Which Routing Protocol Is Best

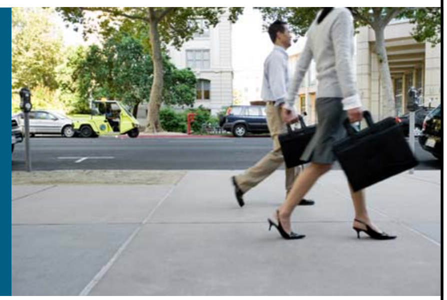

Presented by Rick Burts

Chesapeake NETCRAFTSMEN

CISCO PARTNER Premier Certified

About the Speaker

- Rick Burts, Senior Consultant Chesapeake NetCraftsmen
  - Cisco CCIE # 4615, CCSI # 94069
  - Highly ranked in Cisco NetPro forum
  - Network Design & Implementation consulting
  - Consultant to Government and Enterprise Customers
  - Taught many of the Cisco courses
  - 14 years experience with Cisco networking

www.netcraftsmen.net

rburts@netcraftsmen.net

Chesapeake NETCRAFTSMEN

“ Is one protocol better than the others?

Which routing protocol should I use in my network?

Should I change from the one I’m using? ”

The Questions

- Is one routing protocol better than any other protocol?
- Define “Better!”

- Converges faster?
- Uses less resources?
- Fits a particular topology better?
- Easier to troubleshoot?
- Easier to configure?
- Scales to a larger number of routers, routes, or neighbors?
- More flexible?
- Degrades more gracefully?

Should I Change?

The Answer Could Be Yes if:

- The network is complex enough to “bring out” a protocol’s specific advantages
- A specific feature (or set of features) can be defined that will benefit your network tremendously

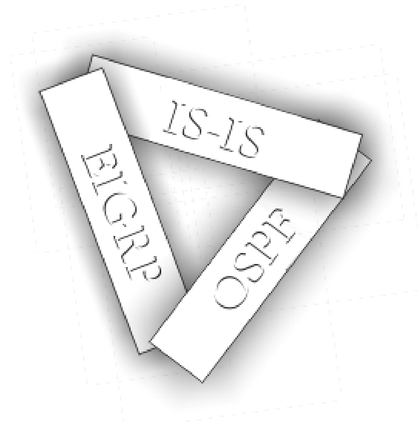

Should I Change?

But, the Answer Could Also Be No if:

- No real benefit would be gained (features/scaling/perform ance). Steep learning curve, complex training etc ..
- Let’s consider some specific topics for each protocol …….

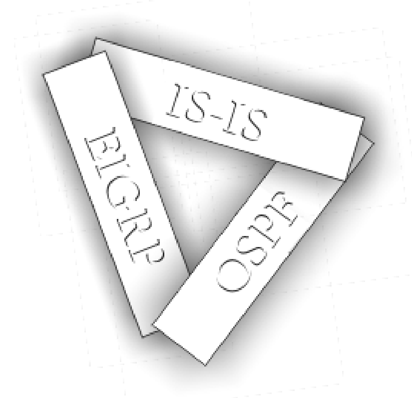

Which Routing Protocol Is Best?

Main Agenda

- Convergence Speed
- Topology Support
- Protocol Features
- Summary

Convergence Speed

Convergence Speed

- Convergence using Equal Cost Paths
- Convergence using Link State
- Convergence using EIGRP
- Convergence Summary

Convergence Speed

- Which protocol converges faster?
- OSPF versus EIGRP
    - OSPF using Dijkstra is a link state protocol
    - EIGRP using DUAL is a hybrid protocol
    - Is DUAL faster, or Dijkstra?
- General Rules of Thumb
    - The more routers involved in convergence, the slower convergence will be
    - The more prefixes involved in convergence, the slower convergence will be

Convergence Speed

- Three main steps to convergence
  - Detect the failure
  - Calculate new routes around the topology change
  - Add changed routing information to the routing table
- The first and third steps are similar for any routing protocol, so we'll focus on the second step
- But, it's important to keep the other two in mind, since they often impact convergence more than the routing protocol does

Convergence Speed—Equal Cost Paths

- B_C_E and B_D_E are equal cost
- If C fails, B and E will send traffic via D only
- Number of nodes whose forwarding behavior is impacted by the change: 2 (B and E)
- Convergence time is in milliseconds for both protocols

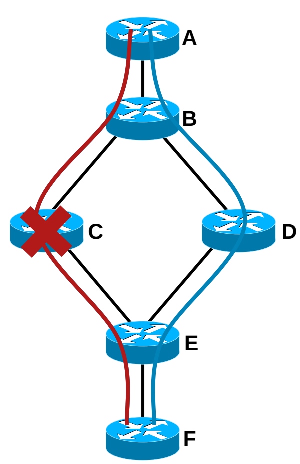

Convergence Speed—Link State

- C fails
- B and E flood new topology information
- All routers run SPF to calculate the new shortest paths through the network
- B and E change their routing tables to reflect the changed topology
- Number of nodes whose routing tables are impacted by the change: 2 (B and E)

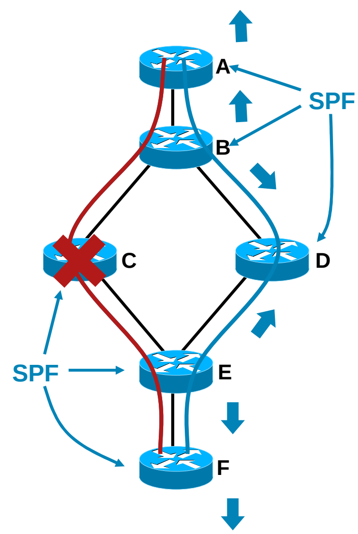

 13

Convergence Speed—Link State

- Within a single flooding domain
  - A single area in OSPF
  - A single flooding domain in IS-IS
- Convergence time depends on flooding timers, SPF timers, and number of nodes/leaves in the SPF tree
- What happens when we cross a flooding domain boundary?

