L43 - IRGP and EIGRP

IGRP and EIGRP

Cisco Routing Protocols

Agenda

- IGRP
- EIGRP

L43 - IRGP and EIGRP

IGRP Background

- IGRP stands for Interior Gateway Routing Protocol
  - designed and deployed successfully in 1986 by Cisco Systems
  - proprietary protocol
- Reason for IGRP?
  - At this time there was no alternative for RIP available
  - RIP disadvantages:
    - Metric limitation
      - max. 15 hops (16 hops = network unreachable)
      - Hop count doesn't reflect the capacity of the transmission medias
      - takes the least hop path instead of the fastest or „best path“
      - => didn't allow flexible routing in complex environments
    - much routing overhead (full routing table every 30seconds)

 IGRP-EIGRP, v3.5 3

IGRP at a glance

- Distance vector protocol (can only be used within an Autonomous System)
- Composite metric
  - Bandwidth
  - Delay
  - Reliability
  - Loading
- Implementation of loop avoidance mechanisms
- Support of multiple unequal-metric paths
- Faster convergence than RIP

 IGRP-EIGRP, v3.5 4

L43 - IRGP and EIGRP

IGRP / EIGRP Metric calculation

- Bandwidth
  - unit: bits/sec
  - default values for LANs
    - corresponding to real bandwidth
  - default values for serial lines
    - corresponding to bandwidth of 1.544Mbps (T1 line)
    - configuration of the real bandwidth on serial lines is a must!!!
    - cisco interface command: *bandwidth <number in kbit/s>*
  - minimum bandwidth along a path is taken
  - BW_{IGRP} is expressed by $(1/bandwidth)*10^{17}$
  - BW_{EIGRP} is expressed by $(1/bandwidth)*10^{17}*256$
  - The range is from a 1200bps line to 10 terabits per second

IGRP / EIGRP Metric calculation

- Bandwidth
  - metric values for bandwidth as part of composite metric:
    - assumption: bandwidth for serial links is configured properly
    - otherwise all serial links will have metric of T1

Bandwidth	BW_{EIGRP}	BW_{IGRP}
Satellite (500 Mbit/s)	5.120	20
Ethernet (100 Mbit/s)	256.000	100
Ethernet (10 Mbit/s)	256.000	1.000
Token Ring (4 Mbit/s)	640.000	2.500
Token Ring (16 Mbit/s)	160.000	625
FDDI (100 Mbit/s)	256.000	100
1.544 Mbps	1.657.856	6.476
128 kbps	20.000.000	78.125
64 kbps	40.000.000	156.250
56 kbps	45.714.176	178.571
10 kbps	256.000.000	1.000.000
1 kbps	2.560.000.000	10.000.000

L43 - IRGP and EIGRP

IGRP / EIGRP Metric calculation

- Delay
  - unit: 10 microseconds
  - it is the sum of the transmission delays along the path and is stored in a 32-bit field, in increments of 39.1 nanoseconds
    - all 1´s indicates an unreachable destination
  - default values for LANs
    - corresponding to real delay
  - default values for serial lines
    - corresponding to delay of 1.544Mbps (T1 line)
    - configuration of the real delay on serial lines is an option
    - cisco interface command: *delay <number in tens of usec>*
  - $Delay_{IGRP}$ is expressed by (delay/10)
  - $Delay_{EIGRP}$ is expressed by (delay/10)***256**

IGRP / EIGRP Metric calculation

- Delay
  - metric values for delay as part of composite metric:
    - assumption: delay for serial links is default value for T1

Delay	**$Delay_{EIGRP}$**	**$Delay_{IGRP}$**
Satellite (2 sec)	51.200.000	200.000
100 Mbit Ethernet (0,1 ms)	25.60	10
10 Mbit Ethernet (1 ms)	25.600	100
Token Ring 4 (2,5ms)	64.000	250
Token Ring 16 (0,6ms)	16.000	62,5
FDDI 100 (0,1ms)	2.560	10
serial links:		
1.544 Mbps (20 ms)	512.000	2.000
128 kbps (20 ms)	512.000	2.000
64 kbps (20 ms)	512.000	2.000
56 kbps (20 ms)	512.000	2.000
10 kbps (20 ms)	512.000	2.000
1 kbps (20 ms)	512.000	2.000

