Chapter 5

EIGRP

- EIGRP is a Cisco-proprietary hybrid routing protocol that contains features of distance-vector and link-state routing protocols. Some of its features are:
  i) **Rapid convergence**. EIGRP uses the **Diffusing Update Algorithm** (DUAL) to achieve rapid convergence. DUAL not only calculates the best loop-free routes, but also calculates backup routes in advanced before they are actually being needed. An EIGRP router stores all available backup routes for fast react upon network topology changes. If no backup route exists in the routing table, an EIGRP router will query its neighbors until an alternative route is found.
  ii) **Reduced bandwidth usage**. EIGRP does not send periodic updates as with DV protocols. It sends **partial updates** upon the route information changes (eg: path, metric). Additionally, the update is propagated only to routers that require it, instead of all routers within an area as with LS routing protocols.
  iii) **Multiple routed protocols support**. EIGRP has been extended from IGRP to be network-layer independent. It supports IP, IPX, and AppleTalk with protocol-dependent modules (PDMs), which are responsible for protocol requirements specific to the corresponding routed protocols. EIGRP offers superior performance and stability when implemented in IPX and AppleTalk networks. EIGRP maintains a **neighbor table**, a **topology table**, and a **routing table** for each running routed protocols (PDMs).
  iv) **Support all LAN and WAN data link protocols and topologies**. EIGRP does not require special configuration across any L2 protocols. OSPF requires different configurations for different L2 protocols, eg: Ethernet and Frame Relay. EIGRP was designed to operate effectively in both LAN and WAN environments. EIGRP supports all multi-access networks, eg: Ethernet, Token Ring, FDDI, and all WAN topologies – leased lines, point-to-point links, and non-broadcast multiaccess (NBMA) topologies, eg: X.25, SMDS, ATM, and Frame Relay.

- EIGRP has its roots as a distance-vector routing protocol (EIGRP is based on IGRP). It is considered an advanced DV routing protocol with traditional DV features, eg: autosummarization, easy configuration; and LS features, eg: dynamic neighbor discovery. Another distance-vector rule is that if a neighbor is advertising a destination, it must also be using that route to forward packets to the particular destination.

- EIGRP (Enhanced IGRP) provides many enhancement features over IGRP, a traditional DV routing protocol, mainly in convergence properties and operating efficiency. Traditional DV routing protocols send periodic full routing updates, which consume unnecessary bandwidth.

- EIGRP utilizes multicasts and unicasts only; broadcasts are not being used. As a result, end systems will not affected by the routing updates and queries.

- EIGRP is a **transport layer** protocol that relies on IP packets to deliver its routing information. EIGRP packets are encapsulated in IP packets with the `Protocol Number` field value 88 (`0x58`) in the IP header. Some EIGRP packets are sent as multicasts (destination IP address `224.0.0.10`), while others are sent as unicasts.

- A significant advantage of EIGRP (and IGRP) over other routing protocols is the support for **unequal-cost load balancing**.

- EIGRP performs autosummarization by default, but this behavior can be disabled with the **`no auto-summary`** router subcommand.

- **Neighbor table** lists the directly connected adjacent EIGRP routers to ensure bidirectional communication with the neighbors. It is similar to the neighborship database in LS routing protocols. It maintains information such as address, hold time, and interface which an adjacent router connected to. An EIGRP router keeps a neighbor table for each running routed protocol. EIGRP routers must form neighbor relationships before exchanging EIGRP updates.

- **Topology table** maintains **all advertised routes to all destinations**, along with the advertising neighbors and advertised metric for each destination. The term "topology table" is confusingly named, as it does not actually store the complete network topology, but rather the routing tables from the directly connected neighbors. All successors and feasible successors to all destinations will be maintained in this table.

- The best routes to a destination will be selected from the EIGRP topology table and placed into the **routing table**. An EIGRP router maintains 1 routing table for each running routed protocol. It contains all best routes selected from the EIGRP topology table and other routing processes. Successors and feasible successors (when unequal-cost load balancing is enabled with the **variance** router subcommand) will be selected from the topology table and stored in this table.

- The **show ip eigrp neighbors**, **show ip eigrp topology**, and **show ip route** EXEC commands display the EIGRP neighbor table, EIGRP topology table, and routing table.

- **Successor** is the **lowest-metric best path** to reach a destination. EIGRP successor routes will be placed into the routing table.

- **Feasible Successor** (FS) is the **best alternative loop-free backup path** to reach a destination. Because it is not the least-cost or lowest-metric path, therefore it is not being selected as the primary path to forward packets and not being inserted into the routing table. Feasible successors are important as they allow an EIGRP router to recover immediately upon network failures and hence reduce the number of DUAL computations and therefore increase performance. The convergence time upon a successor failure with a feasible successor exists is in the range of 2 to 4 seconds (1 ping drop). Feasible successor routes are maintained in the topology table only.

- EIGRP routers use the following procedures to populate their routing tables:
  - i) Each router advertises its IP routing table to all adjacent neighbors in the neighbor table.
  - ii) Each router stores the routing tables of the adjacent neighbors in the topology table.
  - iii) Each router examines its topology database to determine the successor and feasible successor routes to every destination network.
  - iv) The best route to a destination network as selected from the EIGRP topology table or other routing processes is inserted into the routing table

yapchinhoong@hotmail.com

EIGRP Packet Format

- EIGRP sends out the following 5 types of packets:

Hello	Used to discovery neighbor before establishing adjacency. EIGRP Hellos are sent as **multicasts** and contain an acknowledgment number of 0. EIGRP routers must form neighbor relationships before exchanging EIGRP updates.
Update	Used to communicate the routes that a particular router has used to converge. EIGRP Updates are sent as **multicasts** when a new route is discovered or when convergence is completed (the route becomes passive); and are sent as **unicasts** when synchronizing topology tables with neighbors upon the EIGRP startup. They are sent **reliably** between EIGRP routers.
Query	Used to query other EIGRP neighbors for a feasible successor when DUAL is re-computing a route in which the router does not have a feasible successor. EIGRP Queries are sent **reliably** as **multicasts**.
Reply	Sent as the response to an EIGRP Query packet. EIGRP Replies are sent **reliably** as **unicasts**.
Acknowledge	Used to acknowledge EIGRP Updates, Queries, and Replies; Hello and ACK packets do not require acknowledgment. ACKs are Hello packets that contain no data and a **non-zero** acknowledgment number and are sent as **unicasts**.

