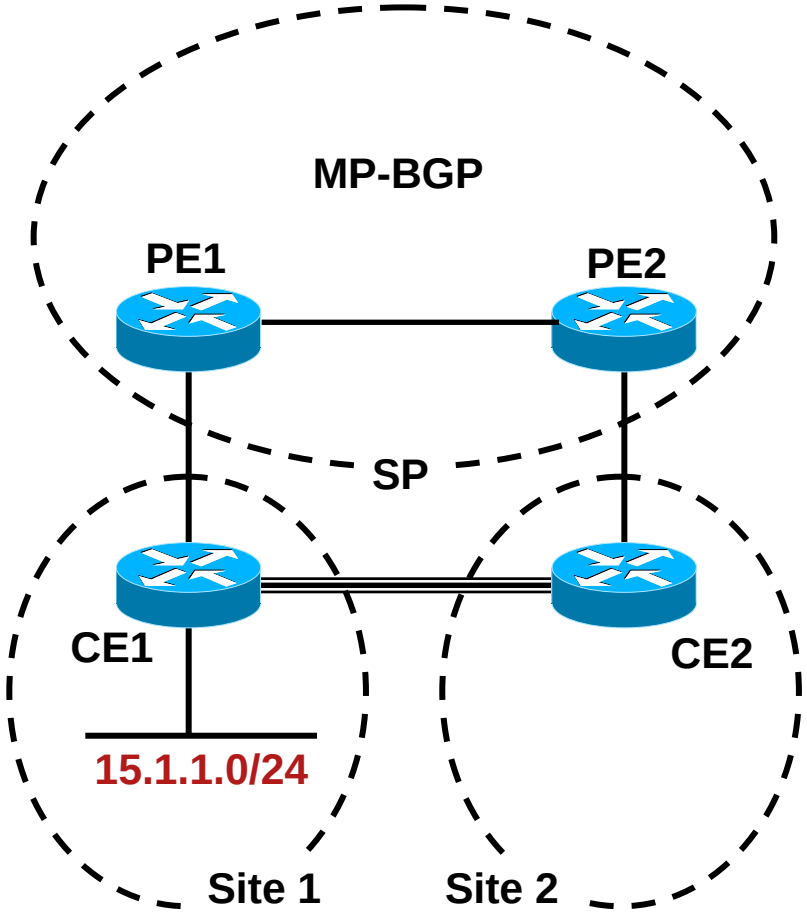

MPLS VPN PE / CE Concept

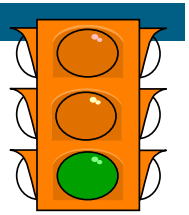

```
PE1(config-if)#ip vrf forwarding site1
PE2(config-if)#ip vrf forwarding site1
```

```
ip vrf site1
 rd 172.16.0.1:20
exit
....
router eigrp 1
 address-family ipv4 vrf site1
 autonomous-system 101
 network 172.16.0.0 255.255.0.0
 redistribute BGP 101 metric 10000 100 255 1 1500
 exit-address-family
```

```
PE1#show ip eigrp vrf site1 topology
...
P 15.1.1.0/24, 1 successors, FD is 307200
        via 10.1.2.1 (307200/281600)...
PE1#
```