L43 - IRGP and EIGRP

IGRP Metric calculation

- Reliability
  - arbitrary number
    - 255 means 100%
    - 1 means 0%
  - worst reliability between source and destination
  - dynamically measured
    - keepalives are sent off the interface every 10 seconds, frame has CRC
    - samples are calculated over 5 minutes
  - time interval needs reconfiguration, when Reliability is used

IGRP Metric calculation

- Loading
  - arbitrary number
    - Load is given as a fraction of 255. A load of 255 indicates a completely saturated link
  - dynamically measured
    - calculated over 5 minutes
    - time interval needs reconfiguration, when Loading is used

L43 - IRGP and EIGRP

IGRP Metric calculation

- Formula for metric calculation
  - defaults:
    - k1=1, k2=0, k3=1, k4=0, k5=0
    - K1 til K5 are arbitrary numbers which can be configured

$$metric1 = k1 * \min BW_{IGRP} + \frac{k2 * \min BW_{IGRP}}{256 - load} + k3 * sumDelay_{IGRP}$$

If k5 doesn´t equal 0, an additional operation is done

$$compositemetric = metric1 * \frac{k5}{reliability + k4}$$

For default values of k parameters:

$$compositemetric = k1 * \min BW_{IGRP} + k3 * sumDelay_{IGRP}$$

 IGRP-EIGRP, v3.5 11

IGRP Metric example 1

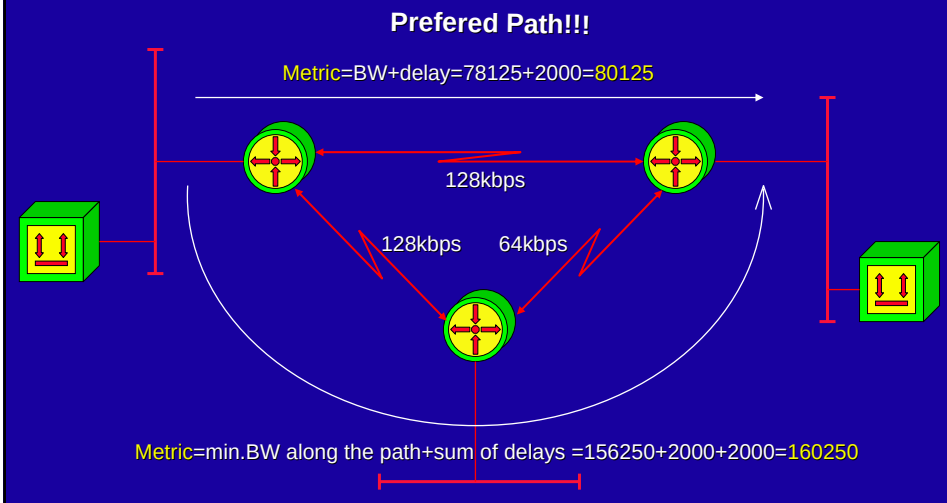

 IGRP-EIGRP, v3.5 12

L43 - IRGP and EIGRP

IGRP Metric example 2

