



# NETWORKERS 2004

## ADVANCES IN EIGRP

RST-4300

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

1

## Advances in EIGRP

Cisco.com

- **Scaling Enhancements**
- **Network Stability Enhancements**
- **Other Enhancements**



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

2

## SCALING ENHANCEMENTS



RST-2310  
9636\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

3

## Scaling Enhancements

Cisco.com

- EIGRP Stubs
- Single Peering Over Parallel Links



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

4

## EIGRP Stubs

Cisco.com

- When a router running EIGRP loses its connection to a network, it first searches for alternate loop free paths
- If it finds none, it then sends queries to each of its neighbors, looking for an alternate path



```
router-a#sho ip eigrp topo
IP-EIGRP Topology Table
....
P 10.1.1.0/24, 1 successors, FD is 281600
   via Connected, Ethernet1/2

router-a#show ip eigrp events
Event information for AS 100:
....
12 Active net/peers: 10.1.1.0/24 1
14 FC not sat Dmin/met: 4294967295 128256
15 Find FS: 10.1.1.0/24 128256
....
18 Conn rt down: 10.1.1.0/24 Ethernet 3/1
```

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

5

## EIGRP Stubs

Cisco.com

- If the neighbor has no path to this destination, it replies
- The router then removes all references to this route from its local tables
- In large hub and spoke networks, the hub routers have to build queries and process replies from each of the spokes
- This limits scaling!



```
router-a#show ip eigrp events
Event information for AS 100:
1 NDB delete: 10.1.1.0/24 1
....
12 Active net/peers: 10.1.1.0/24 1
14 FC not sat Dmin/met: 4294967295 128256
15 Find FS: 10.1.1.0/24 128256
....
18 Conn rt down: 10.1.1.0/24 Ethernet 3/1
```

RST-4300  
9859\_05\_2004\_X

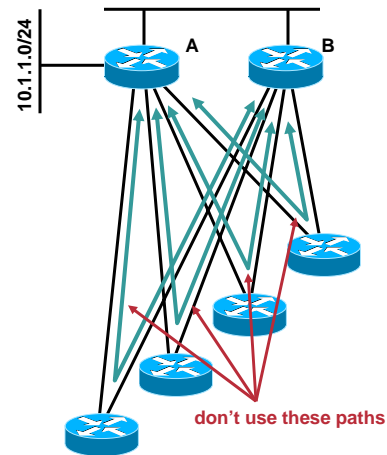
© 2004 Cisco Systems, Inc. All rights reserved.

6

## EIGRP Stubs

Cisco.com

- If these spokes are remote sites, they have two connections for resiliency, not so they can transit traffic between A and B
- A should never use the spokes as a path to anything, so there's no reason to learn about, or query for, routes through these spokes



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

7

## EIGRP Stubs

Cisco.com

- To signal A and B that the paths through the spokes should not be used, the spoke routers can be configured as stubs



```
router#config t#  
router(config)#router eigrp 100  
router(config-router)#EIGRP stub connected  
router(config-router)#
```

RST-4300  
9859\_05\_2004\_X

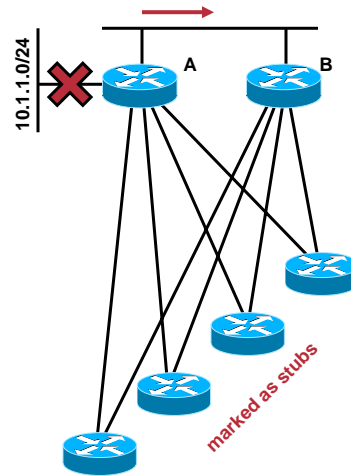
© 2004 Cisco Systems, Inc. All rights reserved.

8

## EIGRP Stubs

Cisco.com

- Marking the spokes as stubs allows them to signal A and B that they are not valid transit paths
- A will not query stubs, reducing the total number of queries in this example to 1
- Marking the remotes as stubs also reduces the complexity of this topology; B now believes it only has 1 path to 10.1.1.0/24, rather than 5



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

9

## EIGRP Stubs

Cisco.com

- If *stub connected* is configured
  - B will advertise 10.1.2.0/24 to A
  - B will not advertise 10.1.2.0/23, 10.1.3.0/23, or 10.1.4.0/24
- If *stub summary* is configured
  - B will advertise 10.1.2.0/23 to A
  - B will not advertise 10.1.2.0/24, 10.1.3.0/24, or 10.1.4.0/24



```

ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
interface serial 0
ip summary-address eigrp 10.1.1.0 255.255.254.0
!
router eigrp 100
 redistribute static 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.1
 network 10.1.2.0 0.0.0.255

```

eigrp stub connected

eigrp stub summary

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

10

## EIGRP Stubs

Cisco.com


- If *stub static* is configured

B will advertise 10.1.4.0/24 to A

B will not advertise 10.1.2.0/24, 10.1.2.0/23, or 10.1.3.0/24

- If *stub receive-only* is configured

B won't advertise anything to A, so A needs to have a static route to the networks behind B to reach them



```
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
interface serial 0
ip summary-address eigrp 10.1.1.0 255.255.254.0
!
router eigrp 100
redistribute static 1000 1 255 1 1500
network 10.2.2.2 0.0.0.1
network 10.1.2.0 0.0.0.255
eigrp stub static
eigrp stub receive-only
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

11

## EIGRP Stubs

Cisco.com

- Any combination of the route types can be specified on the *eigrp stub* statement, **except receive-only**, which cannot be used with any other option
- For example:
  - eigrp stub connected summary redistributed*
- If *eigrp stub* is specified without any options, it will send connected and summary routes

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

12

## EIGRP Stubs

Cisco.com

- At A, you can tell B is a stub using `show ip eigrp neighbor detail`.



