Troubleshooting and Maintaining Routing Based Solutions

Chapter 5

CertificationKits

Routing Protocols

- Each routing protocol has its own data structures used to keep track of routing
- Subsequent slides will look more closely at table / data structures used by the following routing protocols:
  - EIGRP
  - OSPF
  - BGP

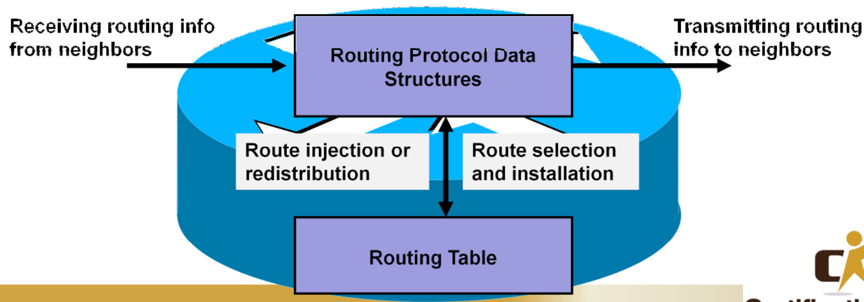

Each routing protocol (e.g. EIGRP, OSPF, BGP) has is own tables used to keep track of routing information.

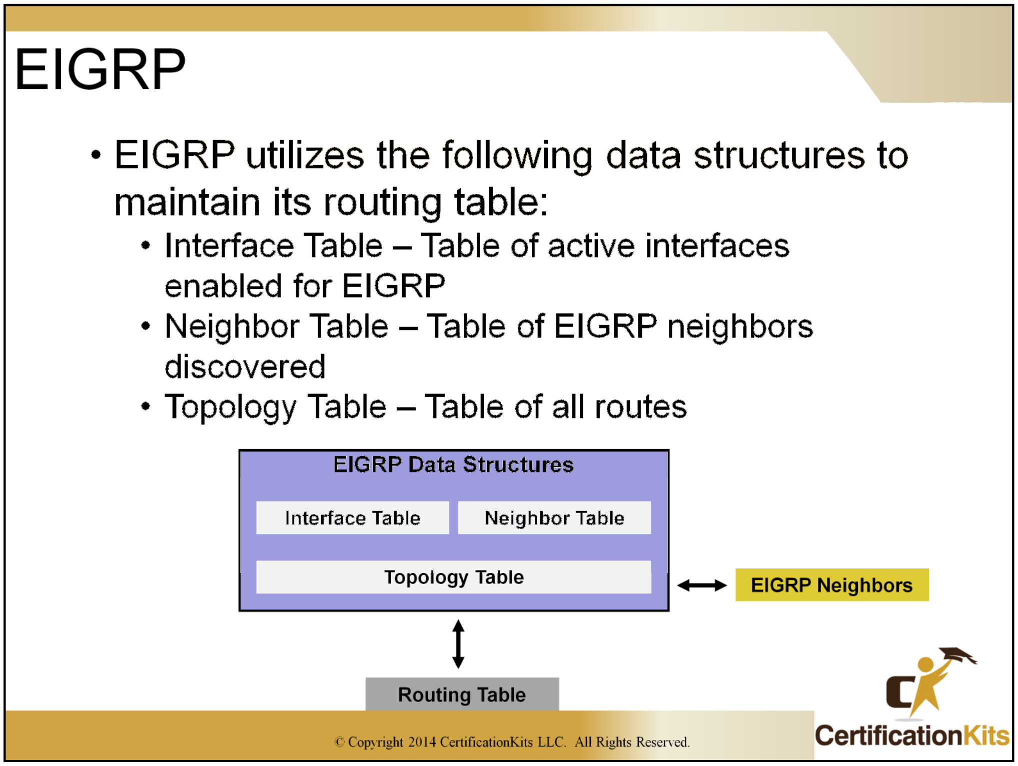

Interface table – Compiles all interfaces that have bee configured to participate in EIGRP.

Neighbor table – Used to keep track of all EIGRP neighbors. Neighbors are added to the table once EIGRP hello packets have been received.

Topology table – Holds all routes, regardless of how they were learned (e.g. via EIGRP, locally injected, redistributed, etc).

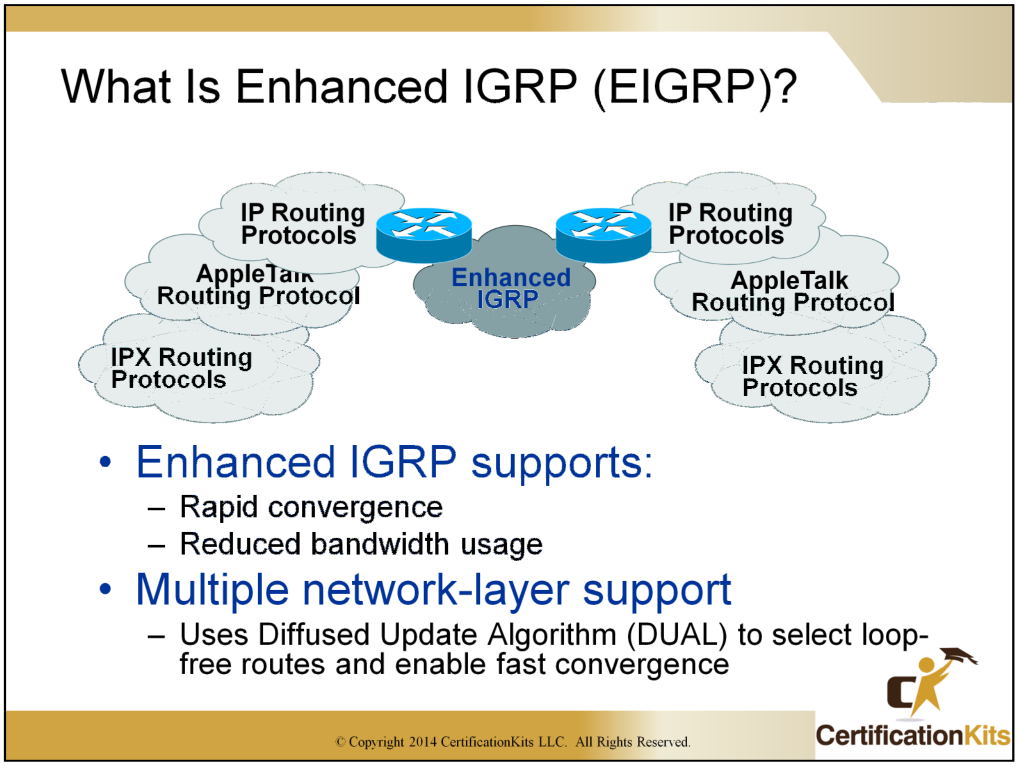

Enhanced Interior Gateway Routing Protocol (EIGRP) is a proprietary Cisco protocol that runs on Cisco routers and internal route processors found in the Cisco Distribution and Core layer 3 switches.

In this section, you'll see the many features of EIGRP and describe how it works, with particular focus on the unique way it discovers, selects, and advertises routes.

EIGRP utilizes Diffused Update Algorithm (DUAL) which is fast to converge and which guarantees loop-free routes.

There are a number of powerful features that make EIGRP a real stand out from IGRP and other protocols. The main ones are listed here:

✓Support for IP, IPX, and AppleTalk via protocol-dependent modules

✓Efficient neighbor discovery

✓Communication via Reliable Transport Protocol (RTP)

✓Best path selection via Diffusing update algorithm (DUAL)

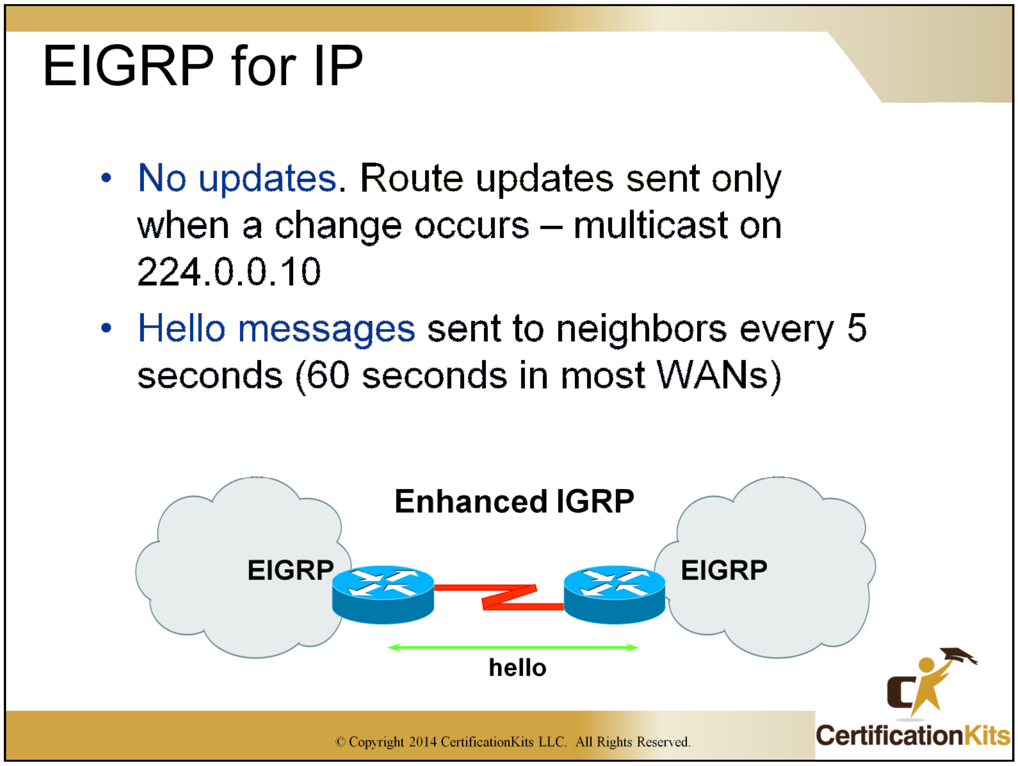

EIGRP doesn't send link-state packets as OSPF does; instead, it sends traditional distance-vector updates containing information about networks plus the cost of reaching them from the perspective of the advertising router.

EIGRP has link-state characteristics as well—it synchronizes routing tables between neighbors at startup, and then sends specific updates only when topology changes occur to multicast address 224.0.0.10. Hello messages are sent to neighbors every 5 seconds in a LAN environment and every 60 seconds in a WAN environment.

EIGRP can be configured to route for IP, IPX, and AppleTalk. This course only focuses on IP.

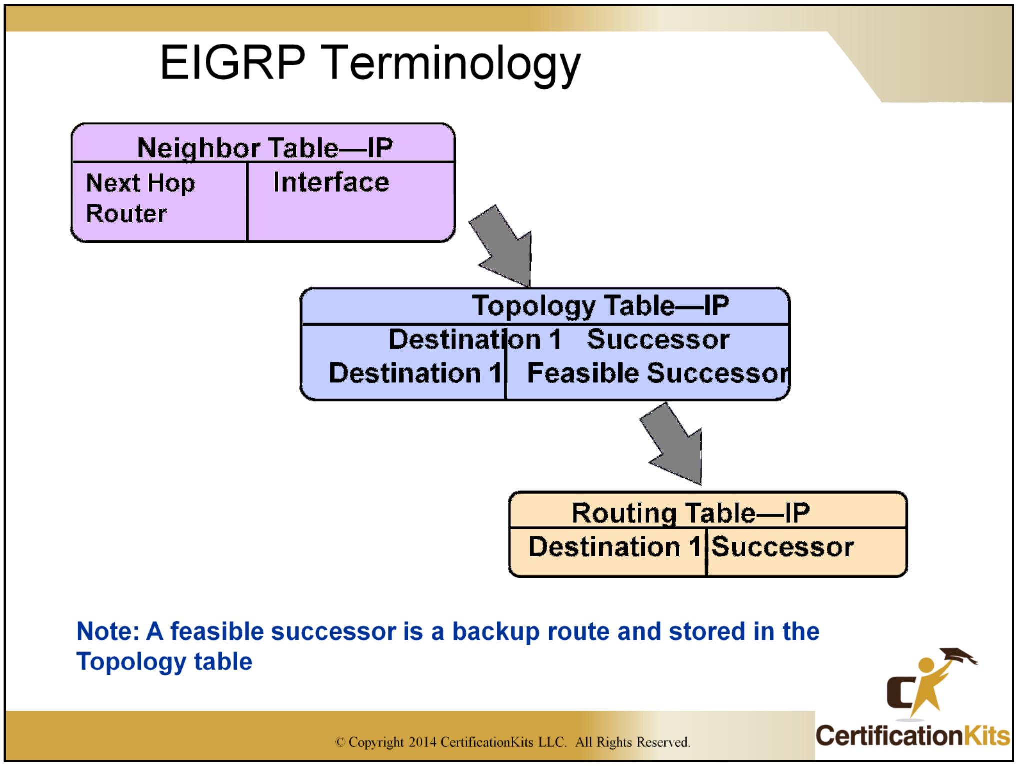

The *neighborship table* (usually referred to as the neighbor table) records information about routers with whom neighborship relationships have been formed.