 14

Convergence Speed—Link State

- E floods topology changes to C and D
- C and D summarize these topology changes (removing the link specific information), and flood it to B
- B then builds a summary itself and floods it into area 2
- A then has the updated information to exit the area and reach other routers in other areas

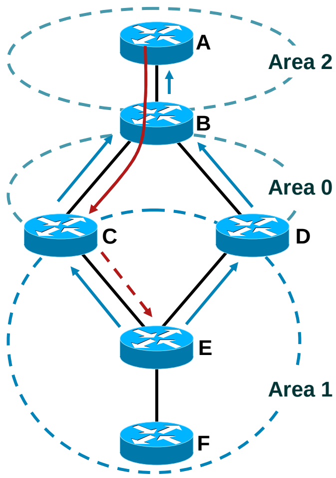

 15

Convergence Speed—Link State

- Between flooding domains, link state protocols have “distance vector” characteristics
- This can have negative or positive impacts on convergence time in a large network
  - Reduces tree size
  - Allows partial SPFs, rather than full SPFs
  - Introduces translation and processing at the flooding domain boundaries
- The impact is primarily dependant on the network design

 16

Convergence Speed—Link State

- Within a flooding domain

  The average convergence time, with default timers, is going to be around three to seven seconds

  With fast timers, the convergence time can be in the milliseconds

  There are operational 200+ node IS-IS and OSPF networks with sub-second convergence times

- Outside the flooding domain

  Network design and route aggregation are the primary determining factors of convergence speed

Convergence Speed—EIGRP

- DUAL works on a simple geometric principle:

  If my neighbor’s cost to reach a given destination is less than my best cost, then the alternate path cannot be a loop

- B_D_E_F is 35
- B_C_E_F is 30
- D_E_F is 20, which is less than the best path, 30, so B_D_E_F cannot be a loop

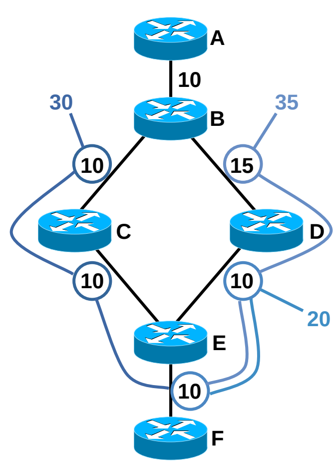

Convergence Speed—EIGRP

- B will install the path through C, and mark the path through D as a feasible successor
- When C fails, B looks for alternate loop free paths
- Finding one, it installs it
- Convergence time is in the milliseconds
- Number of nodes whose routing tables are impacted by the change: 2 (B and E)

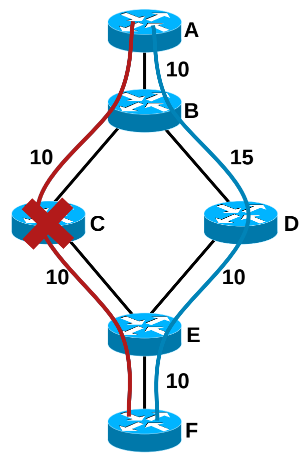

 19

Convergence Speed—EIGRP

- If the second path cannot be proven loop free
- B and E detect the failure, and have no alternate path
- B queries A and D
  - A replies that it has no path
  - D replies with its alternate path
- E queries D and F
  - F replies that it has no path
  - D replies with its alternate path

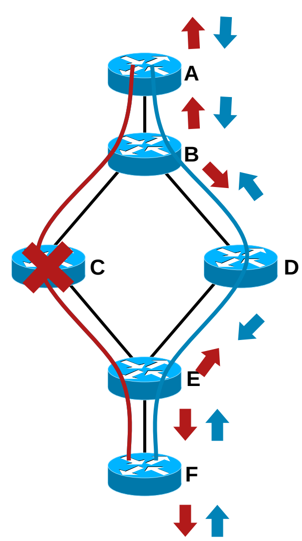

 20

Convergence Speed—EIGRP

- For paths with feasible successors, convergence time is in the milliseconds
  - The existence of feasible successors is dependant on the network design
- For paths without feasible successors, convergence time is dependant on the number of routers that have to handle and reply to the query
  - Queries are blocked one hop beyond aggregation and route filters
  - Query range is dependant on network design
- Good design is the key to fast convergence in an EIGRP network

 21

Convergence Summary

- We can sort typical convergence times into three groups:
  - EIGRP with a feasible successor
  - Link State with modified timers
  - EIGRP without a feasible successor and good design
  - Link State inter flooding domain
  - Link State with default timers
  - EIGRP without a feasible successor and bad design
- Each pair could flip with the grouping

 22

Convergence Summary

- The RP Scaling Team did a series of tests comparing failover times under varying conditions
- Essentially, generate routes on the route generator from A to D
- Fail either B or C, and see how long it takes A to switch paths

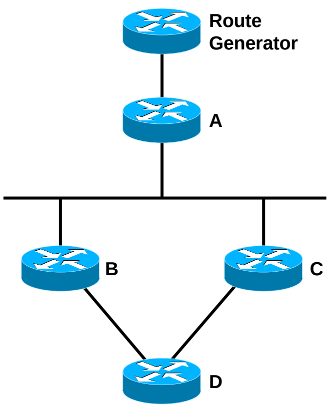

Convergence Summary

- IS-IS with default timers
- OSPF with default timers
- EIGRP without feasible successors
- OSPF with tuned timers
- IS-IS with tuned timers
- EIGRP with feasible successors

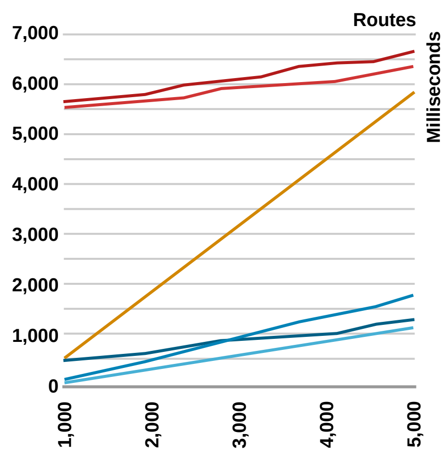

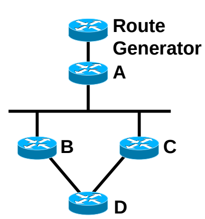

Convergence Summary

- It's possible to converge in under one second using any protocol, with good network design
- Rules of thumb:

  More aggregation tends towards better performance for EIGRP

  Less aggregation tends towards better performance for link state protocols

  If you're going to use link state protocols, tune the timers; but if you tune the timers, be careful with HA features, like GR/NSF

 25

Topology Support

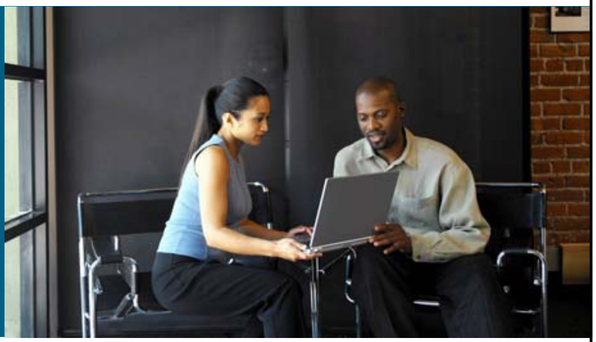

 26

Topology

- Hub and spoke
- Full Mesh
- Support for hierarchy
- Topology summary

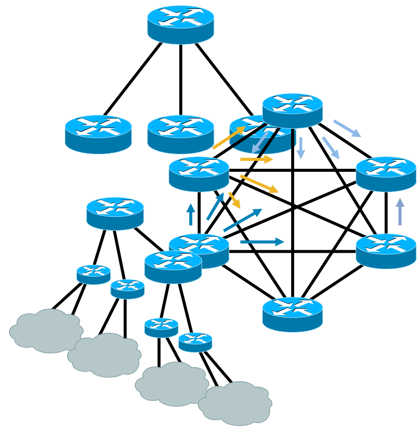

Link State Hub and Spoke

- OSPF and IS-IS are similar when designing for hub and spoke topologies, so we'll look at them together
- Link state protocols rely on every router within a flooding domain having the same view of the network's topology to calculate loop free paths
- Link state flooding rules have implications for scaling and design in hub and spoke networks

