
NETWORKERS
2000
Cisco Systems

Deploying IGRP/E-IGRP

Session 2208

Understanding E-IGRP

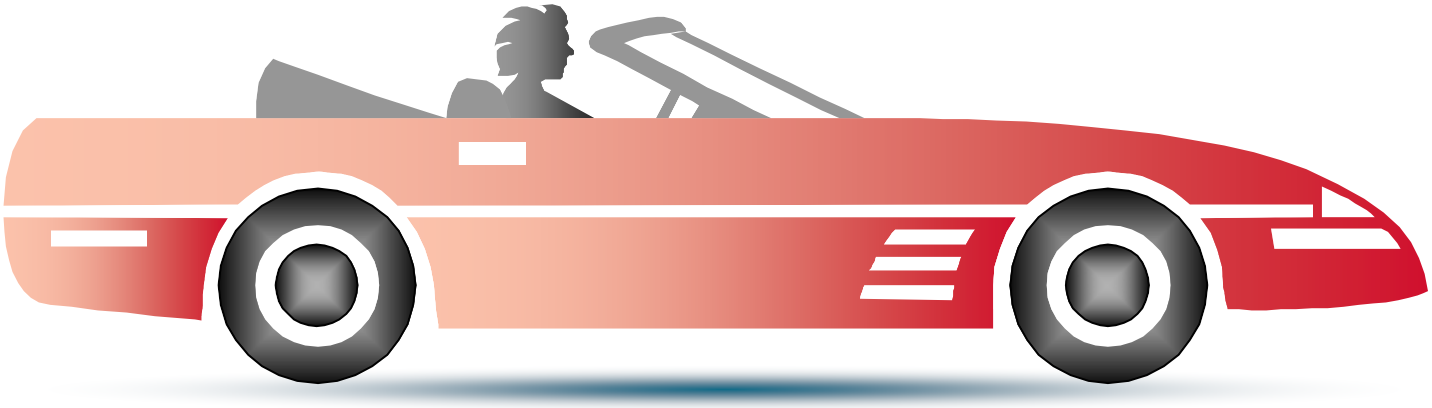

Understanding and deploying E-IGRP is like driving a car

Agenda

- Fundamentals of E-IGRP
- DUAL
- Summarization and Load Balancing
- Query Process
- Deployment Guidelines with E-IGRP
- Summary

IGRP: Interior Gateway Routing Protocol

- Cisco proprietary
- Distance vector
- Broadcast based
- Utilizes link bandwidth and delay
  - 15 hops is no longer the limit
- 90 seconds updates (RIP is 30 sec.)
- Load balance over unequal cost paths

IGRP/E-IGRP Metrics Calculation

- Metric = [K1 x BW + (K2 x BW) / (256 - Load) + K3 x Delay] x [K5 / (Reliability + K4)]

  By Default: K1 = 1, K2 = 0, K3 = 1, K4 = K5 = 0

- Delay is sum of all the delays of the link along the paths

  Delay = Delay/10

- Bandwidth is the lowest bandwidth of the link along the paths

  Bandwidth = 10000000/Bandwidth

Problems with RIP and IGRP

- Slow convergence
- Not 100% loop free
- Don't support VLSM and discontiguous network
- Periodic full routing updates
- RIP has hop count limitation

Advantages of E-IGRP

- Advanced distance vector
- 100% loop free
- Fast convergence
- Easy configuration
- Less network design constraints than OSPF
- Incremental update
- Supports VLSM and discontiguous network
- Classless routing
- Compatible with existing IGRP network
- Protocol independent (support IPX and AppleTalk)

Advantages of E-IGRP

- Uses multicast instead of broadcast
- Utilize link bandwidth and delay

  E-IGRP Metric = IGRP Metric x 256
  (32 bit Vs. 24 bit)

- Unequal cost paths load balancing
- More flexible than OSPF

  Full support of distribute list

  Manual summarization can be done in any interface at any router within network

E-IGRP Packets

- Hello: Establish neighbor relationships
- Update: Send routing updates
- Query: Ask neighbors about routing information
- Reply: Response to query about routing information
- Ack: Acknowledgement of a reliable packet

E-IGRP Neighbor Relationship

- **Two routers become neighbors when they see each other's hello packet**
  - **Hello address = 224.0.0.10**
- **Hellos sent once every five seconds on the following links:**
  - **Broadcast Media: Ethernet, Token Ring, FDDI, etc.**
  - **Point-to-point serial links: PPP, HDLC, point-to-point frame relay/ATM subinterfaces**
  - **Multipoint circuits with bandwidth greater than T1: ISDN PRI, SMDS, Frame Relay**

E-IGRP Neighbor Relationship

- Hellos sent once every 60 seconds on the following links:

  Multipoint circuits with bandwidth less than or equal to T1: ISDN BRI, Frame Relay, SMDS, etc.

- Neighbor declared dead when no E-IGRP packets are received within hold interval

  Not only Hello can reset the hold timer

- Hold time by default is three times the hello time

E-IGRP Neighbor Relationship

- E-IGRP will form neighbors even though hello time and hold time don’t match
- E-IGRP sources hello packets from primary address of the interface
- E-IGRP will not form neighbor if K-values are mismatched
- E-IGRP will not form neighbor if AS numbers are mismatched
- Passive interface (IGRP vs. E-IGRP)