The *topology table* stores the route advertisements about every route in the internetwork received from each neighbor.

The *route table* stores the routes that are currently used to make routing decision. There would be separate copies of each of these tables for each protocol that is actively being supported by EIGRP, whether it's IP, IPX, or AppleTalk.

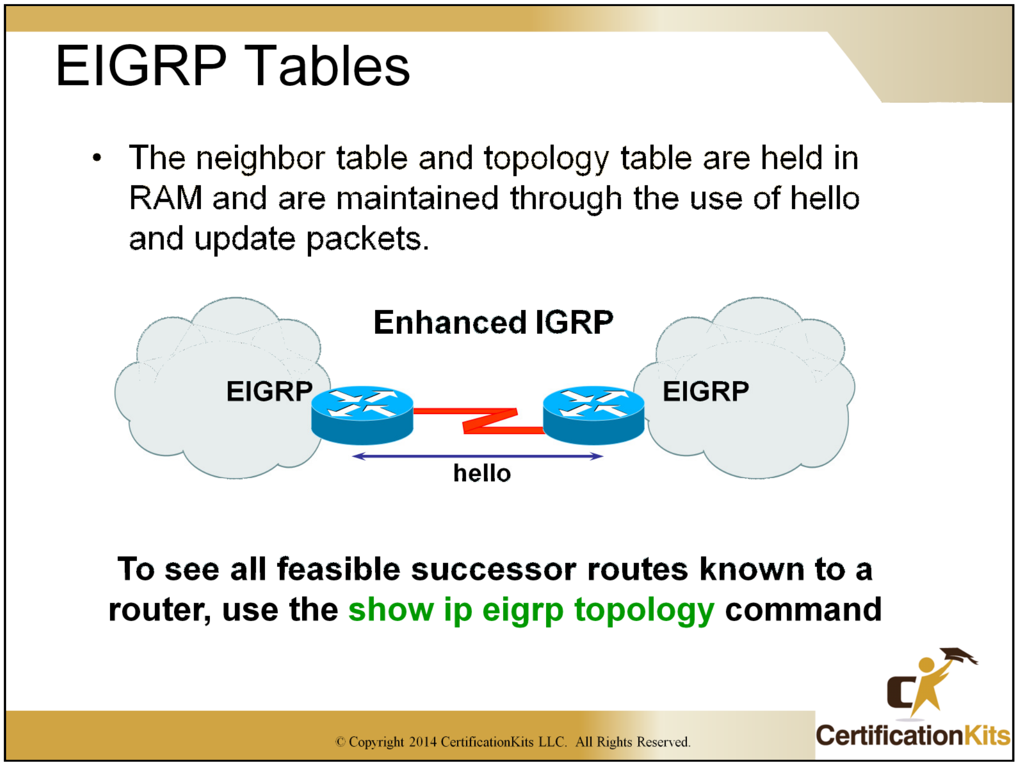

The neighbor table and topology table are held in RAM and are maintained through the use of hello and update packets.

Command syntax:

show ip eigrp topology [*as-number* | [[*ip-address*] *mask*]] [**active** | **all-links** | **pending** | **summary** | **zero-successors**]

as-number - (Optional) Autonomous system number.

ip-address - (Optional) IP address. When specified with a mask, a detailed description of the entry is provided.

mask - (Optional) Subnet mask.

active - (Optional) Displays only active entries in the EIGRP topology table.

all-links - (Optional) Displays all entries in the EIGRP topology table.

pending - (Optional) Displays all entries in the EIGRP topology table that are waiting for an update from a neighbor or are waiting to reply to a neighbor.

summary - (Optional) Displays a summary of the EIGRP topology table.

zero-successors - (Optional) Displays available routes in the EIGRP topology table.

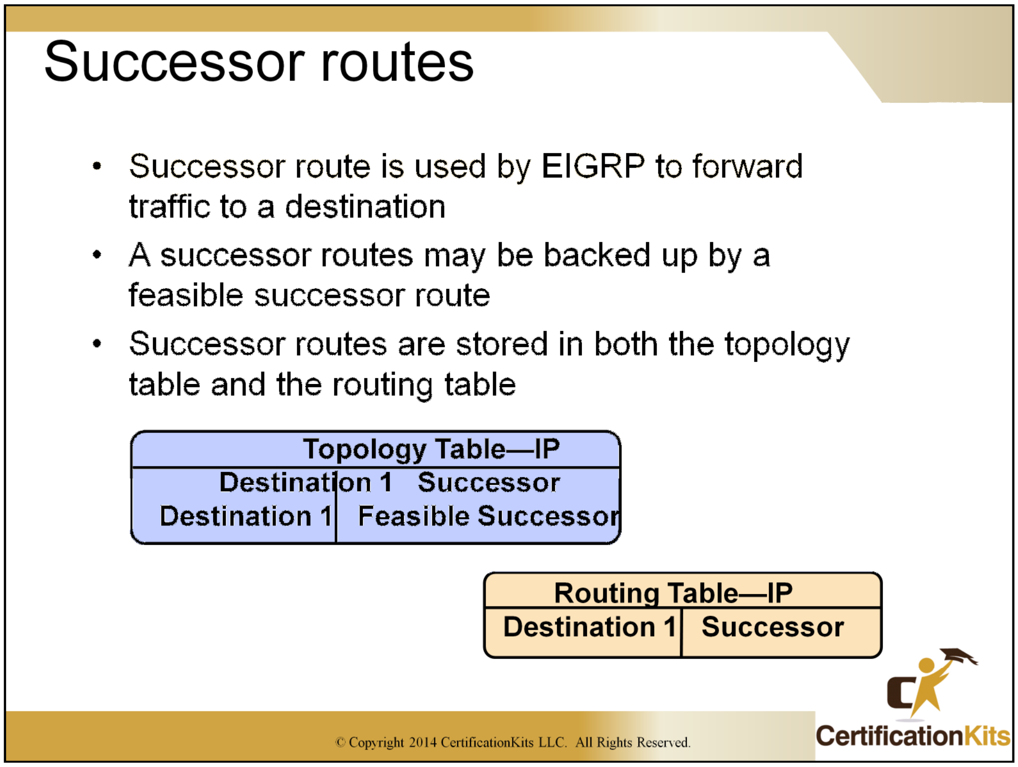

Successor route is used by EIGRP to forward traffic to a destination.

A successor routes may be backed up by a feasible successor route.

Successor routes are stored in both the topology table and the routing table.

EIGRP relies on four fundamental concepts: neighbor tables, topology tables, route states, and route tagging.

When a router discovers a new neighbor, it records the neighbor's address and interface as an entry in the *neighbor table*. One neighbor table exists for each protocol-dependent module. When a neighbor sends a hello packet, it advertises a hold time, which is the amount of time that a router treats a neighbor as reachable and operational. If a hello packet is not received within the hold time, the hold time expires and DUAL is informed of the topology change.

The *topology table* contains all destinations advertised by neighboring routers. The protocol-dependent modules populate the table, and the table is acted on by the DUAL finite-state machine. A topology-table entry for a destination can exist in one of two states: active or passive. A destination is in the *passive state* when the router is not performing a recomputation; it is in the *active state* when the router is performing a recomputation. If feasible successors are always available, a destination never has to go into the active state, thereby avoiding a recomputation.

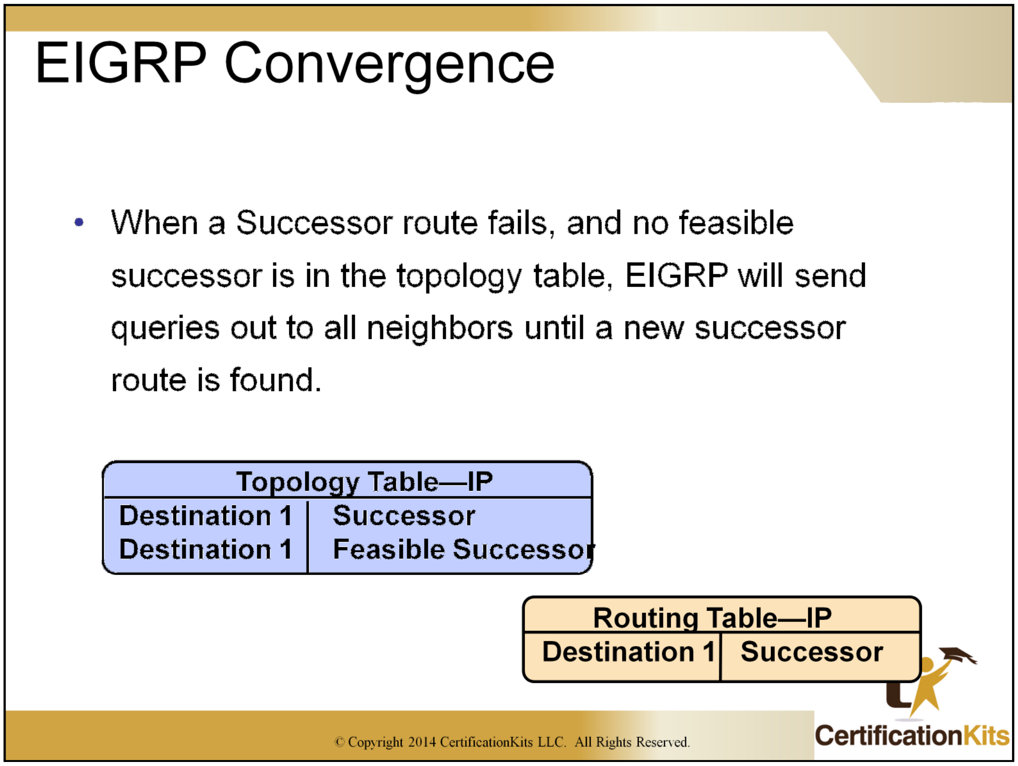

A recomputation occurs when a destination has no feasible successors. The router initiates the recomputation by sending a query packet to each of its neighboring routers. The neighboring router can send a reply packet, indicating that it has a feasible successor for the destination, or it can send a query packet, indicating that it is participating in the recomputation. While a destination is in the active state, a router cannot change the destination's routing-table information. After the router has received a reply from each neighboring router, the topology-table entry for the destination returns to the passive state, and the router can select a successor.

EIGRP supports internal and external routes. Internal routes originate within an EIGRP Autonomous System (AS). Therefore, a directly attached network that is configured to run EIGRP is considered an internal route and is propagated with this information throughout the EIGRP AS. External routes are learned by another routing protocol or reside in the routing table as static routes. These routes are tagged individually with the identity of their origin.