Link State Hub and Spoke

- Although B can only reach C through A, it still receives all of C's routing information
- As the number of remote sites increases, the amount of information each remote site must process and store also increases
- This limits scaling in link state hub and spoke networks

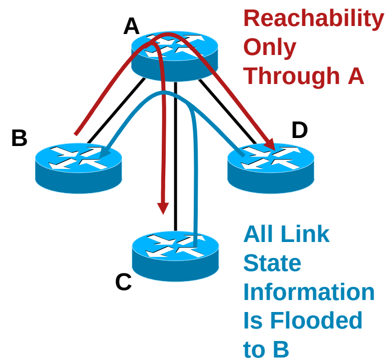

Link State Hub and Spoke

- Controlling route distribution
- There's no way to allow C and D to receive information about 10.1.1.0/24, and not E and F

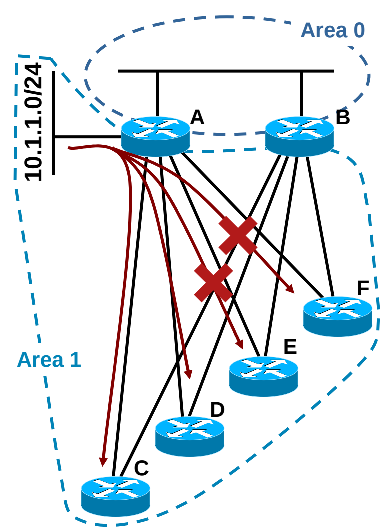

Link State Hub and Spoke

- Transiting remote sites
- C and D issue summaries containing 10.1.1.0/24
- A chooses D as it's best path to the summary
- The D to E link fails
- How can we prevent D from using the link through F to reach 10.1.1.0/24?

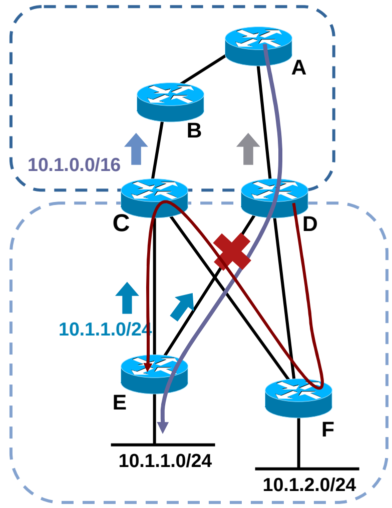

Link State Hub and Spoke

- Place a link between C and D within the same area as the hub and spoke network
- The link cost between C and D should be lower than the link cost through F, causing D to route through this new link

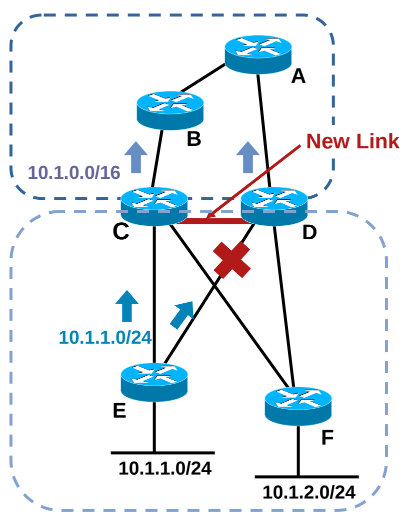

Link State Hub and Spoke

- For each hub and spoke flooding domain you add to the hub routers, you need an additional link between the hub routers in that domain
- You can use virtual links, such as Ethernet VLANS
- This can become difficult to manage in a large scale hub and spoke network

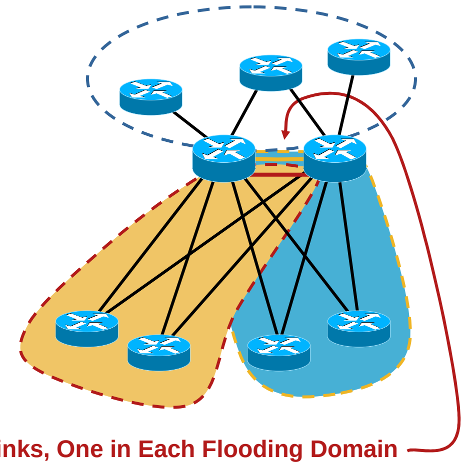

EIGRP Hub and Spoke

- Summary black holes
- C and D are both summarizing 10.1.0.0/16 towards A and B
- C and D are advertising a default only to E and F
- A chooses D's path
- When the D to E link fails, D is still advertising 10.1.0.0/16 (based on 10.1.2.0/24 from F)
- Traffic forwarded to 10.1.1.1 from A will be dropped

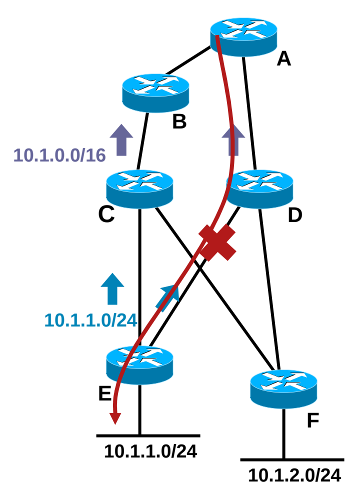

EIGRP Hub and Spoke

- You can resolve this by placing a link between C and D, without summarization

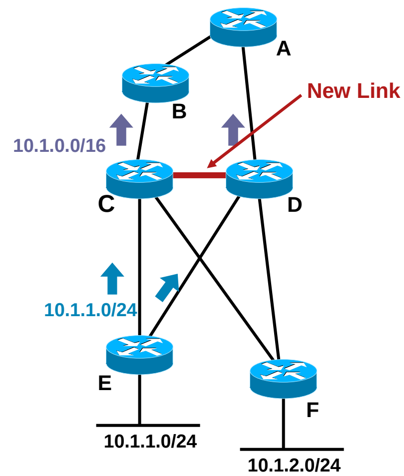

EIGRP Hub and Spoke

- Controlling query range
- If A loses its connection to 10.1.1.0/24, it builds and transmits five queries: one to each remote, and one to B
- Each of the remote sites will query B
- B must process and reply to five queries