```
router-a#show ip eigrp neighbor detail
IP-EIGRP neighbors for process 100
H  Address      Interface      Hold Uptime    SRTT  RTO  Q  Seq Type
   (sec)                (ms)          Cnt Num
0  10.2.2.3      Et1/2         10 00:00:50   320  1920  0  7
Version 12.2/1.2, Retrans: 0, Retries: 0
Stub Peer Advertising ( CONNECTED ) Routes
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

13

## EIGRP Stubs

Cisco.com

- EIGRP stubs are available in 12.0(6.3)T 12.0(6.1)PI 12.0(6.3)XE01(0.35)
- [http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1830/products\\_feature\\_guide09186a00800ab721.html](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1830/products_feature_guide09186a00800ab721.html)

RST-4300  
9859\_05\_2004\_X

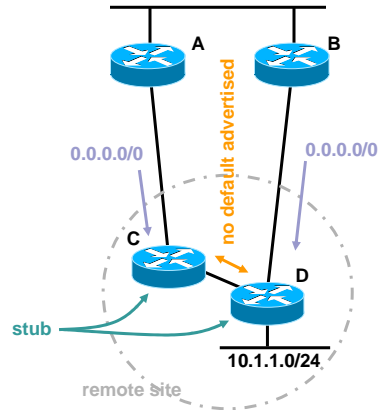
© 2004 Cisco Systems, Inc. All rights reserved.

14

# EIGRP Stubs

Cisco.com

- **Future Enhancement**
- **Dual router remotes are not supported with EIGRP stubs today**
  - C and D are a single remote site
- **If we mark C and D as stub routers**
  - C won't advertise a default learned from A to D
  - D won't advertise a default learned from B to C
- **If the B to D link fails, routing from 10.1.1.0/24 fails**
  - D doesn't have any route towards the core



RST-4300  
9859\_05\_2004\_X

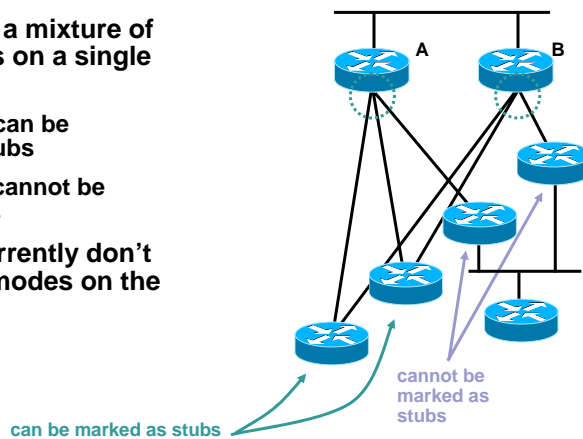
© 2004 Cisco Systems, Inc. All rights reserved.

15

# EIGRP Stubs

Cisco.com

- **Future Enhancement**
- **What if we have a mixture of stub capabilities on a single interface?**
  - Some routers can be declared as stubs
  - Other routers cannot be declared stubs
- **EIGRP stubs currently don't support mixed modes on the same interface**



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

16

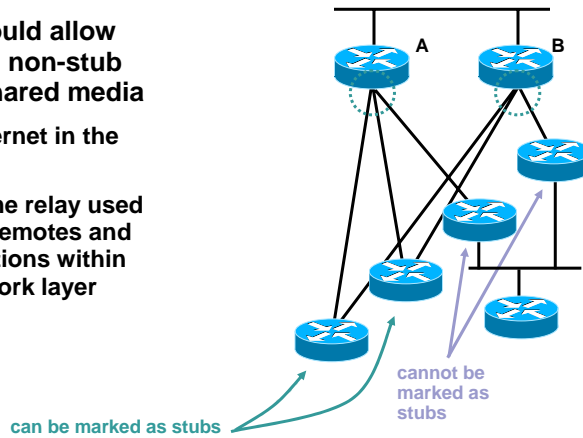
# EIGRP Stubs

Cisco.com

- Future Enhancement
- CSCdx74716 would allow mixing stub and non-stub neighbors on shared media

Useful for Ethernet in the data center

Multipoint frame relay used to connect to remotes and lateral connections within the same network layer



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

17

# EIGRP Stubs

Cisco.com

- Future Enhancement
- EIGRP Stub Site

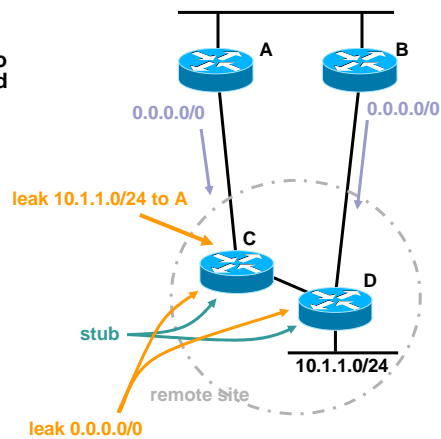
Route-maps used on C and D to identify which routes are leaked through which interfaces

C would allow advertising 10.1.1.0/24 on interface toward A, and 0.0.0.0/0 on interface toward D

C would not advertise 0.0.0.0/0 toward A, since it's not on the allowed list

- This allows both routers to be marked as stubs

The site won't ever be used to transit traffic between A and B



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

18


## Single Peering

Cisco.com

- **Future Enhancement**
- **EIGRP currently peers over every link between a pair of routers**

Each route learned at A will be advertised to B over every neighbor relationship

Each route placed in the active state will cause a query through every neighbor relationship



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

19

## Single Peering

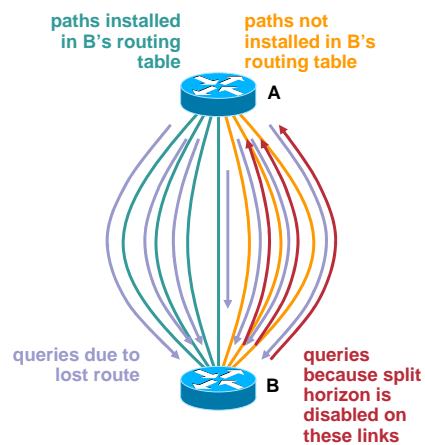
Cisco.com

- **Future Enhancement**
- **Slows down network convergence**

EIGRP only split horizons on routes actually installed in the routing table

If there are more paths than the maximum paths EIGRP can install in the routing table....