- An EIGRP router sends Hello packets out all EIGRP-enabled interfaces. The EIGRP multicast address is 224.0.0.10. An EIGRP router only establishes neighbor relationships (adjacencies) with other routers within the **same autonomous system**.

- EIGRP Hello packets are sent every **5 seconds** on LANs (eg: Ethernet, Token Ring, and FDDI) and point-to-point links (eg: PPP, HDLC, Frame Relay and ATM point-to-point subinterfaces, and multipoint circuits with bandwidth greater than T1 (eg: ISDN PRI, ATM, and Frame Relay), and **60 seconds** on T1 or low-speed interfaces (eg: ISDN BRI, X.25, ATM, and Frame Relay). The **ip hello-interval eigrp** {*as-num*} {*sec*} interface subcommand configures the Hello interval for an EIGRP routing process running upon an interface.

- The EIGRP neighbor table also maintains the **hold time** – the amount of time a router considers a neighbor is up without receiving a Hello or any EIGRP packet from the particular neighbor. The **ip hold-time eigrp** {*as-num*} {*sec*} interface subcommand configures the hold time interval for an EIGRP routing process. The hold time interval is recommended to be at least 3 times the Hello interval. In fact, the hold time interval is 3 times the Hello interval by default. The hold time interval is not automatically adjusted upon the change of the Hello interval. Once the Hello interval is changed, the hold time interval must be manually configured according to the new Hello interval.
  Note: The newly configured hold time value affects neighbor routers instead of local router! Verify the newly configured hold time interval with the **show ip eigrp neighbors** EXEC command on neighbor routers. The hold time value is a parameter in the Hello packet; therefore a neighbor router which receives the Hello packet will be in effect.

- If an EIGRP packet is not received before the expiration of the hold time interval, the neighbor adjacency is deleted, and all topology table entries learnt from the neighbor will be removed, as well as send out an update indicating that the routes are unreachable. If the neighbor is a successor for some destination networks, those networks are removed from the routing table, and alternative paths will be recomputed via DUAL.

- A route is considered **passive** when the router is not performing recomputation for that route; while a route is considered **active** when the router is performing recomputation to seek for a new successor when the existing successor has become invalid.

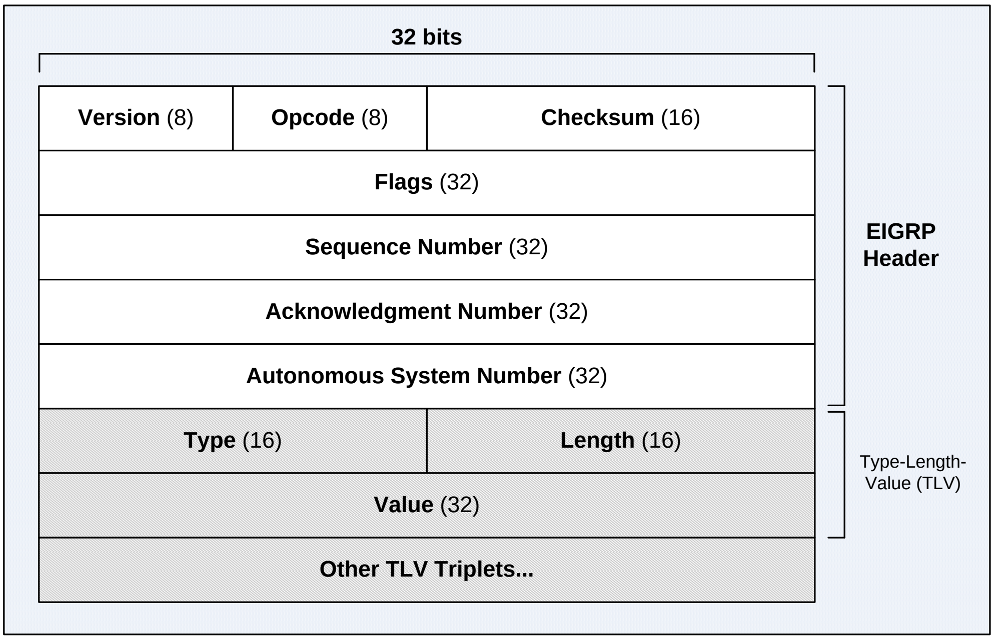

Figure 5-1: EIGRP Packet Format

- The EIGRP header comprises of the following fields:

Field	Description
Version	Identifies the EIGRP process version. The current EIGRP version is 2.
Opcode	Identifies the EIGRP packet type – Update (0x01), Query (0x03), Reply (0x04), Hello (0x05). It determines the TLVs that follow the EIGRP header. **Note:** ACKs are Hello packets that contain a **non-zero** ACK number.
Checksum	The checksum of the **entire EIGRP packet**, excluding the IP header.
Flags	1st LSB bit (0x00000001) – Init bit, used indicate the first set of routing updates upon establishing a new neighbor relationship. 2nd LSB bit (0x00000002) – Conditional Receive bit, used in the Cisco-proprietary reliable multicast protocol – **Reliable Transport Protocol** (RTP). Other bits are not being used.
SEQ and ACK	Used by RTP for reliable EIGRP message exchange.
AS Number	Identifies the autonomous system of an EIGRP packet. An EIGRP process only process EIGRP packets within an EIGRP domain (same AS number).
Type / Length / Value (TLV)	TLVs are comprise of a 16-bit Type field, a 16-bit Length field, and a vary number of fields depends on the type of TLV. **General TLVs:** 0x0001 – EIGRP parameters – K values and hold time. Size of 12 bytes. 0x0002 – Message Digest 5 (MD5) authentication data. Size of 40 bytes. 0x0003 – Sequence. Used by RTP. 0x0004 – Software versions – IOS and EIGRP release versions. Size 8 bytes. 0x0005 – Next Multicast Sequence. Used by RTP. 0x0006 – EIGRP stub parameters. **IP TLVs:** 0x0102 – IP internal route. Size of 28 bytes. 0x0103 – IP external route. Size of 48 bytes.

- An **internal route** contains a destination network within an EIGRP domain; while an **external route** contains a destination network outside an EIGRP domain, eg: redistributed routes from other routing processes into an EIGRP domain.