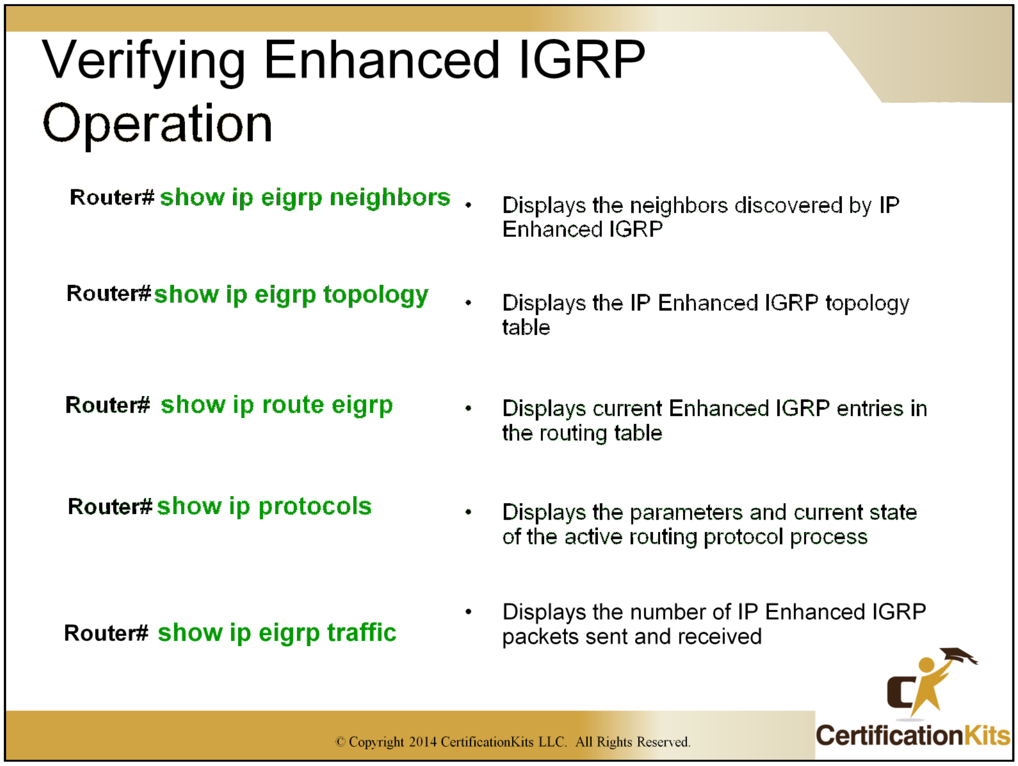

Below are several useful commands to view information about the operation of EIGRP and how it is configured.

show ip eigrp neighbors: Displays all EIGRP neighbors.

show ip eigrp topology: Displays entries in the EIGRP topology table.

show ip route eigrp: Displays only EIGRP entries in the routing table.

show ip protocols: Displays parameters and current state of active routing protocols

show ip eigrp traffic: Displays statistics pertaining to EIGRP packets sent and received.

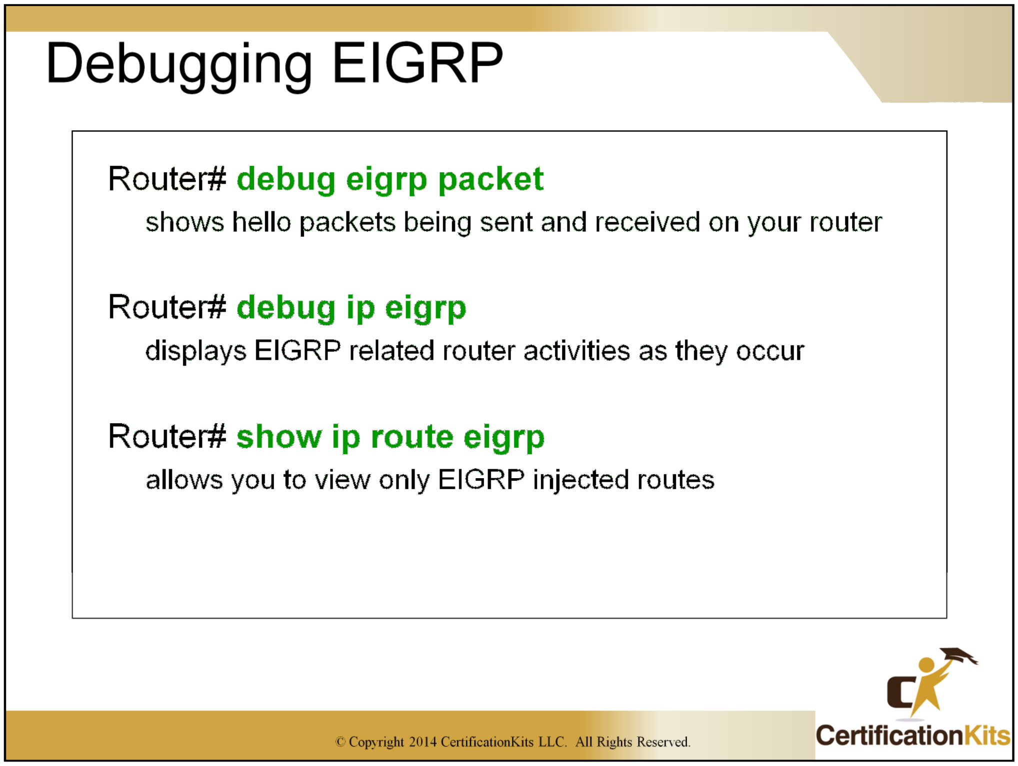

The use of debugging commands should always be used with caution in a production network as debugging is CPU intensive. With that being said the "**debug eigrp packet**" command will display hello packets being sent and received by the router and the "**debug ip eigrp**" command displays EIGRP related activity on the router.

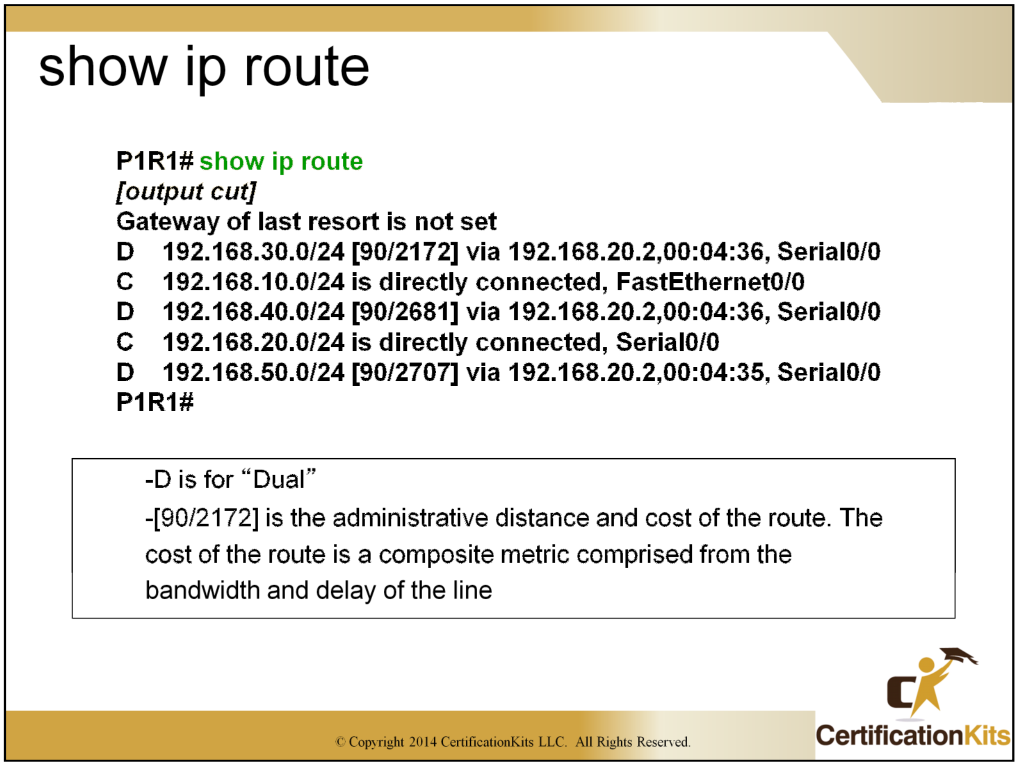

The "**show ip route**" command, or "**show ip route eigrp**" command, displays the routing table. The use of the "**eigrp**" option limits the routes being display to only EIGRP routes which are denoted with the "D" in front of the route which is for "Dual", the algoritm EIGRP uses to provide loop free routing.

-[90/2172] is the Administrative Distance (90) and cost (2172) of the route. By default, the cost of the route is a composite metric comprised from the bandwidth and delay of the line.

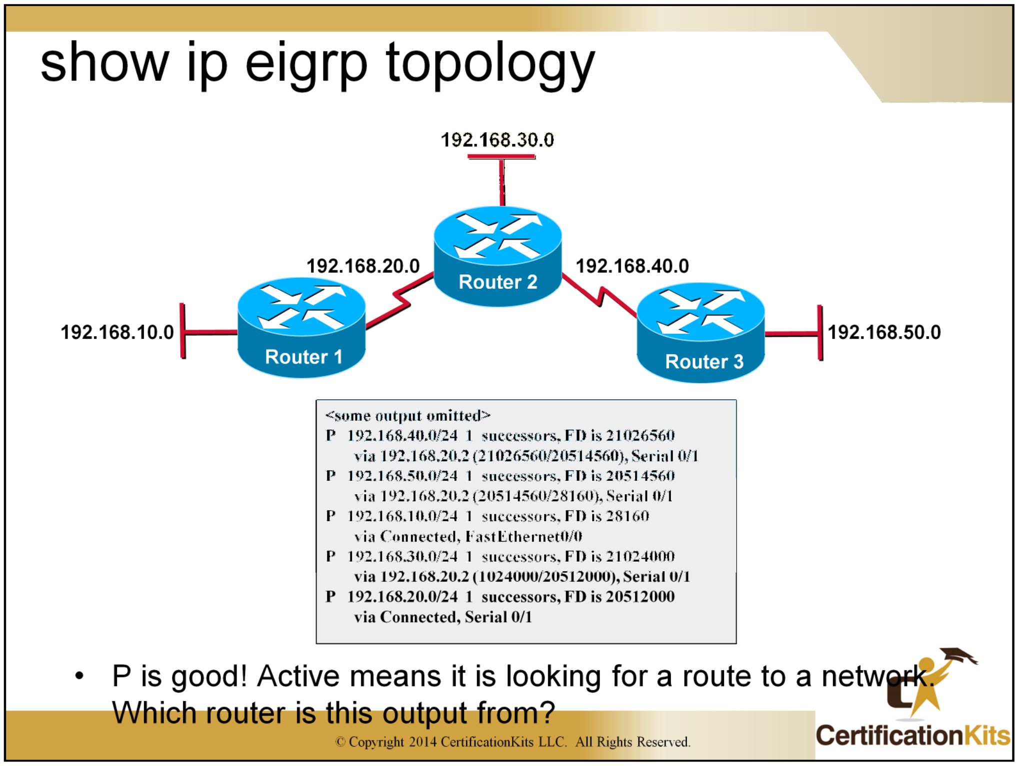

Remember this command, which shows the topology table.

Command syntax:

show ip eigrp topology [*as-number* | [[*ip-address*] *mask*]] [**active** | **all-links** | **pending** | **summary** | **zero-successors**]

as-number - (Optional) Autonomous system number.

ip-address - (Optional) IP address. When specified with a mask, a detailed description of the entry is provided.

mask - (Optional) Subnet mask.

active - (Optional) Displays only active entries in the EIGRP topology table.

all-links - (Optional) Displays all entries in the EIGRP topology table.

pending - (Optional) Displays all entries in the EIGRP topology table that are waiting for an update from a neighbor or are waiting to reply to a neighbor.

summary - (Optional) Displays a summary of the EIGRP topology table.

zero-successors - (Optional) Displays available routes in the EIGRP topology table.

show ip protocols

```
PODxBB# show ip protocols
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
      192.168.255.16/30
      192.168.255.20/30
  Routing Information Sources:
      Gateway          Distance       Last Update
      192.168.255.22      90            00:28:16
      192.168.255.18      90            00:28:16
  Distance: internal 90 external 170
```

How many equal costs routes to the same destination will be allowed in the routing table?

Answer:

A maximum of 4 equal cost paths indicated by the following line in the output:

Maximum path: 4 (default for EIGRP)

Note the "K" values where K1=1 and K3=1 and all other "K" values are 0. This indicates that only bandwidth and delay are taken into account when calculating the cost.