The routers end up querying each other, which can cause major problems in network convergence times



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

20

# Single Peering

Cisco.com

- **Future Enhancement**
- **EIGRP will eventually peer based on router IDs rather than interface addresses**
  - Only one neighbor relationship no matter how many links between the routers
- **Reduces routing traffic**
- **Increases convergence speed**
- **Decreases active times throughout the network**



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

21

## NETWORK STABILITY ENHANCEMENTS



RST-2310  
9859\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

22

## Network Stability Enhancements

Cisco.com

- EIGRP SIA Rewrite
- Multiple Autonomous System Support
- Neighbor Continuity Enhancements
- EIGRP 3-way Handshake
- EIGRP Graceful Shutdown
- EIGRP Graceful Restart



RST-4300  
9859\_05\_2004\_X

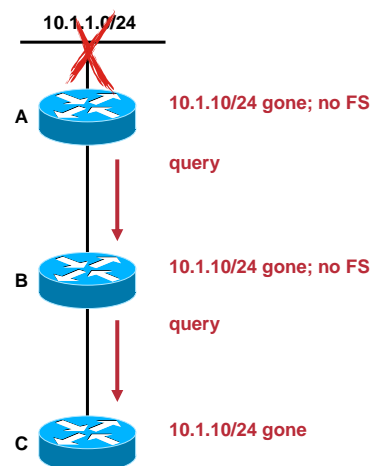
© 2004 Cisco Systems, Inc. All rights reserved.

23

## EIGRP SIA Rewrite

Cisco.com

- A loses its route to 10.1.1.0/24
  - No FS, mark route active
  - Set a 3 minute active timer
  - Query all neighbors (B)
- B receives A's query
  - No FS, mark route active
  - Set 3 minute active timer
  - Query all neighbors (C)
- C receives B's query
  - Examine local topology table
  - No feasible successors
  - No neighbors to query!



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

24

## EIGRP SIA Rewrite

Cisco.com

- C has no alternate path to 10.1.1.0/24
  - Remove from local tables
  - Reply to querying neighbors
- B receives C's reply
  - No outstanding queries
  - Remove from local tables
  - Reply to querying neighbors
- A receives B's reply
  - No outstanding queries
  - Remove from local tables



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

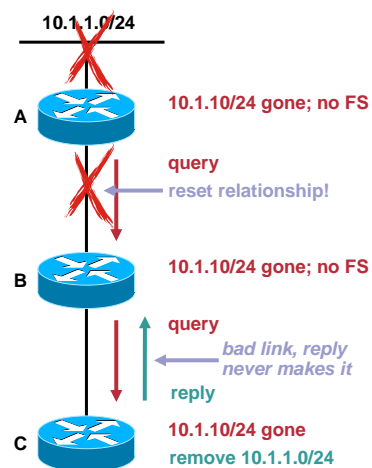
25

## EIGRP SIA Rewrite

Cisco.com

- If C sends the reply, and B never receives it, what happens?
- A's active timer (3 minutes) is still counting down while B and C are trying to get the reply back
- When this timer expires, A declares an SIA
  - The A/B neighbor relationship is reset

Why reset A/B when B/C is the problem??



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

26

## EIGRP SIA Rewrite

Cisco.com

- After the SIA rewrite  
12.1(4.0.3)T and 12.1(4.1)
- A sets its active timer to half  
the configured active time (1.5  
minutes, normally)
- After this time has passed, A  
sends an SIA Query  
  - If B acknowledges this query, A  
resets its timer, and the A/B  
neighbor relationship stays up
- B's relationship with C will fail  
at some point  
  - This clears the query from B's  
point of view
  - B replies to A



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

27

## Multiple AS Support

Cisco.com

Do You Really Want to Do This?



Okay, maybe it's not **that** bad...  
But we still wouldn't recommend it

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

28

## Multiple AS Support

Cisco.com

- A route is redistributed from RIP into AS 200
- At A, it is redistributed into AS 100
- B receives this route as well; which of the two externals will it prefer?
- There are two routes learned through separate routing processes with the same administrative distance, so the route installed first wins



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

29

## Multiple AS Support

Cisco.com

- If router B prefers the route through AS 100, it will redistribute the route back into AS 200
- If the redistribution metric at B is lower than the redistribution metric at C, A will prefer the path through B
- We have a permanent loop!



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

30

## Multiple AS Support

Cisco.com

- External routes can also carry administrative tags; as the external route is redistributed into AS 100 at A, it can be tagged
- This tag can then be used to block the redistribution of the route back into AS 200 at B



```
route-map settag permit 10
  set tag 200
!
router eigrp 100
  redistribute eigrp 200 route-map settag
  ....
```

```
route-map settag deny 10
  match tag 200
route-map settag permit 20
!
router eigrp 200
  redistribute eigrp 200 route-map settag
```

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

31

## Multiple AS Support

Cisco.com

- This blocks the formation of the loop, since A will no longer receive the redistributed from B through AS 200
- B still receives both routes, however, and could still choose the path through AS 100, resulting in suboptimal routing



```
route-map settag permit 10
  set tag 200
!
router eigrp 100
  redistribute eigrp 100 route-map settag
  ....
```

```
route-map filtertag deny 10
  match tag 200
route-map filtertag permit 20
!
router eigrp 200
  redistribute eigrp 100 route-map filtertag
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

32

## Multiple AS Support

Cisco.com

- CSCdm47037 resolves the routing loop and the suboptimal routing
- If two routes with the same administrative distances are compared, and the process type is the same (both EIGRP), then compare the metrics of the routes as well
- <http://www.cisco.com/cgi-bin/Support/Bugtool/onebug.pl?bugid=CSCdm47037>

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

33

## Multiple AS Support

Cisco.com

- If the redistribution metric is not manually set at A, it will be carried from AS 200 into 100
- The cost of the path between A and B is then added at B
- At B, the route through AS 200 wins; it has the lower metric