- EIGRP IP internal routes have a `Type` field of `0x0102`. EIGRP metric information is similar to IGRP, with 2 new fields – `Next Hop`, which can specify a different next-hop router than the advertising router to send packets destined to the `Destination/Prefix Length`; and `Prefix Length` (for VLSM support). The TLV triplets `Value` fields contain the information as in Figure 5-2.

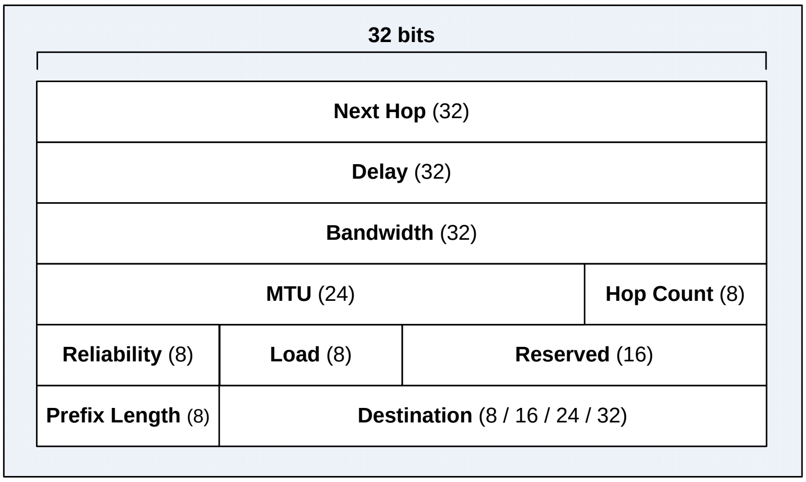

Figure 5-2: EIGRP IP Internal Route Packet Format

- In IGRP, when `RT2` sends an update to `RT1` with a destination network number `10.0.0.0`, the next-hop of `RT1` to the `10.0.0.0` network is `RT2`. With EIGRP, `RT2` can send an update to `RT1` with a `Next Hop` field of `RT3`. Ex: `RT2` and `RT3` are running RIP and `RT2` is redistributing RIP routes into EIGRP. When `RT2` sends an update to its neighbors on a shared network, `RT2` can tell them (`RT1`) to send traffic directly to `RT3` instead of `RT2`. This saves `RT2` from having to accept and reroute the packets to `RT3`. **Note:** This setup is not tested successfully yet!

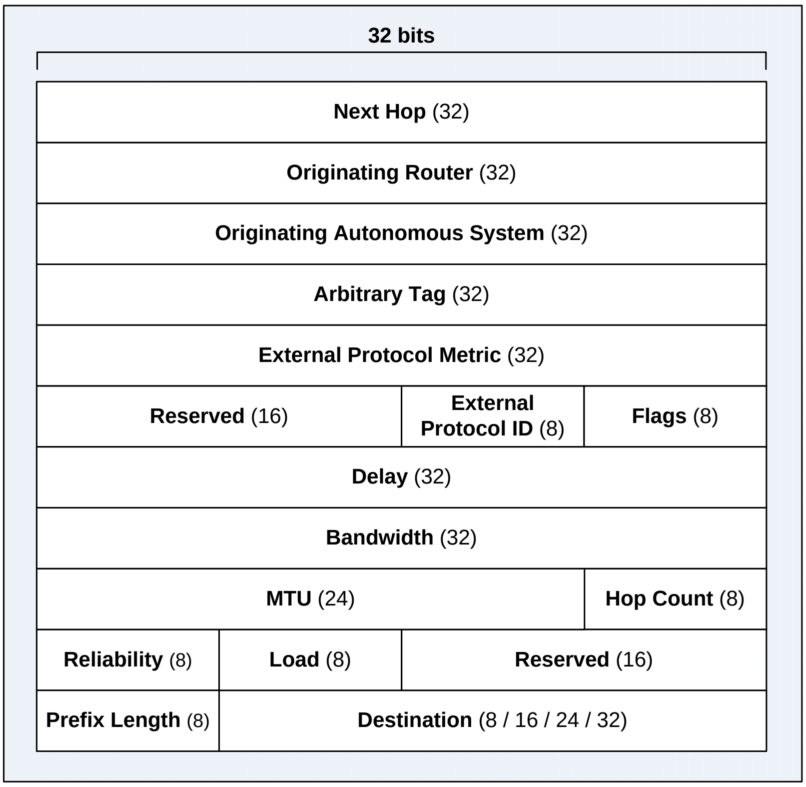

Figure 5-3: EIGRP IP External Route Packet Format

- The Originating Router, Originating AS, External Protocol Metric, and External Protocol ID fields specify the information about the source (router and routing protocol) from which an external route is derived. The common external routing protocol identifiers are 0x01 – IGRP, 0x02 – EIGRP (from different AS), 0x03 – Static Route, 0x04 – RIP, 0x06 – OSPF, 0x07 – IS-IS, 0x08 – EGP, 0x09 – BGP, 0x0A – ISO IDRP, and 0x0B – directly connected link. **Note:** The **Inter-Domain Routing Protocol** (IDRP) which is defined for OSI-based environments is similar to **Border Gateway Protocol** (BGP) in the TCP/IP environments. Cisco IOS Software does not support IDRP!

- The Arbitrary Tag field is used to carry a tag set by route maps – **route tagging**.

- A delay of 0xFFFFFFFF or 4294967295 (infinite distance) indicates an unreachable route.

EIGRP Neighbors

- EIGRP routers can become neighbors even if the Hello and hold time values do not match, in which the Hello and hold time intervals can be set independently on different routers. The EIGRP hold time has the same function as the OSPF dead time.

- EIGRP cannot form neighbor relationships using secondary addresses, as only primary addresses are used as the source IP addresses of all EIGRP packets. A neighbor relationship is formed between EIGRP routers when their primary addresses reside in the same subnet, they reside in the same AS domain, and the metric calculation constants – **K values** for the link are same.

- EIGRP routers send out multicast Hello packets to discover neighbors and form neighborships. An EIGRP router builds a neighbor table with Hello packets received from adjacent EIGRP routers running the same routed protocol. Only adjacent routers will exchange routing updates.

- When an EIGRP router discovers a new neighbor, it will send an update about its known routes to the new neighbor and receives the same information from the new neighbor. These updates are used to populate the topology table, which maintains the advertised metric from neighbors for all destinations and the metric that the local router uses to reach the destinations – **feasible distance**. The feasible distance metric values will be advertised to other neighbors.