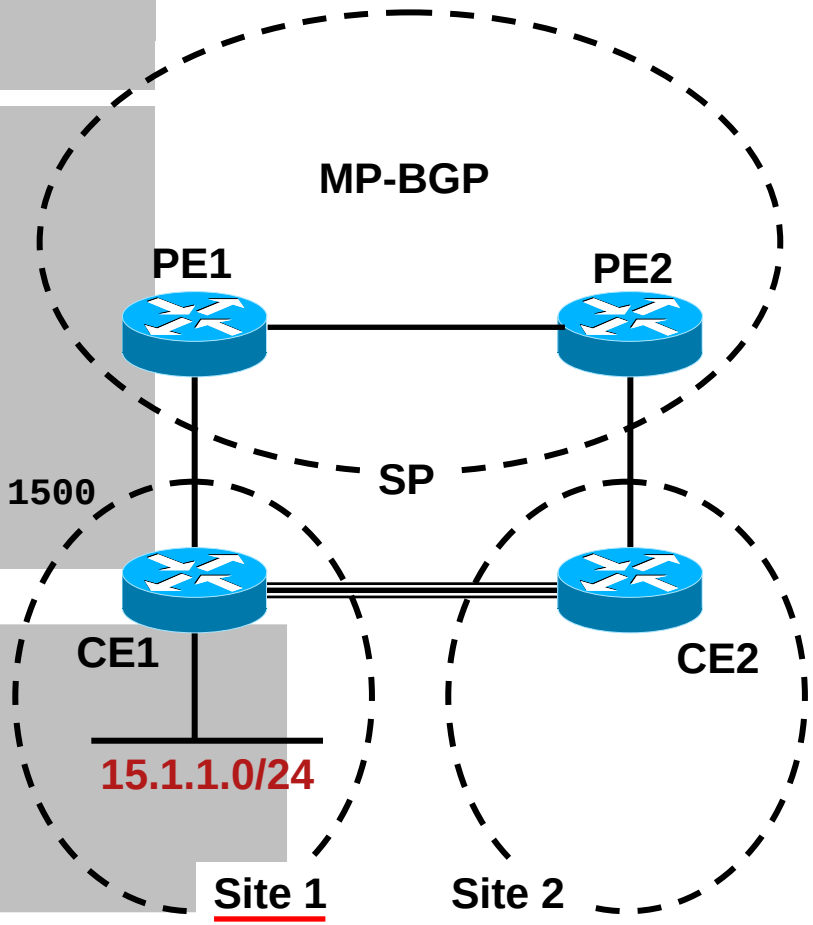

MPLS VPN - PE / CE
Backdoor Links

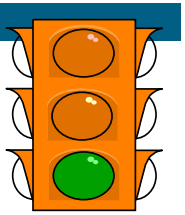

- If backdoor links exist between sites Count-to-Infinity issues can occur when a route goes down
  - Site 1 advertises a route through the backdoor to site 2
  - If the route goes down in site 1 it can be advertised from site 2 via MP-BGP and back into site 1
  - When PE CE was released in 12.0(22)S 12.2(15)T 12.2(18)S we did not cater for this

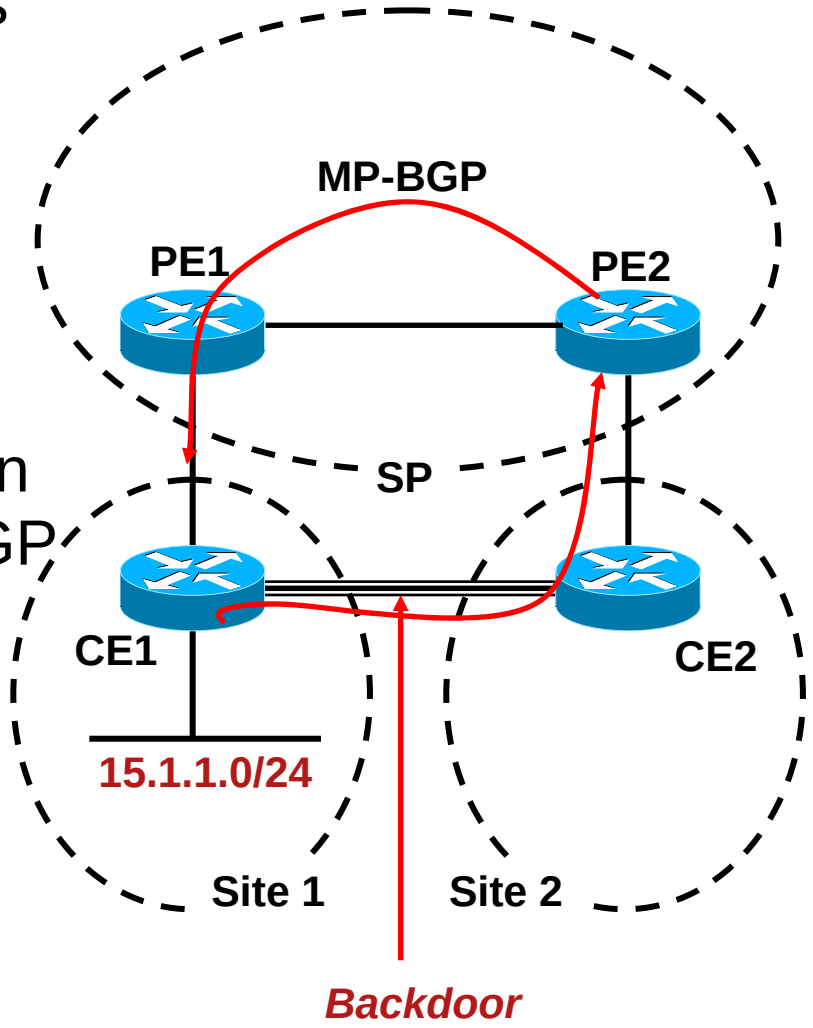

MPLS VPN - PE / CE
Backdoor Links

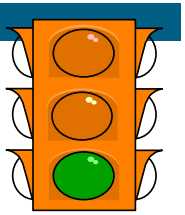

- The solution is to tag routes originating in a site using Site of Origin (SoO)
- EIGRP on the PE will reject (filter out) routes redistributed from MP-BGP if they contain the SoO value for that site

```
route-map SoOrigin permit 10
  set extcommunity soo 100:1
..
interface FastEthernet 0/0
  ip vrf sitemap SoOrigin
```