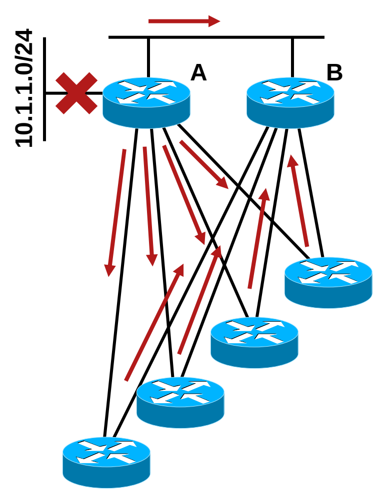

EIGRP Hub and Spoke

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

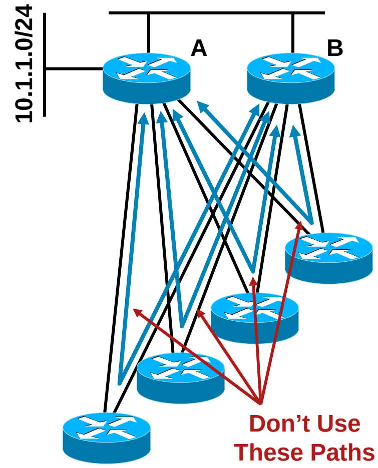

EIGRP Hub and Spoke

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

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

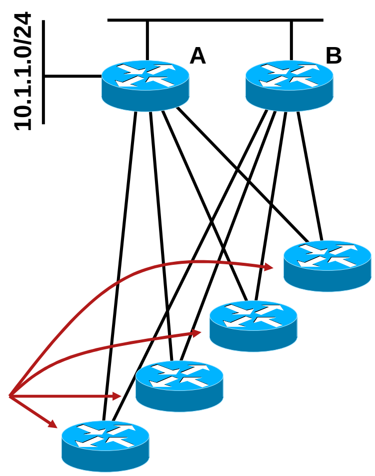

EIGRP Hub and Spoke

- Marking the spokes as stubs allows them to signal A and B that they are not valid transit paths
- A - simply will not query the remotes, reducing the total number of queries in this example to one

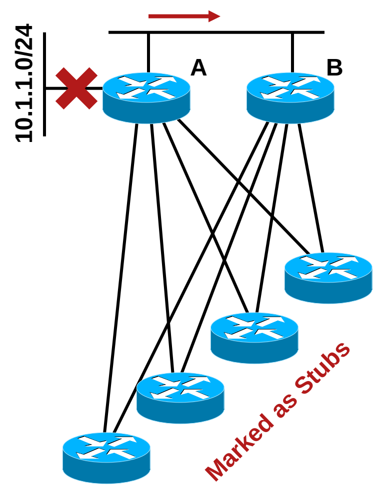

EIGRP Hub and Spoke

- Marking these remotes as stubs also reduces the topological complexity of the network
- Without stubs, B believes it has five paths to 10.1.1.0/24, so it has to maintain five topology table entries

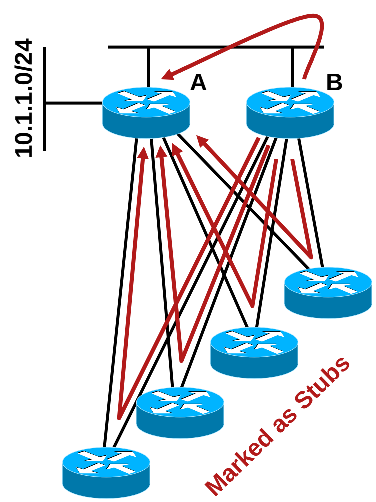

EIGRP Hub and Spoke

- Routers which are configured as stubs will only advertise locally connected or redistributed destinations
- These remotes will not pass A's advertisement of 10.1.1.0/24 to B
- B will only have one path to 10.1.1.0/24

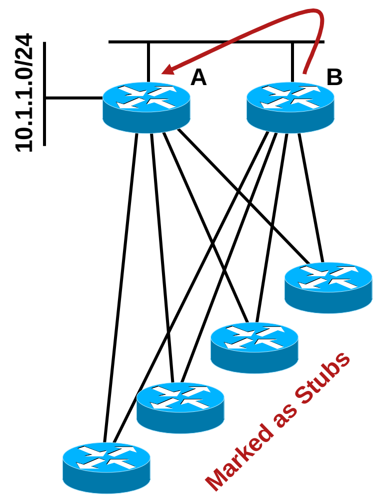

Hub and Spoke

- Testing by the RP Scaling Team

  EIGRP has been tested up to 1400 neighbors in a lab environment (dual homed remote configuration)

  OSPF has been tested up to 700 neighbors in a lab environment (dual homed remote configuration)

- In the field, we see up to 800 dual homed remotes with EIGRP, and up to about 200 with OSPF

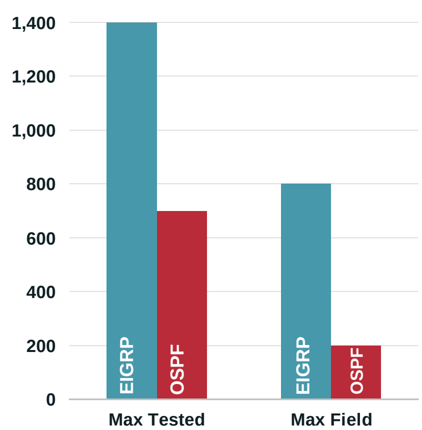

Hub and Spoke

- The RP Scaling Team has also tested initial convergence and hard failover times
  - 600 dual homed remote sites
  - For hard failover, primary hub was powered down

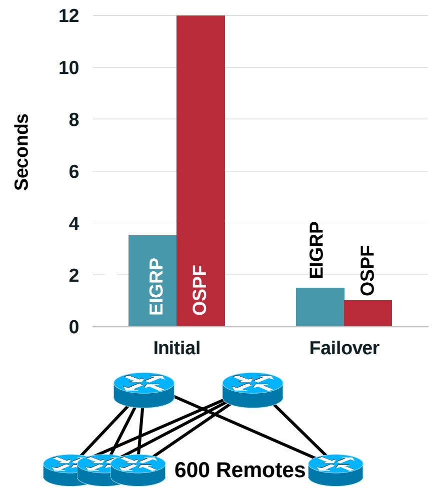

Hub and Spoke

	Scaling	Issues
Link State	All Remote Sites Receive All Other Remote Site Link State Information; Moderate Scaling Capability	No Effective Means to Control Distribution of Routing Information
		Care Must Be Taken to Prevent Transiting Traffic Through Remote Sites
EIGRP	Stub Remote Routers with Filtering and Aggregation; Excellent Scaling Capability	Care Must Be Taken with Summary Black Holes

Full Mesh

- Full mesh topologies are complex:

  2 routers = 1 link

  3 routers = 3 links

  4 routers = 6 links

  5 routers = 10 links

  6 routers = 15 links

  …

- Adjacencies =

  routers * (routers-1) / 2

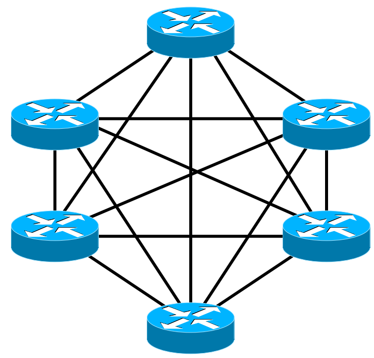

Link State Full Mesh

- Flooding routing information through a full mesh topology is also complicated
- Each router will, with optimal timing, receive
  at least one copy of every new piece of information from each neighbor on
  the full mesh
- There are several techniques you can use
  to reduce the amount of flooding in a full mesh

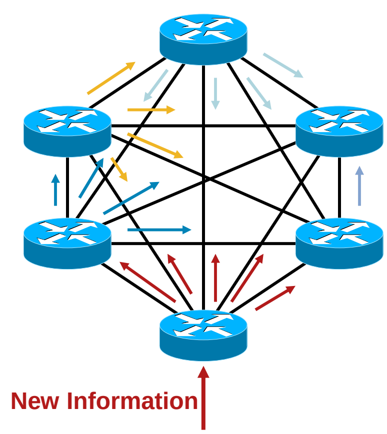

Link State Full Mesh

- OSPF and IS-IS can both use mesh groups to reduce the flooding in a full mesh network
- Mesh groups are manually configured “designated routers” on the full mesh

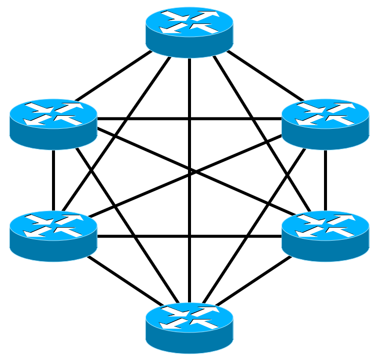

Link State Full Mesh

- Pick one or two routers to flood into the mesh, and block flooding on the remainder
- This will reduce the number of times information is flooded over a full mesh topology
- This isn’t a commonly used configuration

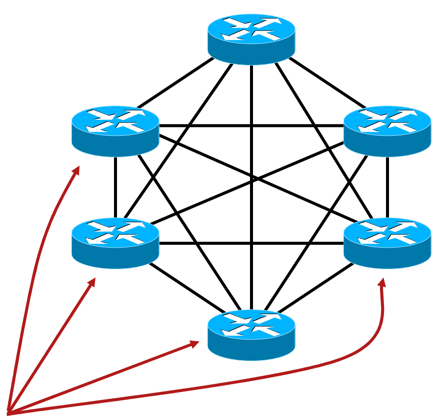

On Each Serial Interface:

```
interface serial x
 ip ospf database-filter all out
 ....
```

EIGRP Full Mesh

- Routes must be advertised between every pair of peers in the mesh so each router has the correct next hop and routing information
- Number the links so they can be summarized to a single advertisement at the edge
- Number the links so the link information can be filtered out at the edge

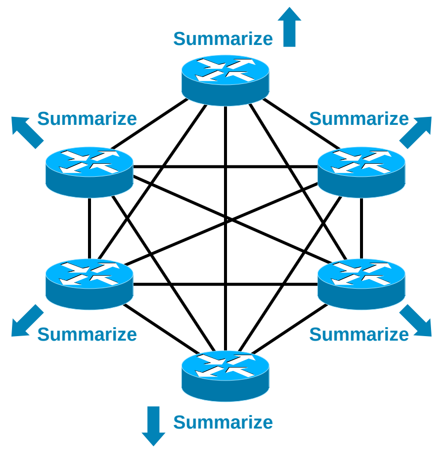

Full Mesh

OSPF	Use `ip ospf database-filter all out` to Manually Designate Flooding Points and Increase Scaling Through a Full Mesh
IS-IS	Use `isis mesh-group` or `isis mesh-group blocked` to Manually Designate Flooding Points and Increase Scaling Through a Full Mesh
EIGRP	Summarize into and out of the Full Mesh

OSPF Support for Hierarchy

- OSPF has a “hard edge” at flooding domain borders
- Summarization and filtering can occur at this border

  Summarization and filtering can also be configured at routers redistributing routes into OSPF
- In a two layer hierarchy, the flooding domain border naturally lies on the aggregation/core boundary

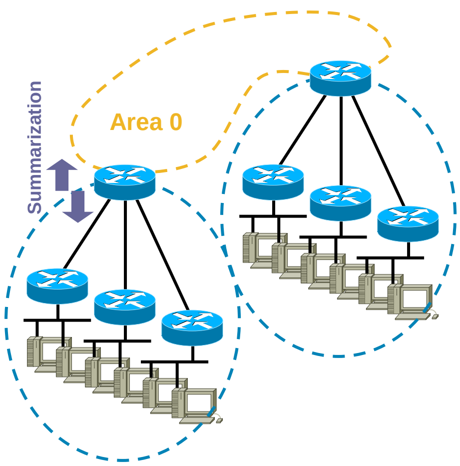

OSPF Support for Hierarchy

- In a three layer hierarchy, the decision of where to place the area border is more difficult
- Typically, the best is to flow around complex areas of the network, attempting to separate them into different areas
- Examples would include full mesh areas, data centers with a large amount of parallelism, and large hub and spoke deployments

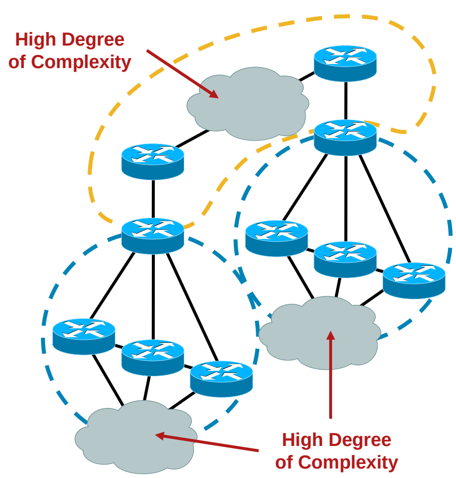

EIGRP Support for Hierarchy

- The depth of the hierarchy doesn't alter the way EIGRP is deployed; there are no "hard edges"
- Summarize at every boundary where possible
- Divide complexity with summarization points