```
IP-EIGRP Topology Table for AS(100)/ID(10.0.17.10)
....
P 10.1.1.0/24, 1 successors, FD is 1500
   via 10.0.6.4 (1500/1000), FastEthernet0/0
....
IP-EIGRP Topology Table for AS(200)/ID(10.2.17.10)
....
P 10.1.1.0/24, 1 successors, FD is 1000
   via 10.2.8.20 (1000/256256), FastEthernet0/1
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

34

## Multiple AS Support

Cisco.com

- CSCdt43016, Support for Incoming Route Filtering Based on Route Maps, makes it possible to filter routes based on any route map condition before it is accepted into the local routing protocol database
- This is listed as an OSPF feature, but it works for all routing protocols
- <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122relnt/xprn122t/122tnewf.htm#33626>

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

35

## Multiple AS Support

Cisco.com

- This blocks the formation of the loop, since B will no longer have the path redistributed from A into AS 100 in its topology table
- This also prevents the suboptimal routing



```
route-map settag permit 10
  set tag 200
!
router eigrp 100
  redistribute eigrp 100 route-map settag
  ....
```

```
route-map settag deny 10
  match tag 200
route-map settag permit 20
!
router eigrp 100
  distribute-list filtertag filter in
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

36

# Neighbor Continuity Enhancements

Cisco.com

- Future Enhancement
- EIGRP currently resets its neighbor relationships for


Filtering configuration change

Summarization configuration

Others....