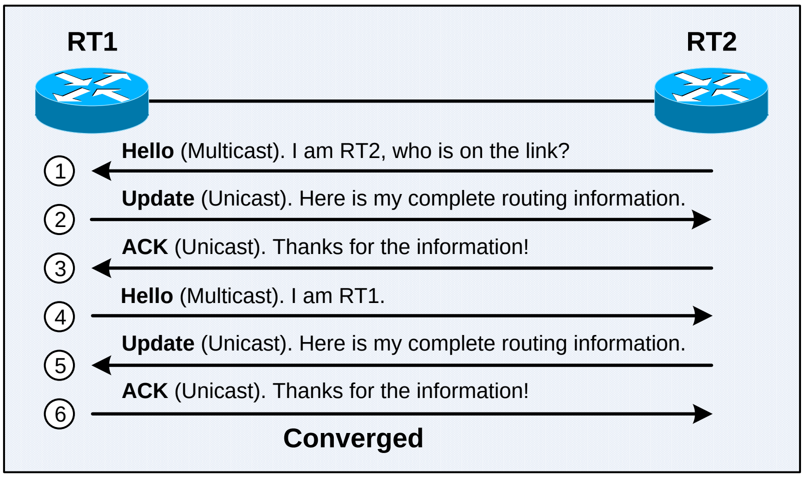

Figure 5-4: EIGRP Initialization – Neighbor and Route Discovery

- Below describes the steps in the EIGRP initial neighbor and route discovery process:
  i) An EIGRP router (RT2) sends out a Hello packet through all its interfaces.
  ii) RT1 which receives the Hello packet reply with a **startup update** packet that contain all its routes in its routing table (RT1 considers RT2 as its neighbor). Even though RT2 received an Update packet from RT1, it does not consider RT1 as its neighbor until it receives a Hello packet from RT1 – the EIGRP: Neighbor not yet found message will be displayed in the output of the **debug eigrp neighbors** privileged command. An EIGRP router considers another router as a neighbor and accepts the update information only after it received a Hello packet from the neighboring router.
  iii) RT2 replies RT1 with an ACK packet, indicating the receipt of the update information.
  iv) RT2 inserts the update information, which includes all destinations along with their associated metrics as advertised by neighboring routers – **advertised distance**.
  v) Step 4, 5, and 6 are similar to Step 1, 2, and 3.

  Note: Above shows a sample EIGRP neighbor establishment scenario. There are very high chances where both routers send Hello, Update, or ACK packets to each other at the same time. However, the sequence of receive Hello, reply with Update, and receive ACK will still follow.

- The **show ip eigrp topology** EXEC command displays only the successor and feasible successor routes, as well as DUAL states which are very useful for debugging EIGRP problems. The **show ip eigrp topology all-links** and **show ip eigrp topology detail-links** EXEC commands display all entries in the topology table.

```
Router#sh ip eigrp topology | b 172.16.1.0
P 172.16.1.0/24, 1 successors, FD is 2172416
        via 10.10.10.2 (2172416/28160), Serial0/0
Router#
Router#sh ip eigrp topology all-links | b 172.16.1.0
P 172.16.1.0/24, 1 successors, FD is 2172416, serno 24
        via 10.10.10.2 (2172416/28160), Serial0/0
        via 10.10.10.6 (2684416/2172416), Serial0/1
Router#
```

  Note: The highlighted path is not a feasible successor as its advertised distance is not less than the feasible distance – both values are same.

- EIGRP distinguishes itself from other routing protocols with its route selection process. EIGRP maintains successor (primary) and feasible successor (backup) routes and inserts them into the topology table (up to 6 per destination). The primary route is then inserted into the routing table.

EIGRP Metric Computation

- EIGRP supports the following types of routes:

Internal	Routes that are originated within an EIGRP autonomous system.
External	Routes that are learnt from another routing protocol or another EIGRP AS.
Summary	Routes that encompass multiple subnets. EIGRP summary routes have an administrative distance value of 5. AD values are **locally significant** and hence are not propagated to other routers.

- EIGRP uses **Diffusing Update Algorithm** (DUAL) to calculate and select **loop-free** primary and backup routes to a destination. When the primary routes fails, EIGRP can immediately uses a backup route without the need for holddown, and hence results in fast convergence.

- The EIGRP composite metric calculation can use up to 5 variables, but only the following 2 are used by default (K1 and K3):

K1 – **Bandwidth**	The minimum or lowest bandwidth between the source and destination.
K3 – **Delay**	The cumulative interface delay values along the path.

- The following variables are not commonly used, as they often cause **frequent recalculation** of the topology table.

K2 – **Load Utilization**	The worst load on a link between the source and destination based on the packet rate and the configured interface bandwidth.
K4 – **Reliability**	The worst reliability between the source and destination.
K5 – **Maximum Transmission Unit**	The smallest MTU along a path. MTU is included in the EIGRP update but is not used in the metric calculation.

- EIGRP calculates the metric to a network by adding weighted values for variables of the links. Below shows the weights attributed to the K variables:
  i) K1 = Bandwidth (1)
  ii) K2 = Load Utilization (0)
  iii) K3 = Delay (1)
  iv) K4 = Reliability (0)
  v) K5 = MTU (0)

- The EIGRP metric calculation formula is as below:

$$metric = \left[\left(\text{K1} \times bandwidth + \frac{\text{K2} \times bandwidth}{256 - load} + \text{K3} \times delay \right) \times \frac{\text{K5}}{\text{K4} + reliability} \right] \times 256$$

- When the weight for K5 as 0, the $\frac{\text{K5}}{\text{K4} + reliability}$ will not be in effect and will be taken as 1.

- The EIGRP metric calculation formula with default weighted K values will be simplied as:

$$metric = \left[\left(1 \times bandwidth + \frac{0 \times bandwidth}{256 - load} + 1 \times delay \right) \times 1 \right] \times 256$$

$$= \left(1 \times bandwidth + \frac{0 \times bandwidth}{256 - load} + 1 \times delay \right) \times 256$$

$$= \boldsymbol{(bandwidth + delay) \times 256}$$

- K values are transmitted in the EIGRP Parameters TLV in the EIGRP Hello packets. They should be modified only after careful planning and must be set identically on all routers within an EIGRP domain. Changing the weight for these constants is not recommended.