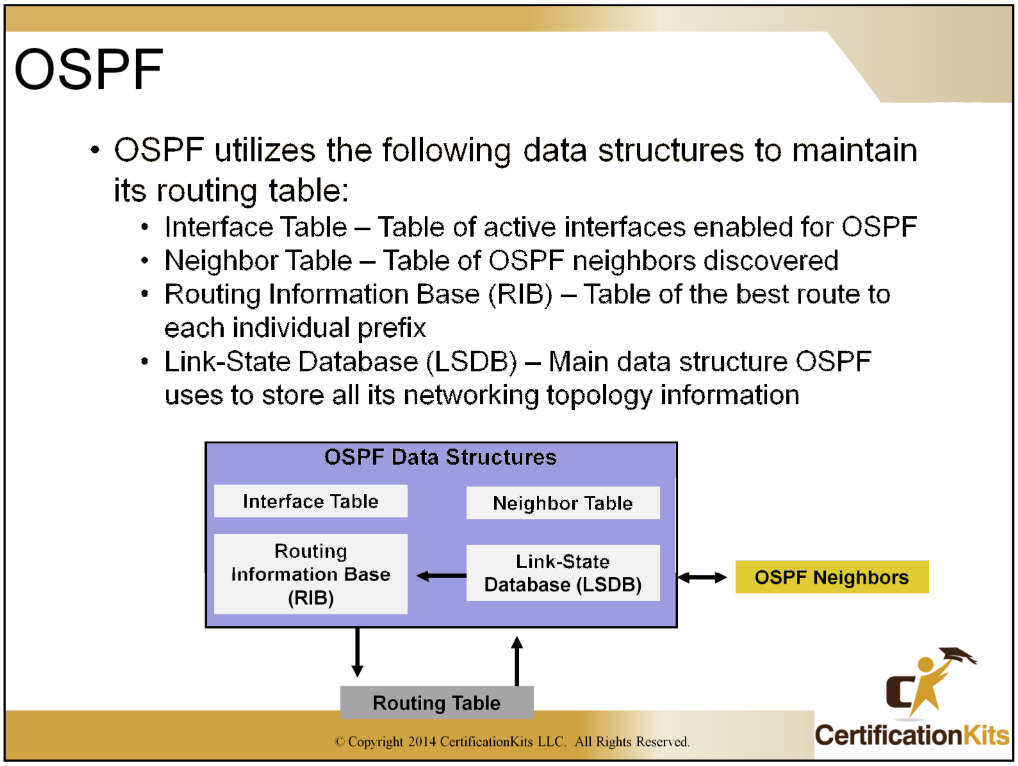

Interface table – Compiles all interfaces that have bee configured to participate in OSPF. The directly connected subnets associated with these interfaces are included in the type 1 router LSA that the router injects into the OSPF LSDB. An interface configured as a passive interface is still listed in the OSPF interface table although no neighbor relationships are established on this interface.

Neighbor table – Used to keep track of all OSPF neighbors. Neighbors are added to the table upon receipt of hello packets whereas they are removed when the OSPF dead timer for a neighbor expires or when the associated interface goes down.

Link-State Database (LSDB) – The main data structure OSPF uses to store all its network topology information. The LSDB contains topology information for the areas that a router participates in. It also contains condensed information about the paths that are available to reach networks and subnets in other areas or other autonomous systems. It is one of the most important data structures to gather information from when troubleshooting OSPF.

Routing Information Base (RIB) – The RIB includes information for the best routes to each individual prefix in the OSPF network with their associated path costs. This information is derived after executing the SPF algorithm. When the information in the LSDB changes, only a partial recalculation might be necessary. Routes might be added or deleted from the RIB without the need for a full SPF recalculation. Based on the RIB, OSPF will attempt to install its routes into the routing table.

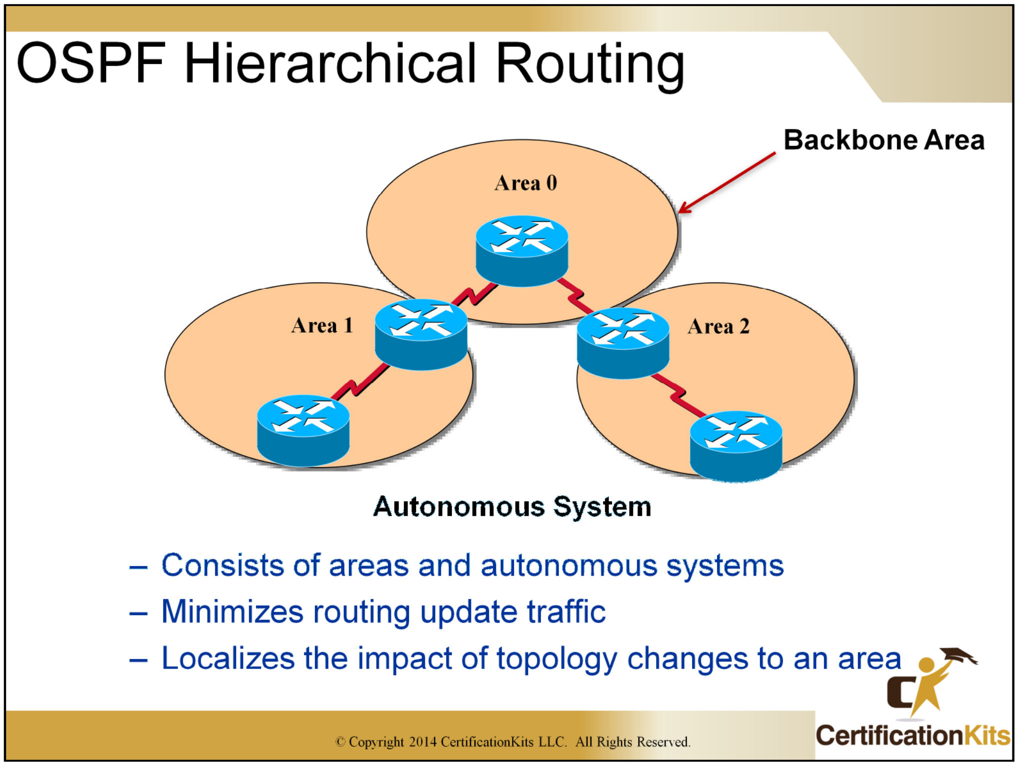

OSPF is designed in a hierarchical fashion, which basically means that you can separate the larger internetwork into smaller internetworks called areas. All OSPF networks MUST contain Area 0 and all areas MUST directly connect to Area 0. This is depicted in the slide where both Area 1 and Area 2 are directly connected to the Backbone Area, Area 0. There is one exception to this rule where a virtual link can be established but it should only be used for a temporary solution (i.e. merging two different company networks). Details of virtual links are not covered in this course.

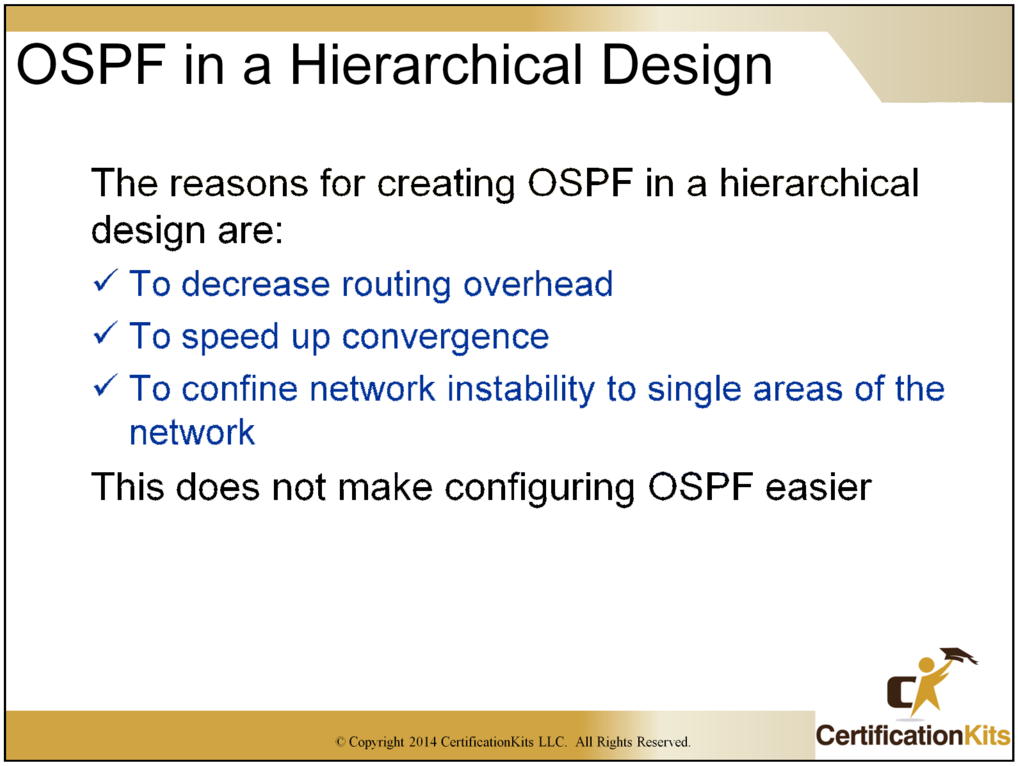

OSPF has special restrictions when multiple areas are involved. If more than one area is configured, one of these areas has be to be area 0. This is called the backbone. When designing networks it is good practice to start with area 0 and then expand into other areas later on.

The backbone has to be at the center of all other areas, i.e. all areas have to be physically connected to the backbone. The reasoning behind this is that OSPF expects all areas to inject routing information into the backbone and in turn the backbone will disseminate that information into other areas.

Utilizing multiple areas is typically a good design for OSPF. Creating OSPF in a hierarchical design has numerous benefits that include:

- Decreasing routing overhead
- Speeding up convergence
- Confines network instability to single areas of the network

Note: While it makes OSPF more efficient, it does not make configuring OSPF easier.

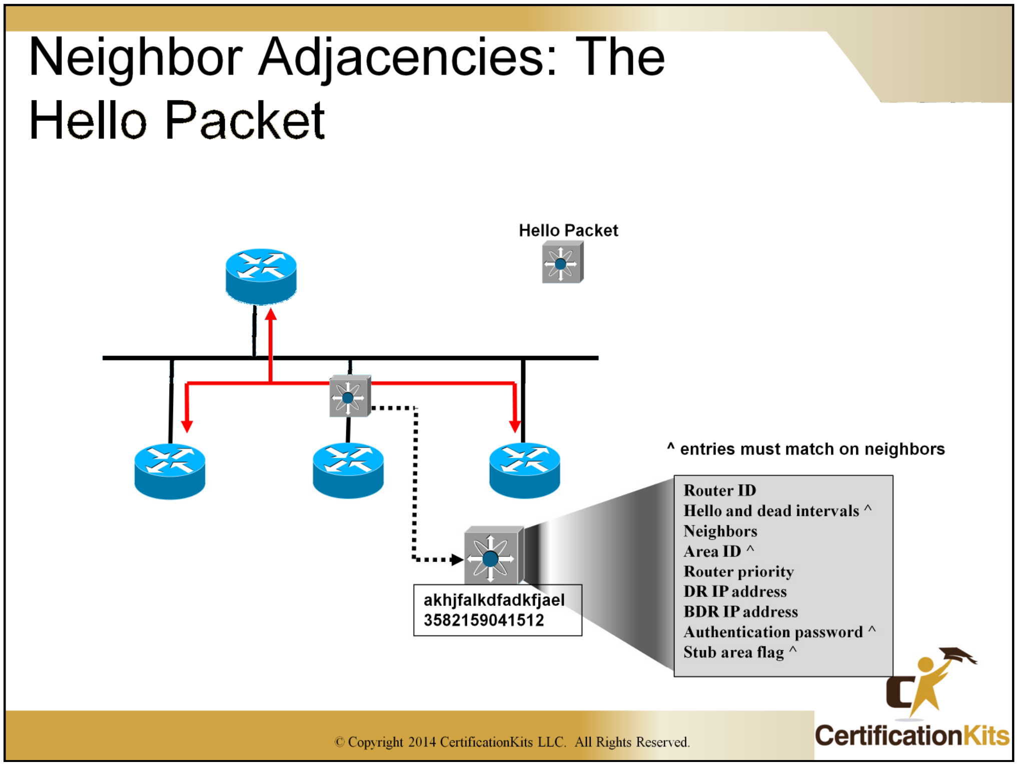

Hello packets are utilized by OSPF routers to recognize each other in order to share information. Each interface participating in OSPF sends hello packets multicast periodically to 224.0.0.5. The entries in the figure with an "^" much match between adjacent routers or a neighbor relationship will NOT be established. These include Hello and dead intervals, Area ID, Authentication password (if used) and Stub area flag.