```
*Aug 27 13:06:26.758: %DUAL-5-NBRCHANGE:
IP-EIGRP(0) 100: Neighbor 10.1.1.1
(Serial0/0) is down: peer restarted
*Aug 27 13:06:27.976: %DUAL-5-NBRCHANGE:
IP-EIGRP(0) 100: Neighbor 10.1.1.1
(Serial0/0) is up: new adjacency
```

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



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

37

# Neighbor Continuity Enhancements

Cisco.com

- Future Enhancement
- EIGRP will use graceful restart techniques in the future to reduce neighbor resets to the minimum



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

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

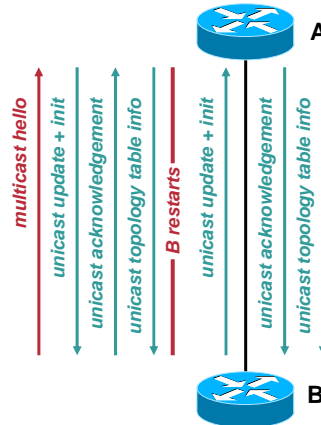
38



## EIGRP 3-Way Handshake

Cisco.com

- A receives a hello, sends the update with the initialization bit set, and B acknowledges it. A begins sending topology table information
- B restarts; once it comes back up, it will send an update with the initialization bit set
- A acknowledges this packet, and continues sending topology information to B



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

41

## EIGRP 3-Way Handshake

Cisco.com

- A and B are out of synchronization, but A doesn't know this
- B ignores the routing information A is sending, since it has lost its neighbor state with A
- Eventually, after A has retransmitted the same packet several times, it will tear down the neighbor relationship



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

42

## EIGRP 3-Way Handshake

Cisco.com

- To prevent EIGRP from forming neighbor relationships under either of these conditions, a new neighbor state is created: the pending state
- When A receives the first multicast hello from B, it places B in the pending state, and transmits a unicast update with the initialization bit set
- While B is in this state, A will not send it any queries or routing information



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

43

## EIGRP 3-Way Handshake

Cisco.com

- When B receives this update with the init bit set, it sends an update with the init bit set as well
- The acknowledgement for A's initial update is piggybacked onto this packet—it is never transmitted by itself
- There is no way for A to receive the acknowledgement for its initial update without also receiving B's initial update



RST-4300  
9859\_05\_2004\_X

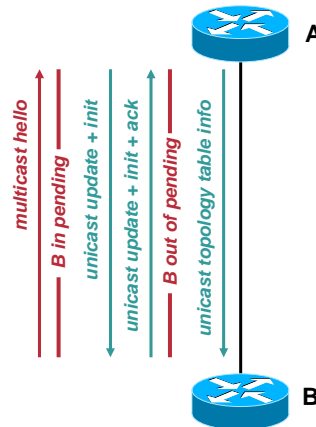
© 2004 Cisco Systems, Inc. All rights reserved.

44

## EIGRP 3-Way Handshake

Cisco.com

- Once the acknowledgement for its initial update is received, A takes B out of the pending state, and begins sending it topology information
- If this acknowledgement isn't ever received, hello's from B are ignored while A attempts to retransmit the initial update
- Eventually, A will time B out, and the process will start over



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

45

## EIGRP 3-way Handshake

Cisco.com

- Committed in 12.2(13.7)T2, 12.2(15.1)S, 12.2(16.1)B as CSCdy45118
- <http://www.cisco.com/cgi-bin/Support/Bugtool/onebug.pl?bugid=CSCdy45118>

RST-4300  
9859\_05\_2004\_X

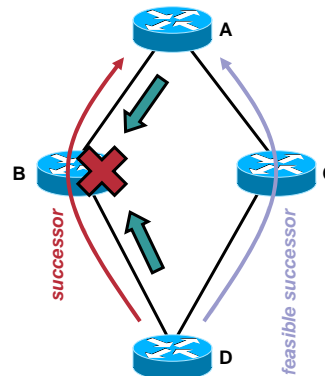
© 2004 Cisco Systems, Inc. All rights reserved.

46

## Graceful Shutdown

Cisco.com

- You want to bring B down for maintenance; the traffic will switch to C because EIGRP will reroute around B when B is taken down
- The packets on the wire will be lost when B is taken off line, though—and this could be a lot of packets, if these are high speed links
- It's better to get A and D to route around B while B can still forward traffic, so it happens gracefully



RST-4300  
9859\_05\_2004\_X

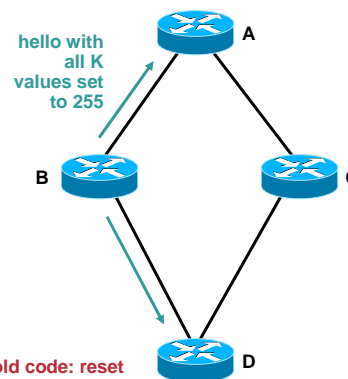
© 2004 Cisco Systems, Inc. All rights reserved.

47

## Graceful Shutdown

Cisco.com

- To shut down routing in B, remove the EIGRP router process
- EIGRP will send a goodbye message, notifying its peers that it is going down, before it cleans up the entries in the routing table
- B can send a goodbye message telling A and D to reset their neighbor relationships
  - Multicast or unicast hello with all K values set to 255
- `12.3(2.3)B 12.3(1.4)T 12.3(1.4)`



- old code: reset neighbor due to K value mismatch
- new code: reset neighbor due to goodbye message

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

48

## GR/NSF Fundamentals

Cisco.com

- **NonStop Forwarding (NSF)** is a way to continue forwarding packets while the control plane is recovering from a failure.
- **Graceful Restart (GR)** is a way to rebuild forwarding information in routing protocols when the control plane has recovered from a failure.
- *The fundamental premise of NSF/GR is to route through temporary failures, rather than around them!*

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

49

## EIGRP Graceful Restart/NSF

Cisco.com

- Router A loses its control plane for some period of time.
- It will take some time for Router B to recognize this failure, and react to it.



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

50

## EIGRP Graceful Restart/NSF

Cisco.com

- During the time that A has failed, and B has not detected the failure, B will continue forwarding traffic through A.
- Once the control plane resets, the data plane will reset as well, and this traffic will be dropped.
- NSF reduces or eliminates the traffic dropped while A's control plane is down.



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

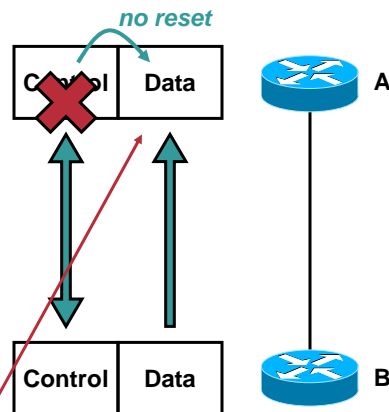
51

## EIGRP Graceful Restart/NSF

Cisco.com

- If A is NSF capable, the control plane will not reset the data plane when it restarts
- Instead, the forwarding information in the data plane is marked as stale.
- Any traffic B sends to A will still be switched based on the last known forwarding information.

*mark forwarding information as stale*



RST-4300  
9859\_05\_2004\_X

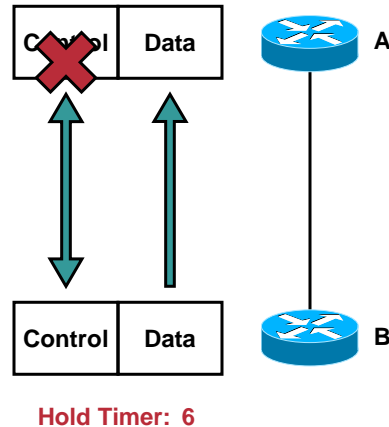
© 2004 Cisco Systems, Inc. All rights reserved.

52

## EIGRP Graceful Restart/NSF

Cisco.com

- While A's control plane is down, the routing protocol hold timer on B counts down....
- A has to come back up and signal B before B's hold timer expires, or B will route around it
- When A comes back up, it signals B that it is still forwarding traffic, and would like to resync
- This is the first step in Graceful Restart (GR)



RST-4300  
9859\_05\_2004\_X

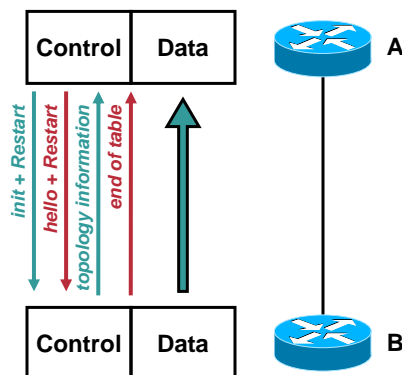
© 2004 Cisco Systems, Inc. All rights reserved.

53

## EIGRP Graceful Restart/NSF

Cisco.com

- The signal in EIGRP is an update with the initialization and restart (RS) bits set
- A sends its hellos with the restart bit set until GR is complete
- B transmits the routing information it knows to A
- When B is finished sending information, it sends a special end of table signal so A knows the table is complete



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

54

## EIGRP Graceful Restart/NSF

Cisco.com

- When A receives this end of table marker, it recalculates its topology table, and updates the local routing table
- When the local routing table is completely updated, EIGRP notifies CEF
- CEF then updates the forwarding tables, and removes all information marked as stale



RST-4300  
9859\_05\_2004\_X

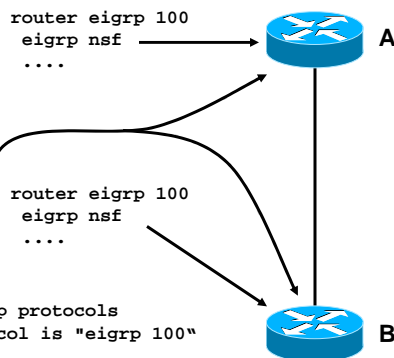
© 2004 Cisco Systems, Inc. All rights reserved.

55

## EIGRP Graceful Restart/NSF

Cisco.com

- *eigrp nsf* enables graceful restart
- *show ip protocols* verifies graceful restart is operational
- [http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products\\_feature\\_guide09186a0080160010.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080160010.html)



```
router eigrp 100
eigrp nsf
....

router eigrp 100
eigrp nsf
....