Discovering Routes

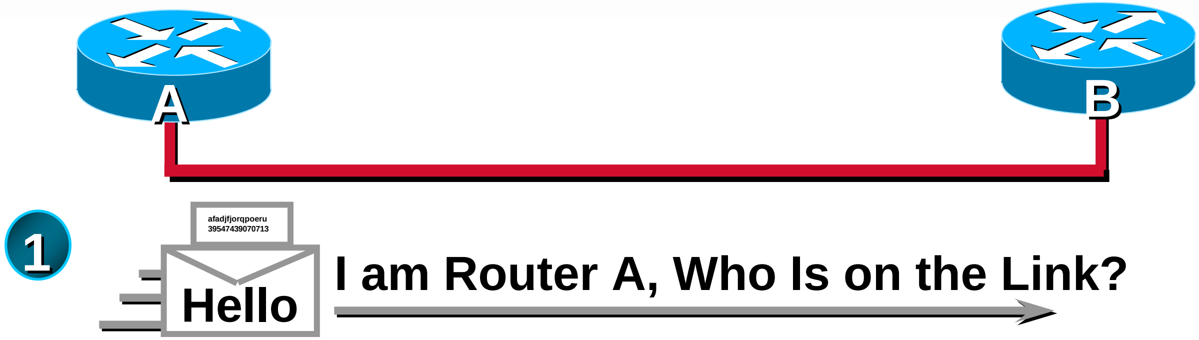

Discovering Routes

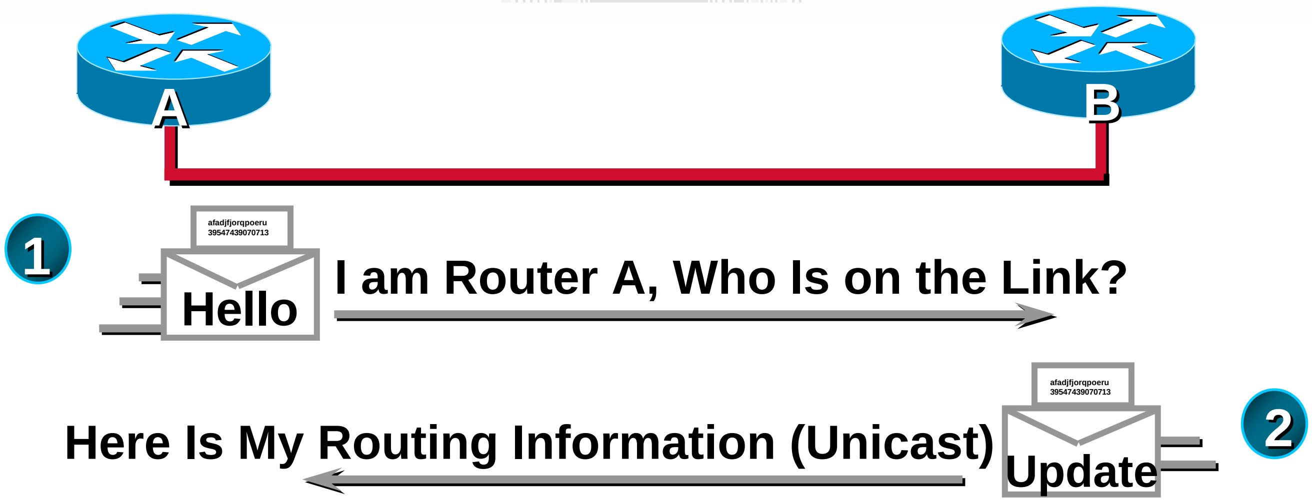

Discovering Routes

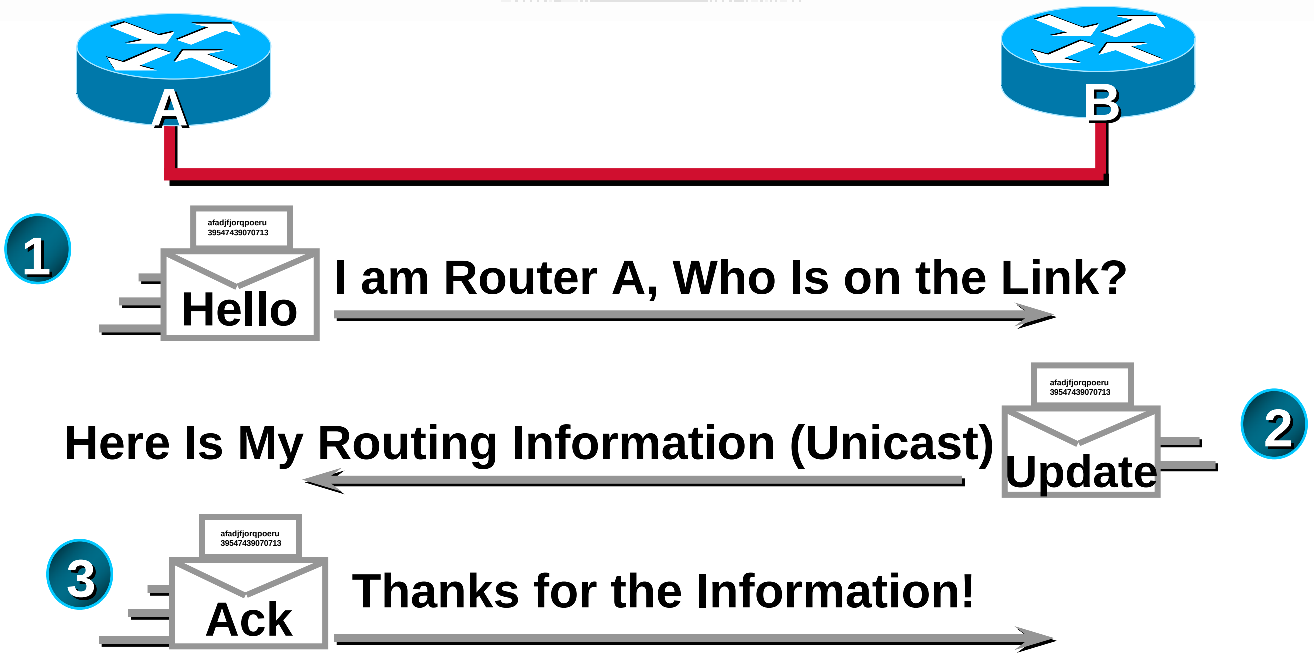

Discovering Routes

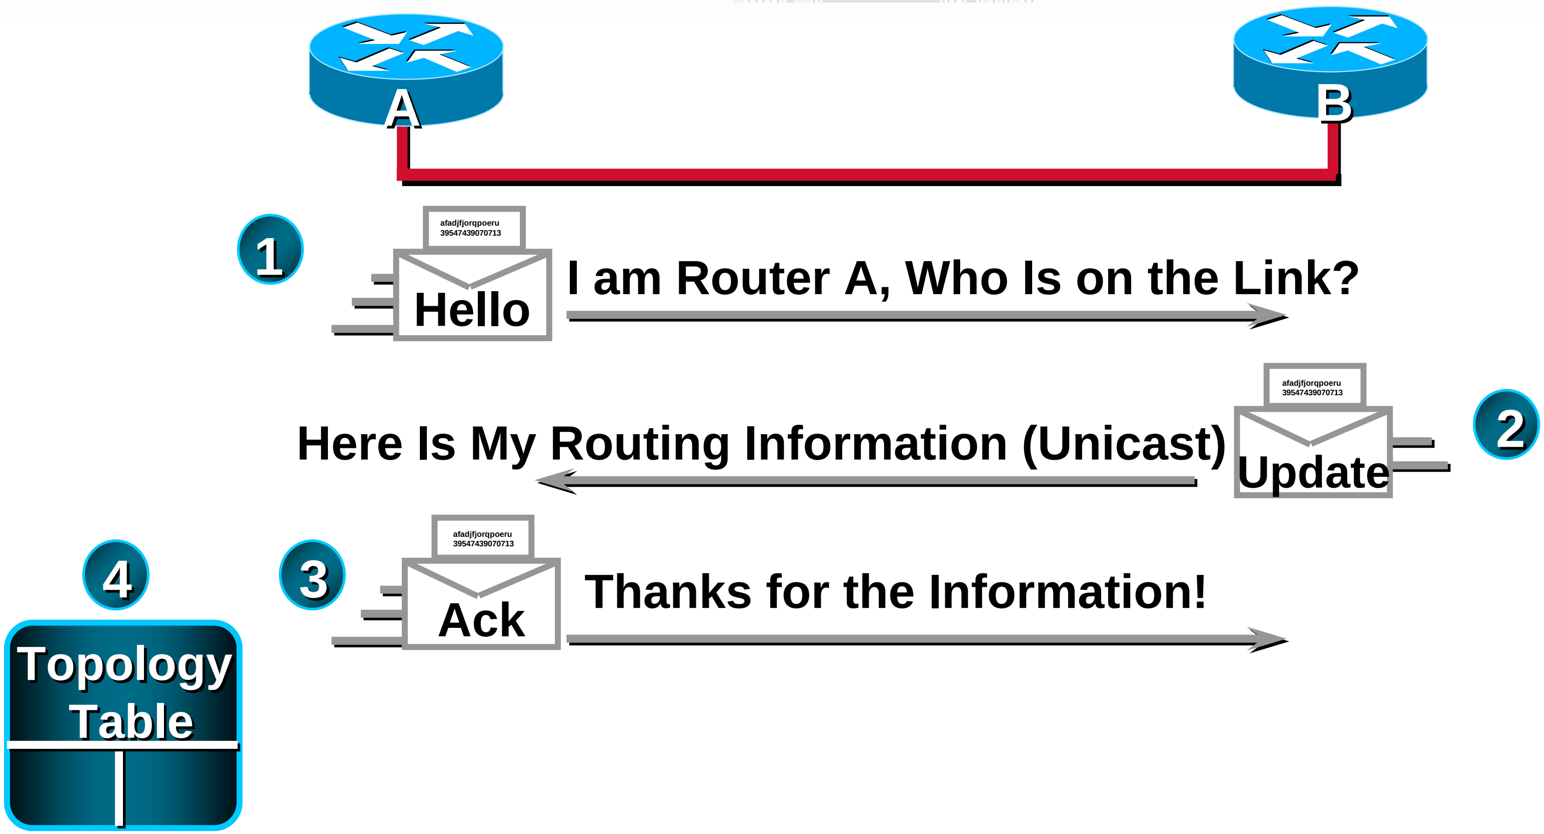

Discovering Routes

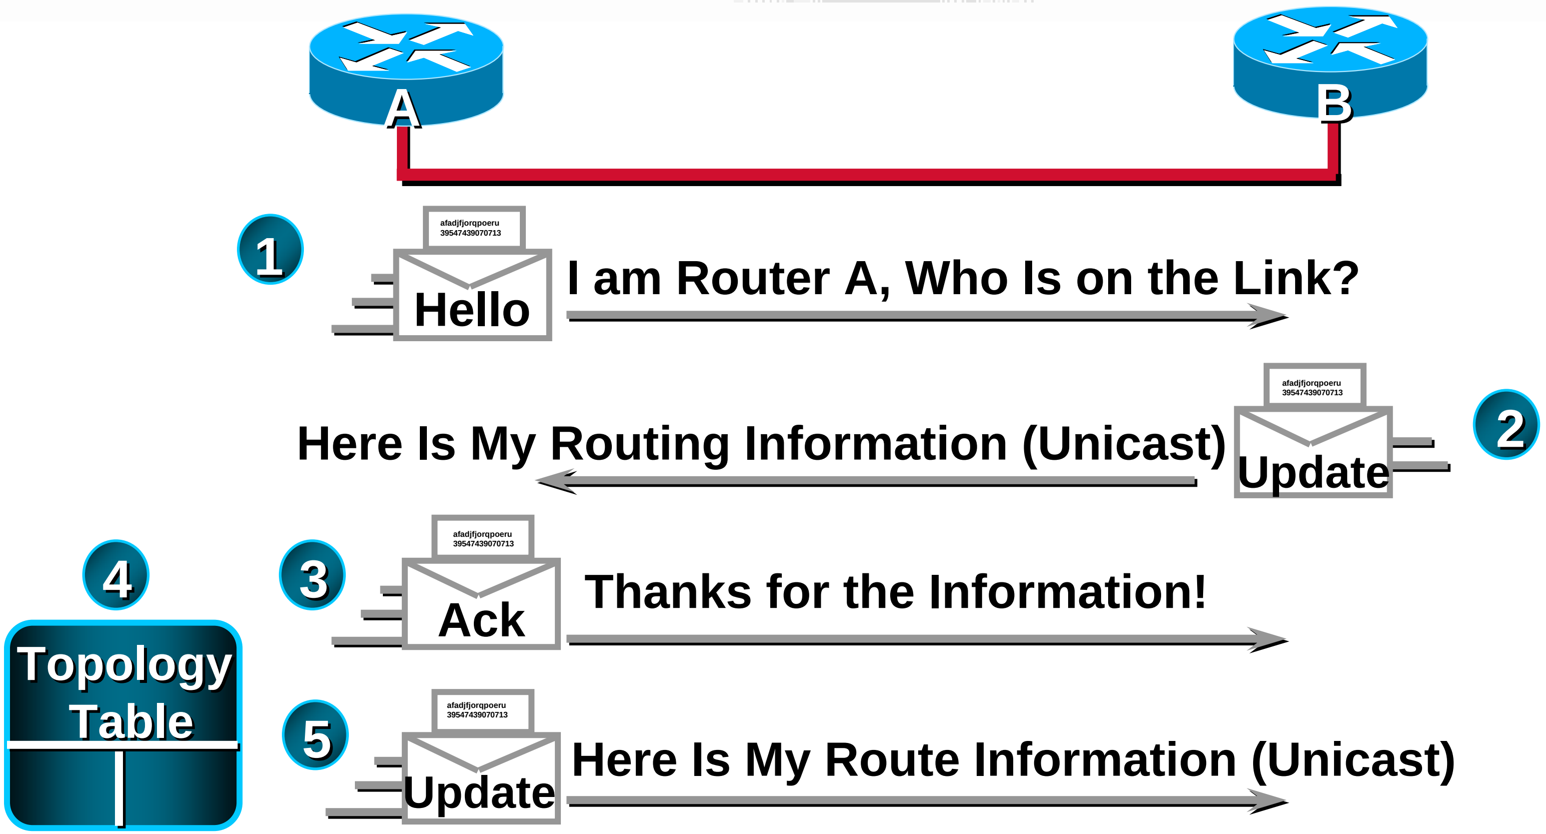

Discovering Routes

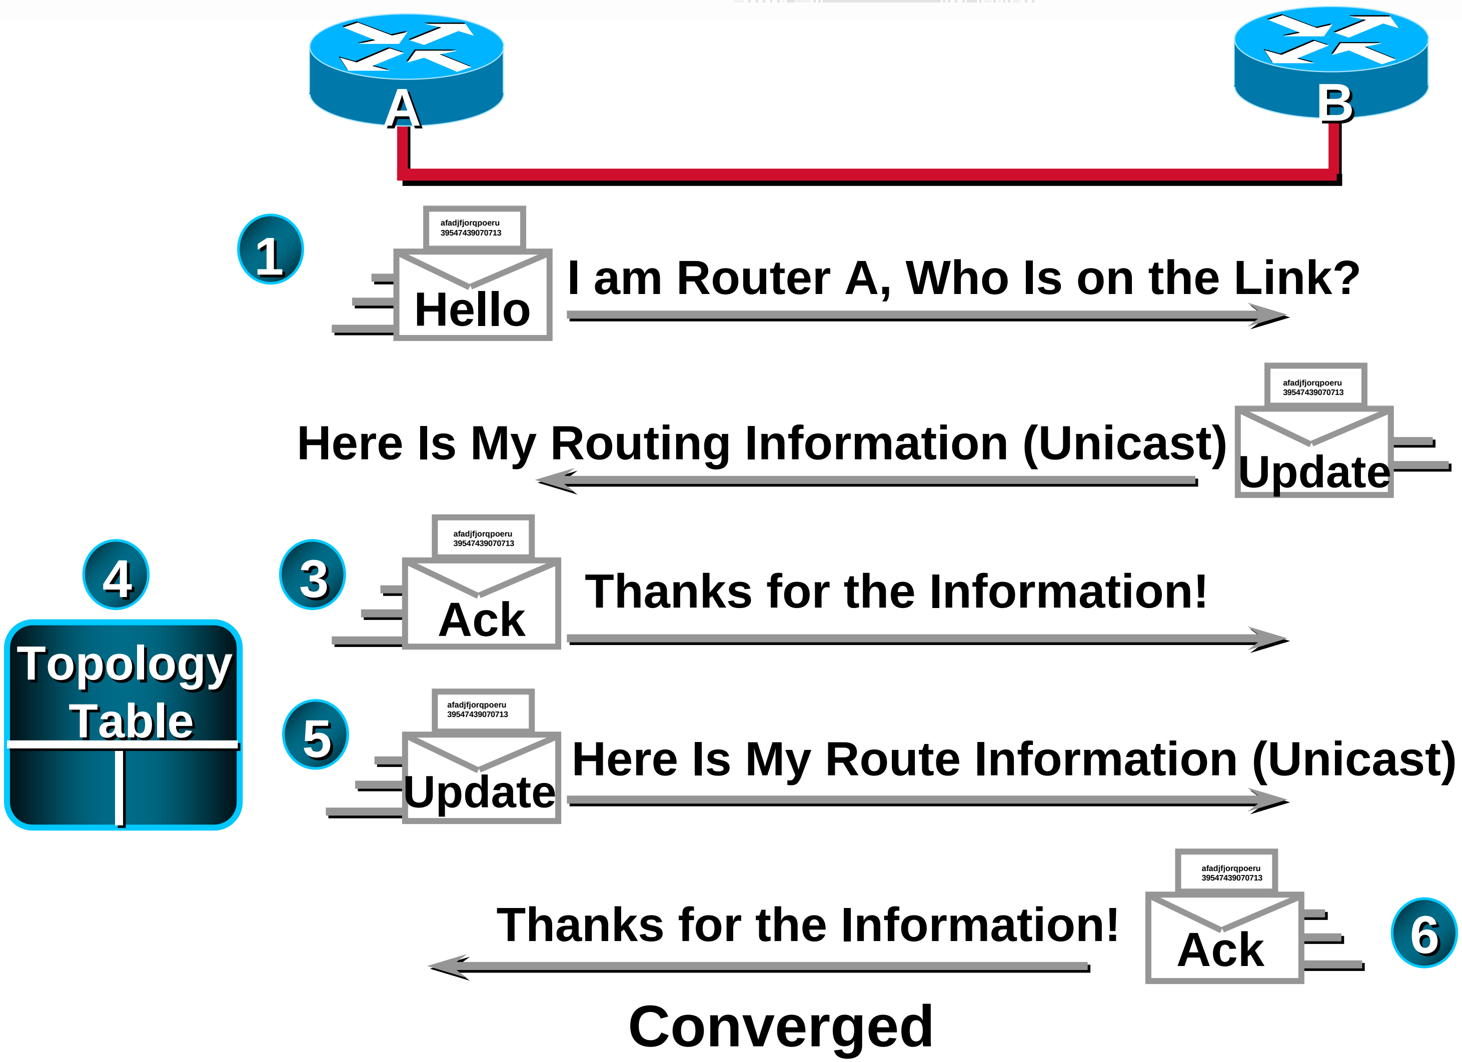

Agenda

- Fundamentals of E-IGRP
- DUAL
- Summarization and Load Balancing
- Query Process
- Deployment Guidelines with E-IGRP
- Summary

E-IGRP DUAL

- Diffusing update algorithm
- Finite-State-Machine
  - Track all routes advertised by neighbors
  - Select loop-free path using a successor and remember any feasible successors
  - If successor lost
    - Use feasible successor
  - If no feasible successor
    - Query neighbors and recompute new successor

E-IGRP Feasible Distance (FD)

- **Feasible distance is the minimum distance (metric) along a path to a destination network**

Feasible Distance Example

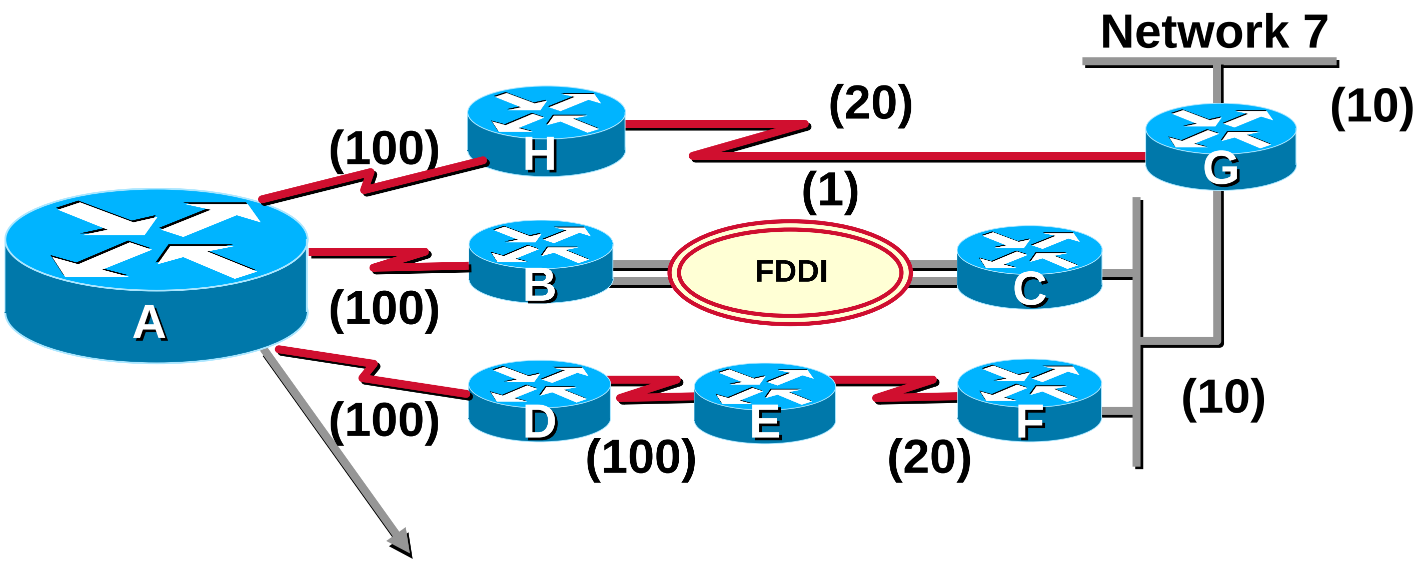

Topology Table

Destination	Feasible Distance (FD)	Neighbor
7	100+20+10=130	H
7	100+1+10+10=121	B
7	100+100+20+10+10=240	D

E-IGRP Reported Distance (RD)

- **Reported distance is the distance (metric) towards a destination as advertised by an upstream neighbor**

  Reported distance is the distance reported in the queries, the replies and the updates

Reported Distance Example

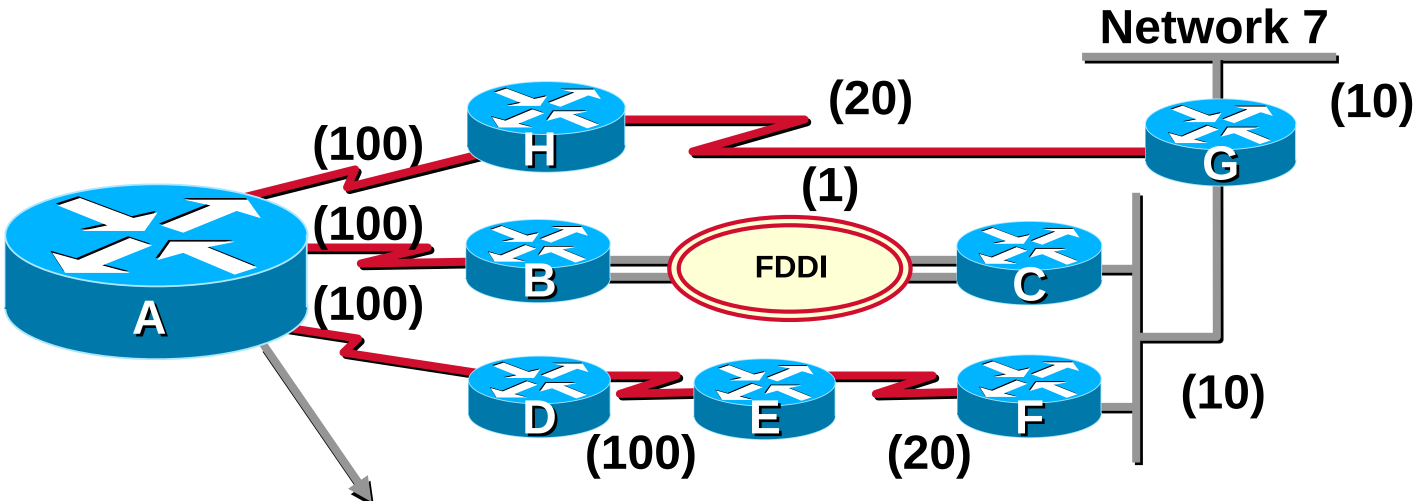

Topology Table

Destination	Reported Distance (RD)	Neighbor
7	20+10=30	H
7	1+10+10=21	B
7	100+20+10+10=140	D