- The format of the bandwidth and delay values is different from those displayed in the **show interfaces** EXEC command. The EIGRP bandwidth metric value is calculated with the formula **10000000 / minimum bandwidth along the path (in kbps)**; while the EIGRP delay metric value is calculated as the **cumulative delays along the path (in tens of microseconds / usec)**.
Notes: tens of microseconds = 10 usec. The **show interfaces** EXEC command displays delay value in microseconds. By issuing the **delay 100** interface subcommand, the delay for an interface will be displayed as 1000 usec in the **show interfaces** EXEC command.

```
Router#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)#int s0/0
Router(config-if)#delay ?
  <1-16777215>  Throughput delay (tens of microseconds)

Router(config-if)#delay 100
Router(config-if)#do sh int s0/0 | in DLY
  MTU 1500 bytes, BW 1544 Kbit, DLY 1000 usec,
Router(config-if)#
```

- EIGRP metrics are represented in a 32-bit format (**1 – 4'294'967'296**) instead of the 24-bit format (**1 – 16'777'216**) as with IGRP. This provides additional granularity when determining the successor and feasible successor. Present-day bandwidth ranges from 9600bps to 10Gbps. EIGRP 32-bit metric format accommodates this range better than IGRP 24-bit metric format.

- EIGRP uses the same algorithm used by IGRP for metric calculation. The EIGRP metric is the IGRP metric multiplied by 256. When redistributing IGRP routes into an EIGRP domain, the IGRP metric is multiplied by 256 to compute the EIGRP-equivalent metric. When redistributing EIGRP routes to an IGRP routing domain, the EIGRP metric is divided by 256 to compute the IGRP-equivalent metric.

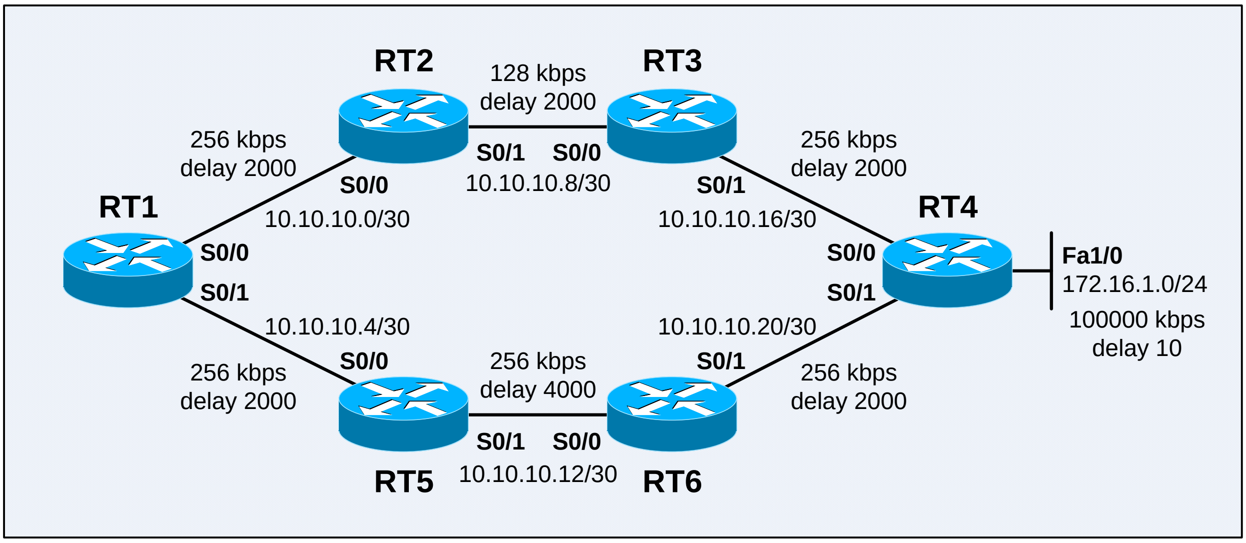

Figure 5-5: Sample Network for EIGRP Metric Calculation

- Figure 5-5 shows `RT1` has 2 paths to reach `172.16.1.0/24`. The bandwidth (in kbps) and delay (in tens of microseconds) of all links are also shown.

- The least bandwidth along the upper path (RT1 – RT2 – RT3 – RT4) is 128kbps.
  The EIGRP bandwidth calculation for this path is:

$$bandwidth = \frac{10000000}{128} = 78125$$

  The delay through this path is:

$$\begin{aligned} delay &= (\text{RT1-RT2}\, delay) + (\text{RT2-RT3}\, delay) + (\text{RT3-RT4}\, delay) + \text{RT4 Fa1/0}\, delay \\ &= 2000 + 2000 + 2000 + 10 \\ &= 6010 \end{aligned}$$

  Finally, the EIGRP metric for this path is:

$$\begin{aligned} metric &= (bandwidth + delay) \times 256 \\ &= (78125 + 6010) \times 256 \\ &= 84115 \times 256 \\ &= 21538560 \end{aligned}$$

- The least bandwidth along the upper path (RT1 – RT5 – RT6 – RT4) is 256kbps.
  The EIGRP bandwidth calculation for this path is:

$$bandwidth = \frac{10000000}{256} = 39062.5 = 39062$$

  The delay through this path is:

$$\begin{aligned} delay &= (\text{RT1-RT5}\, delay) + (\text{RT5-RT6}\, delay) + (\text{RT6-RT4}\, delay) + \text{RT4 Fa1/0}\, delay \\ &= 2000 + 4000 + 2000 + 10 \\ &= 8010 \end{aligned}$$

  Finally, the EIGRP metric for this path is:

$$\begin{aligned} metric &= (bandwidth + delay) \times 256 \\ &= (39062 + 8010) \times 256 \\ &= 47072 \times 256 \\ &= 12050432 \end{aligned}$$

- RT1 chooses the lower path over the upper path due to lower metric value. RT1 installs the route through the lower path with a next-hop of RT5 and a metric of 12050432 in its IP routing table.

```
RT1#sh ip eigrp topo 172.16.1.0/24
IP-EIGRP (AS 100): Topology entry for 172.16.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 12050432
  Routing Descriptor Blocks:
  10.10.10.6 (Serial0/1), from 10.10.10.6, Send flag is 0x0
      Composite metric is (12050432/11538432), Route is Internal
      Vector metric:
        Minimum bandwidth is 256 Kbit
        Total delay is 80100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 3
```

- The bottleneck along the upper path is the 128kbps link between RT2 and RT3, as this means that the throughput rate through this path would be at a maximum of 128kbps. While the lowest speed along the lower path is 256kbps, allowing the throughput rate up to that speed.