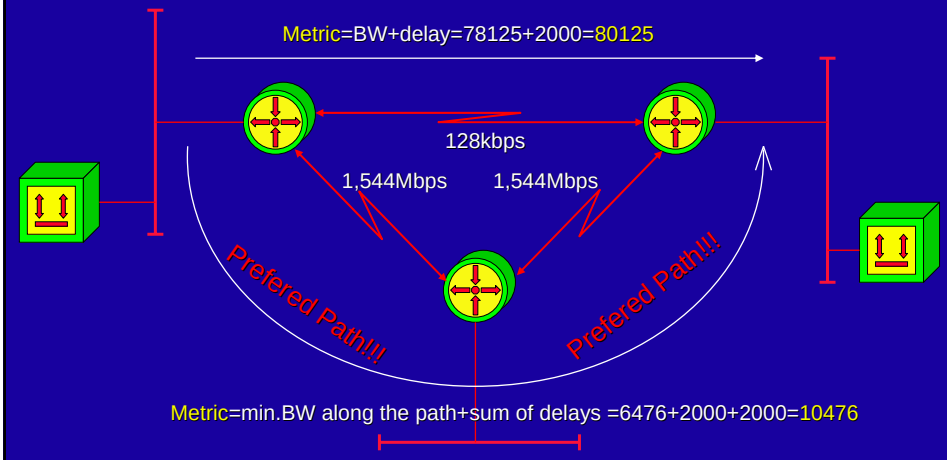

IGRP multiple paths

- Support of up to 6 parallel paths (default:4) to the destination for load balancing
  - default: same metric for parallel paths is necessary
- Support of unequal-metric load balancing
  - Prerequisite: configuration of a variance factor
  - The alternative path metric must be within the specified variance of the best local metric

L43 - IRGP and EIGRP

IGRP multiple paths

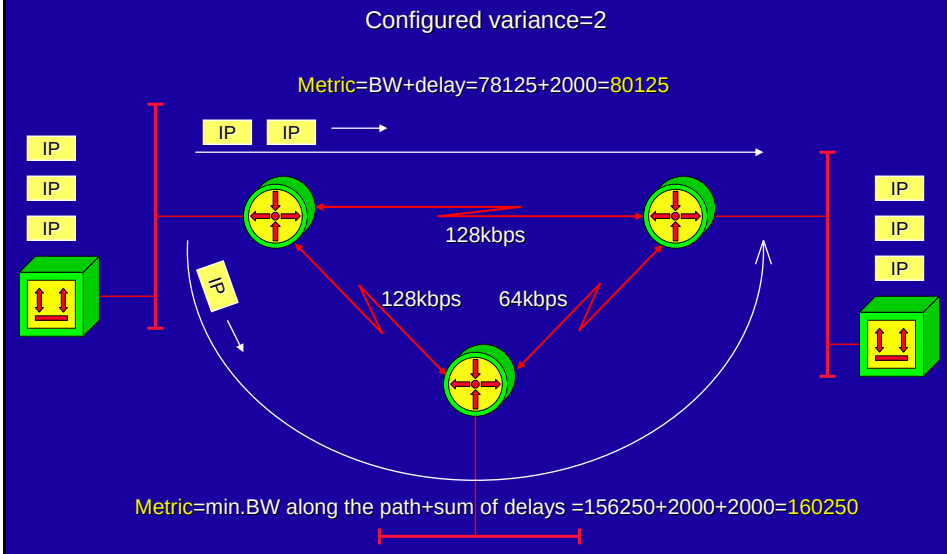

IGRP Metric

- Routing updates also include a count of hops and a computation of the path MTU
  - => max. network diameter 255 hops (IP TTL-field)
  - default: 100 hops

IGRP & default routes

- RIP and OSPF are using 0.0.0.0 as their default route (=>metric is not related to a distance)
- IGRP allows to mark real networks as „candidates for being default“
  - e.g. if multiple exit points to the Internet exist, then this implementation allows to choose the optimal „border router“

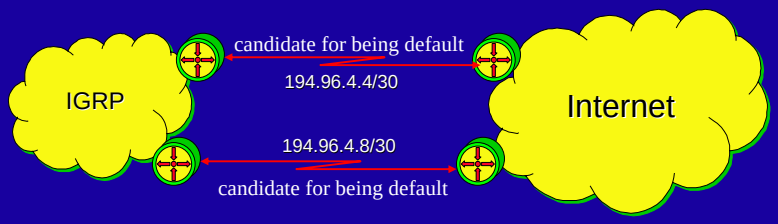

IGRP Loop avoidance & timers

- Periodical Routing updates
  - destination address: 255.255.255.255
- Topology changes are reported with triggered updates
- Hold-downs
- Split horizon
- Route Poisoning Updates
  - are intended to defeat larger routing loops
  - are sent if a route metric has increased by a factor of 1.1 or greater