E-IGRP Feasibility Condition (FC)

- **A neighbor meets the feasibility condition (FC) if the reported distance by the neighbor is smaller than the feasible distance (FD) of this router**

E-IGRP Successor

- A successor is a neighbor that has met the feasibility condition and has the least cost path towards the destination
- It is the next hop for forwarding packets
- Multiple successors are possible (load balancing)

E-IGRP Feasible Successor (FS)

- **A feasible successor is a neighbor whose reported distance (RD) is less than the feasible distance (FD)**

Successor Example

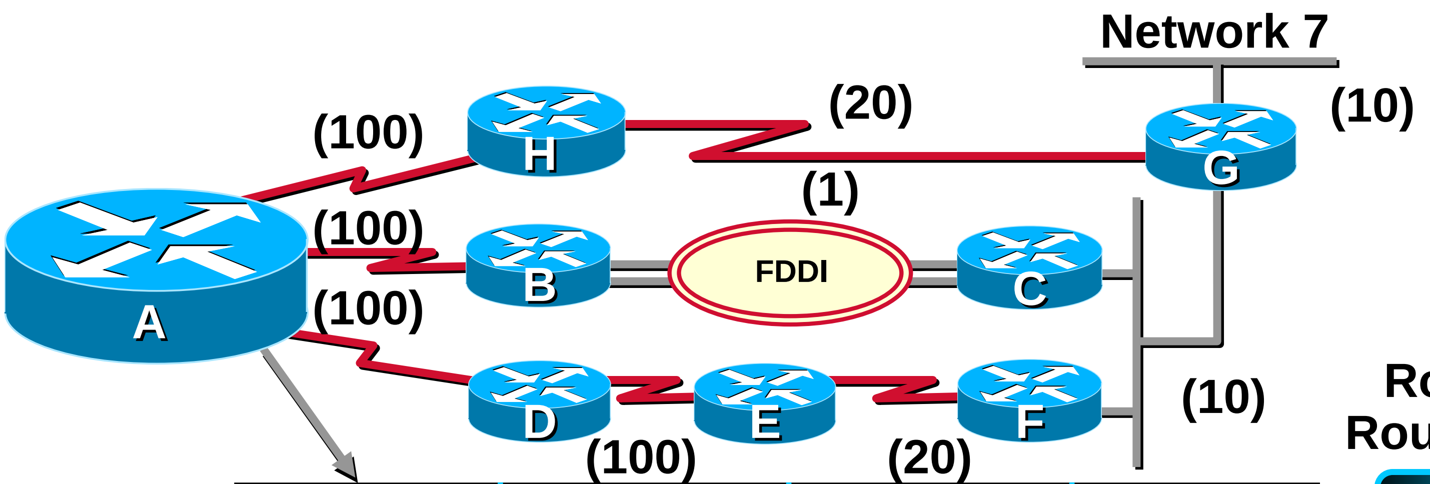

Topology Table

Destination	FD	RD	Neighbor
7	130	30	H
7	121	21	B
7	240	140	D

Router A's Routing Table

7	121	B

- B is current successor (FD = 121)
- H is the feasible successor (30 < 121)

Passive, Active, and Stuck in Active (SIA)

- **Passive routes are routes that have successor information**

  Passive route = Good

- **Active routes are routes that have lost their successors and no feasible successors are available. The router is actively looking for alternative paths**

  Active route = Bad

- **Stuck in Active means the neighbor still has not replied to the original query within three minutes**

  Stuck in active = Ugly

Dual Algorithm

- **Local computation**
  - When a route is no longer available via the current successor, the router checks its topology table
  - Router can switch from successor to feasible successor without involving other routers in the computation
  - Router stays passive
  - Updates are sent

DUAL: Local Computation

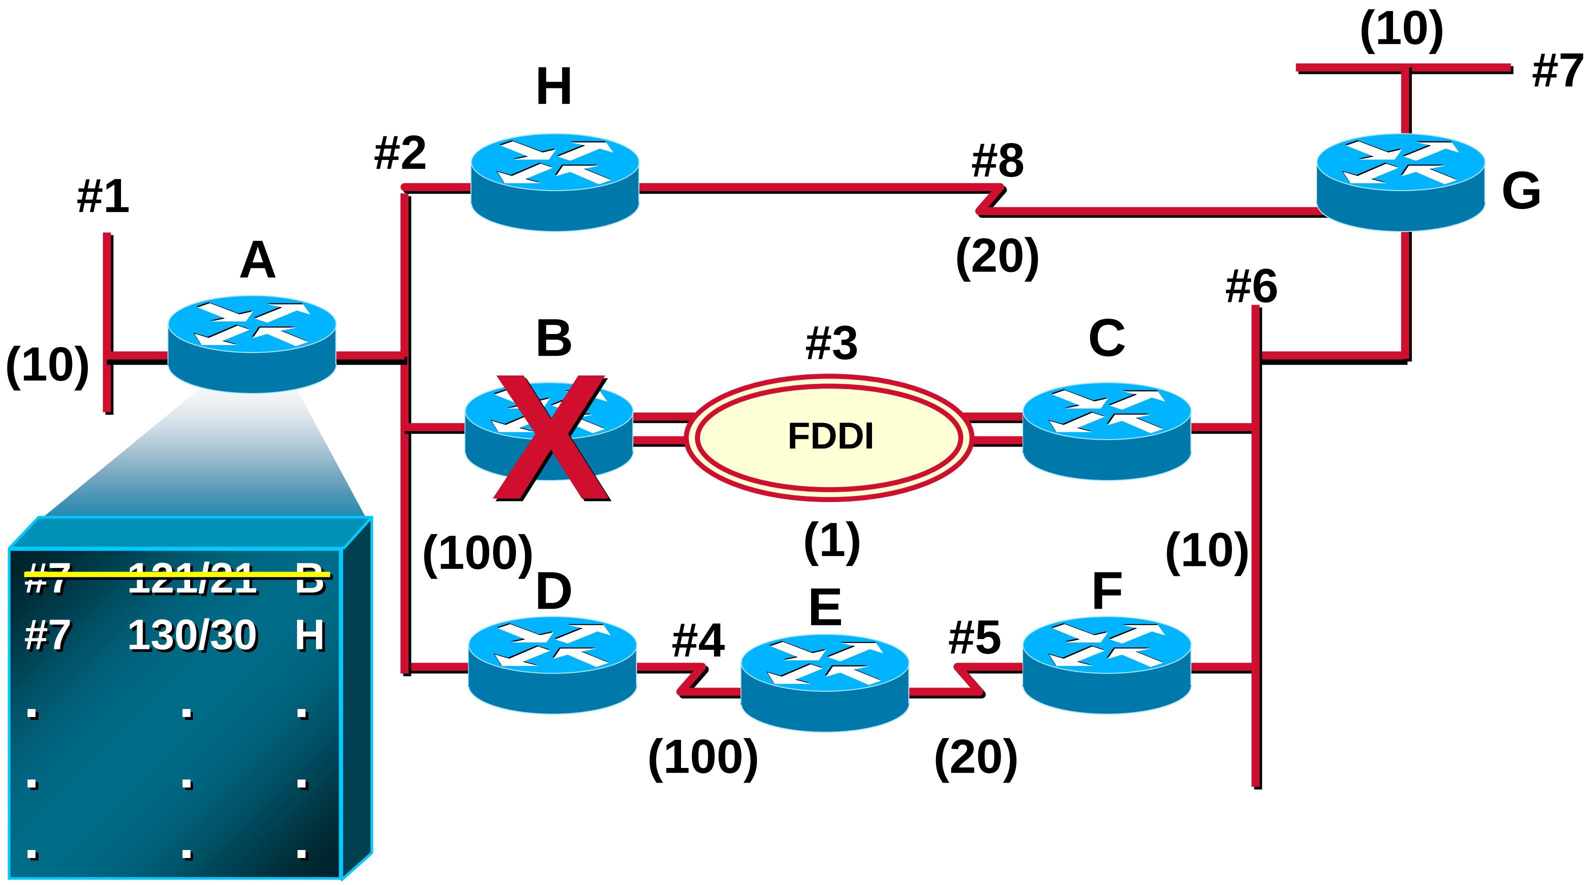

Dual Algorithm

- **Diffused Computation**

  When a route is no longer available via its current successor and no feasible successor is available, queries are sent out to neighbors asking about the lost route

  The route is said to be in active state

  Neighbors reply to the query if they have information about the lost route. If not, queries are sent out to all of their neighbors.

  The router sending out the query waits for all of the replies from its neighbors and will make routing decision based on the replies

DUAL: Diffused Computation

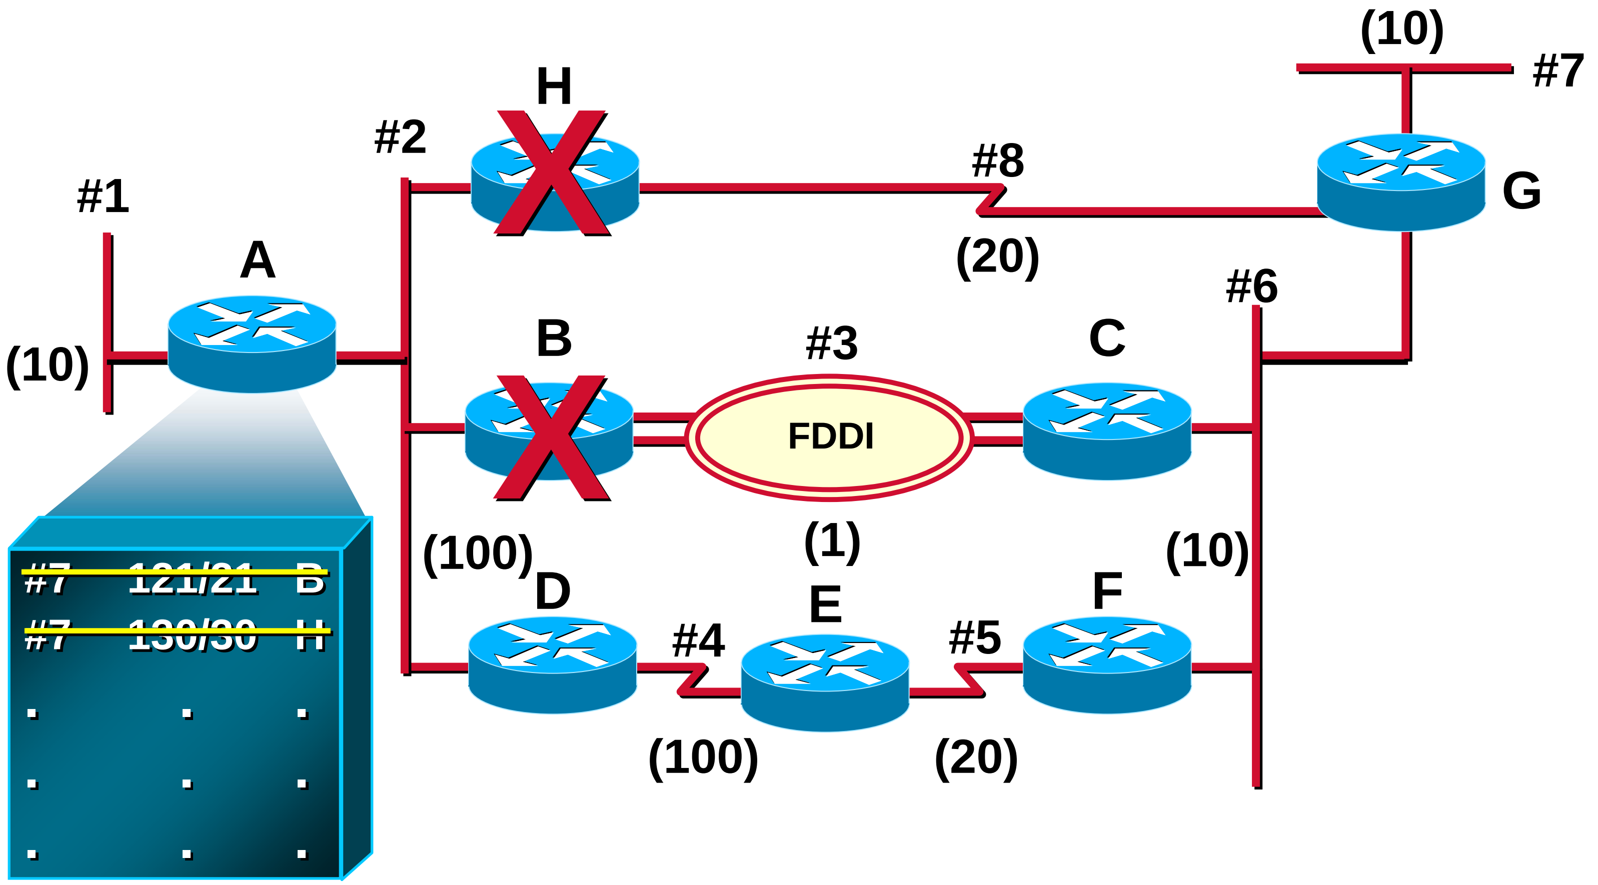

cisco.com

DUAL Example

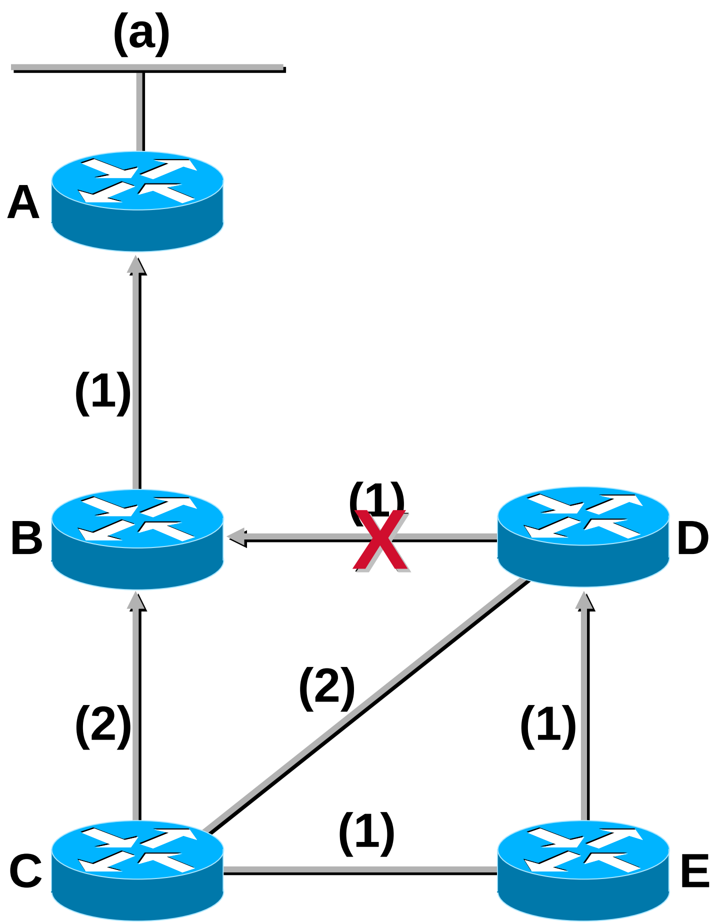

C	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D	Cost (4/2)	(fs)
	via E	Cost (4/3)	

D	E-IGRP Topology		
(a)		Cost (2)	(fd)
	~~via B~~	~~Cost (2/1)~~	~~(Successor)~~
	via C	Cost (5/3)	