EIGRP DUAL – Diffusing Update Algorithm

- EIGRP uses the DUAL finite-state machine to tracks all routes advertised by all neighbors with the topology table, performs route computation on all routes to select an efficient and loop-free path to all destinations, and inserts the lowest metric route into the routing table.

- EIGRP uses the **advertised distance** and **feasible distance** to determine the **successor** (best route) and **feasible successor** (backup route) to a destination network.

Advertised Distance	The EIGRP metric for a next-hop EIGRP neighboring router to reach a destination network. Also known as **Reported Distance**.
Feasible Distance	The EIGRP metric for the local router to reach a destination network. Theoretically it is the sum of the advertise distance of an EIGRP neighbor and the metric to reach the neighbor, but actually the local router would recalculate the EIGRP metric to the destination network.

The xxx and yyy in the via A.B.C.D (xxx/yyy), *interface* entry in the **show ip eigrp topology** EXEC command represent **feasible distance** and **advertised distance** respectively.

- An EIGRP router compares all feasible distances to reach a destination network in its topology table. The lowest-metric route(s) will be placed in its IP routing table as the successor route(s).

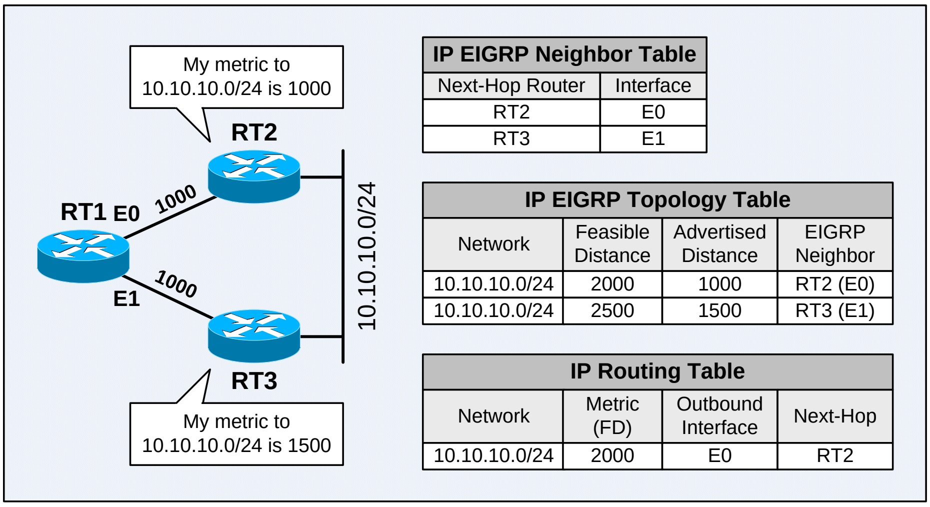

Figure 5-6: EIGRP Neighbor Table, Topology Table, and Routing Table

- RT1's topology table has 2 paths to reach 10.10.10.0/24. The EIGRP metric for RT1 to reach RT2 and RT3 is 1000. The metric (1000) will be added to the respective AD from each router, which represents the FD for RT1 to reach network 10.10.10.0/24 via a particular neighbor. RT1 chooses the least-cost FD (2000) as the best route and inserts it into its IP routing table. An EIGRP route in the routing table is the best FD selected from the EIGRP topology table.
  Note: Advertised distance is used only to calculate the feasible distance (the least-cost path); it does not affect the selection of the best route. Because there could be cases where the advertised distance from a neighbor is promising, but the metric to the neighbor is disappointing. Hence the path via the neighbor would not be selected as the feasible distance.

- **Successor** is a neighboring router that has the least-cost or lowest-metric (best) path of all possible paths to reach a destination network. After the best path to reach a network is selected, EIGRP inserts the destination network, metric to reach the network, outbound interface to reach the next-hop router, and the IP address of the next-hop router into the IP routing table. If the EIGRP topology table has multiple equal-cost FD entries to a destination network, all successors (4 by default, max 6) for the destination network will be inserted into the routing table.

- **Feasible Successor** is a neighboring router that does not does not provide the least-cost path but provides a backup route. The route through the feasible successor must be **loop-free**. Successors and feasible successors are selected at the same time during the metric computation process. Feasible successor routes are maintained in the topology table, as well as in the routing table when unequal-cost load balancing is enabled with the **variance** router subcommand.

  To be qualified as a feasible successor, the AD from a neighboring router must be **less than** (and not equal to) the FD of the successor route to a destination network. This ensures an FS is **loop-free** and will never route packets to the destination network through the local router, or else this would result in a routing loop. The AD of the local router (RT1) would often be greater than the FD of the neighboring feasible successor router (RT2) and will eventually notify the neighboring feasible successor router (RT2) that the local router is not a feasible successor for the destination network and never route packets to the destination network via the local router (RT1). When the AD from a neighboring router is greater than or same as the FD of the local router to the destination network, it means that the neighboring router could route packets to the destination network via the local router. Considering the path via the neighboring router as a feasible successor (backup route) and immediately uses the path to forward packets upon the failure of the successor route would result in routing loop.

- Every neighbor that satisfies the relation of AD < FD for a particular destination network is a loop-free route to that destination. This feasibility condition is proven in Loop-free Routing using Diffusing Computations by Dr. J. J. Garcia-Lunes-Aceves published by IEEE on Feb 1993.

- Below shows an excerpt from the output of the **show ip eigrp topology** EXEC command. The FD (successor) to the 10.10.10.12/30 network is via Fa0/0. The route via S1/0 is considered a FS as the AD from 10.10.10.1 (30720) is less than the FD (284160) to the network.

```
P 10.10.10.12/30, 1 successors, FD is 284160
         via 10.10.10.10 (284160/281600), FastEthernet0/0
         via 10.10.10.1 (2174976/30720), Serial1/0
```

- Refer to Figure 5-6, RT3 is an FS to 10.10.10.0/24, as the AD through RT3 (1500) is less than the FD of the current successor, RT2 (2000).