IGRP Loop avoidance & timers

- Routing update timer: 90 seconds
- Hold-down timer: (3*90sec)+10sec=280sec
- Invalid timer: 3*90sec=270sec
  - (starts in the absence of routing information about a specific route)
- Flush timer: 7*90sec=630sec
  - How much time should pass before a route should be flushed from the routing table

IGRP Protocol stack

OSI stack

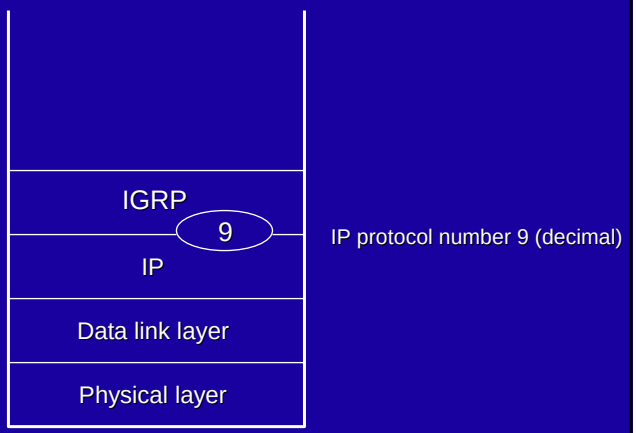

L43 - IRGP and EIGRP

Agenda

- IGRP
- EIGRP

EIGRP Background

- EIGRP stands for Enhanced Interior Gateway Routing Protocol
  - proprietary protocol
  - tries to combine the advantages of the distance vector and the link state protocol world without their specific problems
  - Distance vector
    - advantages:
      - less CPU power and memory usage, simple configuration
    - disadvantages:
      - slow convergence, lot of routing overhead, possibility of loops
  - Link state
    - advantages:
      - fast convergence, no loops, better metric (cost factor)
    - disadvantages:
      - high CPU power and memory usage (SPF algorithm, LS database), not so easy to configure (area concept)

EIGRP at a glance

- Advanced distance vector protocol (can only be used within an AS)
- uses exactly the same metric as IGRP
- Loop avoidance by DUAL
  - DUAL: Diffusing Update Algorithm
    - evolved by Mr. J.J. Garcia-Luna-Aceves
- Event triggered updates (Multicasts)
- Fast convergence

EIGRP at a glance

- Support of multiple unequal-metric paths
- Classless routing protocol (supports route summarization)
  - Update contains network plus prefix
- It supports IP, IPX and Appletalk

EIGRP Concepts

- 1st : Every EIGRP router has to discover its neighbors
  - it´s done with a Hello protocol
  - a Router doesn´t expect an acknowledge for its Hello
  - network type dependent
    - Point-to-point network
    - Multiaccess with broadcast/multicast support (BMA)
    - NBMAs nonbroadcast multiaccess network
- 2nd: based on this Hellos it builds a neighbor table
  - When a newly discovered neighbor is learned, the address and interface of the neighbor is recorded.
  - The HoldTime is the amount of time a router treats a neighbor as reachable and operational: 3* Hello interval

 IGRP-EIGRP, v3.5 25

EIGRP Concepts

Point-to-Point: Neighbor relationship is formed with the router on the other end
Hello time interval: 5 sec

Broadcast multiaccess: Neighbor relationships are formed dynamically using multicast hellos
Hello time interval: 5 sec
Multicast address: 224.0.0.10

Nonbroadcast multiaccess: Neighbor relationships are formed by manual configuration
Hello time interval: 60sec
Dest. Address: Unicast address of the neighbor