E	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via D	Cost (3/2)	(Successor)
	via C	Cost (4/3)	

DUAL Example

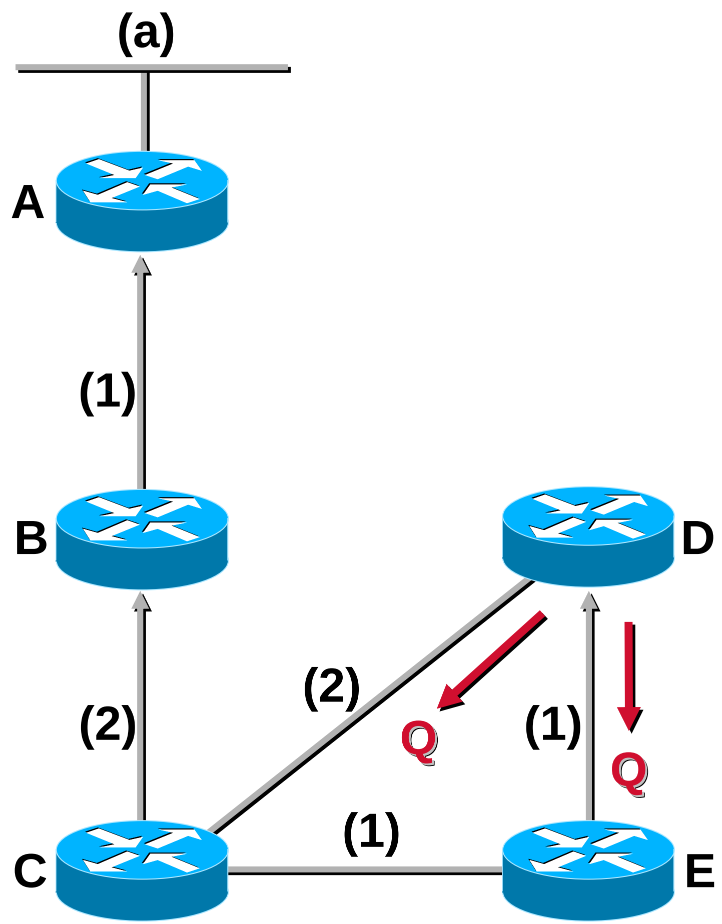

C	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D		
	via E	Cost (4/3)	

D	E-IGRP Topology		
(a)	**ACTIVE**	Cost (-1)	(fd)
	via E		(q)
	via C	Cost (5/3)	(q)

E	E-IGRP Topology		
(a)		Cost (3)	(fd)
	~~via D~~	~~Cost (3/2)~~	~~(Successor)~~
	via C	Cost (4/3)	

DUAL Example

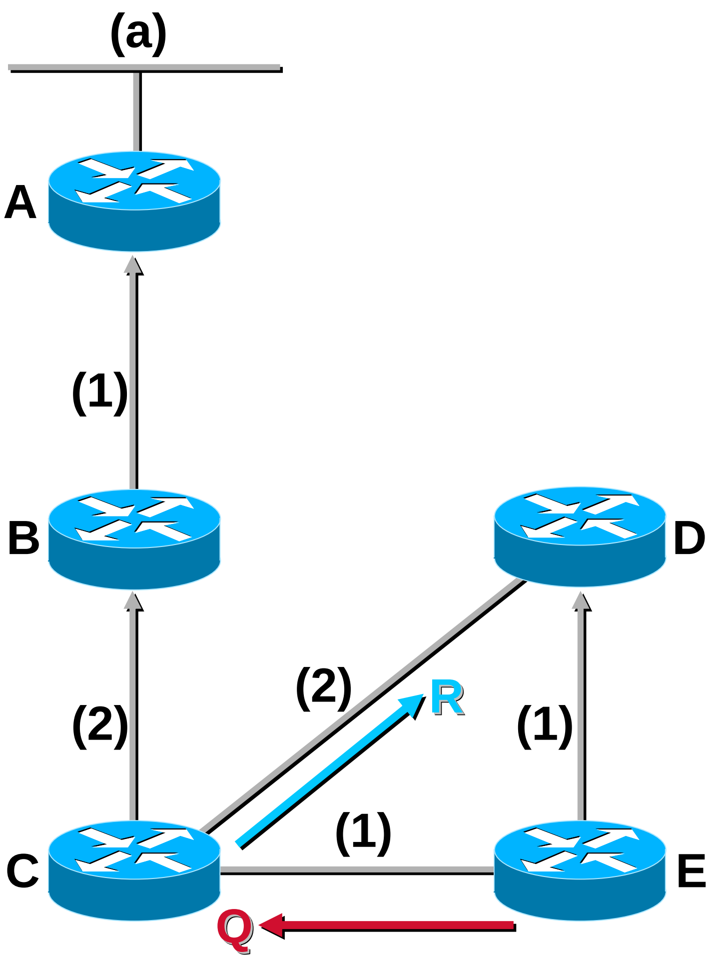

C E-IGRP Topology
(a) Cost (3) (fd)
via B Cost (3/1) (Successor)
via D
via E

D E-IGRP Topology
(a) **ACTIVE** Cost (-1) (fd)
via E (q)
via C Cost (5/3)

E E-IGRP Topology
(a) **ACTIVE** Cost (-1) (fd)
via D
via C Cost (4/3) (q)

DUAL Example

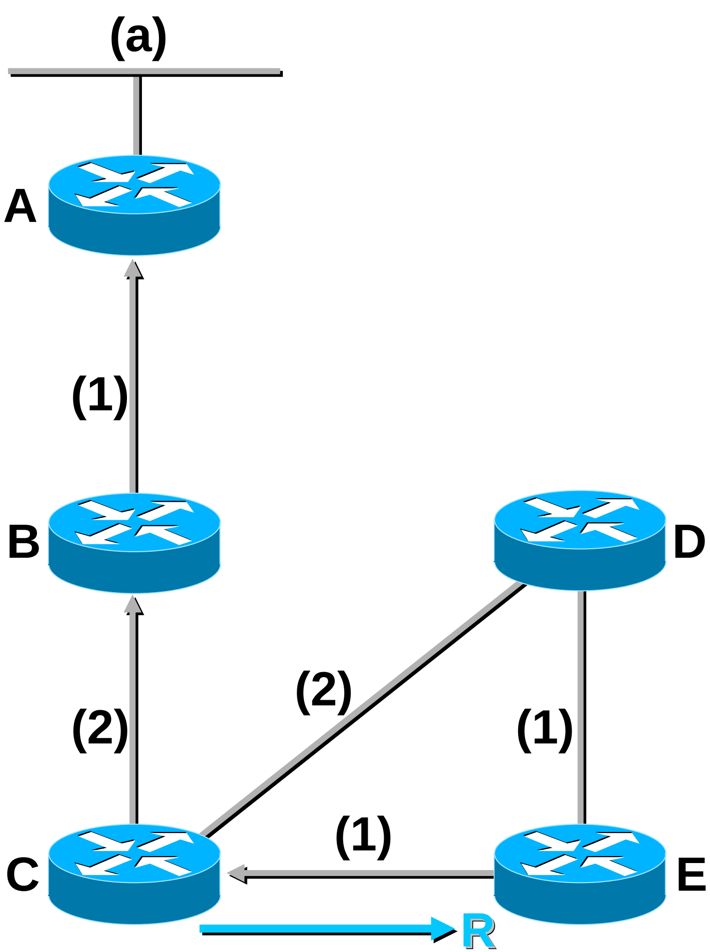

C	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D		
	via E		

D	E-IGRP Topology		
(a)	**ACTIVE**	Cost (-1)	(fd)
	via E		(q)
	via C	Cost (5/3)	

E	E-IGRP Topology		
(a)		Cost (4)	(fd)
	via C	Cost (4/3)	(Successor)
	via D		

DUAL Example

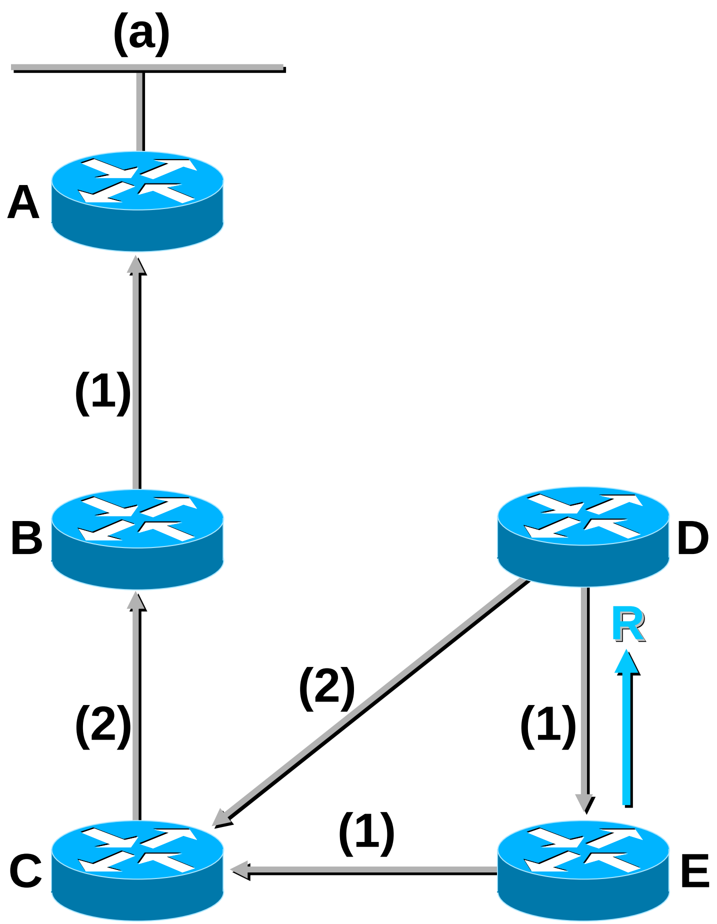

C E-IGRP Topology

(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D		
	via E		

D E-IGRP Topology

(a)		Cost (5)	(fd)
	via C	Cost (5/3)	(Successor)
	via E	Cost (5/4)	(Successor)

E E-IGRP Topology

(a)		Cost (4)	(fd)
	via C	Cost (4/3)	(Successor)
	via D		

DUAL Example

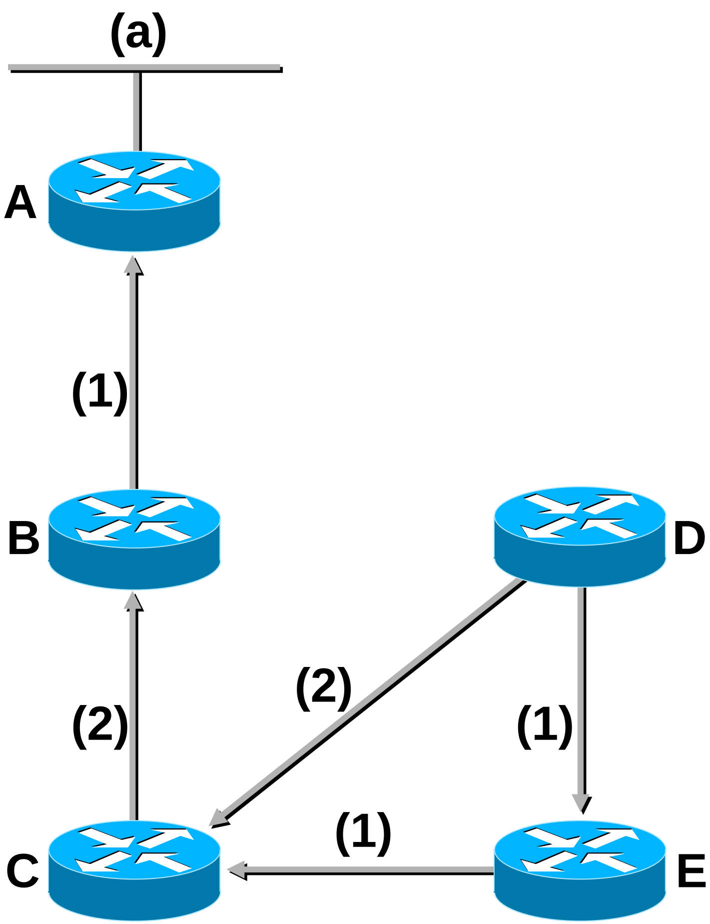

C	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D		
	via E		

D	E-IGRP Topology		
(a)		Cost (5)	(fd)
	via C	Cost (5/3)	(Successor)
	via E	Cost (5/4)	(Successor)

E	E-IGRP Topology		
(a)		Cost (4)	(fd)
	via C	Cost (4/3)	(Successor)
	via D		

DUAL Example (Start)

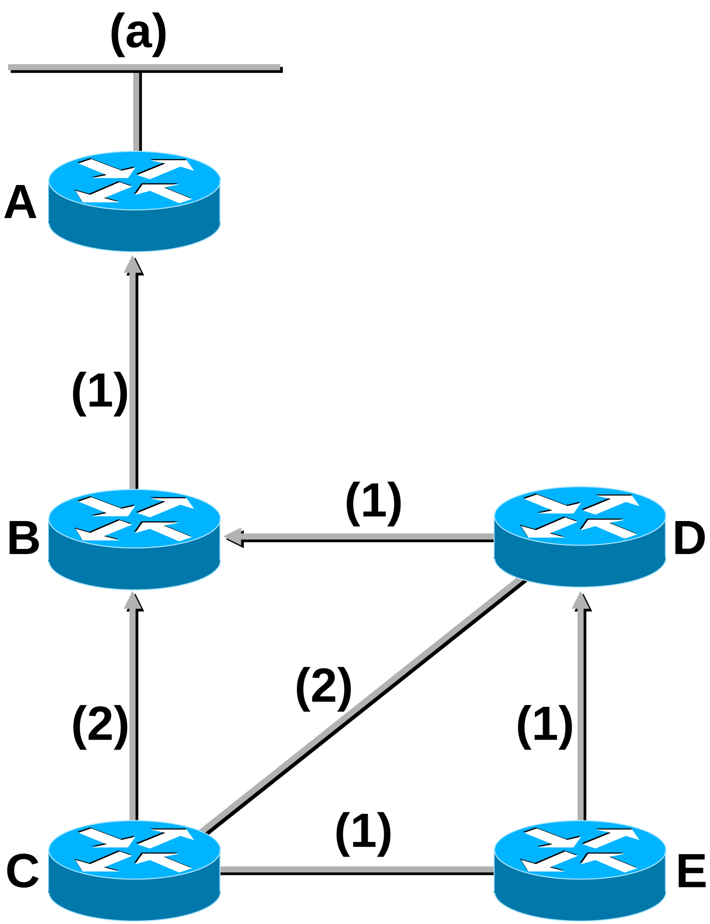

C E-IGRP Topology

(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D	Cost (4/2)	(fs)
	via E	Cost (4/3)	

D E-IGRP Topology

(a)		Cost (2)	(fd)
	via B	Cost (2/1)	(Successor)
	via C	Cost (5/3)	

E E-IGRP Topology

(a)		Cost (3)	(fd)
	via D	Cost (3/2)	(Successor)
	via C	Cost (4/3)	

DUAL Example (End)