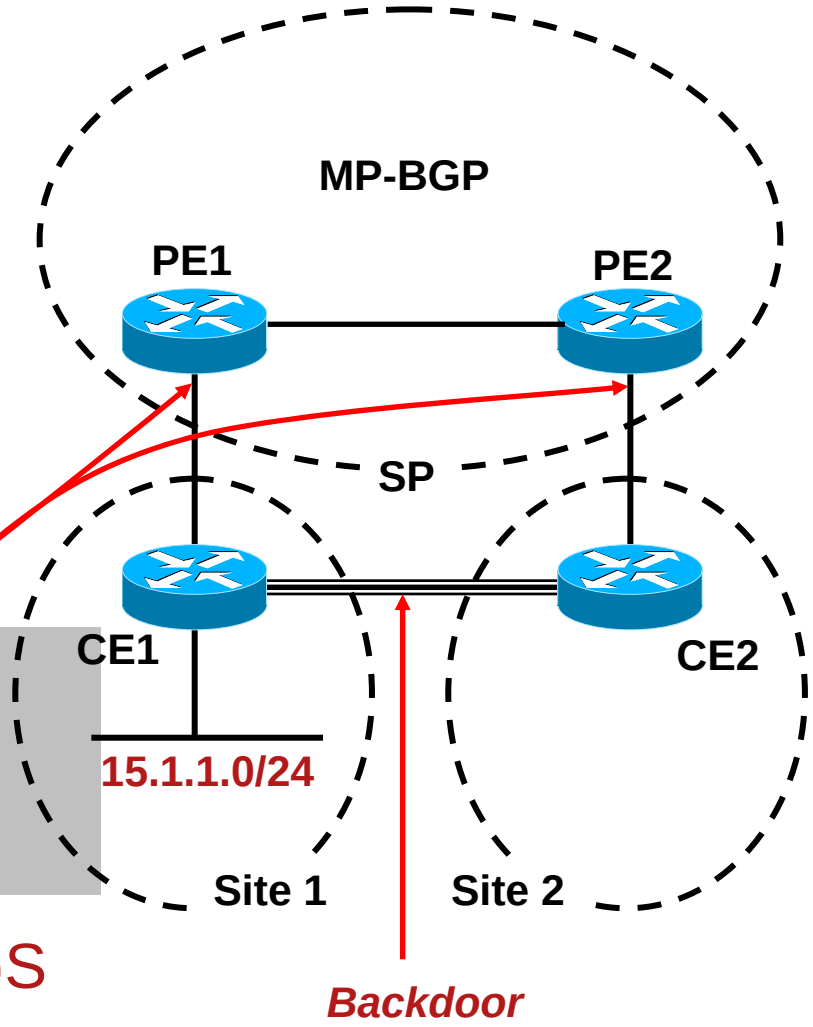

- 12.0(27)SV 12.0(26)SZ 12.0(26.1)S

MPLS VPN - PE / CE
PE Route Preference

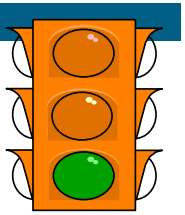

- CE2 receives 15.1.1.0/24 through the backdoor and advertises the route to PE2
- CE1 then advertises 15.1.1.0/24 and it is received at PE2 via MBGP
- The BGP best path algorithm at PE2 prefers the locally originated route from CE2
- To resolve this situation we need BGP to be aware of EIGRP metrics

