9. EIGRP Concepts

ROUTING PROTOCOL CHOICE

RIPv2 isn't used much in corporate environments, thanks to its slow convergence time (adaptation to network topology changes) and less useful hop-count metric. Thanks to its metric, RIP may well choose a path with fewer hops (routers between source and destination) even if those fewer links are slower.

EIGRP advantages

- Metric based on bandwidth and delay (OSPF based on bandwidth alone)
- Fast Convergence

EIGRP used to be proprietary. To this day, few non-Cisco vendors implement it on their routers.

DISTANCE-VECTOR CONCEPTS

Distance Vector (DV) Protocols—RIP is truly a DV protocol. EIGRP has some optimizations that blur the line. In DV, each router knows only:

- A destination's direction AKA vector (exit interface, next hop address)
- A destination's distance (metric)

This differs from link-state protocols, which store a map of the entire network area

Full Update Messages—Traditional DV protocols periodically send their entire routing table to their neighbors, whether or not changes have occurred. RIP does this every 30 seconds, in lieu of a hello message.

Split Horizon—Never advertise a route out that route's exit interface. Useful in DV protocols to avoid routing loops. Always safe because a route's next hop would already know about the destination and in a more reliable way since it's closer. In other words, don't pretend to be an expert when you're talking to a *real* expert. This is only true for DV; link-state protocols share a map instead of conclusions, and that's harmless.

Route Poisoning—When a route goes down, it is advertised with an infinity metric (unreachable). In RIP, infinity is implemented as 16, hence RIP's maximum hopcount of 15.

EIGRP NEIGHBORS AND UPDATES

Partial Updates—After initially sending route information, only updated or new information is sent to neighbors, unlike OSPF, which floods the topology every 30 minutes, or RIP, which tells about every route it knows every 30 seconds, in lieu of a hello.

Tables and how they are filled

Neighbor Table	Hello messages are sent out all involved interfaces that aren't passive.
Topology Table	At startup, neighbors exchange everything they know of the topology
Routing Table	From the topology table, EIGRP chooses and offers the best routes to the routing table

Neighbor Requirements—Hellos are multicast to 224.0.0.10. Routers can be neighbors if

- Their interfaces are on the same LAN and aren't passive (no hellos). In other words, the source IP address in the hello needs to be in the same subnet as the receiving interface.
- Same Autonomous System (AS) number
- If authentication is used, they need to pass
- The K values must match. These are the weightings in the metric calculation given to things like interface speed, delay, etc. Cisco recommends using the defaults, so this shouldn't matter.

Update Messages

- If sent to multiple routers on the same LAN, the multicast address is used. Otherwise, unicast.
- RTP (Reliable Transport Protocol) is used instead of TCP or UDP. This is not the same RTP that's used for IP telephony
- When the neighbor relationship is first formed, a full update (all known routes) is sent; thereafter, partials (just changes)
- The topology data doesn't cover the entire AS, but each router knows both its own metric *and* the next hop's metric for a route (used to avoid loops)

METRIC CALCULATION

By Default uses static items to reduce recalculation frequency

- Bandwidth—Minimum along the path in Kb/s
- Delay—Summed along the path. Default delays: Gigabit = 10 µsec, 100 Mb/s fast Ethernet = 100 µsec. This is automatically adapted by IOS, i.e. a gigabit port operated at 100 Mb/s will have a 100 µsec delay, just like a fast Ethernet port would.

Can also use (Cisco recommends against)

- Interface Load
- Reliability
- MTU (Max Transmission Unit)—Advertised, but not usable in calculation

Calculation

- Receive subnet #, min BW, cumulative delay
- Replace min BW if receiving interface is slower
- Add own delay to received cumulative delay

$$metric = \left(\frac{10^7}{bandwidth} + delay \right) x\, 256$$

Bandwidth is least along path in Kb/s.

Delay is summed along the path in 10s of µsecs.

Note that (10^7 / bw) is rounded down (truncated) before delay is added.

Tweaking—Serial links default to T1 speed (1.544 Mbps) bandwidth and 20,000 µsec delay (=0.02 secs) no matter what the real clock rate is.

```
R1(config-if)# bandwidth 64
                    1000's of bits per second
R1(config-if)# delay 1000
                    10's of microseconds (µsec)
```

Units Recap

Parameter	Calculation	Update Msgs	IOS Setting	IOS Show Interface
Bandwidth	1000s of bps		1000s of bps	1000s of bps (labeled)
Delay	10s of μsecs	μsec	10s of μsecs	μsec (labeled)

CONVERGENCE

Successor Route—The chosen best route from EIGRP.

Feasible Distance (FD)—Successor route metric = Distance to neighbor + their Reported Distance.

Reported Distance (RD)—The next-hop neighbor's metric for a route. This used to be called the Advertised Distance (AD). Perhaps AD was dropped because it could be confused with a route's Administrative Distance.

Feasible Successor—A non-best route that can be immediately used f the successor fails because it is known to be loop-free. This allows fast convergence. Defined as having an RD less that the successor's FD. In other words, it's impossible for that second-best route to be coming back through us. Link-State (OSPF) stores whole topology, so loops not a concern.

DUAL (Diffusing Update Algorithm)—The EIGRP process. When a route is lost and a feasible successor is unavailable, DUAL will send out queries to ensure that a non-feasible successor is still functioning via an EIGRP reply message from the next hop. Convergence should be possible in <10 seconds.

ROUTER ID

Like with OSPF, a router's ID is set in one of three ways. In order of preference:

- The “eigrp router-id x.x.x.x” command (note “eigrp router-id” vs. just “router-id” for OSPF)
- The highest “up/up” loopback
- The highest “up/up” interface