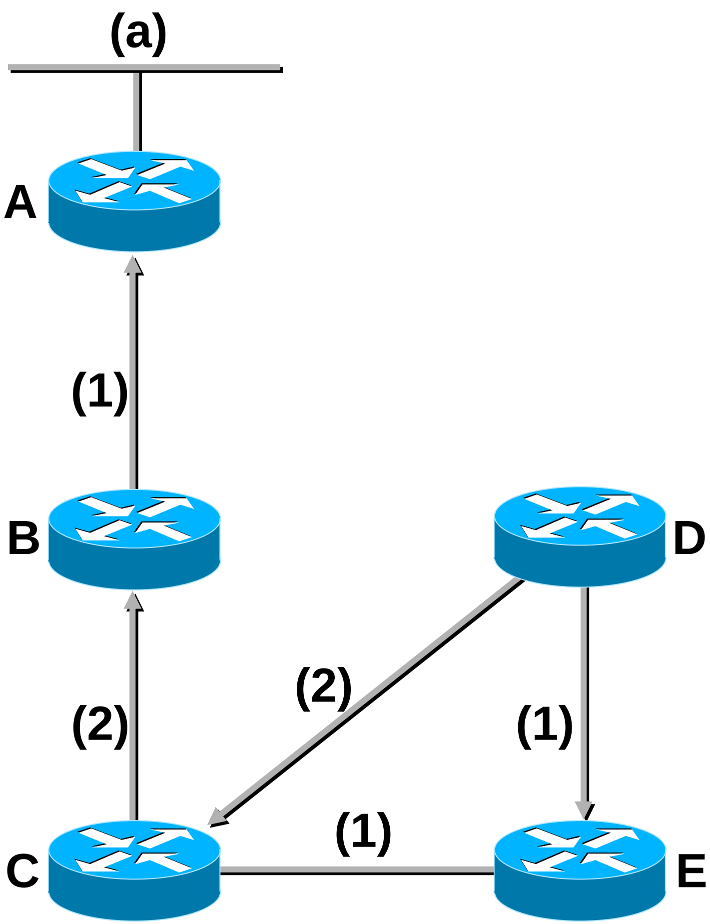

C	E-IGRP Topology		
(a)		Cost (3)	(fd)
	via B	Cost (3/1)	(Successor)
	via D		
	via E		

D	E-IGRP Topology		
(a)		Cost (5)	(fd)
	via C	Cost (5/3)	(Successor)
	via E	Cost (5/4)	(Successor)

E	E-IGRP Topology		
(a)		Cost (4)	(fd)
	via C	Cost (4/3)	(Successor)
	via D		

E-IGRP Reliable Transport Protocol

- E-IGRP **reliable packets** are packets that requires explicit acknowledgement:

  Update

  Query

  Reply

- E-IGRP **unreliable packets** are packets that do not require explicit acknowledgement:

  Hello

  Ack

cisco.com

E-IGRP Reliable Transport Protocol

- The router keeps a neighbor list and a retransmission list for every neighbor
- Each reliable packet (Update, Query, Reply) will be retransmitted when packet is not acked
- E-IGRP transport has window size of one (stop and wait mechanism)

  Every single reliable packet needs to be acknowledged before the next sequenced packet can be sent

E-IGRP Reliable Transport Protocol

- **With reliable multicast traffic, one must wait to transmit the next reliable multicast packets, until all peers have acknowledged the previous multicast**
- **If one or more peers are slow in acknowledging, all other peers suffer from this**
- **Solution: The nonacknowledged multicast packet will be retransmitted as a unicast to the slow neighbor**

E-IGRP Reliable Transport Protocol

- **Per neighbor, retransmission limit is 16**
- **Neighbor relationship is reset when retry limit (limit = 16) for reliable packets is reached**

Agenda

- Fundamentals of E-IGRP
- DUAL
- Summarization and Load Balancing
- Query Process
- Deployment Guidelines with E-IGRP
- Summary

E-IGRP Summarization

- **Purpose: Smaller routing tables, smaller updates, query boundary**
- **Auto summarization:**

  On major network boundaries, networks are summarized to the major networks

  Auto summarization is turned on by default

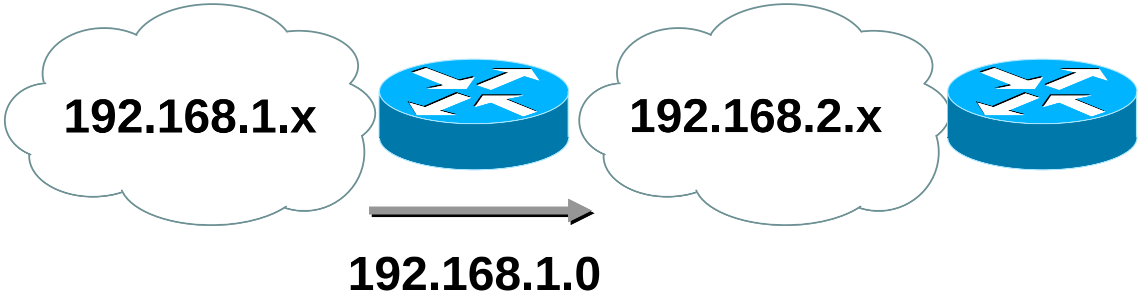

E-IGRP Summarization

- **Manual summarization**
  - **Configurable on per interface basis in any router within network**
  - **When summarization is configured on an interface, the router immediate creates a route pointing to null zero with administrative distance of five**
    - **Loop prevention mechanism**
  - **When the last specific route of the summary goes away, the summary is deleted**
  - **The minimum metric of the specific routes is used as the metric of the summary route**

E-IGRP Summarization

- **Manual summarization command:**

  ip summary-address E-IGRP <as number> <address> <mask>

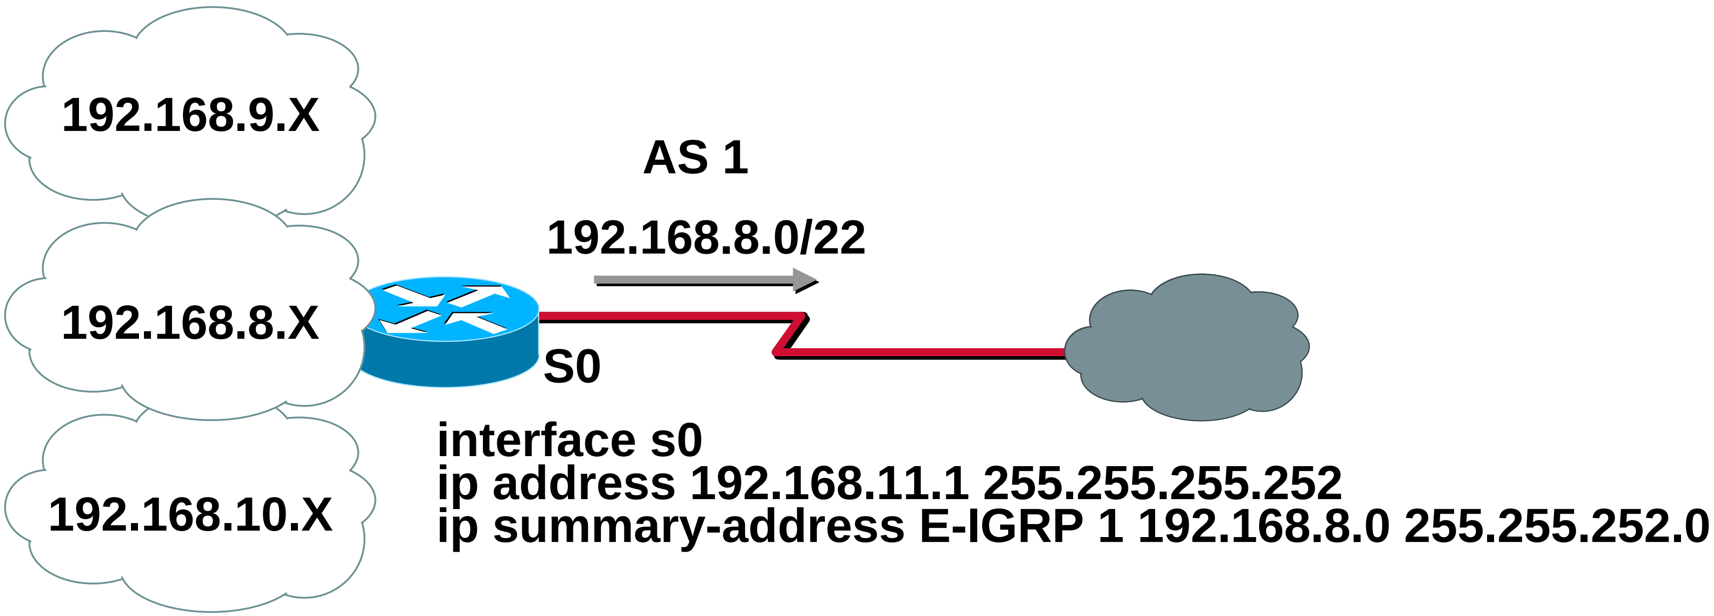

cisco.com

E-IGRP Load Balancing

- Routes with equal metric to the minimum metric, will be installed in the routing table (Equal Cost Load Balancing)
- There can be up to six entries in the routing table for the same destination (default = 4)

  ip maximum-paths <1-6>

E-IGRP Unequal Cost Load Balancing

- **E-IGRP offers unequal cost load balancing feature with the command:**

  Variance <multiplier>

- **Variance command will allow the router to include routes with a metric smaller than multiplier times the minimum metric route for that destination, where multiplier is the number specified by the variance command**

Variance Example

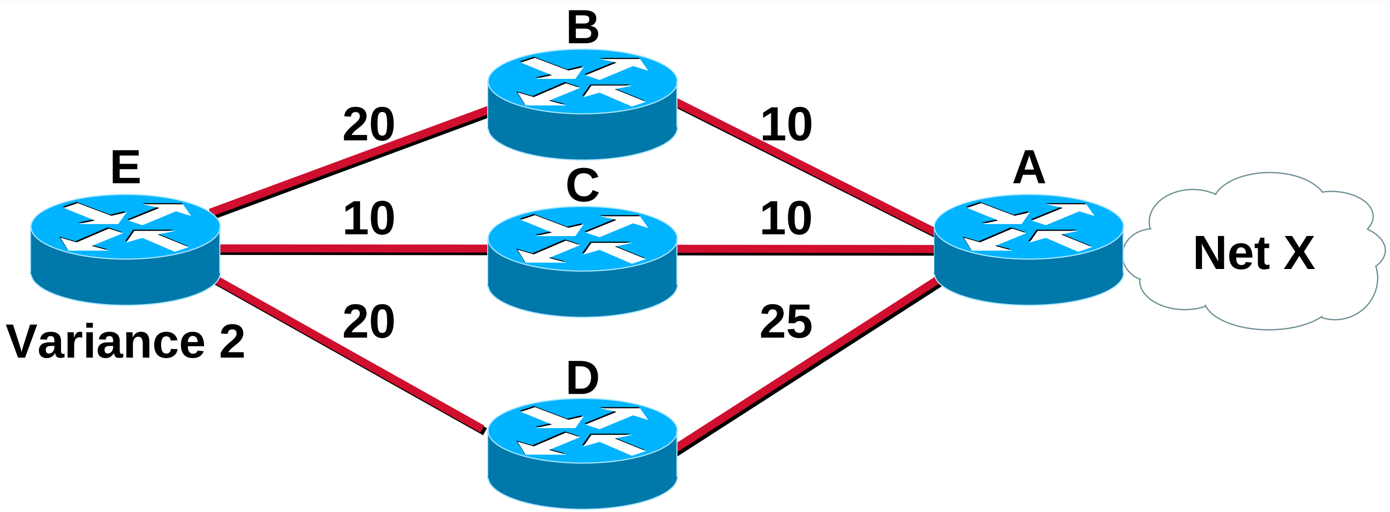

- **Router E will choose router C to get to net X FD=20**
- **With variance of 2, router E will also choose router B to get to net X**
- **Router D will not be used to get to net X**

Agenda

- Fundamentals of E-IGRP
- DUAL
- Summarization and Load Balancing
- Query Process
- Deployment Guidelines with E-IGRP
- Summary

E-IGRP Query Process

- **E-IGRP is Advanced Distant Vector. It relies on its neighbor to provide routing information**
- **If a route is lost and no feasible successor is available, E-IGRP needs to converge fast, its only mechanism for fast convergence is to actively query for the lost route to its neighbors**

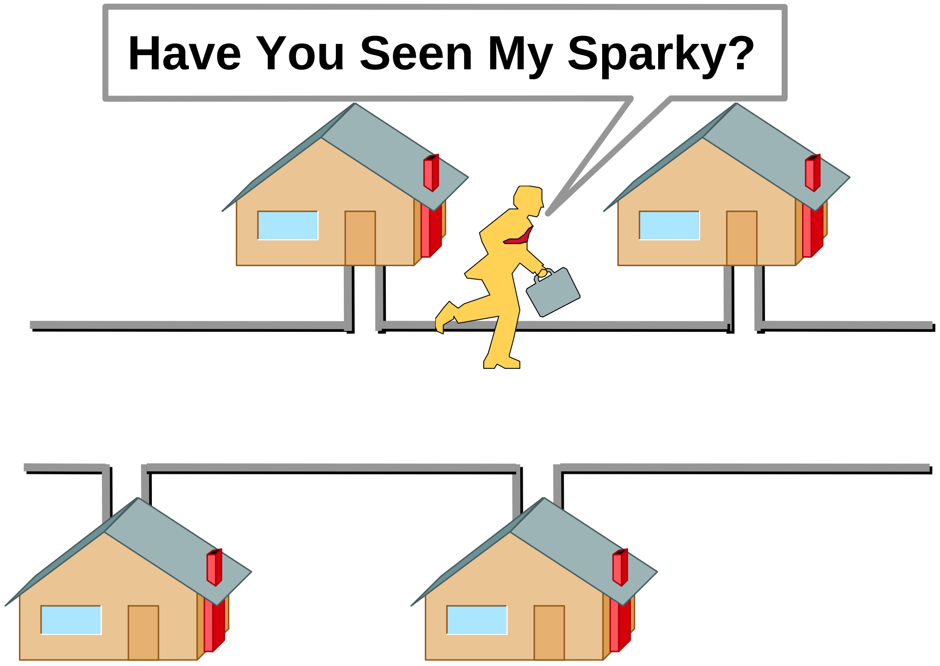

E-IGRP Query Process

- Queries are sent out when a route is lost and no feasible successor is available
- The lost route is now in active state
- Queries are sent out to all of its neighbors on all interfaces except the interface to the successor
- If the neighbor does not have the lost route information, queries are sent out to their neighbors

E-IGRP Query Process

- The router will have to get ALL of the replies from the neighbors before the router calculates the successor information
- If any neighbor fails to reply the query in three minutes, this route is stuck in active and the router resets the neighbor that fails to reply
- Solution is to limit query range to be covered later in presentation

E-IGRP Query Range

- **Autonomous System Boundaries**

  Contrary to popular belief, queries are not bounded by AS boundaries. Queries from AS 1 will be propagated to AS 2