- All routing protocols can insert only the next-hop router information in the routing table, information about the subsequent routers along the path can never appear in the routing table. Each router counts on the next-hop router to reach a destination network. Each router makes a path selection to reach a network and installs the best next-hop address to reach the network. A router trusts a successor (the best next-hop router) to forward traffic to the destination network.

- The routing table is a subset of the topology table. The topology table contains more detailed information about each route, backup routers, and information used exclusively by DUAL.

- When a router loses a route, it first looks at the topology table for an FS. If there is an FS, the route does not go into active state, and the best FS is promoted as the successor and inserted into the routing table – an FS is used immediately without any route computation. If there is no FS, a route goes into active state. A new successor will be determined through the route computation process, in which an EIGRP router will begin the process by sending queries to all its neighbors. The amount of route computation time affects the convergence time.

- An EIGRP router would perform the following tasks upon the receipt of a query for a route:
  i) If the EIGRP topology table of the router does not contain an exact entry for the queried route, it immediately replies with a **network is unreachable** message to state that there is no path for the queried route through this neighboring router.
  ii) If the EIGRP topology table of the router lists the querying router as the successor for the queried route and contains one or more FSs, the FS with the lowest metric will be installed into the routing table and the router immediately replies the information of the FS to the querying router.
  iii) If the EIGRP topology table of the router lists the querying router as the successor for the queried route and does not contain a FS, the router propagates the query to all its neighboring routers except the querying router to seek for another alternative successor. The router will not reply to the querying router until it has received a reply for each query it propagated to its neighboring routers.
  iv) If the querying router is not the successor for the queried route, the router immediately replies the information of its successor to the neighboring router.