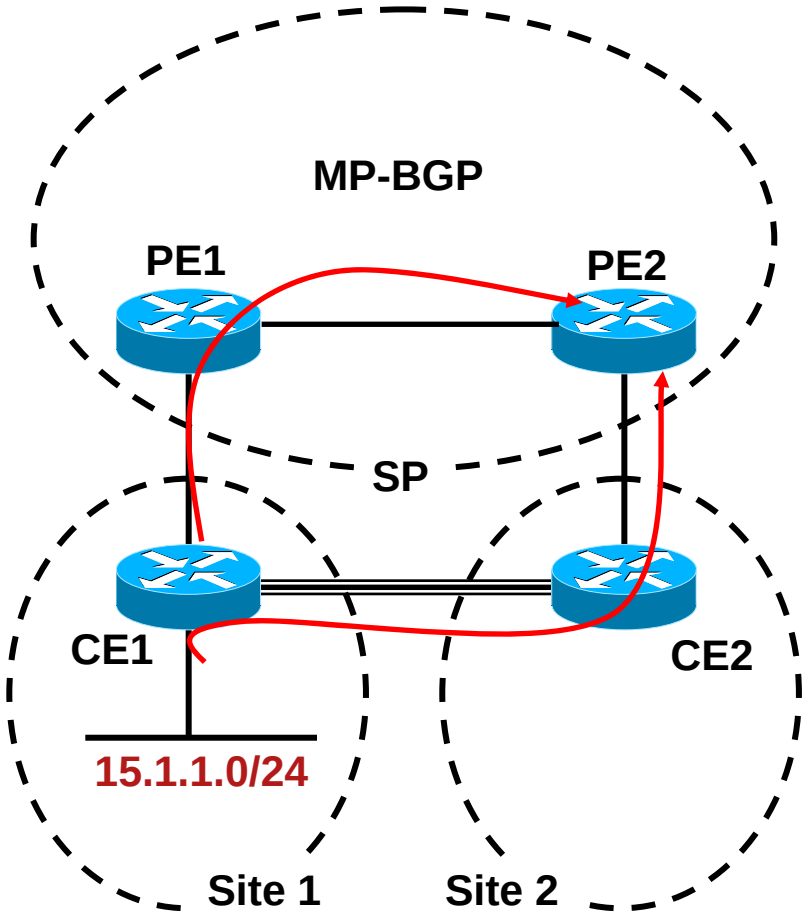

MPLS VPN - PE / CE
PE Route Preference

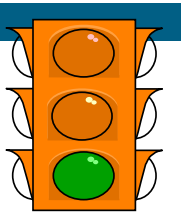

```
PE1#sh ip eigrp vrf site1 topology 15.0.0.0/8
IP-EIGRP (AS 65535): Topology entry for 15.0.0.0/8
...
Routing Descriptor Blocks:
  10.1.2.1 (Ethernet0/0), from 10.1.2.1, Send flag is
0x0
      Composite metric is (307200/281600), Route is..
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 2000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
PE1#
```

```
PE1#sh ip bgp vpnv4 vrf site1 15.0.0.0
BGP routing table entry for 1:1:15.0.0.0/8, version 150
...
Extended Community: RT:1:1 Cost:pre-bestpath:128:307200
0x8800:32768:0 0x8801:65535:51200 0x8802:65281:256000
0x8803:65281:1500
      mpls labels in/out 23/nolabel
PE1#
```

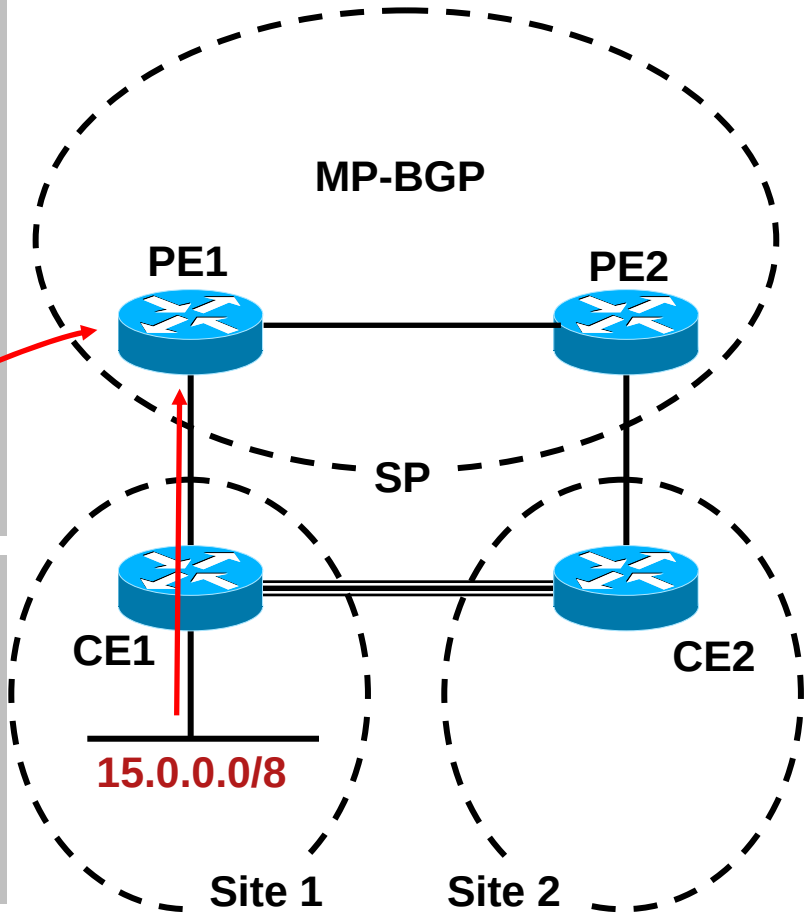

- **Pre-bestpath is automatic 12.0(27)S, 12.3(8)T, and 12.2(25)S**

MPLS VPN - PE / CE
Prefix Limits

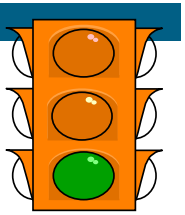

- This feature is primarily about protection E.g. Full BGP table being accidentally redistributed into EIGRP (or for monitoring)
- It can be done on specific redistribution statements or using a per-process scope

```
nw2005(config-router-af)#redistribute maximum-prefix
```

```
nw2005(config-router-af)#maximum-prefix 5000 ?
  <1-100>         Threshold value (%) at which to generate a warning message
  dampened        Exponentially increase restart time interval
  reset-time      Duration after which restart history is cleared
  restart         Duration for which a prefix source is ignored
  restart-count   Number of times sessions are auto-restarted
  warning-only    Only give warning message when limit is exceeded
```