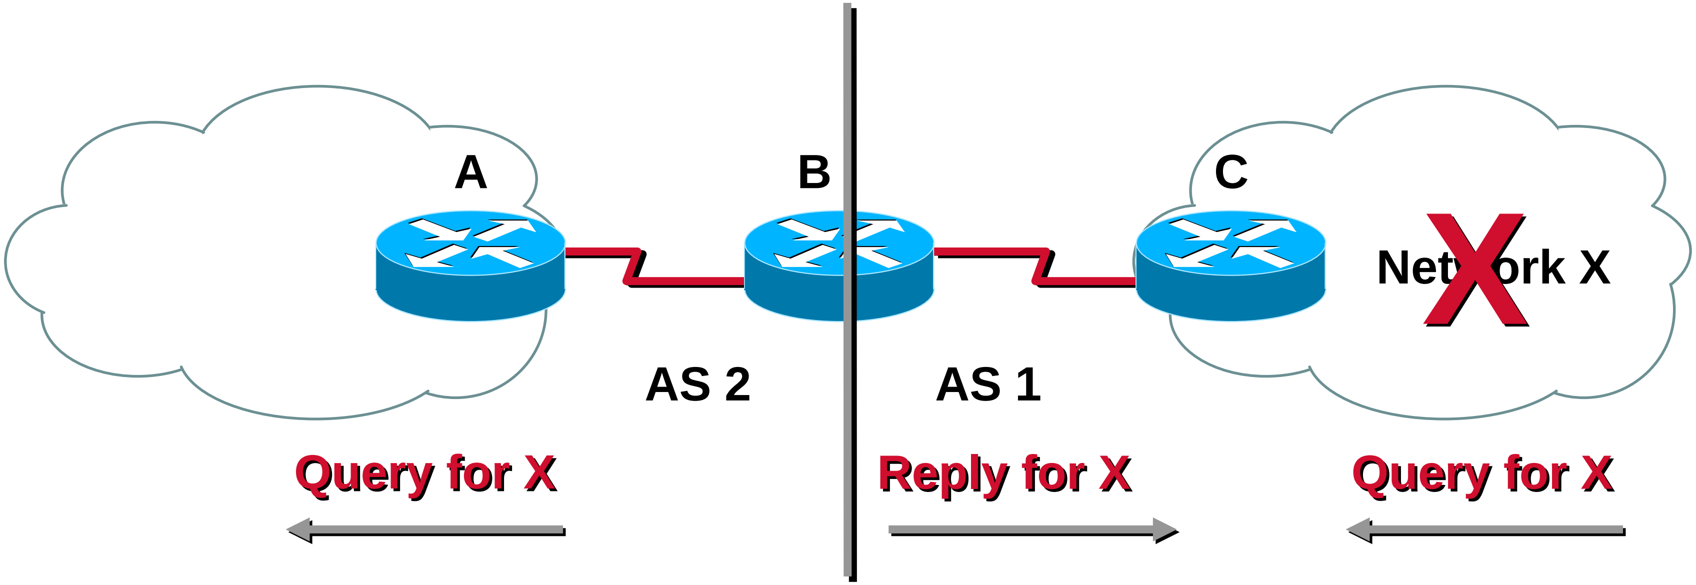

E-IGRP Query Range

- **Summarization point**
  - Auto or manual summarization bound queries
  - Requires a good address allocation scheme

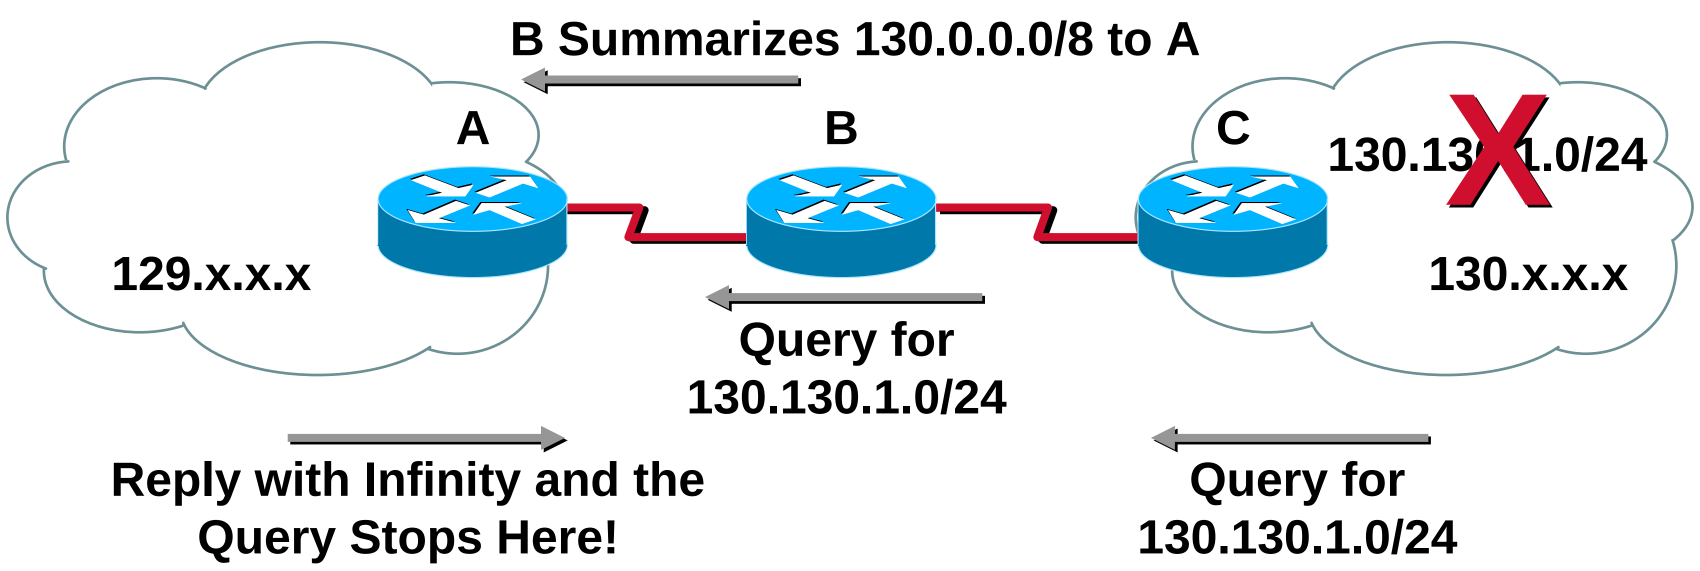

E-IGRP Bandwidth Utilization

- E-IGRP by default will use up to 50% of the link bandwidth for E-IGRP packets
- This parameter is manually configurable by using the command:

  ip bandwidth-percent E-IGRP <AS-number> <nnn>

- Use for greater E-IGRP load control

Bandwidth over WAN Interfaces

- **Bandwidth utilization over point-to-point subinterface Frame Relay**
  - **Treats bandwidth as T1 by default**
  - **Best practice is to manually configure bandwidth as the CIR of the PVC**

Bandwidth over WAN Interfaces

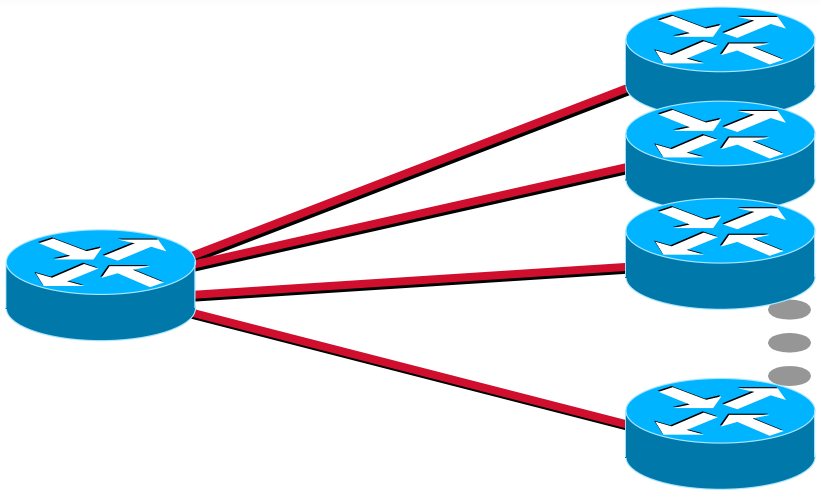

- **Bandwidth over multipoint Frame Relay, ATM, SMDS, and ISDN PRI:**

  E-IGRP uses the bandwidth on the main interface divided by the number of neighbors on that interface to get the bandwidth information per neighbor

Bandwidth over WAN Interfaces

- **Each PVC might have different CIR, this might create E-IGRP packet pacing problem**
  - **Multipoint interfaces:**
    - **Convert to point-to-point**
    - **Bandwidth configured = (lowest CIR x number of PVC)**
  - **ISDN PRI:**
    - **Use Dialer Profile (treat as point-to-point link)**

Agenda

- Fundamentals of E-IGRP
- DUAL
- Summarization and Load Balancing
- Query Process
- Deployment Guidelines with E-IGRP
- Summary

Factors That Influence E-IGRP Scalability

- **Keep in mind that E-IGRP is not plug and play for large networks**
- **Limit E-IGRP query range!**
- **Quantity of routing information exchanged between peers**

Limiting Updates/Queries—Example

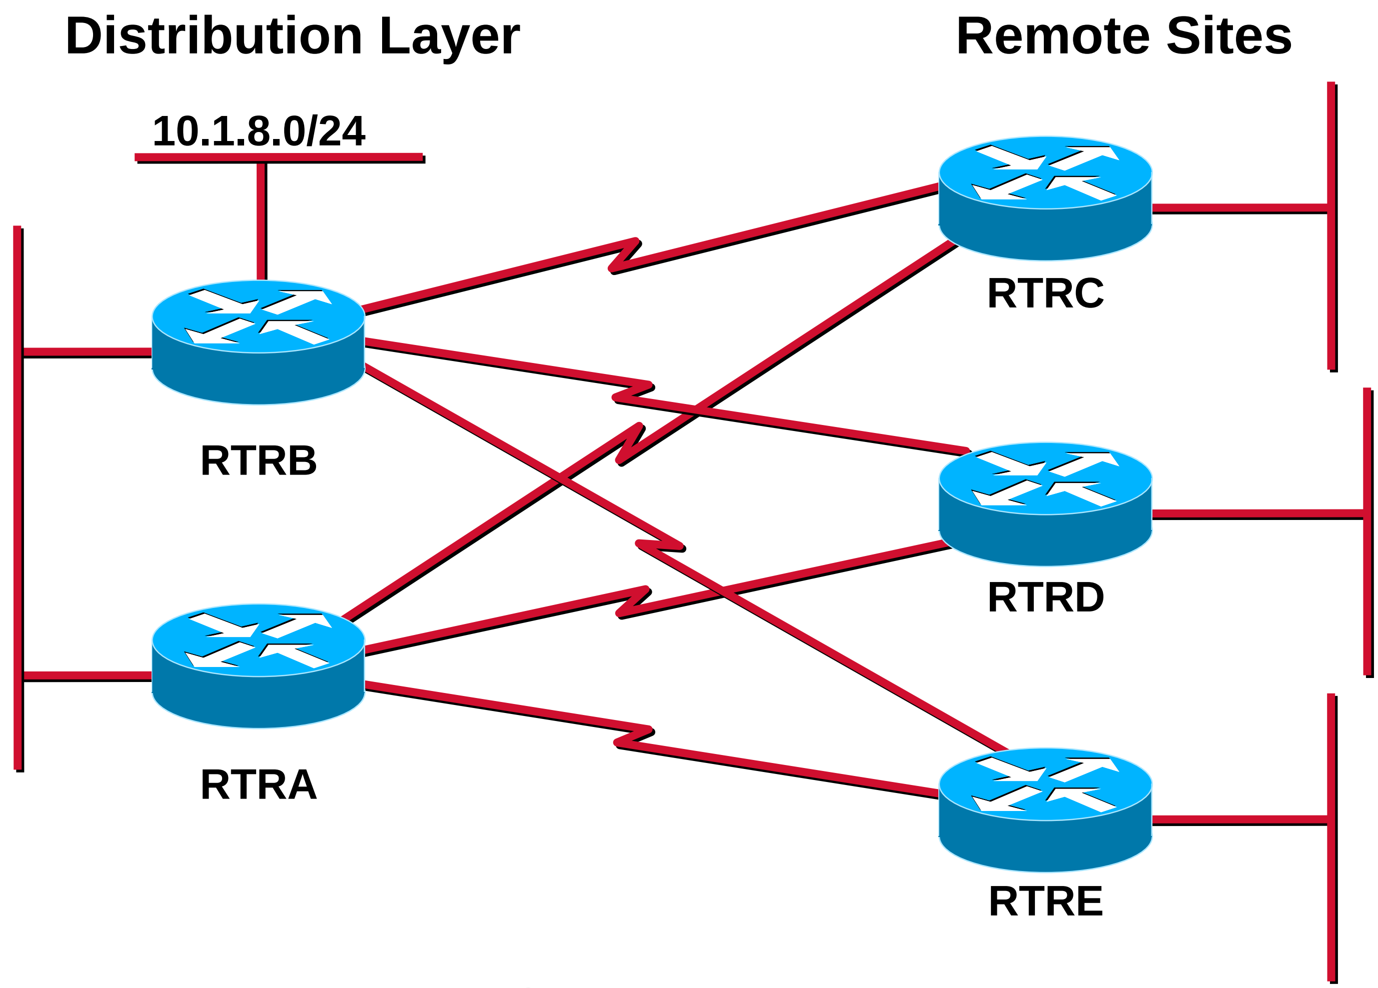

Limiting Size/Scope of Updates/Queries

- **Evaluate routing requirements**

  What routes are needed where?