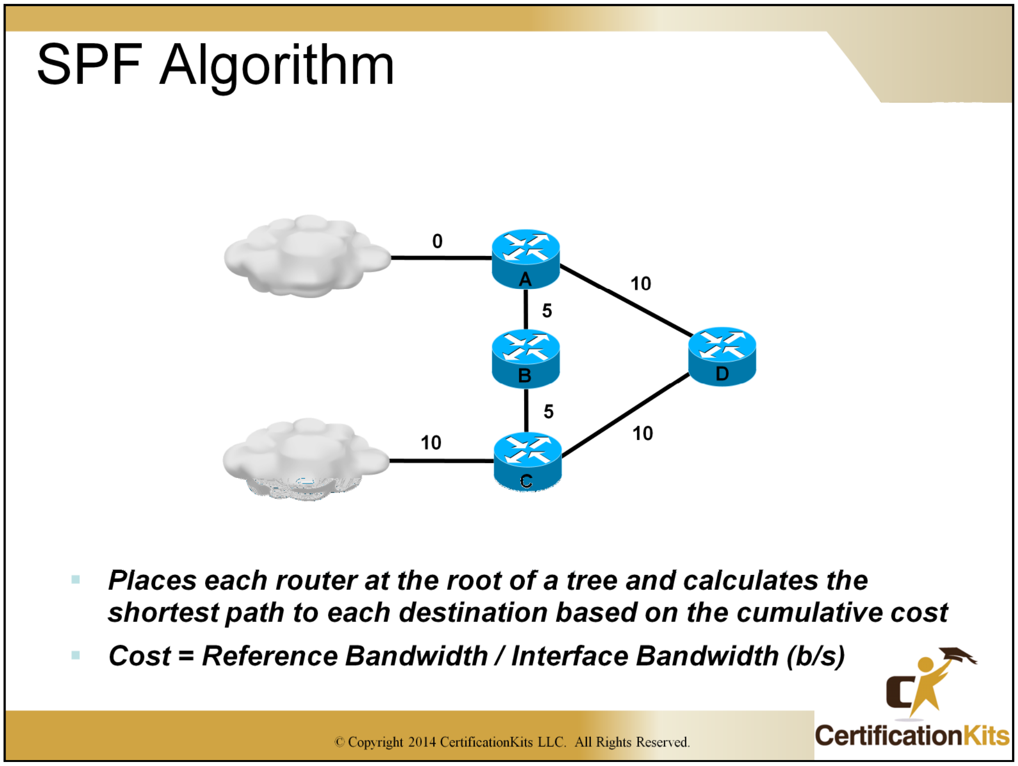

The SPF algorithm places each router at the root of the tree and calculates the shortest path to each node utilizing the Dijkstra's algorithm. It is based on the cumulative cost required to reach each node. The cost is based on bandwidth. The default reference bandwidth is 10 to the 8th, which is 100,000,000 or the equivalent of the fast ethernet which means a fast ethernet has a default cost of 1. A 10Mb/s ethernet has a default cost of 10 and so on. Due to the increased speeds of networks today (i.e. 1Gbs and 10Gbs) the reference bandwidth can be changed so that a 1Gbs link would be more desirable than a 100Mbs link.

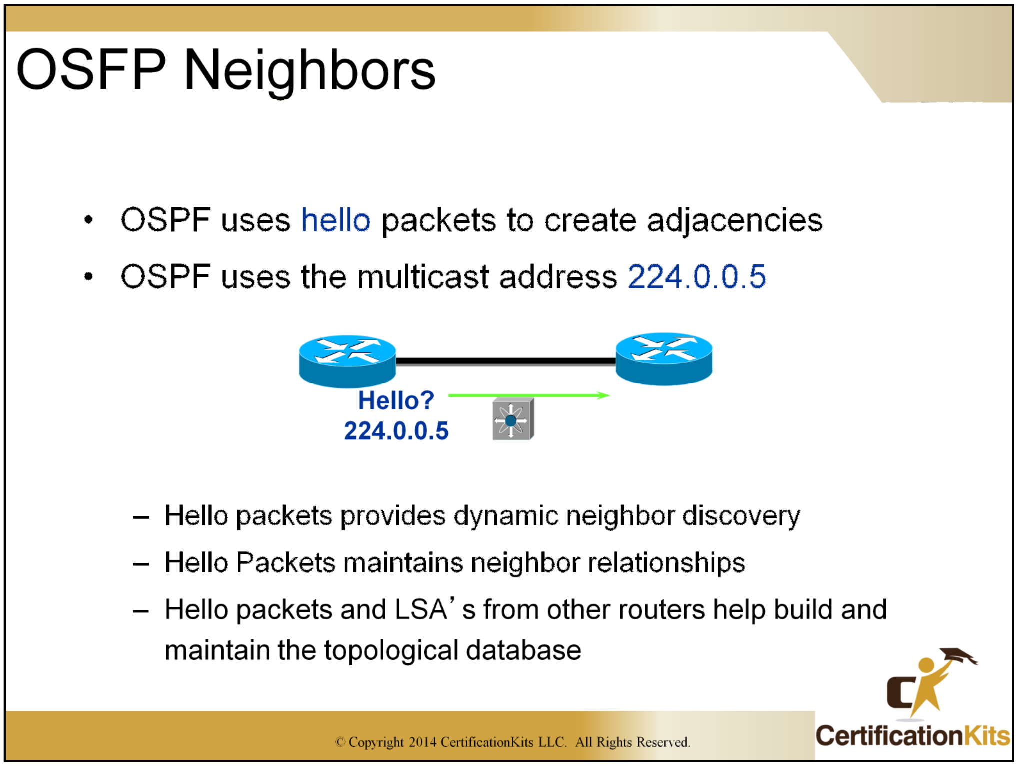

Neighbors

Neighbors are two or more routers that have an interface in a common network, such as two routers connected on a point-to-point serial link.

Adjacency

An *adjacency* is a relationship between two OSPF routers that permits the direct exchange of route updates. OSPF is really picky about sharing routing information, unlike EIGRP that directly shares routes with all of its neighbors.

Instead, OSPF directly shares routes only with neighbors that have also established adjacencies.

Link State Advertisement

A *Link State Advertisement (LSA)* is an OSPF data packet containing link-state and routing information that's shared among OSPF routers.

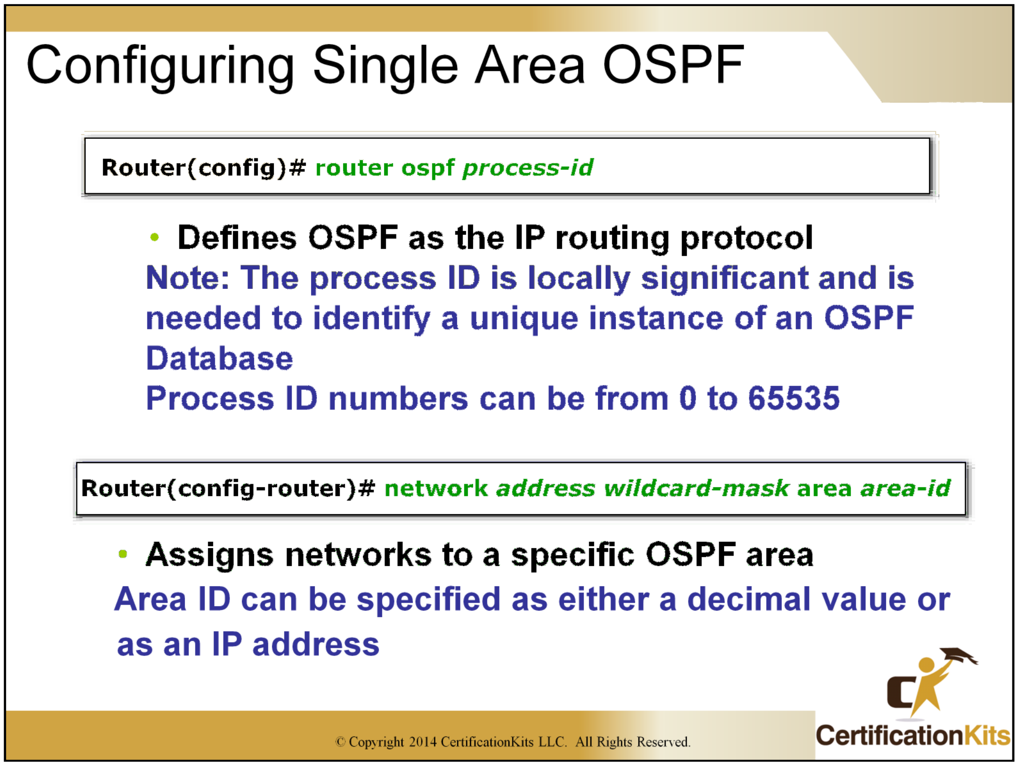

Configuring basic OSPF isn't as simple as RIP, IGRP and EIGRP, and it can get can really complex once the many options that are available within OSPF are factored in.

Two basic elements of OSPF configuration are:

-Enabling OSPF

This is accomplished with the "**router ospf process-id**" command.

-Configuring OSPF areas

This is accomplished with the "**network address wildcard-mask area area-id**" command.

The easiest, and also least scalable way to configure OSPF is to just use a single area. Doing this requires a minimum of two commands.

router ospf <process-id>

A value in the range 0 – 65535 identifies the OSPF Process ID.

network address mask area area-id

The area-id should be 0 if only a single area is configured.

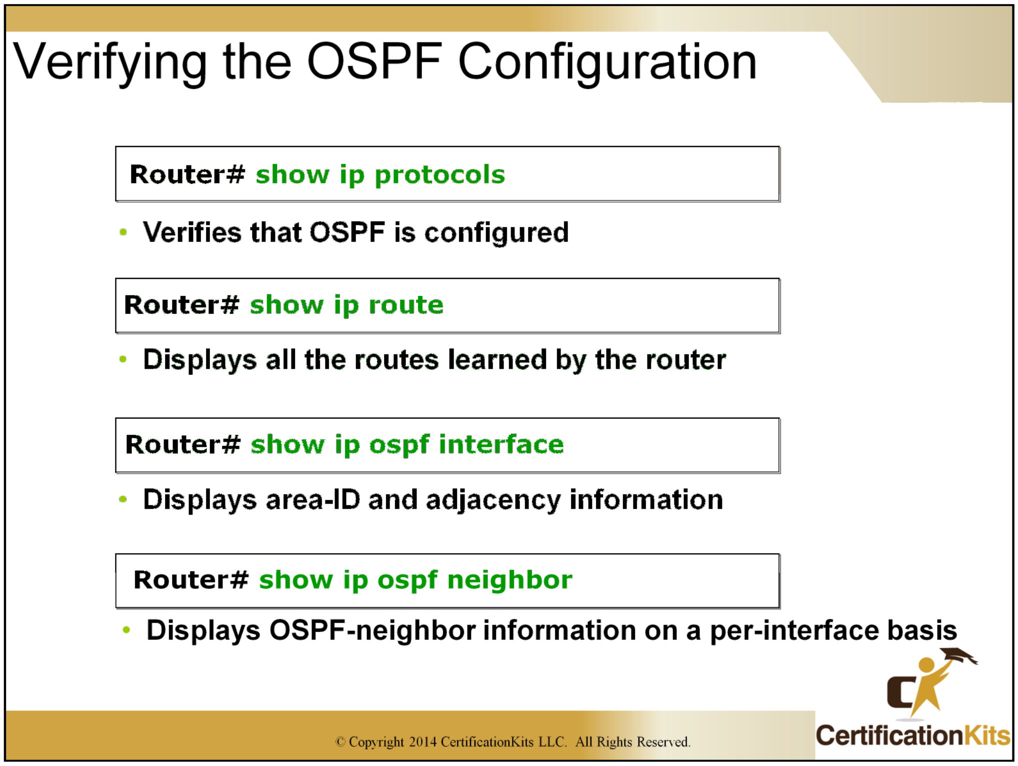

There are several ways to verify proper OSPF configuration and operation. Some of the commands are as follows:

show ip protocols: This command is useful in debugging routing operations. Information in the Routing Information Sources field of the **show ip protocols** output can help you identify a router suspected of delivering bad routing information.

show ip route or show ip route ospf: Displays the current state of the routing table

show ip ospf interface: Displays general information about Open Shortest Path First (OSPF) routing processes.

show ip ospf neighbor: Displays OSPF-neighbor information on a per-interface basis.

OSPF is great if you configure it properly. Now let's look at other ways to do that.