- For each neighbor to which a query is sent, an EIGRP router will set a **reply status flag** (r) to keep track of all outstanding queries that are waiting for replies. The DUAL computation is completed when the router has received a reply for every query sent out earlier.

```
Router>sh ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(10.10.10.1)

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

Passive (P)	Indicates an available and stable network which can be installed into the routing table.
Active (A)	Indicates an unavailable network which has outstanding queries and cannot be installed into the routing table.
Update (U)	Indicates that a network is being updated and the router is waiting for an acknowledgment for the update packet.
Query (Q)	Indicates that a network has an outstanding query packet or the router is waiting for an acknowledgment for a query packet.
Reply (R)	Indicates that the router is generating a reply for the network or is waiting for an acknowledgment for the reply packet.
Reply Status (r)	Indicates that the router is waiting for the reply to a query packet from a neighboring router.

- At the beginning of the DUAL computation, an Active timer will be set for 3 minutes. If the expected reply for each outstanding query is not received before the Active timer expires, the route will enter into the **Stuck-in-Active** (SIA) state. The router will reset the neighbor relationships for a neighbor that failed to reply, and cause the router to go active on all routes known through the neighbor and re-exchange routing information with the neighbor.

- SIA should not occur in well-designed and stable networks. If a neighboring router or a link fails, it should be detected by the expiry of the hold timer instead of the expiry of the active timer. Such problems are often not the root cause of SIA.

- SIA may occur due to looping of queries, heavily congested and/or low bandwidth links, and/or overloaded routers in a large network. SIA may lead to flushing of valid neighbors and routes which would affect the stability of a network.

- The **EIGRP SIA Rewrite** or **Active Process Enhancement** feature was introduced to provide improved network reliability by preventing unintended resets of neighbor relationships between querying and queried routers due to the SIA conditions.

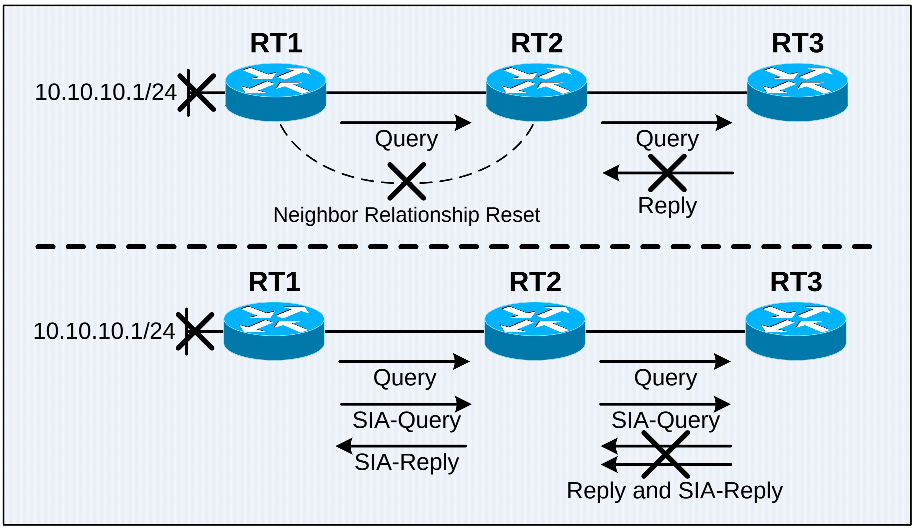

Figure 5-7: EIGRP SIA Rewrite

- Figure 5-7 shows the differences of the SIA conditions with and without the SIA Rewrite feature. When RT1 loss the route to 10.10.10.0/24, it would send a query for the network to RT2. RT2 would propagate the query to RT3 as it has no feasible successor for the network. As RT3 has no other neighbors to query, it would reply RT2 that it has no feasible successor. If RT2 never receives the reply from RT3 due to an errored link between RT2 and RT3, RT1 would reset the neighbor relationship between RT1 and RT2 after its Active timer expired. The question is why interrupts RT1 and RT2 when RT2 and RT3 are experiencing problems?

- With the SIA Rewrite feature, RT1 would query RT2 about the status of the route with a SIA-Query when the SIA-retransmit timer expires (90 seconds). RT2 would respond to the SIA-Query from RT1 with a SIA-Reply indicates that it is still searching for a new successor. Upon the receipt of SIA-Reply from RT2, RT1 would acknowledge the status of RT2 and reset the Active and SIA-retransmit timers. RT2 would also send a SIA-Query to RT3. If no SIA-Reply or reply to the original query for the active route is received from RT3 within the SIA-retransmit timer due to network link problem, RT2 would terminate the neighbor relationship with RT3 and inform RT1 that 10.10.10.0/24 is unreachable. RT1 and RT2 would remove the active route from their topology and routing tables and the neighbor relationship between them remains intact.
Note: When a router starts querying for an active route, it would activate the Active timer (3 minutes) and SIA-retransmit timer, which is half of the value of the Active timer (90 seconds).

- Assuming that no reply to the original query is received, an EIGRP router will send up to 3 SIA-Query messages as long as SIA-Reply messages are received before resetting a neighbor. Hence as long as a neighbor router responds to every SIA-Query within the SIA-transmit timer, it won't be declared SIA and reset for 6 minutes, assuming a default Active timer of 180 seconds. This gives more time for large networks to respond to the query-reply processes for active routes.

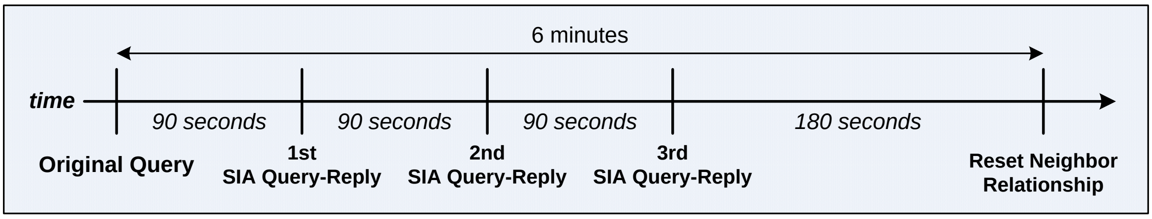

Figure 5-8: 3 SIA Query-Reply Cycles Maintain Neighbor Relationship up to 6 Minutes

EIGRP Reliability

- EIGRP reliability mechanism ensures the delivery of important routing information – Update, Query, and Reply packets, to neighboring routers in order to maintain a loop-free topology. A sequence number is assigned to every packet and requires an explicit acknowledgment for the sequence number. Acknowledgments are not necessary sent via ACK packets, as the ACK field in any RTP unicast packet is sufficient to acknowledge the received EIGRP packets. For efficiency purpose, only certain EIGRP packets are being transmitted reliably.

- **Reliable Transport Protocol** (RTP) is responsible for guaranteed and ordered delivery of EIGRP packets with the use of sequence and acknowledge numbers, but without any fancy windowing or congestion control mechanism (because only one packet will be sent at a time). It supports transmission of both multicast and unicast packets.

- Sending Hello packets to all neighbors individually is inefficient on a multi-access networks that provide multicast capabilities (eg: Ethernet), therefore Hello packets do not require acknowledgement are sent as **unreliable multicasts**. All packets that carry routing information – Update, Query, and Reply packets are sent reliably and require explicit acknowledgment.
  Note: Hello and ACK packets which are not being transmitted reliably have no sequence number.

- RTP ensures ongoing communication is maintained between neighboring routers by maintaining/a retransmission list for each neighbor. The list is used to track all the reliable packets that were sent but not acknowledged within the **Retransmission Time Out** (RTO). If the RTO timer expires before an ACK packet is received, EIGRP will transmit another copy of the reliable packet until the hold time expires and terminate the neighbor relationship.

- The use of reliable multicast packets is efficient. However, delays are potential to exist on multi-access media with multiple neighbors, as a reliable multicast packet cannot be transmitted until all peers have acknowledged the previous multicast packet. If a router is slow to respond, it would delay the transmission of next packet and affects all other routers. RTP is designed to handle such situation – neighbors that respond slow to multicasts would have the unacknowledged multicast packets retransmitted as unicasts when the RTO timer expires. This allows the reliable multicast operation to proceed without delaying communications with others and ensure low convergence time in the environments with variable-speed links.

- The multicast flow timer determines how long to wait for an ACK packet before switching from multicast to unicast; while the RTO determines how long to wait between subsequent unicasts. The EIGRP process for each neighbor calculates both the multicast flow timer and RTO timer based on the **Smooth Round-Trip Time** (SRTT). The formulas for the SRTT, RTO, and multicast flow timer are Cisco-proprietary.

- RTO is a dynamically adjusted over time. It is based on the SRTT, which specifies the average time in milliseconds between the transmission of a packet and the receipt of an acknowledgment. As more unacknowledged updates are sent, the SRTT would get higher and higher, which causes the RTO to increase exponentially. The maximum RTO value is 5000 ms (5 seconds).

- In a steady-state network where no routes are flapping, EIGRP waits for the specified hold-time interval to expire before determining that an EIGRP neighbor adjacency is down. By default, EIGRP waits up to 15 seconds for high-speed links and up to 180 seconds for low-speed links. When EIGRP determines that a neighbor is down and the router cannot reestablish the adjacency, the router will remove all reachable networks through that neighbor from the routing table. The router will attempt to find alternative paths to those networks when the convergence occurs.

- The 180-second hold time interval for low-speed links seems excessive, but it accommodates the links which are generally connected to less-critical remote sites. Additionally, 15 seconds is considered too long for some networks with high-speed links and serving mission-critical and time-sensitive applications. Always remember that there are situations and reasons in which modifying the default hold time interval is necessary in order to achieve faster convergence.

- When a local router sends an update packet to a remote router and the remote router does not acknowledge the packet, the local router would retransmit the update, every time the RTO expires, until the hold time expires and terminate the neighbor relationship. The statement which mentioned by other sources that the neighbor relationship will be terminated after 16 retransmissions rather than waiting for the hold timer to expire is inaccurate!

- EIGRP packets are only being generated at the moment of transmission. The transmit queues contain small and fixed-size structures that indicate which parts of the topology table to include in an EIGRP packet. As a result, the queues do not consume large amounts of memory and only the latest information will be transmitted. Ex: If a route changes state several times, only the last state is transmitted in the packet. This approach reduces bandwidth utilization.

- The **show ip eigrp traffic** EXEC command shows the numbers EIGRP packets sent and received on a router.

```
Router#sh ip eigrp traffic
IP-EIGRP Traffic Statistics for AS 100
  Hellos sent/received: 26/22
  Updates sent/received: 3/3
  Queries sent/received: 0/0
  Replies sent/received: 0/0
  Acks sent/received: 2/0
  Input queue high water mark 2, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 72
  PDM Process ID: 71

Router#
```