- **Once needs are determined**

  Use summary address

  Use new E-IGRP Stub feature
  (To be discussed later)

  Use distribute lists

Limiting Updates/Queries—Example

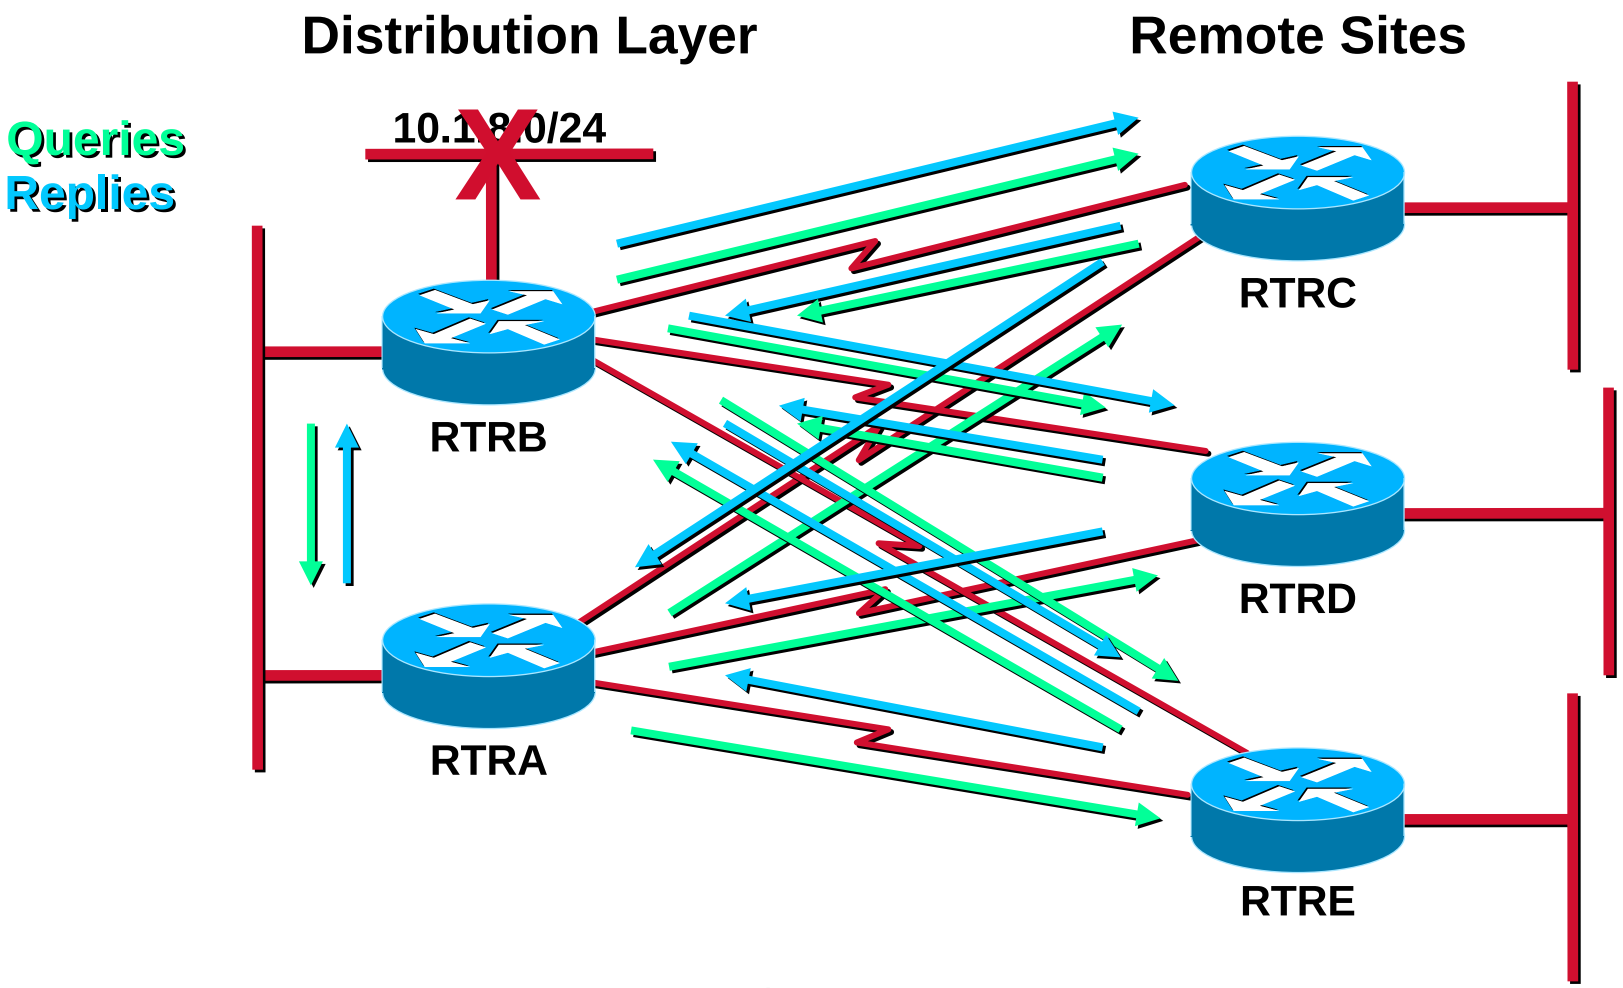

Limiting Updates/Queries—Summary

- **Remote routers fully involved in convergence**

  Most remotes are never intended to be transit

  Convergence complicated through lack of information hiding

Limiting Updates/Queries—Better

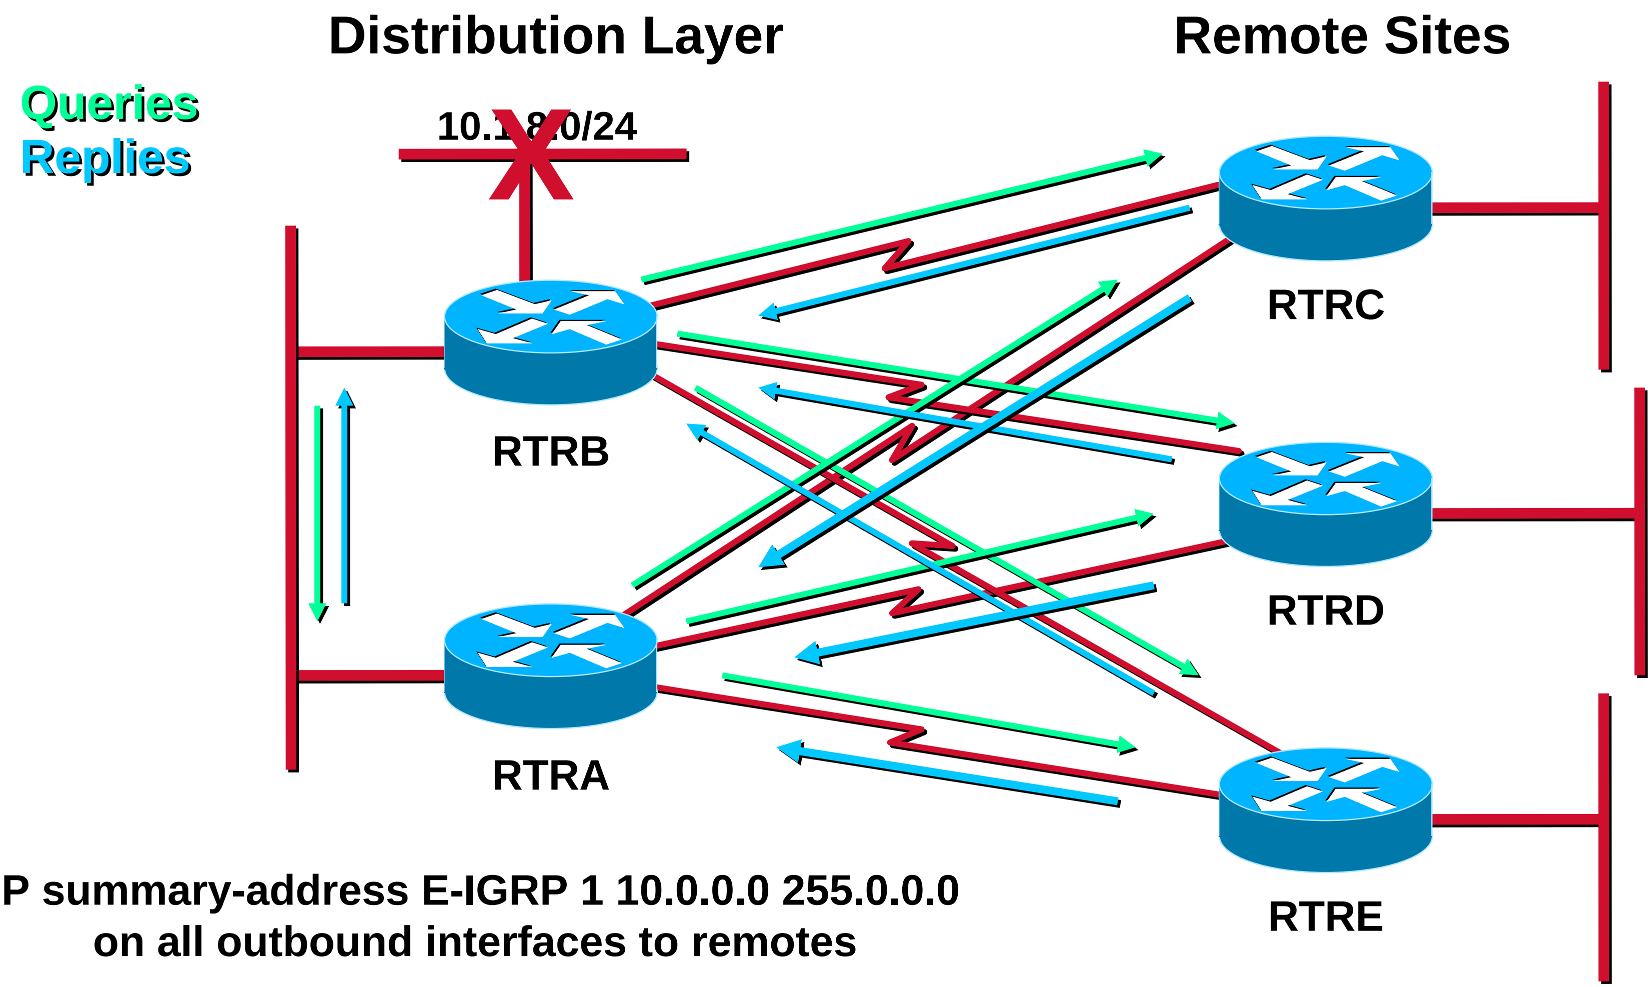

IP summary-address E-IGRP 1 10.0.0.0 255.0.0.0 on all outbound interfaces to remotes

Limiting Updates/Queries— Summary

- Convergence simplified by adding the summary-address statements
- Remote routers just reply when queried

Limiting Updates/Queries New Feature

- New E-IGRP STUB command is now available (12.0.7T and higher)
- [no] E-IGRP stub [receive-only] [connected] [static] [summary]

  Only specified routes are advertised.

  Any neighbor receiving "stub" information from a neighbor will not query those routers for any routes

Limiting Updates/Queries—Best

- Best practice is to combine Summarization and E-IGRP STUB command

Limiting Updates/Queries—Best

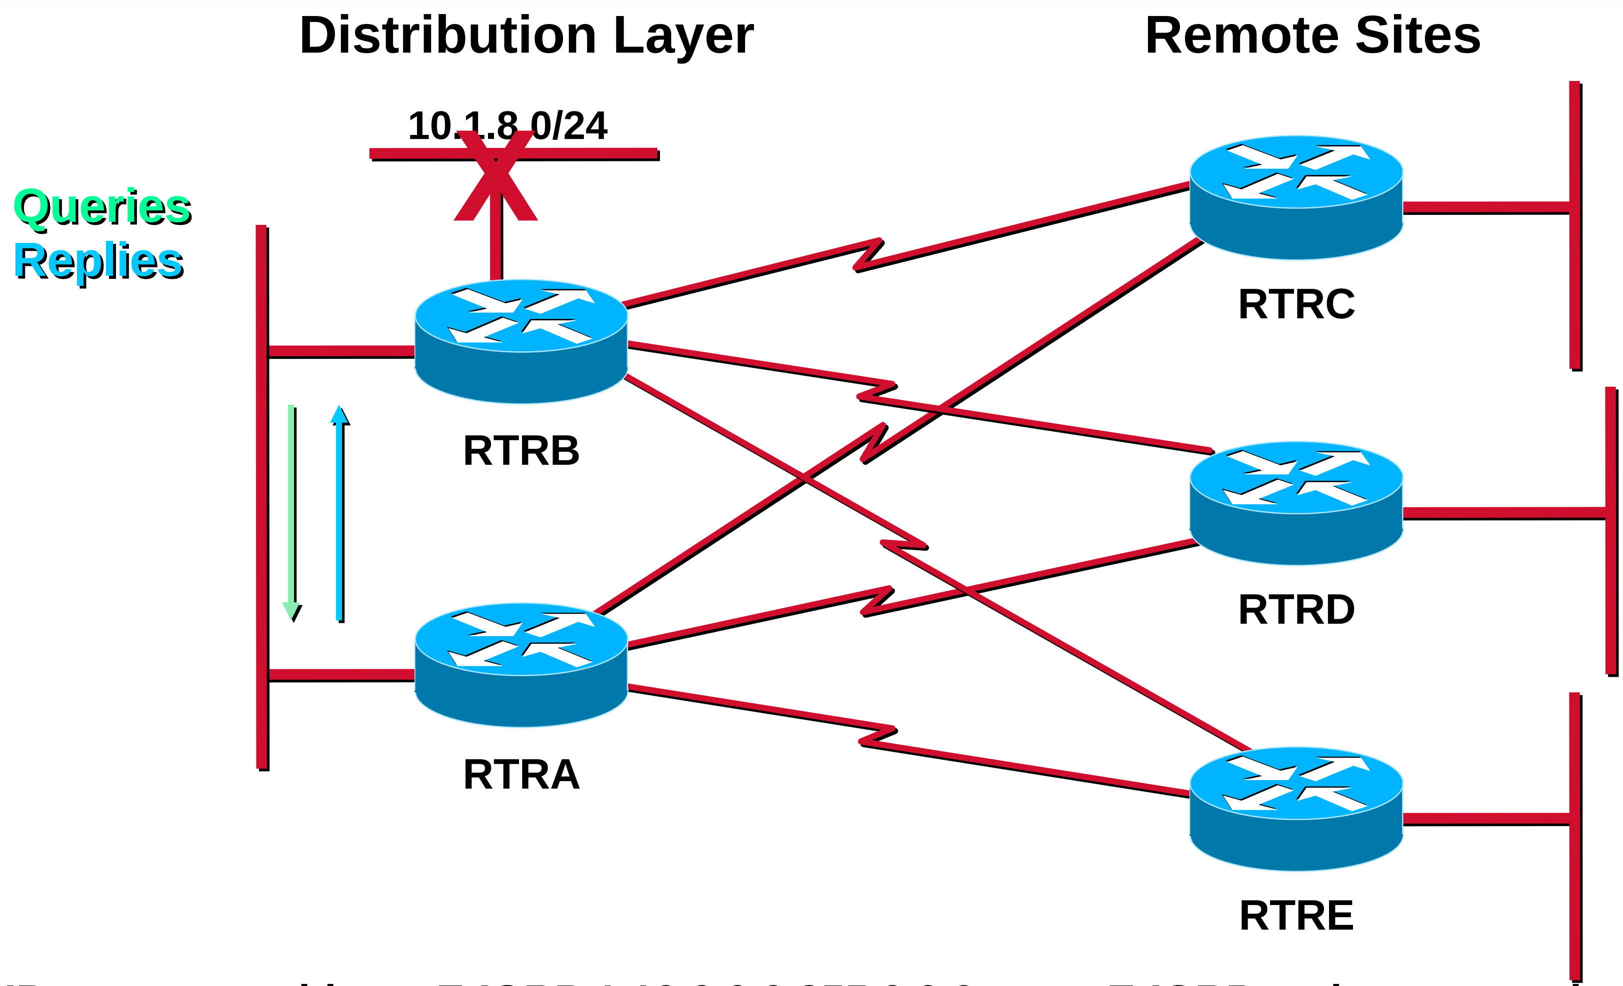

IP summary-address E-IGRP 1 10.0.0.0 255.0.0.0
on all outbound interfaces to remotes

E-IGRP stub connected
on all remote routers C, D, and E

Hierarchy/Addressing

- Permits maximum information hiding
- Advertise major net or default route to regions or remotes
- Provides adequate redundancy

E-IGRP Scalability

- E-IGRP is a very scalable routing protocol if proper design methods are used:
  - Good allocation of address space
    - Each region should have an unique address space so route summarization is possible
  - Have a tiered network design model (Core, Distribution, Access)

E-IGRP Scalability

- **Use of E-IGRP Stub command if possible**
- **Proper network resources**
  - **Sufficient memory on the router**
  - **Sufficient bandwidth on WAN interfaces**
- **Proper configuration of the "bandwidth" statement over WAN interfaces, especially over Frame Relay**
- **Avoid blind mutual redistribution between two routing protocols or two E-IGRP processes**