- Available 12.0(29)S and 12.3(14)T

Multi Topology Routing Overview

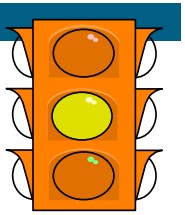

- A service-topology is a logical path that the traffic will take across the given network
- A service-topology will route/forward a sub-set of the traffic as defined by classification criteria
- Mapping traffic to a service-topology to determining which traffic (based on a classification criteria E.g. DSCP) gets the service-topology specific forwarding treatment
- QoS provides per-hop service differentiation within a single path and MTR provides PATH-BASED service differentiation within a single domain
- *EIGRP MTR is under current development so any output shown here is subject to change*

Multi Topology Routing

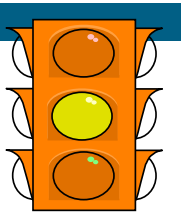

Base Topology
Voice Service-Topology
Video Service-Topology

DiffServ QoS Policies Apply on a Per Hop Basis

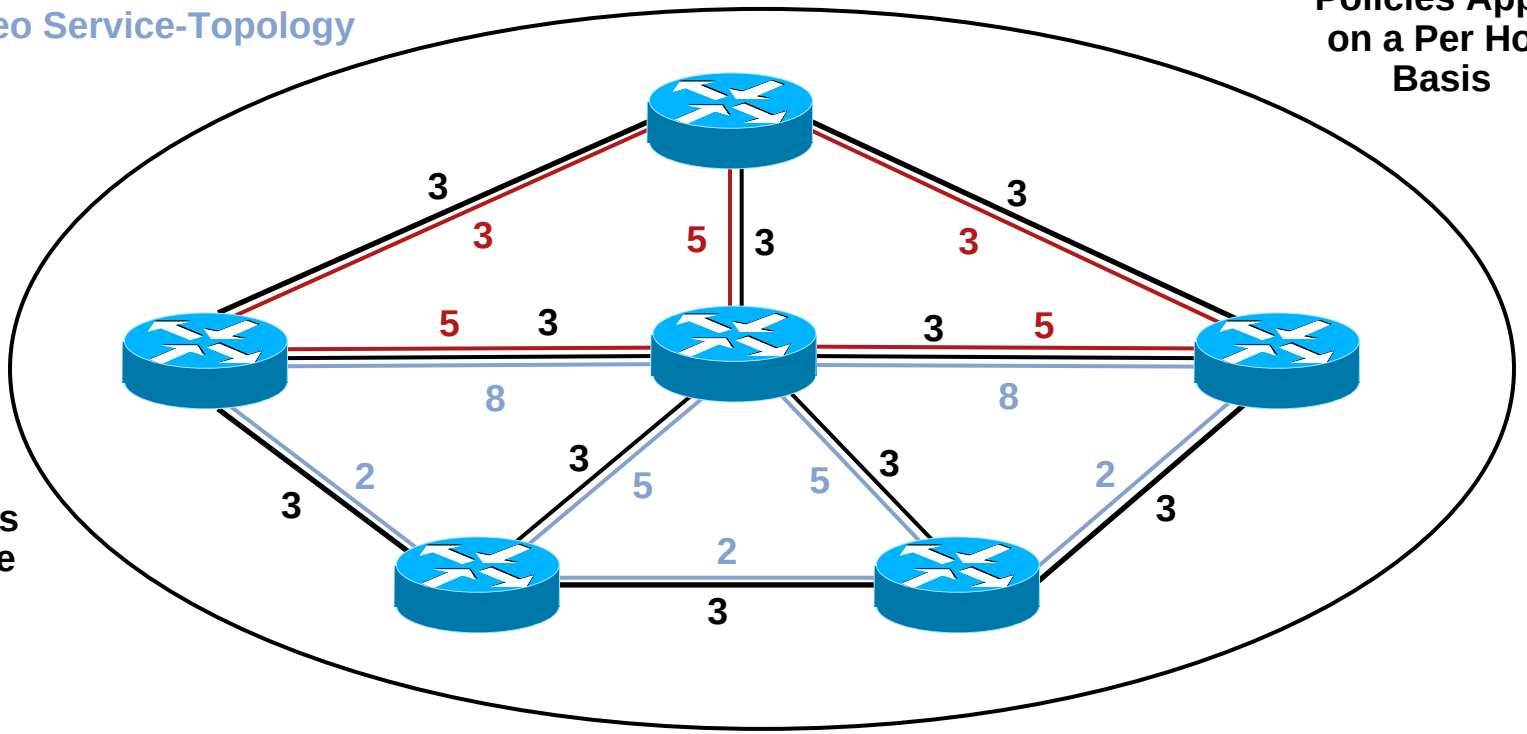

Classify and Mark the Packets at the Edge

Goal: Destination Based Routing Based on Traffic Types…

Service-Topologies Are Independent Entity in Itself

Each Link Can Have Multiple Service-Topologies

Multiple DSCP Values Can Be Mapped to a Single Service-Topology

MTR Topology Routing
Router Mode Configuration

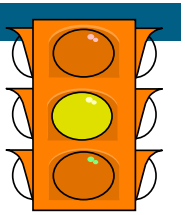

These slides focus only on some EIGRP specifics

```
!
router eigrp EMEA
 !
 address-family ipv4 unicast autonomous-system 5
  network 10.0.0.0
  passive-interface Loopback0
  !
  topology base
  !
  topology DATA tid 20
  !
  topology VOICE tid 10
!
```