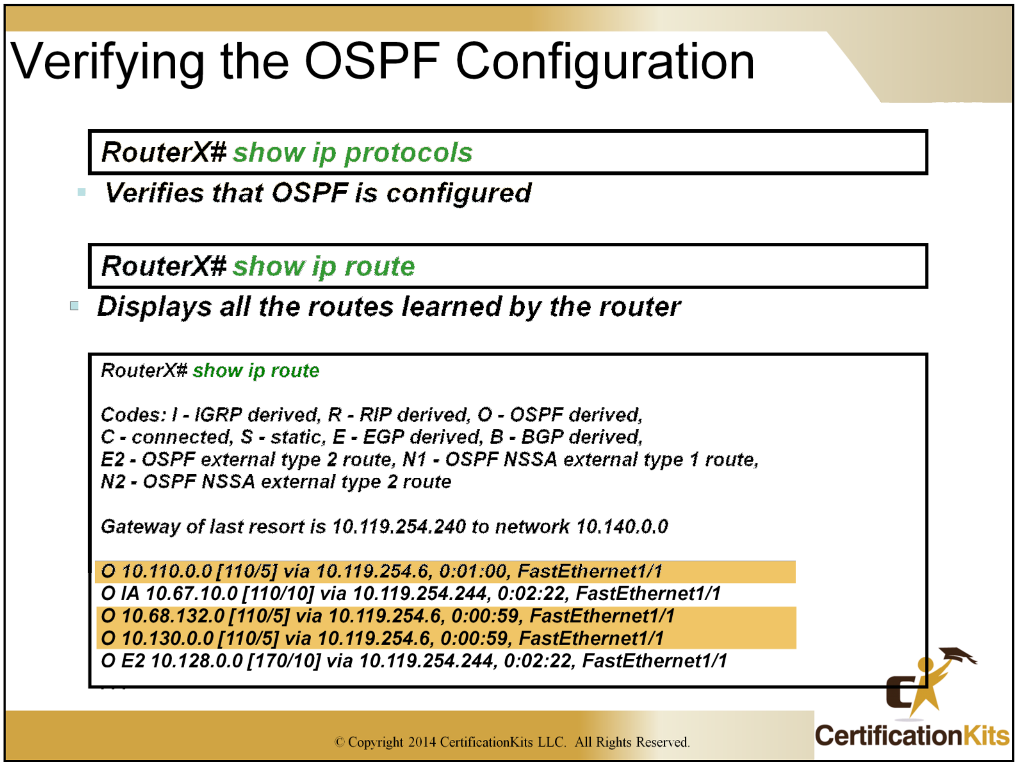

Notice from the output of the "**show ip route**" command, the default administrative distance (AD) of OSPF is 110. The AD for an external (E2) OSPF route is 170 as shown in the output of the **show ip route** command.

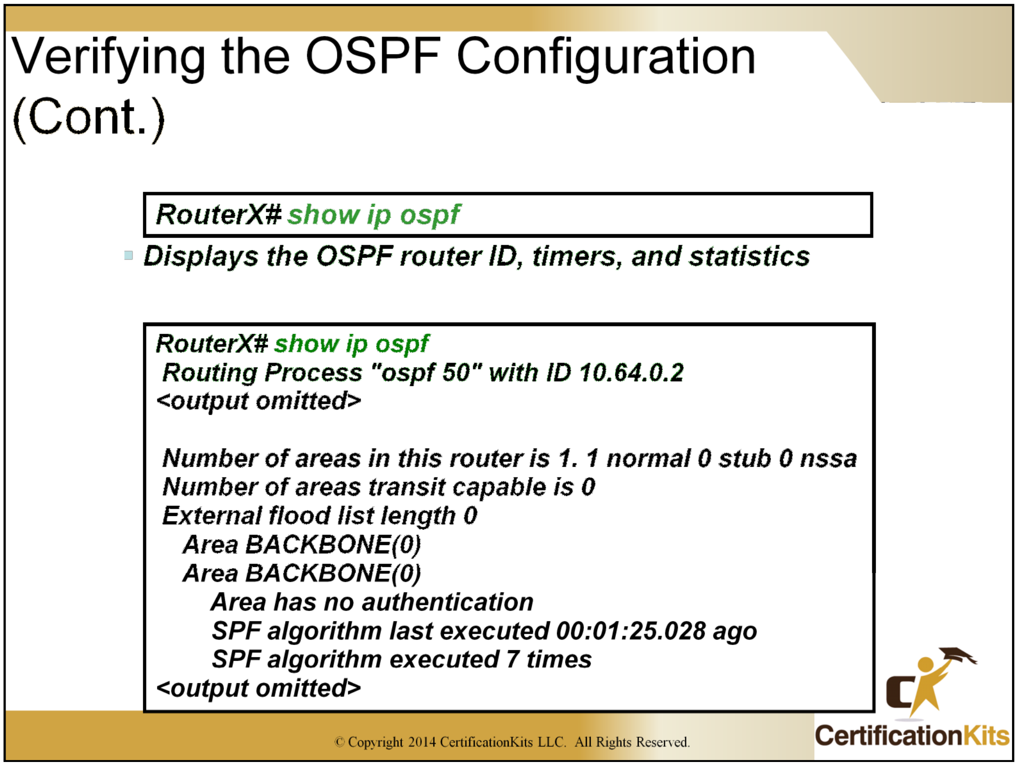

The **show ip ospf** command is great for troubleshooting OSPF as it has a significant amount of valuable information in the output. Things such as timer settings, Link State Advertisement (LSA) information, area information, etc.

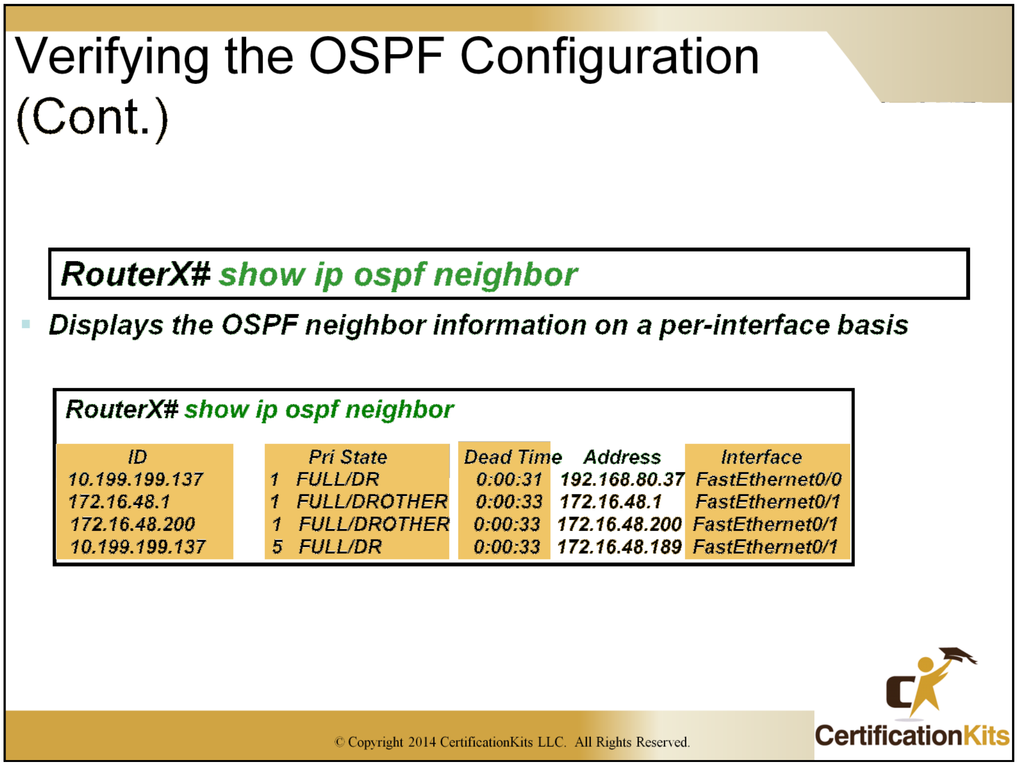

Another good command for troubleshooting OSPF is the **show ip ospf neighbor** command. As you can see from the output on the slide, fields that are displayed are as follows:

ID – Router ID of the OSPF neighbor

Priority – OSPF priority of the OSPF neighbor

State – Displays the state of the OSPF router. Is it a DR, BDR, etc.

Dead Time – Expected time before Cisco IOS software will declare the neighbor dead.

Address – Actual address of the interface the OSPF router is assigned.

Interface – Interface the neighbor is learned on.

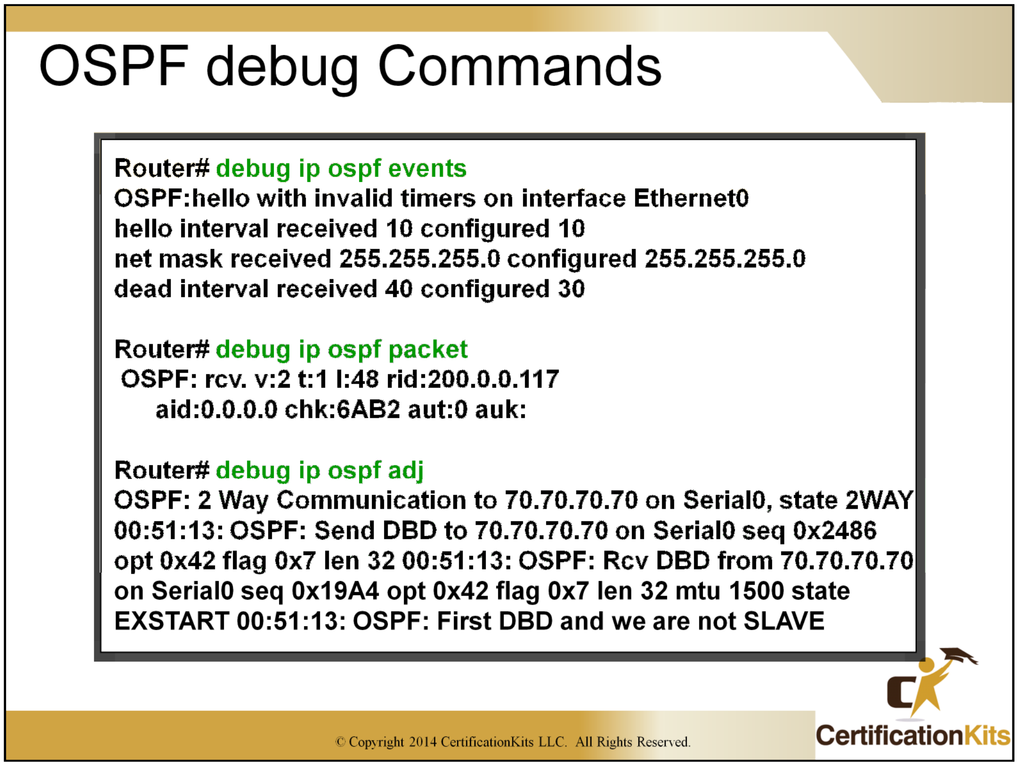

To display information on Open Shortest Path First (OSPF)-related events, such as adjacencies, flooding information, designated router selection, and shortest path first (SPF) calculation, use the **debug ip ospf events** command in privileged EXEC mode.

To display general IP debugging information and IP security option (IPSO) security transactions, use the **debug ip packet** command in privileged EXEC mode.

By issuing the **debug ip ospf adj** command, you can capture the authentication process and DR/BDR entries.