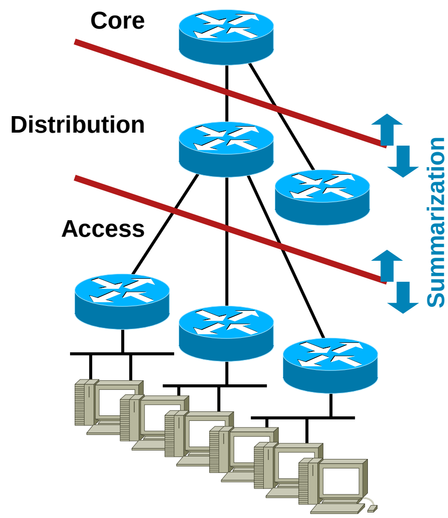

Hierarchical Division Points

OSPF	▪ Normal: All Routing Information Is Flooded into the Area ▪ Stub: Only Internal Summarized (Type 3) Information Is Flooded into the Area ▪ Totally Stub: Only a Default Is Flooded into the Area ▪ Not-So-Stubby: Only Internal Summarized (Type 3) Information Is Flooded into the Area; Routers Within the Area Can Originate Type 7's, Which Are Converted to Type 5's at the ABR ▪ Totally Not-So-Stubby: Only a Default Is Flooded Into the Area; Routers Within the Area Can Originate Type 7's, Which Are Converted to Type 5's Are the ABR
IS-IS	▪ L1/L2 Borders Set the Attached Bit, Equivalent to a Default Route ▪ L2 Routes Are Not Propagated into L1 Areas, Except Through Manually Configured Route Leaking

Hierarchical Division Points

OSPF	"Hard" Flooding Domain, Summarization, and Filter Border; Area Borders Need to Be Considered When Designing or Modifying the Network
IS-IS	"Softer" Flooding Domain, Summarization, and Filtering Border; L2 Overlaps L1 Domains, Providing Some Flexibility; Network Design Needs to Consider Flooding Domain Border
EIGRP	Summarization and Filtering Where Configured, No Hard or Soft Borders Other Than What the Network Dictates (But This Doesn't Imply the Network Doesn't Need to Be Designed!)

Topology Summary

Rules of Thumb

- EIGRP performs better in large scale hub and spoke environments
- Link state protocols perform better in full mesh environments, if tuned correctly
- EIGRP tends to perform better in more strongly hierarchical network models, link state protocols in flatter networks

Protocol Features

Protocol Features

- Routing policy
- On the wire efficiency
- Other features (mixed bag)
- Future work
- Summary

Routing Policy

- What is routing policy?
- Marking routes using tags for further processing elsewhere in the network
- Filtering routes learned from adjacent routers or external protocols
- Setting metrics for routes learned from an adjacent router
- Summarizing routing information

Routing Policy

OSPF Can Filter Prefixes Out of Type 3 Summary LSAs at an ABR (Cscdi43518)	router ospf 1 log-adjacency-changes area 1 filter-list prefix AREA_1_OUT out area 3 filter-list prefix AREA_3_IN in .... ! ip prefix-list AREA_1_OUT seq 10 permit 10.25.0.0/8 ge 16 ip prefix-list AREA_1_OUT seq 20 permit 172.20.20.0/24 ! ip prefix-list AREA_3_IN seq 10 permit 172.31.0.0/16
OSPF Can Summarize at ABR's	router ospf 100 .... area 0 range 10.1.0.0 255.255.0.0
OSPF Can Filter Routes, Set Their Metric, Type, Tag, and Next Hop When Redistributing Using a Route Map	access-list 100 deny 10.1.1.0 0.0.0.255 access-list 100 permit any ! route-map filter-server permit 10 match ip address 100 ! router ospf 100 .... redistribute static route-map filter-server

Routing Policy

OSPF Can Filter Routes Between the Local Database and the Routing Table Using Route Maps to Match Tags, Prefix Lists to Match Prefixes, and Access Lists to Match Prefixes; This Doesn't Impact the Routes Advertised to Adjacent Routers

(CSCdt43016)

```
hostname router-a
!
access-list 10 permit 10.1.1.0 0.0.0.255
access-list 10 deny any
!
route-map settag permit 10
 match ip address 10
 set tag 1000
!
router ospf 1
 ....
 network 10.1.2.0 0.0.0.255 area 0
 redistribute connected route-map settag

hostname router-b
!
route-map matchtag deny 10
 match tag 1000
route-map matchtag permit 20
!
router ospf 1
 network 10.1.2.0 0.0.0.255 area 0
 distribute-list route-map matchtag in
```

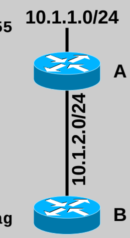

Routing Policy

EIGRP Can Mark Routes with Tags During Redistribution or Using a Route Map; These Tags Can Be Used for Filtering or Other Policy Implementations

```
route-map settag permit 10
 set tag 1000
!
router eigrp 100
 redistribute static route-map settag
 default-metric 10000 1 255 1 1500
 ....
```

EIGRP Can Set the Metrics for Any Route Using a Route Map (CSCdw22585)

```
route-map setmetric permit 10
 set metric <bandwidth> <delay> <reliability> <load>
<MTU>
!
router eigrp 100
 distribute-list route-map setmetric in
```

Routing Policy

EIGRP Can Set the Next Hop for Any Route Using Either Route Maps or no ip next-hop-self
(CSCdk23784)
	`route-map setnh permit 10`
	` set next-hop`
	`!`
	`router eigrp 100`
	` redistribute static route-map setnh`
	` default-metric 10000 1 255 1 1500`
	` ....`

EIGRP Can Filter Routes at Any Point in the Network on a per Interface Basis	`access-list 10 permit 10.1.1.0 0.0.0.255`
	`!`
	`router eigrp 100`
	` distribute-list 10 in serial 0/0`

EIGRP Can Summarize Routes at Any Point in the Network on a per i/f Basis	`interface serial 0/0` ` ip summary-address eigrp 100 10.1.0.0 255.255.0.0`

Routing Policy

	OSPF	IS-IS	EIGRP
Route Marking	Tags for Externals at Redistribution	Tags for All Routes	Tags for All Routes
Metrics	Can Be Changed for Externals at Redistribution	Can Be Set at Redistribution or Summary Creation	Can Be Set Using Route Maps
Next Hop	Can Be Changed for Externals at Redistribution	Cannot Be Changed or Set	Can Be Set for All Routes Under Varying Conditions
Filtering	Summary Information Can Be Filtered at ABR's and ASBR's	Filtering and Leaking Can Be Done at L1/L2 Borders	Anyplace
Summarization	At ABR's and ASBR's	Only at L1/L2 Borders & Redistribution Points	Anyplace

On the Wire Efficiency

How Much Information Is Transmitted on the Wire?