"Named" or "Classic" Mode

Address-family

Service-Topologies

- *Classic Mode* is entered by specifying the "AS number after 'router eigrp'
- *Virtual Router, or "Named" Mode* is entered by providing a user definable string after 'router eigrp'

Multi Topology Routing
Interface Configuration

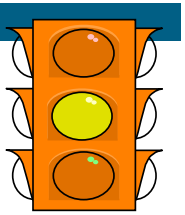

- A service-topology can have interface specific qualities such as the delay metric

```
EMEA1(config-if)#topology ipv4 BETA
EMEA1(config-if-topology)#eigrp 65535 ?
  delay            Set delay for EIGRP metric calculations
  next-hop-self    Configures EIGRP-IPv4 next-hop-self
  shutdown         Shutdown EIGRP for topology on this interface
  split-horizon    Perform split horizon
  summary-address  Perform address summarization
```

Multi Topology Routing
Routing Information Base

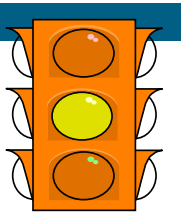

```
EMEA1#show ip route topology red 10.7.0.1
Routing entry for 10.7.0.0/16
  Known via "bgp 5", distance 200, metric 0
  Tag 7, type internal
  Last update from 10.5.0.4 00:46:05 ago
  Routing Descriptor Blocks:
  * 10.5.0.4, from 10.5.0.4, 00:46:05 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 7
```

Multi Topology Routing
EIGRP Service-Topology Aware CLI

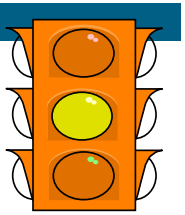

```
EMEA1#show ip eigrp topology red 10.5.0.4/32
EIGRP-IPv4 (AS 5): Topology(10) entry for 10.5.0.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2809856
  Routing Descriptor Blocks:
  10.5.2.2 (Serial1/0), from 10.5.2.2, Send flag is 0x0
      Composite metric is (2809856/2297856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 45000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
  10.5.1.2 (Serial2/0), from 10.5.1.2, Send flag is 0x0
      Composite metric is (4345856/2297856), Route is Internal
      ..
EMEA1#
```

PIX EIGRP Support

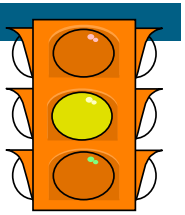

- The Cisco® PIX® EIGRP support is targeted for IOS® version 7.3
- Common portable EIGRP core code with a platform dependent OS-shim
- EIGRP stub and other key features
- MTR, IPv6 coming soon..
- Newer platforms supported (Note: 506, 515 & 520 are EOS/EOL)
- Additional CCO information

  http://www.cisco.com/go/pix

EIGRP MANET
Dynamic Cost Routing

CSCek40468

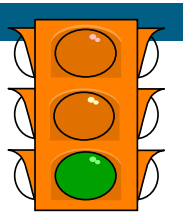

- Support for Mobile Ad-hoc Network deployments

- The fundamental requirement for MANET applications is effective integration of routing and radio technologies
- Effective routing requires immediate recognition of topology changes, the ability to respond to radio link quality fluctuations, and a means by which routers can receive and act upon feedback from a radio network
- New Virtual Multipoint Interface (VMI) and L2L3 API connects Layer 2 RF network with layer 3