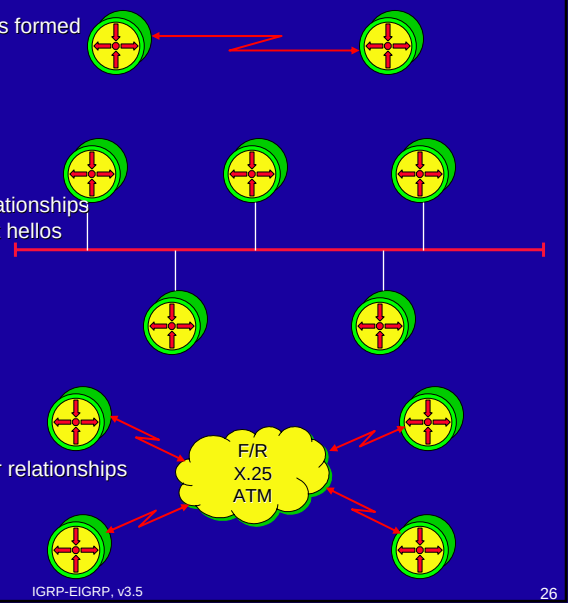

 IGRP-EIGRP, v3.5 26

L43 - IRGP and EIGRP

EIGRP Concepts

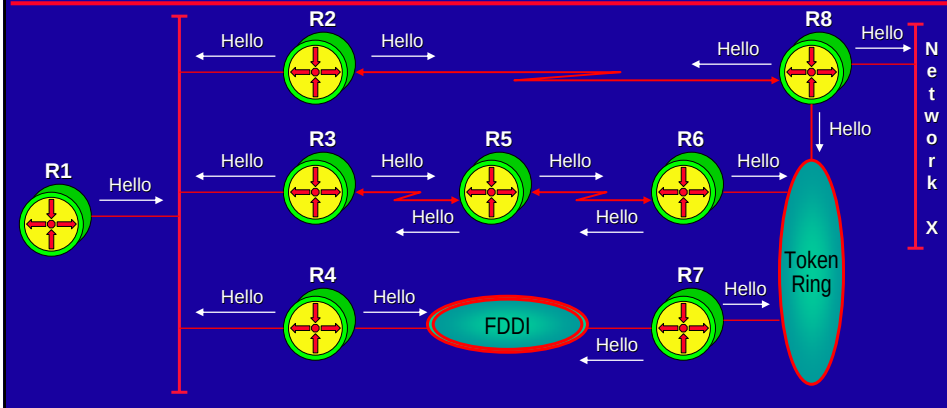

Neighbor table R1

R2
R3
R4

Neighbor table R2

R1
R3
R4
R8

Neighbor table R3

R1
R2
R4
R5

Neighbor table R4

R1
R2
R3
R7

 IGRP-EIGRP, v3.5 27

EIGRP Concepts

- 3rd:After the neighbor discovery a Topology table is built
  - Neighbor routers exchanging their complete routing tables and store these informations in a Topology table
  - information exchange through Update packets
    - Update packets
    - contain a sequence number field in the header and must be acknowledged by the receiver (reliable transmission)
    - are sent in the following instances:
      - when a neighbor first comes up (packet´s dest. addr is an unicast)
        when a network has failed (packet´s dest. addr. is 224.0.0.10)
        when there is a metric change for a certain destination (packet´s dest. addr. is 224.0.0.10)

 IGRP-EIGRP, v3.5 28

EIGRP Concepts

- Update packets
  - In contrast to OSPF every EIGRP router is modifying any received update packet. It adds its own local distance to the information and sends the packet with an own sequence number to its neighbors.