- This is a complex question; the answer is primarily dependant on the network design
- But, there are some characterizations we can observe by examining routing protocol packet formats

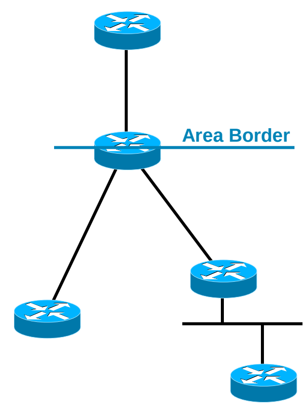

OSPF On the Wire Efficiency

OSPF Uses a Series of Packets:

- Type 1: connected links and adjacent routers
- Type 2: Designated Router (DR) representing a broadcast link
- Type 3: summary information
- Type 4: border router
- Type 5: external routes

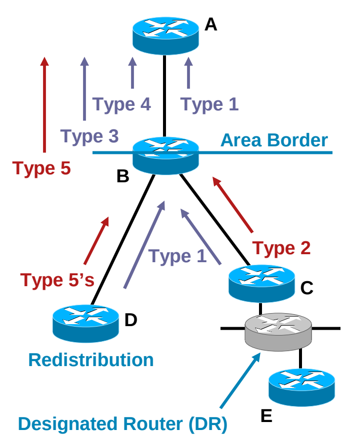

OSPF on the Wire Efficiency

- Single route changes require a single type of LSA to be transmitted, reducing on the wire overhead
- For instance, a change in the routes learned from some other protocol will cause just the type 5 containing external routes
  to be reflooded, rather than all the routing information known at A

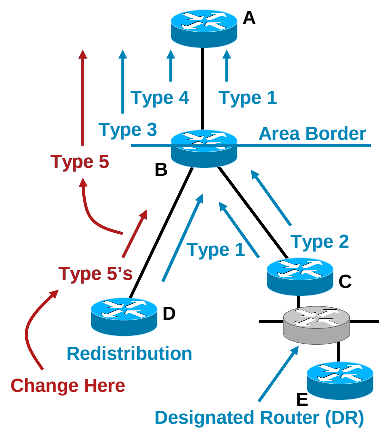

OSPF On the Wire Efficiency

- Each OSPF LSA has an LS age, set to 0 by the originating router
- Receiving routers increment the LS age over time, so they know when the LSA should be removed from
  the local database
- The originator must reoriginate the LSA before this timer reaches the maximum (60 minutes by default)

```
14. Aging The Link State Database

Each link state advertisement has
  an LS age field. The LS age is
  expressed in seconds. An
  advertisement's LS age field is
  incremented while it is
  contained in a router's
  database. Also, when copied into
  a Link State Update Packet for
  flooding out a particular
  interface, the advertisement's
  LS age is incremented by
  InfTransDelay.
....
RFC 2328
```

On the Wire Efficiency

- When a router originates an LSA, it sets a timer as well; when the timer reaches 30 minutes, it re-originates the LSA
- Because timer counts up, all the routers in a network must be reconfigured to slow these retransmissions
- OSPF flooding reduction removes these periodic retransmissions

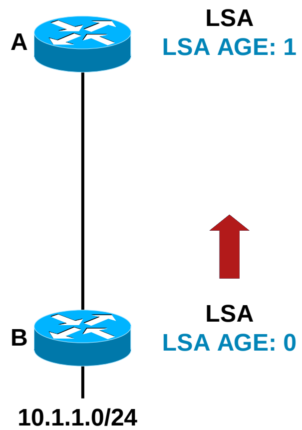

OSPF On the Wire Efficiency

- OSPF Flooding Reduction essentially uses the “do not age” bit set in all LSAs within an area to keep the LSA age timer from aging
- http://www.cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a008008011e.html

OSPF On the Wire Efficiency

- OSPF paces packets to one every 33 milliseconds by default
  - This can be configured to lower or higher numbers manually (CSCds86112)
  - This is not primarily aimed at link utilization, but at buffer utilization on the inbound side of the router
- Pacing can be useful when dealing with lower speed links to reduce the amount of traffic OSPF puts on the wire
- Pacing will slow down network convergence

EIGRP On the Wire Efficiency

- EIGRP transmits only reachable routes to neighbors
- As long as the neighbor is up, any routes received from that neighbor are assumed valid and operational

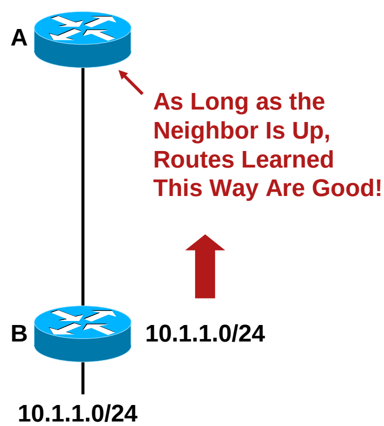

EIGRP On the Wire Efficiency

- EIGRP paces packets based on the bandwidth on links below T1 speed
- (1/bandwidth) x (packet size in bits) x bandwidth-percentage
- For a 1,500 byte packet at:

  56k Link: (1/56000) * 1500 * 8 * .5 = One packet every 107msecs

  64k Link: (1/64000) * 1500 * 8 * .5 = One packet every 94msecs

  128k Link: (1/128000) * 1500 *8 * .5 = One packet every 47msecs

  256k Link: (1/256000) * 1500 *8 * .5 = One packet every 23msecs

  512k Link: (1/512000) * 1500 *8 * .5 = One packet every 12msecs

On the Wire Efficiency

- Wire efficiency is the least of our problems …
- Few modern protocols can, or will, exhaust network resources except in unusual situations

  All routing protocols operate equally well over higher speed links, 64k and above

  A large number of lower speed links may push you towards optimizations reducing the protocol's burden, or to a protocol that adjusts to fewer resources more readily

MPLS and Traffic Engineering

- MPLS/TE:
  - Complex, not common or practical in most enterprise networks
  - OSPF and IS-IS support
  - EIGRP has no immediate plans to support
- MTR:
  - New technology
  - Will probably be more useful and commonly deployed in enterprise networks
  - Supported by OSPF, IS-IS, EIGRP, and BGP

Unequal Cost Load Sharing

- Can we load share over these two paths?
  - The path B_C_E costs 10
  - The path B_D_E costs 20
- EIGRP: As long as the B_D_E link is proven to be loop free, we can load share in proportion to the link costs
- OSPF: No
- IS-IS: No

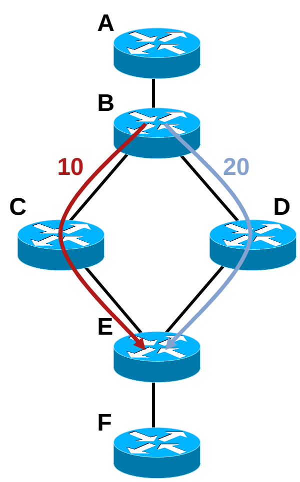

Metrics

OSPF	▪ Metric Based on the Bandwidth of the Interface: 10^8/Bandwidth ▪ You Can Change the Number in the Numerator Using ospf auto-cost (CSCdi73355)
IS-IS	▪ Metric Set to 10 on All Interfaces by Default (Default Metric Can Be Set in the Global Configuration Mode) ▪ Narrow Metrics from 1 to 63 for Any Link; Wide Metrics 1 to 2^{32}
EIGRP	▪ $\left[\left(\frac{10^7}{\text{min bandwidth}}\right)+\sum\left(\text{delays}\right)\right] \times 256$ ▪ Support for Gigabit and Higher Metrics Is Planned (CSCdx36932)