router#show ip protocols
Routing Protocol is "eigrp 100"
....
Redistributing: eigrp 100
EIGRP NSF-aware route hold timer is 240s
Automatic network summarization is in effect
Maximum path: 4
....
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

56

## EIGRP Graceful Restart/NSF

Cisco.com

- Routing protocol graceful restart is supported in IOS 12.2(15)T
- NonStop Forwarding is Supported on the:
  - Cisco 10000 and Cisco 12000 12.0(22)S
  - Cisco 7500 in 12.0(22)S, with the caveat that inserting a new standby RSP will cause some traffic loss, and switching from the primary to standby RSP will cause a microcode reload on the line cards
- [http://www.cisco.com/en/US/partner/tech/tk826/tk364/technologies\\_white\\_paper09186a008016317c.shtml](http://www.cisco.com/en/US/partner/tech/tk826/tk364/technologies_white_paper09186a008016317c.shtml)

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

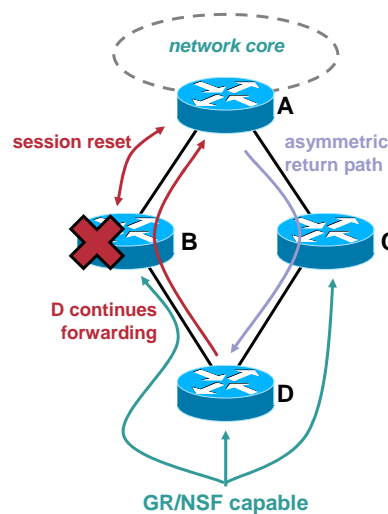
2 January 2003

57

## EIGRP Graceful Restart/NSF

Cisco.com

- Be careful with partial deployments of GR/NSF capability
- If B restarts, A will reset its session, removing all the routing information it learned from B. However, D will continue to forward traffic through B
- This will, at best, cause asymmetric routing. At worst, it could cause a routing loop
- A must be GR capable



RST-4300  
9859\_05\_2004\_X

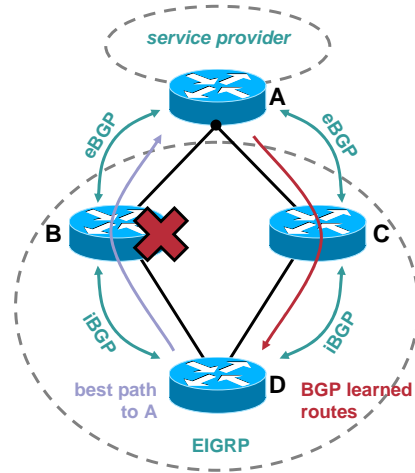
© 2004 Cisco Systems, Inc. All rights reserved.

58

## EIGRP Graceful Restart/NSF

Cisco.com

- EIGRP is configured for GR/NSF, while BGP is not
- D's next hop for all routes is A; the path to A is learned via EIGRP
- If the control plane on B restarts, D will continue learning BGP routes with a next hop of A through C; it will also maintain the best path to that next hop through B



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

59

## OTHER CURRENT ENHANCEMENTS



RST-2310  
9636\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

60

## Other Current Enhancements

Cisco.com

- EIGRP PE/CE Deployment
- EIGRP PE/CE Backdoor Links
- EIGRP Third Party Next Hop



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.


61

## EIGRP PE/CE Deployment

Cisco.com

- In this network, we have two corporate sites, connected by a leased line and VPN through a service provider
- EIGRP routes redistributed into BGP at B, and back into EIGRP at C, appear as external routes at Site 2

We want them to appear as internal routes



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

62

## EIGRP PE/CE Deployment

Cisco.com

- As routes are redistributed into BGP as B, extended communities containing the EIGRP metrics are attached to them
- As routes are redistributed back into EIGRP at C, these extended communities are used to reconstruct the routes as internals
- The VPN is considered a 0 cost link in this configuration



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.


63

## EIGRP PE/CE Deployment

Cisco.com

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

```
router-c#show ip eigrp vrf VRF-RED topology
IP-EIGRP Topology Table for AS(1)/ID(192.168.10.1)
 Routing Table:VRF-PINK
P 10.17.17.0/24, 1 successors, FD is 409600
 via 50.10.10.2 (409600/128256), Ethernet3/0
P 172.16.19.0/24, 1 successors, FD is 409600
```



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

64

## EIGRP PE/CE Deployment

Cisco.com

- 12.0(27)SV 12.0(21.1)SY2  
12.0(21.1)S2
- Backdoor links are not supported
- [http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products\\_feature\\_guide09186a0080154db3.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080154db3.html)



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

65

## EIGRP PE/CE Backdoor Links

Cisco.com

- The biggest danger with backdoor links is possible routing loops
  - Site1 advertises a network through the back door to site 2
  - C prefers this route, and redistributes it into BGP
  - B prefers the BGP route, and redistributes it into EIGRP, forming a loop
- The solution is to automatically tag all the routes originating in site 1 so they will be rejected by C
- This tag is called the Site of Origin (SoO)



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

66

## EIGRP PE/CE Backdoor Links

Cisco.com

- The SoO is set on all PE routers on the interface connecting to the PE, and on backdoor link routers
- The CE will always reject the marked EIGRP learned routes, and prefer the BGP learned routes.
- You can then set the backdoor link so the path through the VPN is always preferred over the backdoor link.



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

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

67

## EIGRP PE/CE Backdoor Links

Cisco.com

- 12.0(27)SV 12.0(26)SZ 12.0(26.1)S
- [http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a00801eff60.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801eff60.html)

RST-4300  
9859\_05\_2004\_X

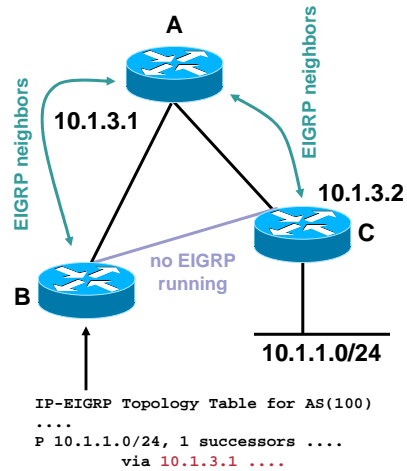
© 2004 Cisco Systems, Inc. All rights reserved.