 IGRP-EIGRP, v3.5 29

EIGRP Concepts

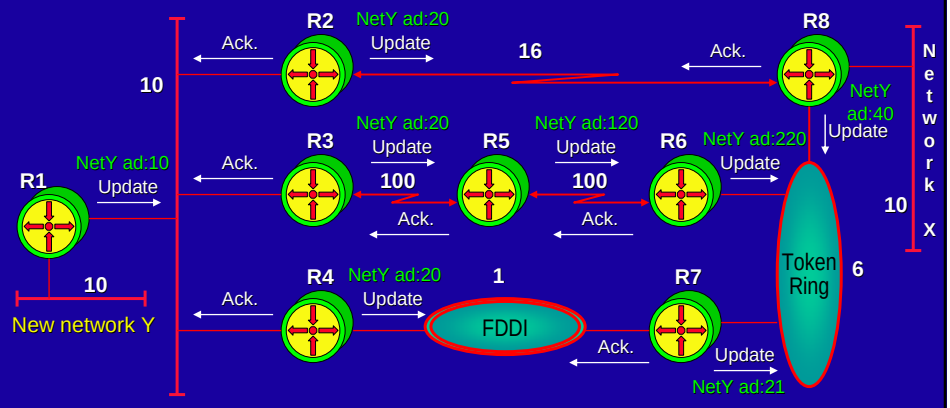

Update propagation

ad = advertised distance

 IGRP-EIGRP, v3.5 30

EIGRP Concepts

- Topology Table
  - Stores routing information that neighbors exchange after the first Hello exchange (=> smaller compared to an OSPF topology table). DUAL acts on the Topology table to determine Successors and FSs.
    - Successor:
      - A neighbor that has been selected as the next hop for a destination, it ends up in the Routing Table
    - Feasible Successor (FS):
      - A neighbor that has satisfied the Feasibility Condition and has a path to the destination (is an alternate route to the current successor)
    - Feasibility Condition (FC):
      - A condition that is met when the lowest of all the neighbors' costs plus the link cost to that neighbor is found, and the neighbor's advertised cost is less than the current successor's cost.

 IGRP-EIGRP, v3.5 31

EIGRP Topology table and feasible successor

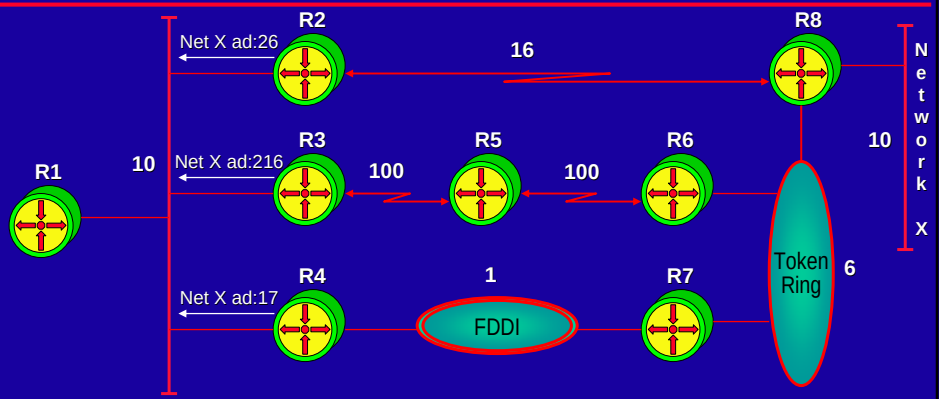

Part of R1´s Topology Table

	Network	Advertised Distance	Feasible Distance	Neighbor
Successor	X	17	27	R4
	X	216	226	R3
Feasible Successor	X	26	36	R2