 77

IPv6 Support

- OSPF

  Two separate protocols, OSPFv2 and OSPFv3, running concurrently

- IS-IS supports IPv6
- EIGRP supports IPv6

 78

PE-CE Support

- Many enterprises now use MPLS VPNs through service providers to replace Frame Relay or their core links
- In either of these environments, it's important to maintain IGP routing information across the VPN
- Backdoor links complicate this situation

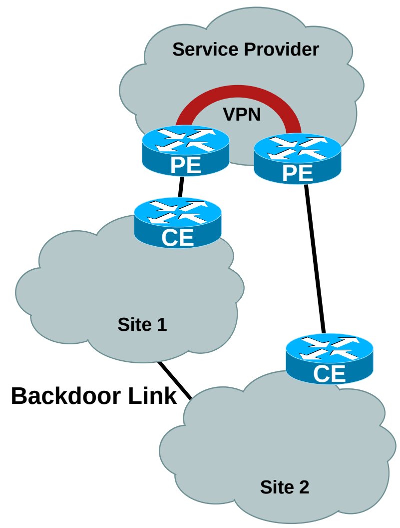

MPLS PE-CE Support

- OSPF

  Supports MPLS PE/CE through sham links in cooperation with the provider's BGP

  Supports back door links through the down bit

  http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087ce2.html

- IS-IS

  No support

MPLS PE-CE Support

EIGRP

- Supports MPLS PE-CE in cooperation with the provider's BGP
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080154db3.html
- Supports back door links through Sight Of Origin (SOO) communities
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801eff60.html
- Supports BGP metric adjustment through cost communities
  http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801eff5f.html

Protocol Management

	Debugs	Event Log	Neighbor Logging	SNMP
OSPF	Neighbor and Protocol Events	Yes, but Not Easy to Read	Yes	RFC 1253
IS-IS	Neighbor and Protocol Events	No	No	RFC 4444
EIGRP	Neighbor and Protocol Events	Yes, Moderately Difficult to Read	Yes	Yes

OSPF Future Developments

- Fast reroute
  - Current work in the IETF
  - draft-ietf-rtgwg-ipfrr-framework and others
- Multitopology routing
  - draft-ietf-ospf-mt (IPv4)
  - draft-ietf-ospf-mt-ospfv3 (IPv6)
- Address families
  - To support IPv4 and IPv6 in OSPFv3
  - draft-ietf-ospfv3-af-alt and others

EIGRP Future Developments

- Multi-topology routing
- Route servers
- Conditional advertisement
- Full community support (improved policy)
- Improved debugging and event logging

Summary

" Is one protocol better than the others?

Which routing protocol should I use in my network?

Should I change from the one I'm using? "

Did We Answer THIS Question?

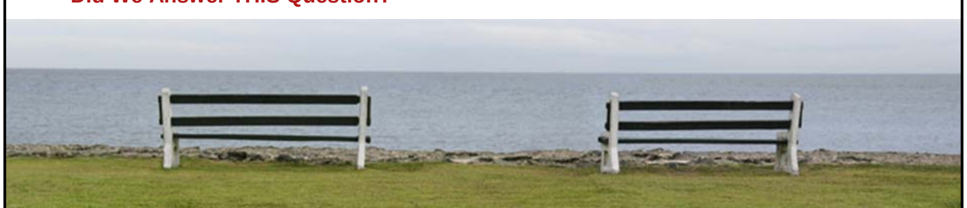

Summary

- There is no “right” answer
- Consider:
  - Your business requirements
  - Your network design
  - Intangibles
- The three advanced IGP’s are generally pretty close in capabilities, development, and other factors

Expertise (Intangible)

- What is your team comfortable with?
- What “escalation resources” and other support avenues are available?
- But remember, this isn’t a popularity contest - you don’t buy your car based on the number of a given models sold, do you?
- An alternate way to look at it: what protocol would you like to learn? ☺

Standardization (Intangible)

- Who's standard?
  - OSPF: standardized by the IETF
  - IS-IS: standardized by the ISO and the IETF
  - EIGRP: "Cisco Standard" ☺
- Standardization is a tradeoff:
  - Promises Interoperability
  - Larger number of eyes looking at problems and finding new features
  - Politics often influence standards
  - New features are often difficult to push through standards committees, slowing their release

Summary

Rules of Thumb

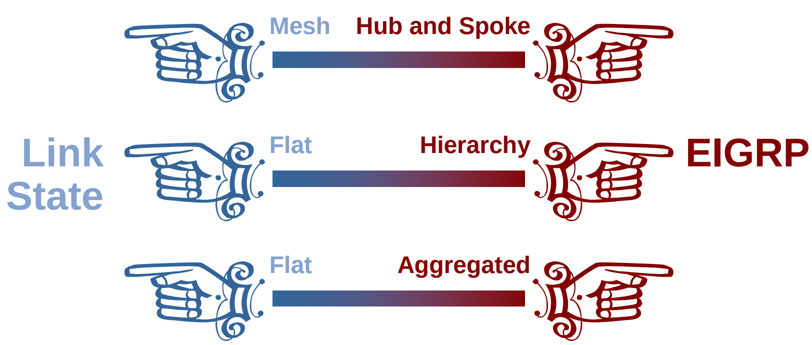

Other References

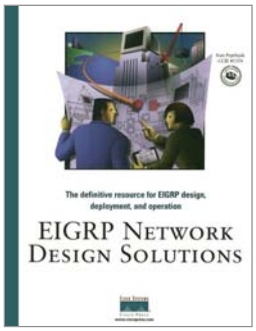

ISBN: 1578701651

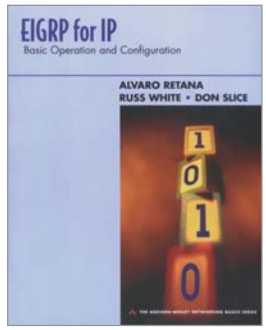

ISBN: 0201657732

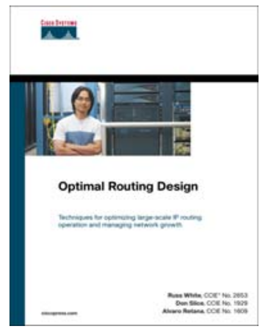

ISBN 1587051877

Other References

ISBN: 1587050323

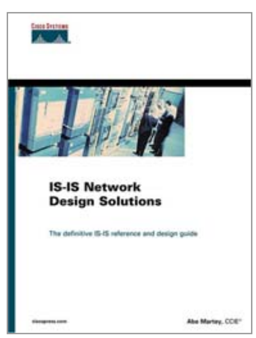

ISBN: 1578702208

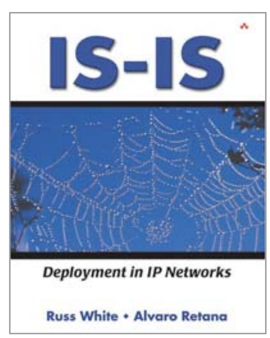

ISBN: 0201657724

Q and A