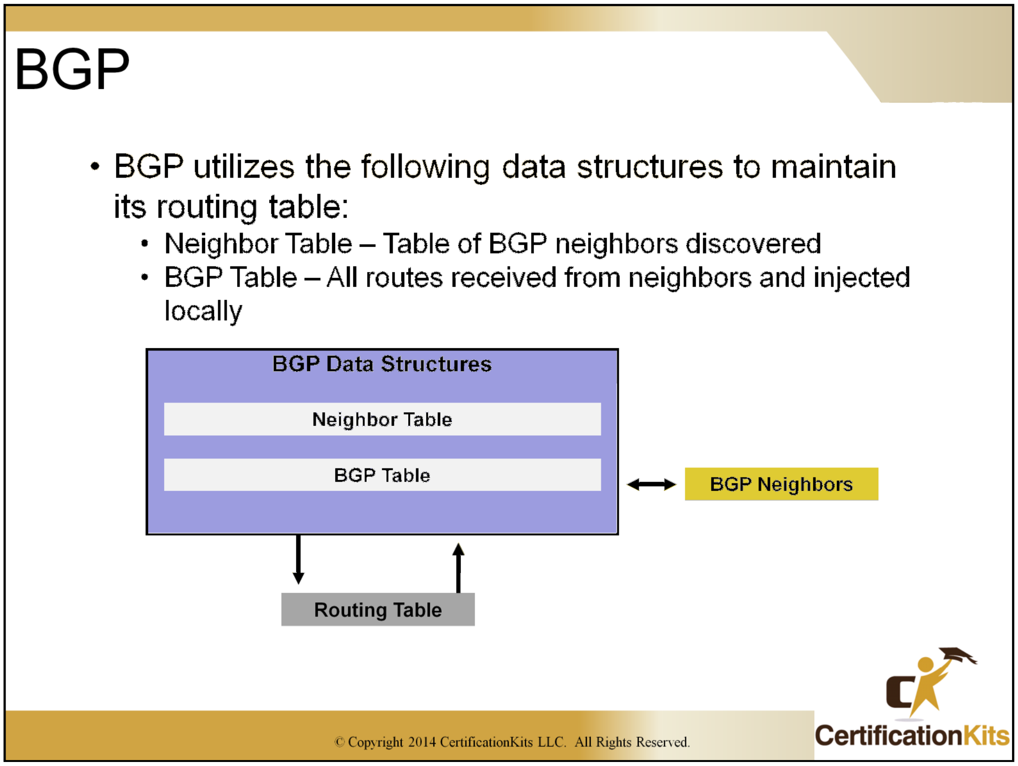

Neighbor table – Used to keep track of all BGP neighbors that have been configured on this router by IP address. Stores the configured autonomous system (AS) number of the neighbor, whether the session is internal or external, the state of the session, the uptime of the neighbor as well as other things.

BGP table – Also called the BGP Routing Information Base (RIB). It stores all routes that were received from all the router's peers, together with all the BGP attributes that are associated with the route.

Implications of High CPU Utilization

- How can you determine when high CPU utilization is normal and when high CPU utilization is a problem?
- What problems can arise due to high CPU utilization on a router?

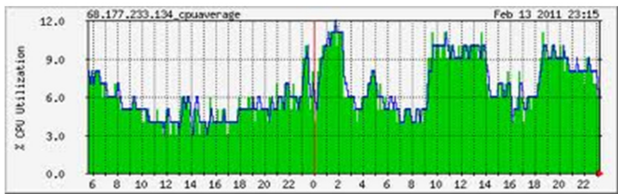

CertificationKits

The main functions a router CPU performs are packet forwarding and executing control plane and management processes. When high CPU utilization occurs it is important to understand when it is a potential problem and when temporary high CPU utilization might be normal. High CPU utilization could be attributable to short network burst which would be considered normal or if CPU utilization is consistently high it will be impacting performance and needs to be investigated.

Router Slowdown

- High CPU utilization can result in things such as packet loss, jitter (variations in delay), high delay, and other network performance related problems

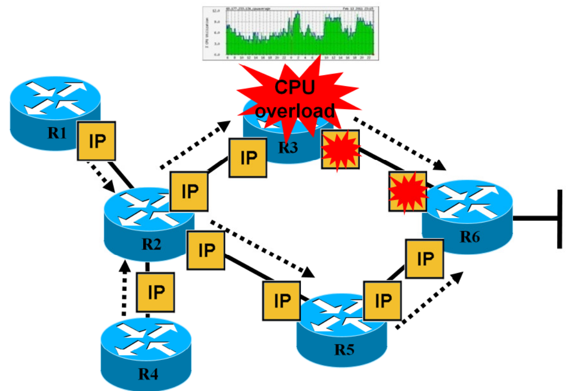

When the CPU is too busy network performance suffers. Things such as packet loss can occur as buffers fill up, jitter (variations in packet delay) as the CPU is not processing packets in a timely manner and higher than normal delays due to the same problem as just mentioned.

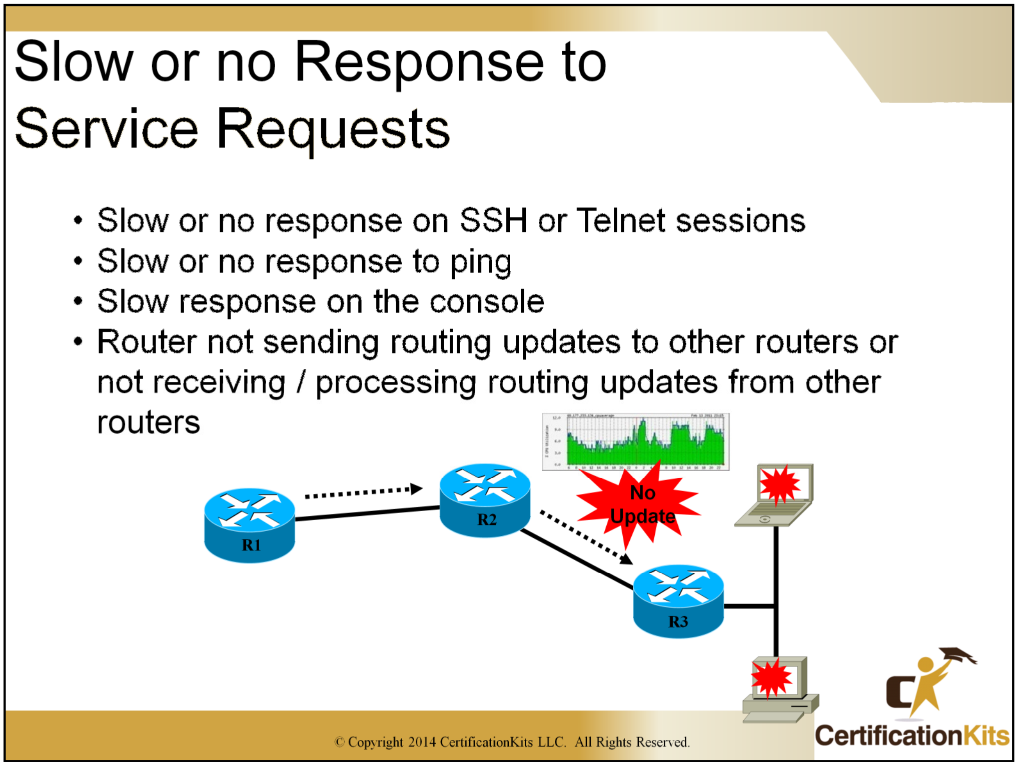

Numerous problems can arise due to high CPU utilization. Some common problems are listed on the slide above.

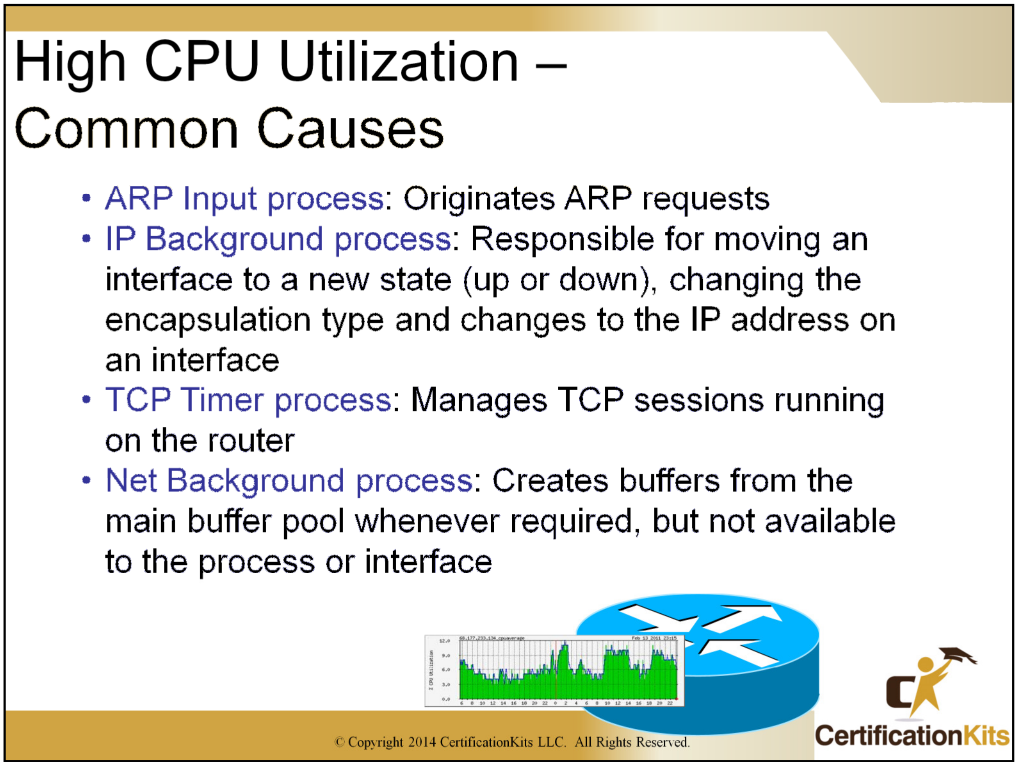

Some of the common causes to high router CPU utilization are listed on the slide above.

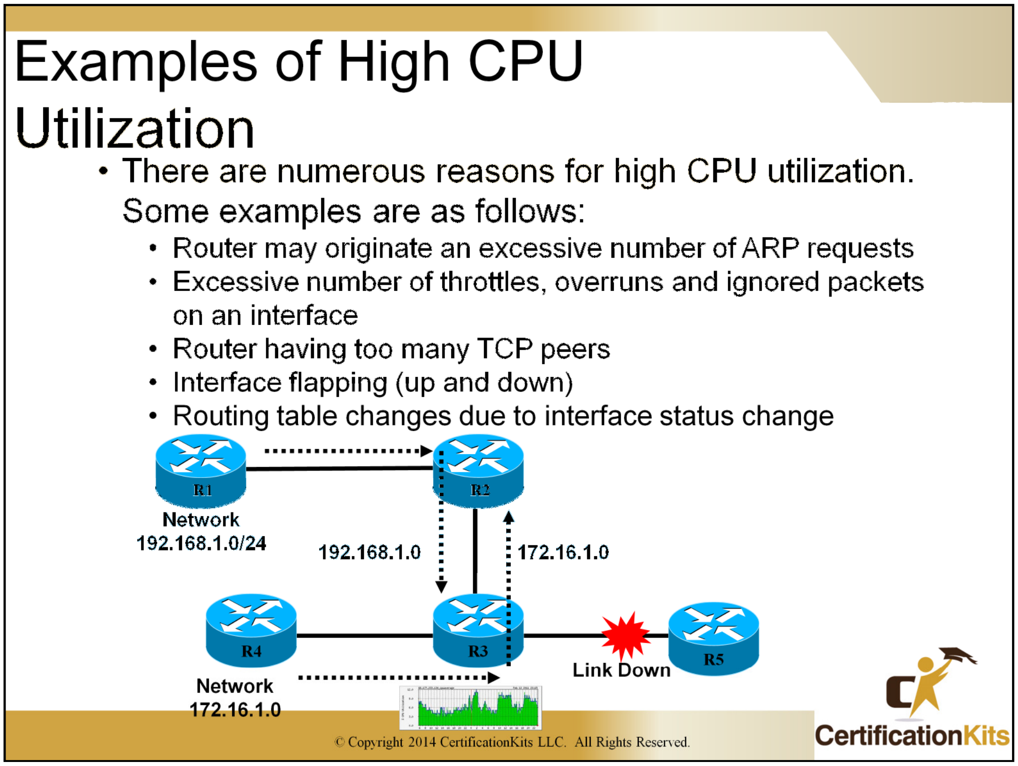

Listed on the slide above are just a handful of examples that can cause high router CPU utilization.