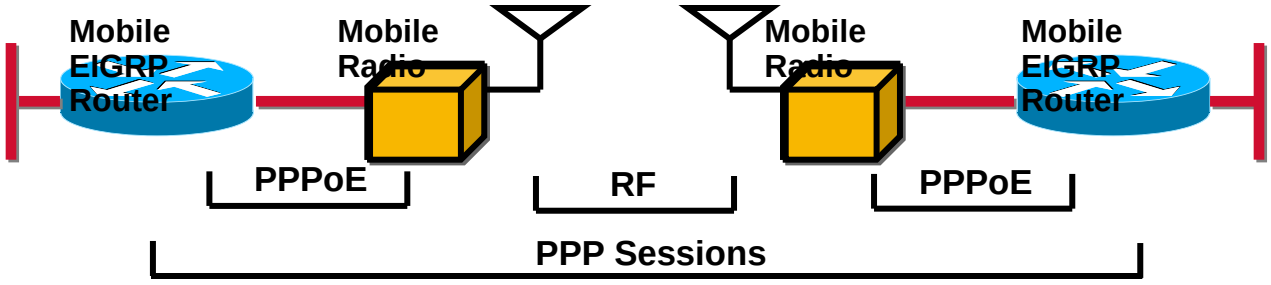

EIGRP MANET
Dynamic Cost Routing

CSCek40468

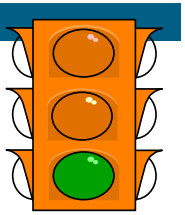

- The VMI interface maps multiple PPPoE sessions into a broadcast-capable multi-access interface
- The quality of a neighbor will vary based on raw radio link characteristics collected dynamically. It is from this that we compute the composite EIGRP metric based on a proprietary formula
- To avoid churn in the network through frequent changes a dampening mechanism is implemented
- Initial platform support for 2800, 3800, and 3200 routers supporting IPv4 and IPv6 in Cisco® IOS® 12.4T

OER EIGRP Support

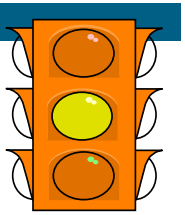

OER Components

- BR—Border Router (Forwarding Path)
- MC—Master Controller (Decision Maker)