ad = advertised distance

 IGRP-EIGRP, v3.5 32

L43 - IRGP and EIGRP

EIGRP Topology table without feasible successor

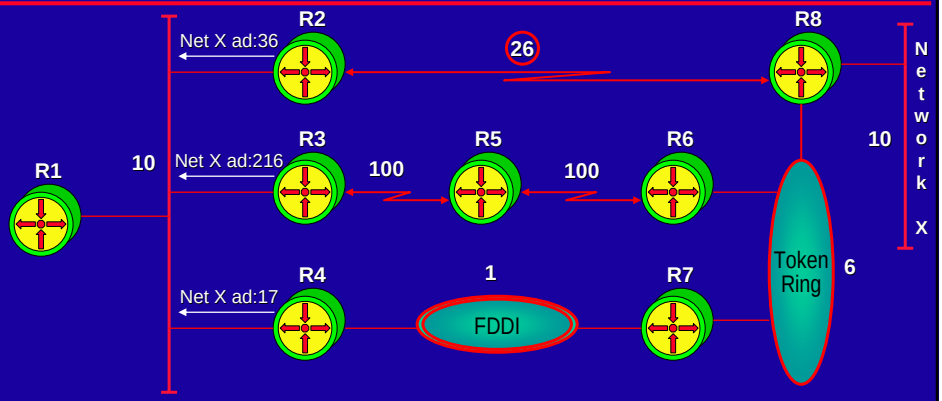

Part of R1´s Topology Table

Network	Advertised Distance	Feasible Distance	Neighbor
X	17	27	R4
X	216	226	R3
X	36	46	R2

Successor

no Feasible Successor!!!

ad = advertised distance

EIGRP Active state

- If the successor disappears from the topology table because of a network change and there is a feasible successor, DUAL keeps the route in a passive state.
  - Passive state
    - A router's state after losing its successor when it has an FS to the destination available in its Topology table
- If the successor disappears from the topology table because of a network change and there is no feasible successor, DUAL puts the route into the active state.

L43 - IRGP and EIGRP

EIGRP Active state

- Active state
  - A router's state for a destination when it has lost its successor to that destination and has no other feasible successor (FS) available. The router is forced to compute a route to the destination.
    It´s sending a query packet to all its neighbors.
  - Query packet (will be acknowledged from the receiver)
    - Sent to all neighbors when a router goes into Active for a destination and is asking for information on that destination. Unless it receives replies back from *all* its neighbors, the router will remain in Active state and not start the computation for a new successor.
  - Reply packet (will be acknowledged from the receiver)
    - Sent by every EIGRP neighbor which receives a query. If the neighbor doesn't have the information, it queries its neighbors indicating that it is also performing route recomputation

 IGRP-EIGRP, v3.5 35

EIGRP Topology table without feasible successor

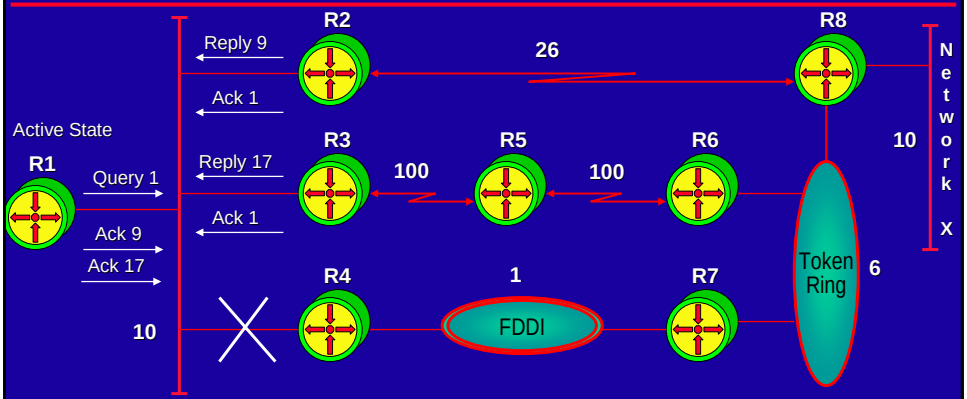

Part of R1´s Topology Table

Network	Advertised Distance	Feasible Distance	Neighbor
X	17	27	R4
X	216	226	R3
X	36	46	R2

Successor

no Feasible Successor!!!