68

## EIGRP Third Party Next Hop

Cisco.com

- C advertises 10.1.1.0/24 to A, which it readvertises it to B, and sets the next hop to 0.0.0.0
- When B receives this route, it sets the next hop to 10.1.3.1
- If B is receiving packets for 10.1.1.1, it will need to forward them through A, even though it has a direct connection to C



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

69

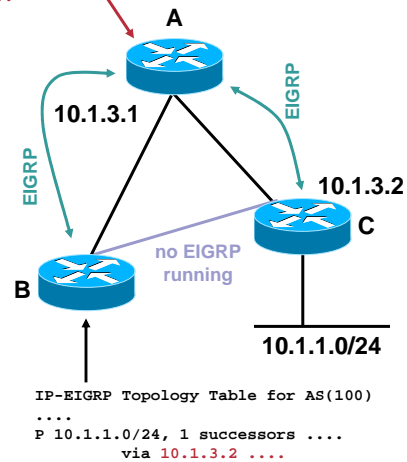
## EIGRP Third Party Next Hop

Cisco.com

- EIGRP third party next hop allows A to leave the next hop at 10.1.3.2
- B can then use the direct link between B and C to forward traffic to 10.1.1.0/24, even though EIGRP isn't running between the two routers

```

interface Serial 0
no ip next-hop-self eigrp 100
....
    
```



RST-4300  
9859\_05\_2004\_X

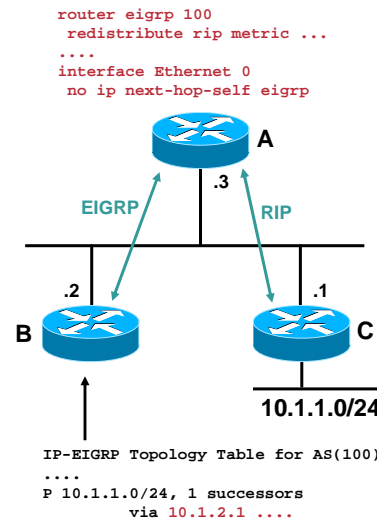
© 2004 Cisco Systems, Inc. All rights reserved.

70

## EIGRP Third Party Next Hop

Cisco.com

- Here A, B, and C are all connected to the same broadcast segment, 10.1.1.0/24
  - A is redistributing Rip into EIGRP
  - B isn't running RIP
  - C isn't running EIGRP
- B would normally show A as the next hop, rather than C, although it can reach C directly
- With `egrp no next hop self` on the Ethernet, A will send its updates to B with C as the next hop



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

71

## EIGRP Third Party Next Hop

Cisco.com

- Applications for third party next hop:
  - Dynamic Multipoint Virtual Private Networks  
[http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products\\_feature\\_guide09186a0080110ba1.html#1039490](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080110ba1.html#1039490)
  - Preserving the next hop in redistribution from broadcast networks
- <http://www.cisco.com/cgi-bin/Support/Bugtool/onebug.pl?bugid=CSCdk23784>

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

72

## EIGRP Enhanced Route Map Support

Cisco.com

- CSCdw22585 provides enhanced support of route maps for EIGRP
- This allows setting the metric 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
    
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

73

## EIGRP Enhanced Route Map Support

Cisco.com

|  |  |
|--|--|
| match tag 100                                  | matches against tags on internal routes                  |
| match tag external 100                         | matches against tags on external routes                  |
| match metric external 1000                     | matches against the external metric of an external route |
| match metric 1000 deviation 100                | matches routes with metrics from 900 to 1100             |
| match route-type external route-type bgp 65000 | matches routes sourced from BGP autonomous system 65500  |
| match route-type external route-type bgp 65000 | matches routes sourced from BGP autonomous system 65500  |

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

74

## EIGRP Enhanced Route Map Support

Cisco.com

|                                     |   |
|-------------------------------------|---|
| <b>match ip next-hop 10.1.1.1</b>   | matches against the next hop listed in the route            |
| <b>match interface serial 0/0</b>   | matches against the interface the route was learned through |
| <b>set metric 1000 1 255 1 1500</b> | sets the component metrics for a route                      |
| <b>set ip next-hop 10.1.1.1</b>     | sets the next hop listed in the route                       |
| <b>set tag 100</b>                  | sets the tag on internal routes                             |
| <b>set tag external 100</b>         | sets the tag on external routes                             |

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

75

## OTHER FUTURE ENHANCEMENTS



RST-2310  
9636\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

76

## Other Future Enhancements

Cisco.com

- EIGRP IPv6
- MIB Support
- Loadable EIGRP
- Dynamic Metrics
- Bundled Metrics
- Summary Only
- Leak Through a Summary
- Default Information Originate



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

77

## EIGRP IPv6

Cisco.com

- **Future Enhancement**
- **Initially, EIGRP IPv6 will**
  - Run EIGRP over an IPv6 multicast address
  - Add additional TLVs to the EIGRP packets to carry IPv6 addresses
  - Interface based configuration
- **In the future, all configuration will be moved under the EIGRP process**
  - Interfaces grouped by range, rather than network statements
  - IPv4 configuration will follow the IPv6 configuration style

### Per-interface configuration

```
router#conf t
Enter configuration commands, one per
line. End with CNTL/Z.
router(config)#ipv6 unicast
router(config)#int fastEthernet 0/0
router(config-if)#ipv6 eigrp 1
router(config-if)#exit
```

### Assign router-id and no shutdown

```
router(config)#!
router(config)#ipv6 router eigrp 1
router(config-rtr)#router-id 1.1.1.1
router(config-rtr)#no shutdown
router(config-rtr)#end
router#
```

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

78

## MIB Support

Cisco.com

- **Future Enhancement**
- **EIGRP will support 68 MIB objects in 4 major tables**
  - EIGRP traffic statistics
  - EIGRP topology data
  - EIGRP neighbor data
  - EIGRP interface data
- **A fifth table, the EIGRP VPN Table, is included for indexing**
- **eigrpRouteSIA and eigrpAuthFailure can trigger SNMP traps**

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

79

## MIB Support

Cisco.com

- **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 Neighbor Data**
  - Peer Address
  - Peer Interface
  - Hold Time
  - Up Time
  - SRTT/RTO
  - Version

And many more....