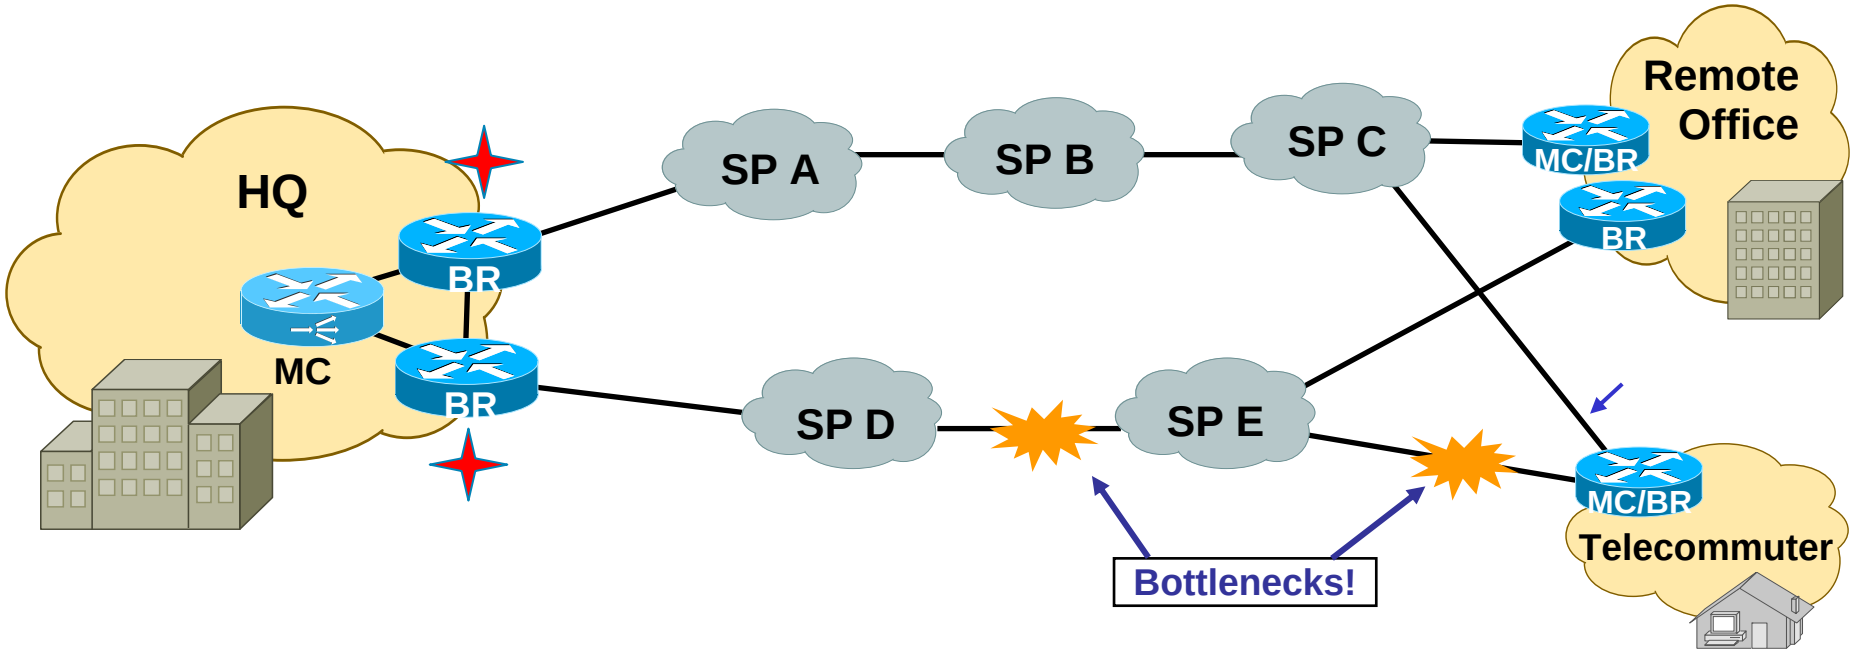

probe

www.cisco.com/go/oer/

OER EIGRP Support

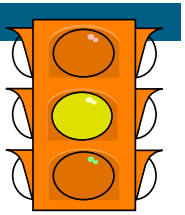

- Cisco® IOS® Optimized Edge Routing will support Route control using EIGRP
- Currently OER supports BGP, static routes and PBR only for route control
- Monitors traffic performance for prefixes passively with NetFlow and/or actively using IP SLA probes
- Chooses best performing path to a given destination
  - Delay, MOS
  - Load Balancing
  - For prefix, traffic-class and application

Meet the Experts
IP and MPLS Infrastructure Evolution

- Andy Kessler
  Technical Leader
- Beau Williamson
  Consulting Engineer
- Benoit Lourdelet
  IP services Product manager
- Bertrand Duvivier
  Consulting Systems Engineer
- Bruce Davie
  Cisco Fellow
- Bruce Pinsky
  Distinguished Support Engineer






Meet the Experts
IP and MPLS Infrastructure Evolution

- Gunter Van de Velde
  Technical Leader
- John Evans
  Distinguished Systems Engineer
- Oliver Boehmer
  Network Consulting Engineer
- Patrice Bellagamba
  Consulting Engineer
- Shannon McFarland
  Technical Leader





Meet the Experts
IP and MPLS Infrastructure Evolution

- Andres Gasson
  Consulting Systems Engineer

- Steve Simlo
  Consulting Engineer

- Toerless Eckert
  Technical Leader

- Dino Farinacci
  Cisco Fellow & Senior Software Engineer

Recommended Reading

BRKIPM -3008

- Routing TCP/IP, Volume I
- Optimal Routing Design
- EIGRP for IP
- EIGRP Network Design Solutions

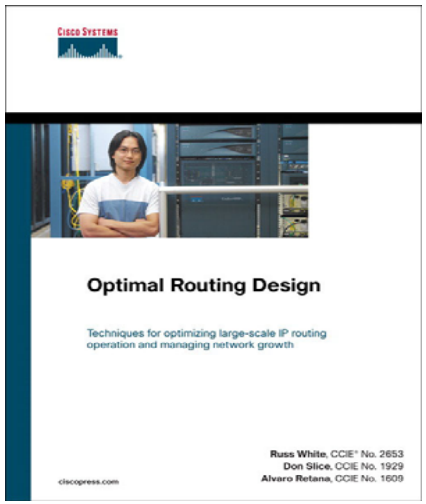

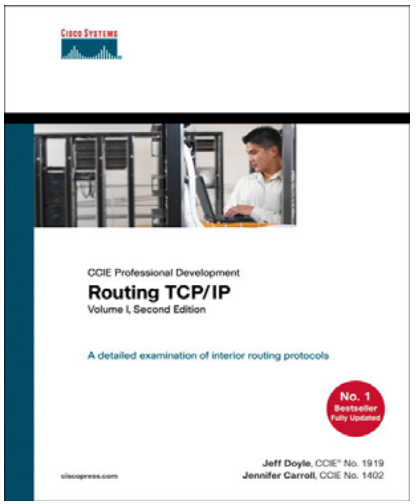

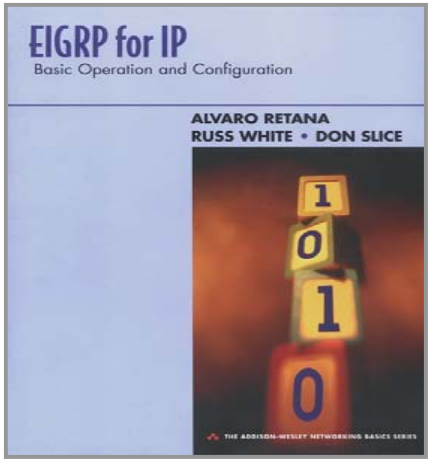

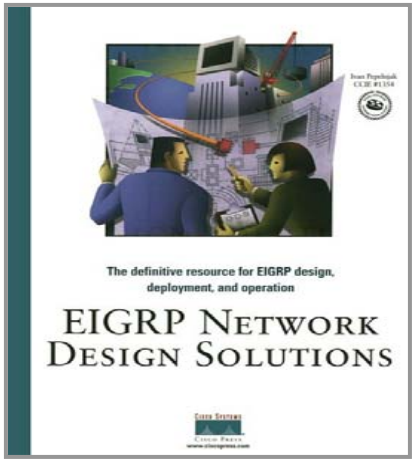

Available in the Cisco Company Store

Q and A