Tiered Network Design

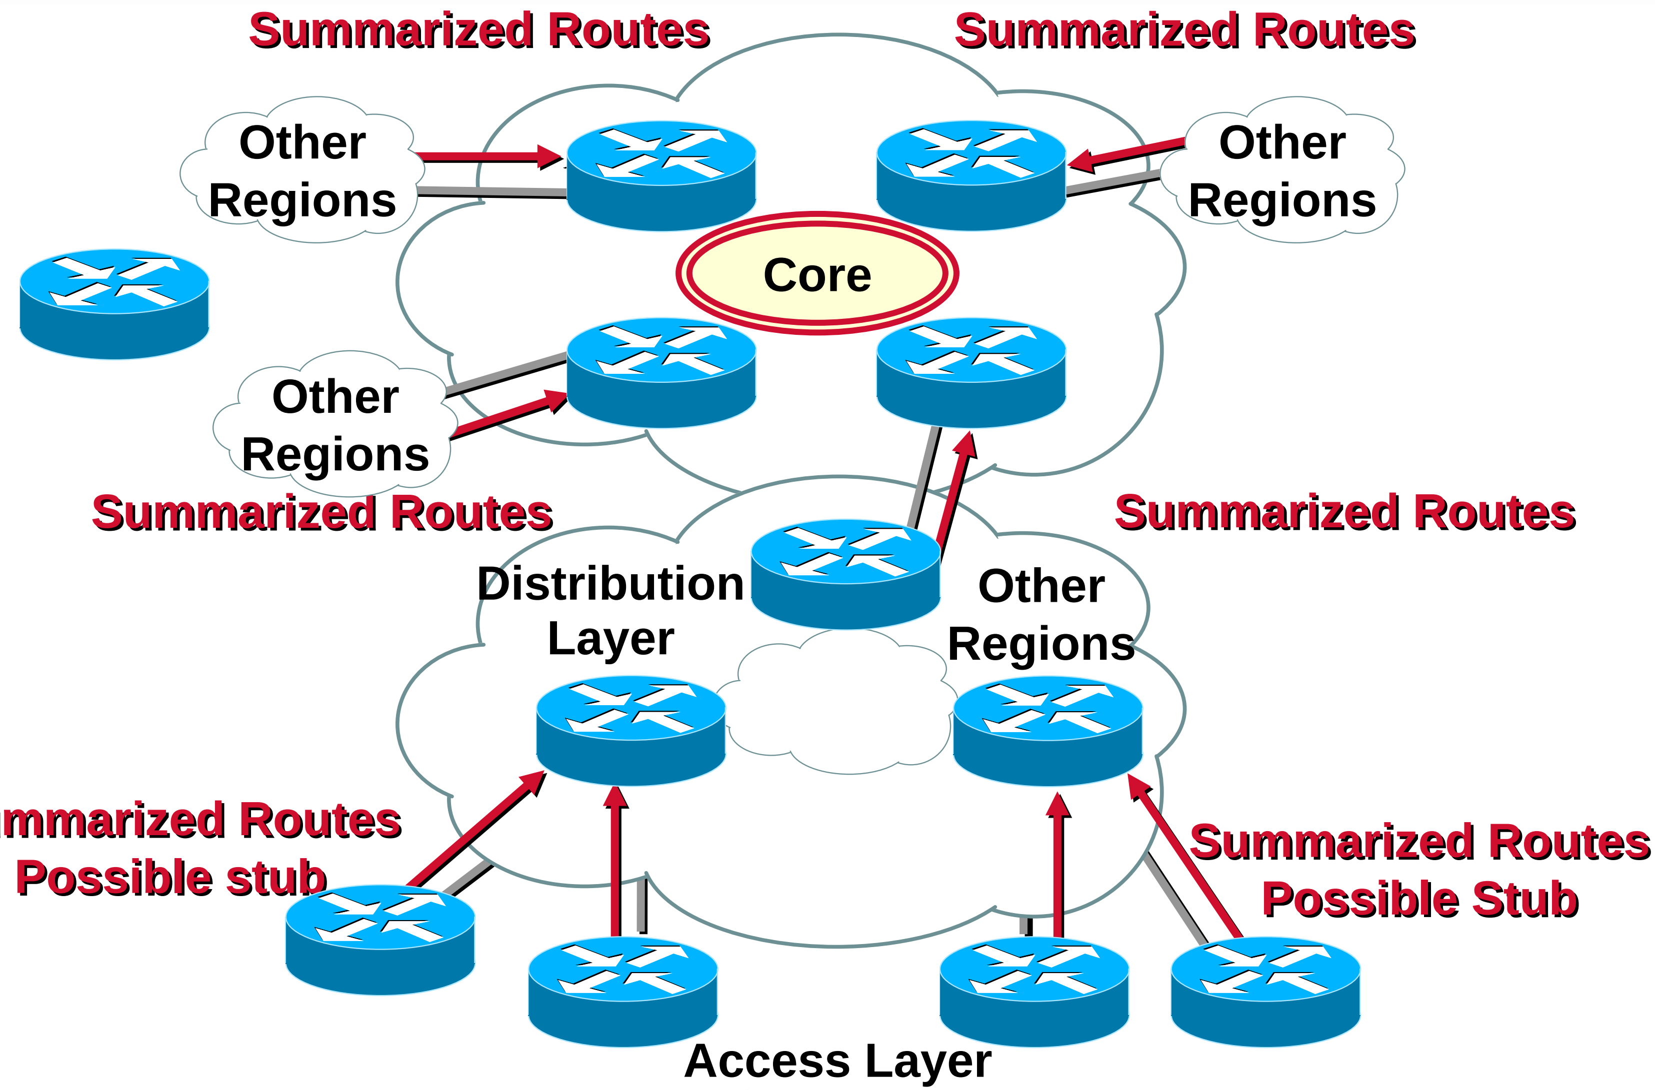

Nonscalable Network

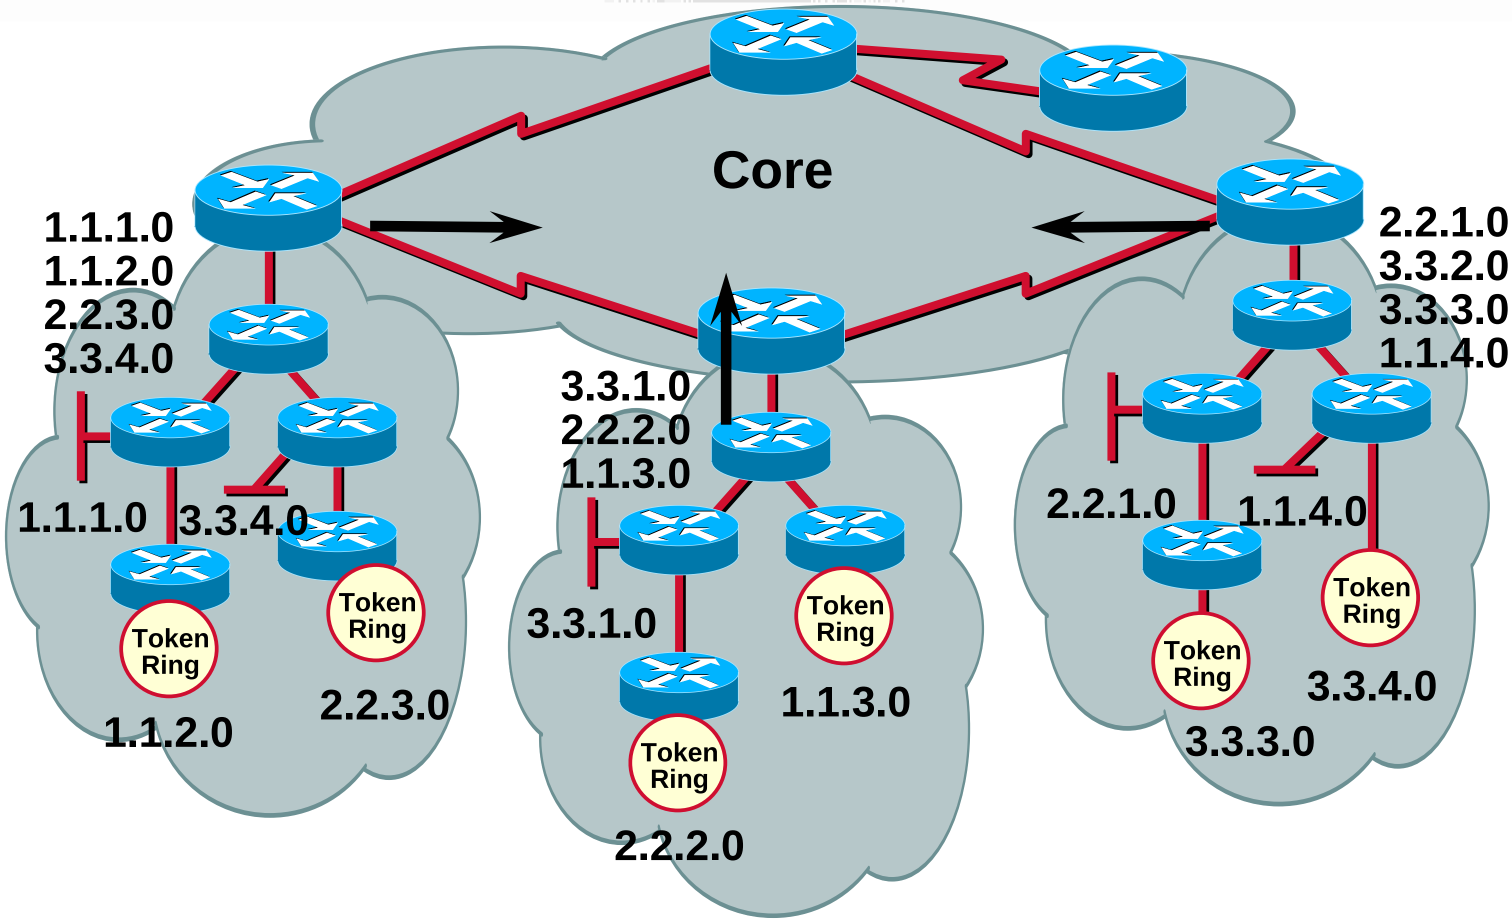

- **Bad addressing scheme**

  Subnets are everywhere throughout entire network

- **Queries not bounded**

Scalable Network

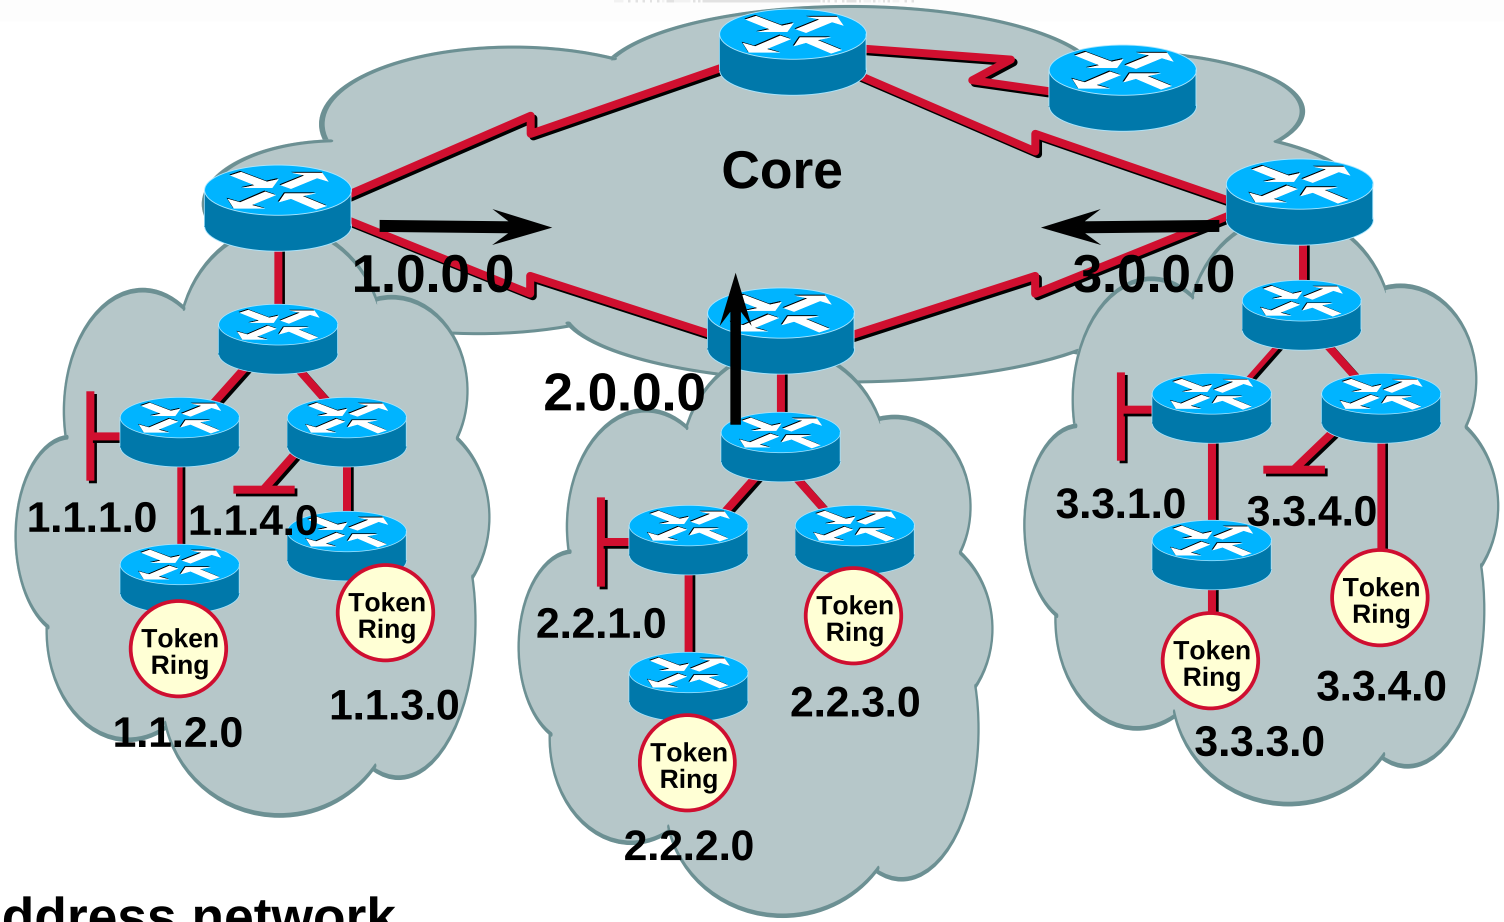

- **Readdress network**
  - **Each region has its own block of address**
- **Queries bounded by using "ip summary-address E-IGRP" command**

Summary

- Query range
  - Best way to limit query is through route summarization and new E-IGRP Stub command
- E-IGRP is not plug and play for large networks
  - It's a very scalable protocol with little design requirement
- Optimizing E-IGRP network
  - Limiting query range
    - Route summarization
    - Tiered network design
    - Use of E-IGRP Stub command
  - Sufficient network resources

Deploying IGRP/E-IGRP

Session 2208

Please Complete Your Evaluation Form

Session 2208

Cisco Systems®

Empowering the Internet Generation℠