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

80

## Loadable EIGRP

Cisco.com

- **Future Enhancement**
- **Cisco IOS Software is currently monolithic**
  - To get a new version of any specific module, you need a new version of IOS
  - When you get a new version of IOS, you get a new version of everything, not just the module you wanted
- **Development is planned to break EIGRP off as a loadable unit**
  - This means you could load a new version of EIGRP into an existing running instance of IOS
  - New features and bug fixes could be targeted
  - You load only EIGRP

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

81

## Dynamic Metrics

Cisco.com

- **Future Enhancement**
- **EIGRP uses manually configured bandwidth and delay metrics**
- **EIGRP only “reads” the load and reliability metrics off an interface**
  - When the neighbor relationship comes up
  - When the bandwidth or delay are manually changed
- **Why can't we make these dynamic?**

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

82

## Dynamic Metrics

Cisco.com


- **Future Enhancement**
- The problem is EIGRP can easily get into positive feedback loops if metrics are read dynamically from the interface

As traffic is shifted to 10.1.1.0/24, the load increases

This drives traffic back to 10.1.2.0/24, increasing its load, decreasing 10.1.1.0/24's load

This, in turn, drives traffic back to 10.1.1.0/24, increasing its load, and decreasing 10.1.2.0/24's load

- **Constant Churn**



RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

83

## Dynamic Metrics

Cisco.com

- **Future Enhancement**
- What we need is a way to balance the traffic properly
  - We don't want any churn
- Cisco has a patented mechanism to allow the shifting of load without the churn
- This capability will be included in a future feature allowing dynamic metrics to be used in EIGRP



RST-4300  
9859\_05\_2004\_X

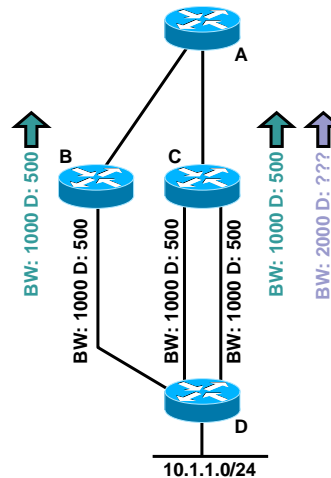
© 2004 Cisco Systems, Inc. All rights reserved.

84

## Bundled Metrics

Cisco.com

- **Future Enhancement**
- If there are multiple links used for load sharing, only a single link's bandwidth and delay will be advertised to neighbors
  - In this network, B and C advertise the same metrics to 10.1.1.0/24
  - C actually has more bandwidth available to reach 10.1.1.0/24
- EIGRP will be able to bundle the metrics of the equal cost links between C and D
- A will get a more true picture of the paths available to 10.1.1.0/24



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

85


## Summary Only

Cisco.com

- **Future Enhancement**
- A would like to advertise just the 10.1.0.0/16 to B and C
  - B and C don't need reachability to 10.2.1.0/24
- To do this, build
  - A summary for 10.1.0.0/16
  - A distribute list blocking everything except the summary towards B and C
- Create a new keyword on the summary statement
  - Just the summary would be advertised, and no other routes
  - Eliminates the distribute list

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

```
interface serial 0/0
 ip summary-address .... summary only
```



RST-4300  
9859\_05\_2004\_X

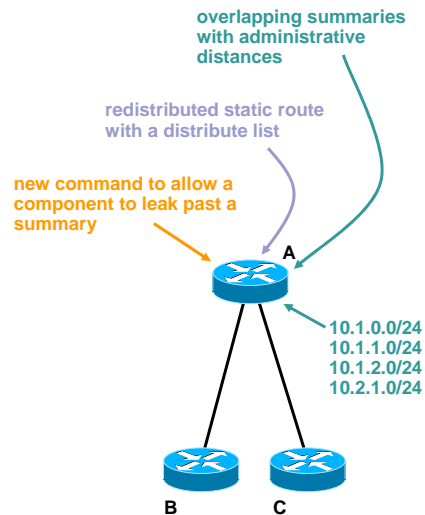
© 2004 Cisco Systems, Inc. All rights reserved.

86

## Leak Through a Summary

Cisco.com

- **Future Enhancement**
- **A wants to advertise the summary 10.1.0.0/16 and 10.1.1.0/24**
  - Use a pair of overlapping summaries, and play with the administrative distance
  - Use a static to null 0 for 10.1.0.0/16, rather than a summary, and then build the correct distribution list
  - There's no real easy way to do this today
- **Create a new command to allow a set of routes matching a route map to leak through a summary (or stub, etc.)**



RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

87

## Default Information Originate

Cisco.com

- **Future Enhancement**
- **The only options for originating a default route into EIGRP today are:**
  - A redistributed static route, which produces an external
  - A summary, which produces an internal, but isn't conditional
- **A new feature will be added to create the command *default-information originate* under *router eigrp***
  - A route map will be able to control when the default route is generated or not
  - The default route generated will be an EIGRP internal

RST-4300  
9859\_05\_2004\_X


© 2004 Cisco Systems, Inc. All rights reserved.

88

## Recommended Reading

Cisco.com

- Continue your Networkers learning experience with further reading for this session from Cisco Press.
- Check the Recommended Reading flyer for suggested books.



Available on-site at the Cisco Company Store

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

89

## Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

RST-4300  
9859\_05\_2004\_X

© 2004 Cisco Systems, Inc. All rights reserved.

90

**CISCO SYSTEMS**