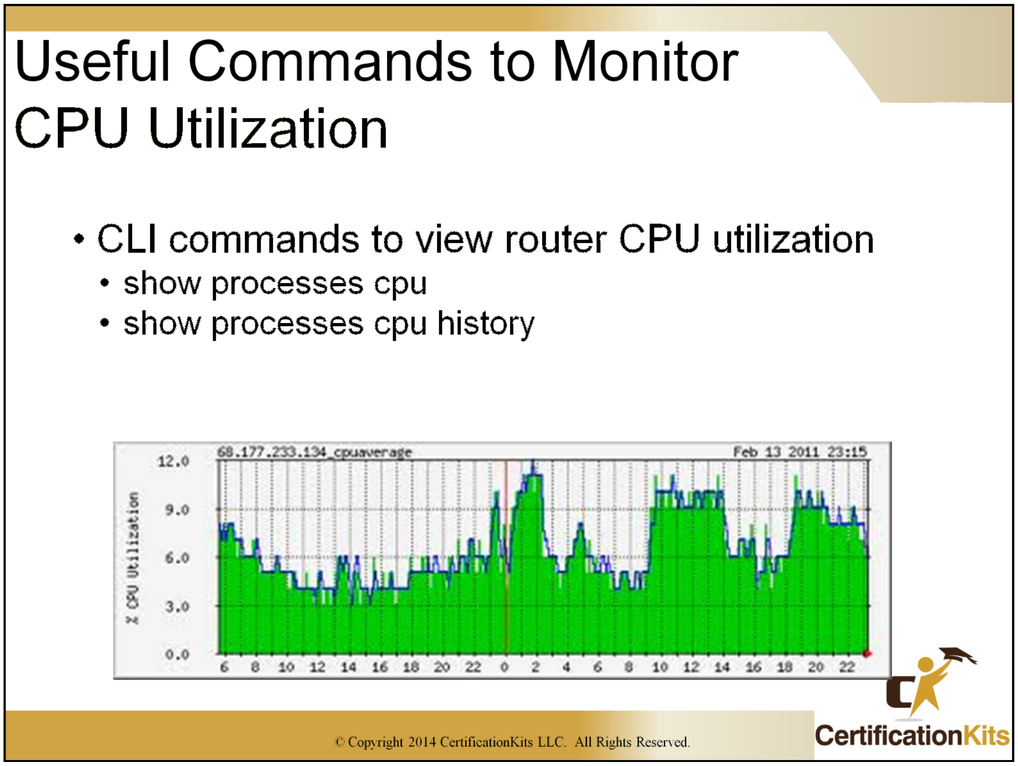

The **show processes cpu** command displays overall router CPU utilization as well as the CPU utilization of each of the processes running on the router. The output CPU utilization for the past 5 seconds, past 1 minute and past 5 minutes.

The **show processes cpu history** command is similar to the **show processes cpu** command although it displays router utilization for the past 1 minute, 60 minutes and 72 hours.

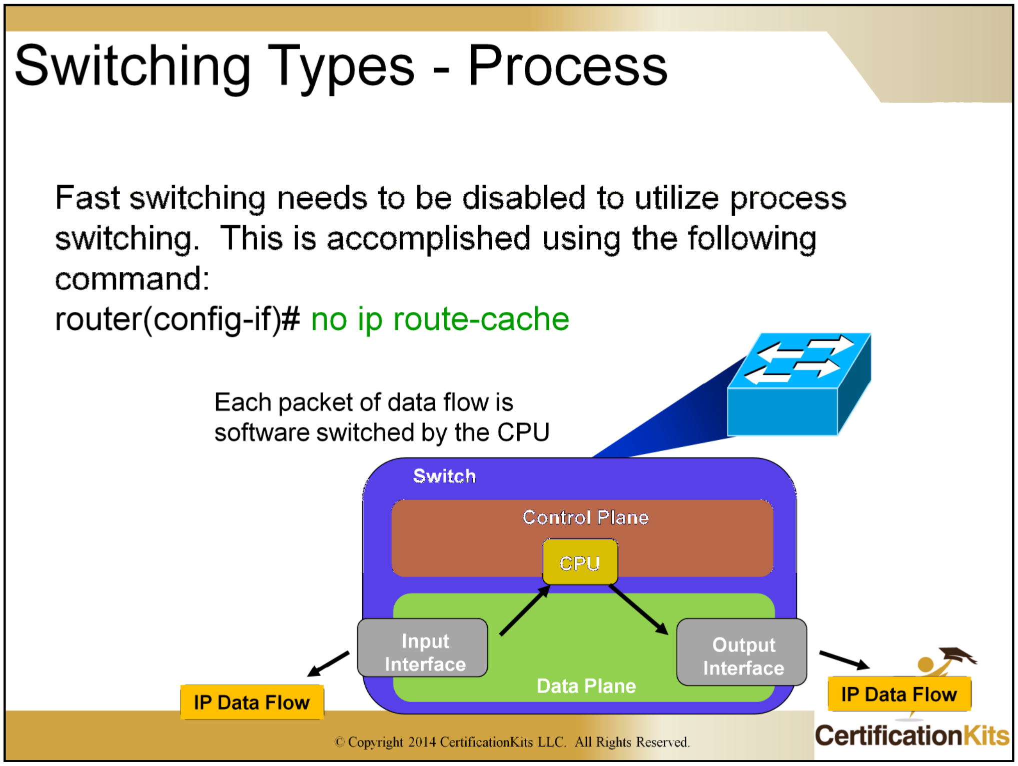

Process switching utilizes the CPU to process packets. Can be CPU intensive.

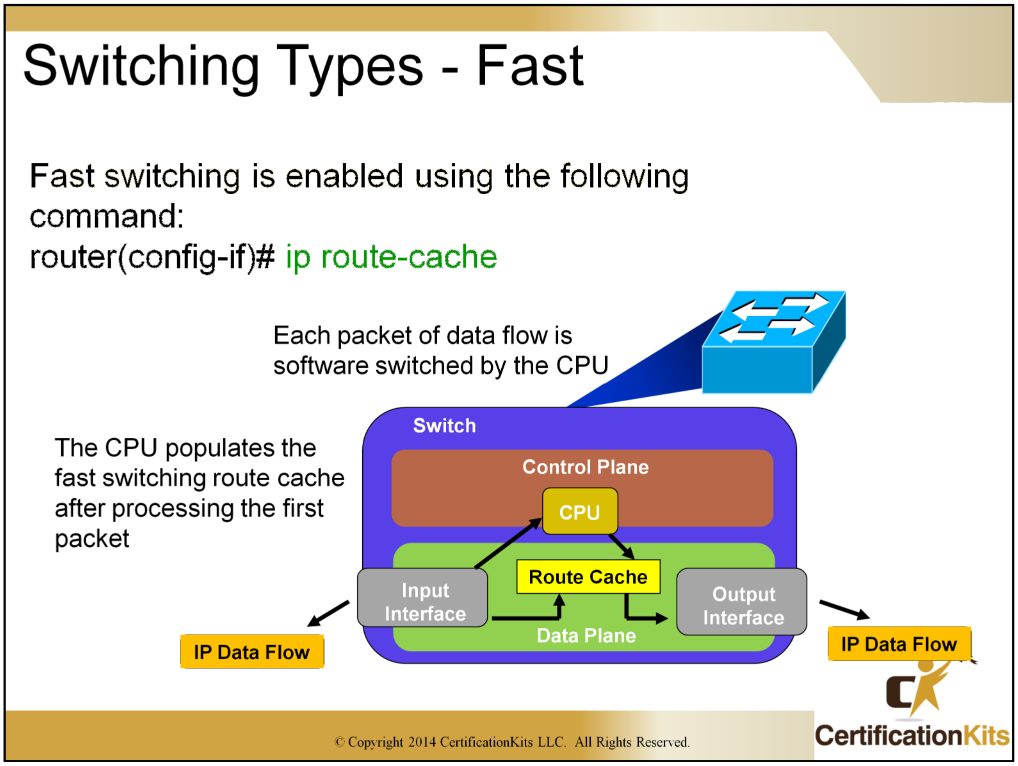

Fast switching is enabled by default on all interfaces that support it. It is less processor intensive that process switching.

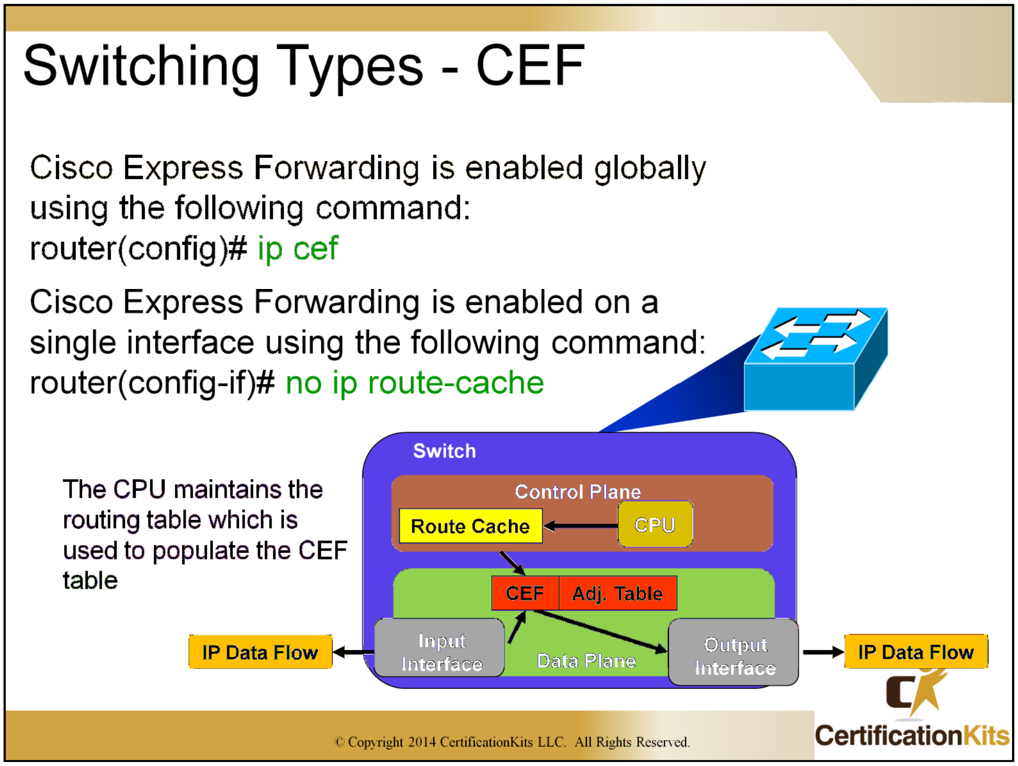

Cisco Express Forwarding utilizes the Forwarding Information Base (FIB) to make IP switching decisions. The FIB is only update after a network change.

Common Router Memory Problems

- The router does not have enough memory to support the Cisco IOS Software image:
  - Prior to downloading an IOS image from Cisco's web site it tells you what the memory requirements are

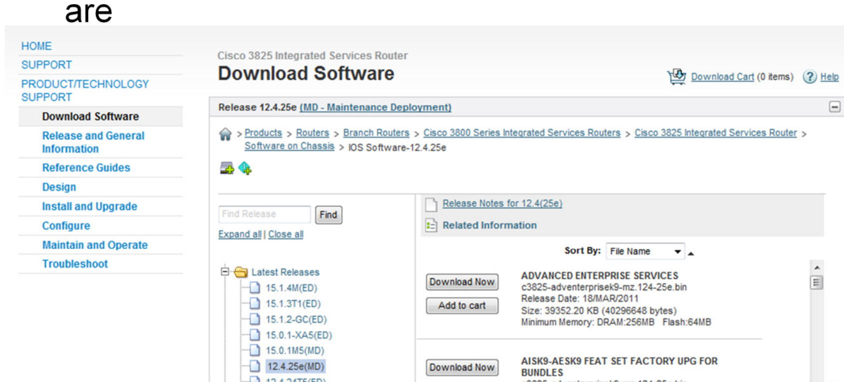

CertificationKits

In the example on the slide, in order to utilize the Advanced Enterprise Services version of IOS your 3800 series router would need to have a minimum of 256MB of DRAM and 64MB of Flash.

BGP and Memory Usage

- BGP is a very memory intensive routing protocol
- To determine the total amount of memory allocated to the processor pool and the I/O pool and how much memory has been allocated to a certain process, use the **show process memory** command

Chapter 5 Summary

- We reviewed the main data structures used by the EIGRP as well as the processes EIGRP uses to exchange routing information
- We learned different IOS commands that can be used to troubleshoot EIGRP
- We reviewed the main data structures used by the OSPF as well as the processes OSPF uses to exchange routing information within a single area as well as between areas
- We learned different IOS commands that can be used to troubleshoot OSPF
- We reviewed the main data structures used by the BGP as well as the processes BGP uses to exchange and process routing information
- We learned different IOS commands that can be used to troubleshoot BGP

CertificationKits

Chapter 5 Summary

- We learned typical causes of high router CPU utilization and different IOS commands that can be used to identify and solve performance problems caused by high CPU utilization
- We learned about the different switching types Cisco IOS has available and how each work
- We learned about common memory allocation problems and how to recognize symptoms associated with them

CertificationKits