 IGRP-EIGRP, v3.5 36

L43 - IRGP and EIGRP

EIGRP Topology table without feasible successor

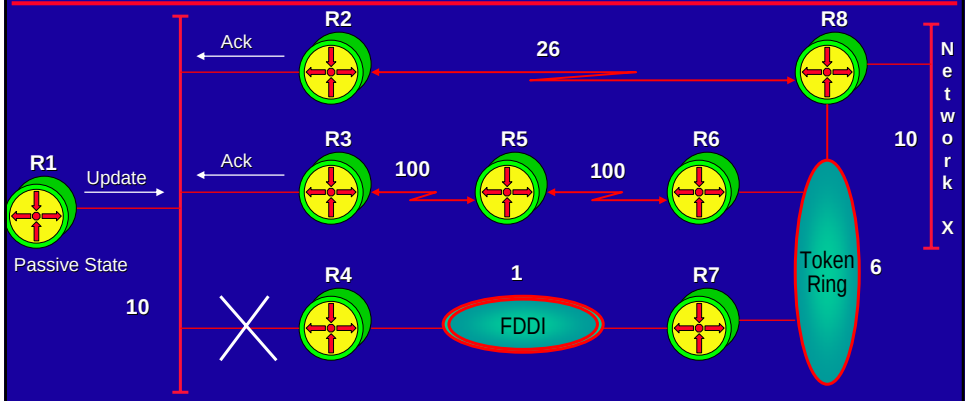

Part of R1´s Topology Table

Network	Advertised Distance	Feasible Distance	Neighbor
X	216	226	R3
X	36	46	R2

Successor

EIGRP Compatibility

- Route tagging
  - EIGRP has the notion of internal and external routes.
    - Internal routes
      are ones that have been originated within an EIGRP autonomous system (AS).
    - External routes
      are ones that have been learned by another routing protocol or reside in the routing table as static routes.
- Route redistribution
  - in the case of IGRP is done automatically, when EIGRP and IGRP are belonging to the same Autonomous System (compatible metric!!!). IGRP derived routes are treated as external routes in EIGRP (also OSPF, RIP, EGP, BGP...)

L43 - IRGP and EIGRP

EIGRP Compatibility

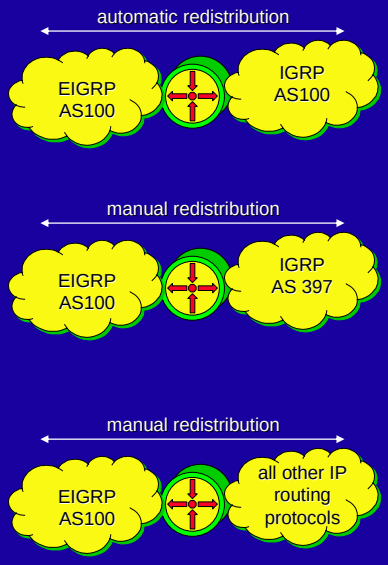

 IGRP-EIGRP, v3.5 39

EIGRP Protocol stack

OSI stack

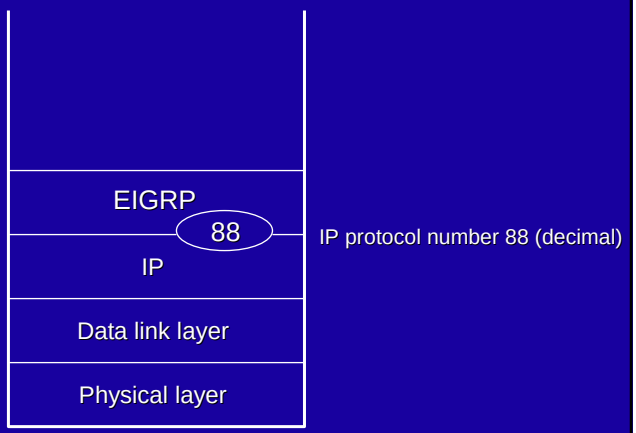

 IGRP-EIGRP, v3.5 40