Planning for Information Network

Lecture 8:

Network Routing Protocols

Routing protocol features

There are many ways to characterize routing protocols, including the following:

- Static versus dynamic routing
- Interior versus exterior routing protocols
- Distance vector versus link-state versus hybrid protocols
- Routing protocol metrics
- Routing protocol convergence
- Flat versus hierarchical routing protocols

Static Routing

The term **static routing** denotes the use of manually configured or injected static routes for traffic forwarding purposes. Using a static route might be appropriate in the following circumstances:

- When it is undesirable to have dynamic routing updates forwarded across slow bandwidth links, such as a dialup link.
- When the administrator needs total control over the routes used by the router.
- When it is necessary to reach a network that is accessible by only one path (a stub network)

A stub network is a computer network, or part of an internetwork, with no knowledge of other networks, that will typically send much or all of its non-local traffic out via a single path, with the network aware only of a default route to non-local destinations.

* Configuring and maintaining static routes is time-consuming.
* Properly implementing static routes requires complete knowledge of the entire network.

Static Routing

The following figure illustrates a stub network scenario in which the use of static routes is favored over a dynamic routing protocol. The right side of the figure shows a stub network with a single entry/exit point over the S0 interface of Router A. On the stub network router (Router A), a static default route is configured so that the S0 link forwards all traffic toward destinations outside the stub network. On Router B, a static route is installed toward the stub network and then is redistributed into the routing protocol so that reachability information for the stub network is available throughout the rest of the network.

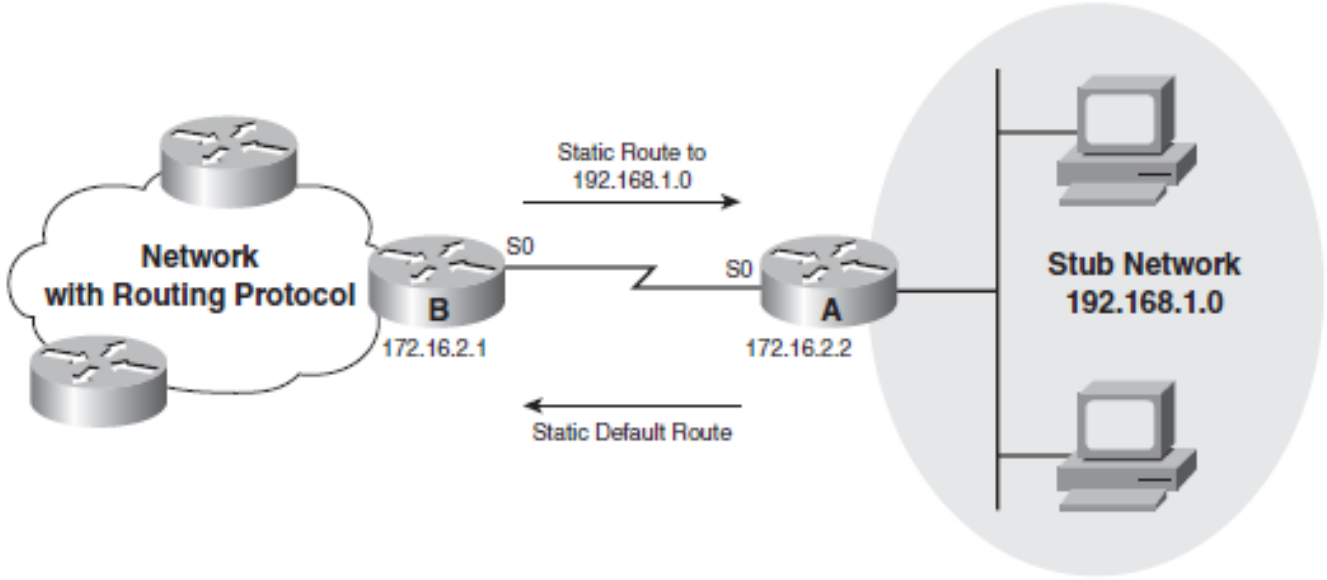

Static Routing

Static routes are unidirectional. A static route configured in one direction via one router must have a corresponding static route configured on the adjacent router, in the opposite direction, for the return path.

Static routes are appropriate in situations such as with stub networks, hub-and-spoke connections (Star connections), and dialup environments.

Dynamic Routing

Dynamic routing allows the network to adjust to changes in the topology automatically, without administrator involvement.

A static route cannot dynamically respond to changes in the network. If a link fails, the static route is no longer valid if it is configured to use that failed link, so a new static route must be configured. If a new router or new link is added, that information must also be configured on every router in the network. In a very large or unstable network, these changes can lead to considerable work for network administrators. It can also take a long time for every router in the network to receive the correct information. In situations such as these, it might be better to have the routers receive information about networks and links from each other using a dynamic routing protocol.

Dynamic Routing

Dynamic routing protocols must do the following:

- Find sources from which routing information can be received (usually neighboring routers).
- Select the best paths toward all reachable destinations, based on received information.
- Maintain this routing information.
- Have a means of verifying routing information (periodic updates or refreshes).

Dynamic Routing

When using a dynamic routing protocol, the administrator configures the routing protocol on each router. The routers then exchange information about the reachable networks and the state of each network. Routers exchange information only with other routers running the same routing protocol.

When the network topology changes, the new information is dynamically propagated throughout the network, and each router updates its routing table to reflect the changes.

Interior Versus Exterior Routing Protocols

An autonomous system (AS), also known as a domain, is a collection of routers that are under a common administration, such as a company's internal network or an Internet service provider's (ISP's) network.

Because the Internet is based on the AS concept, two types of routing protocols are required:

- **Interior gateway protocols (IGP)** are **intra-AS** (inside an AS) routing protocols. Examples of IGPs include Routing Information Protocol (RIP) version 1 (RIPv1), RIP version 2 (RIPv2), Open Shortest Path First (OSPF), Integrated Intermediate System-to-Intermediate System (IS-IS), and Enhanced Interior Gateway Routing Protocol (EIGRP).

- **Exterior gateway protocols (EGP)** are **inter-AS** (between autonomous systems) routing protocols. Border Gateway Protocol (BGP) is the only widely used EGP protocol on the Internet. BGP version 4 (BGP-4) is considered the acceptable version of BGP on the Internet.

Interior Versus Exterior Routing Protocols

Different types of protocols are required for the following reasons:

- **Inter-AS** connections require more options for manual selection of routing characteristics. EGPs should be able to implement various policies.
- The speed of convergence (distribution of routing information) and finding the best path to the destination are crucial for **intra-AS** routing protocols.

Therefore, EGP routing protocol metrics (used to measure paths to a destination) include more parameters to allow the administrator to influence the selection of certain routing paths.

*EGPs are slower to converge and more complex to configure.

*IGPs use less-complicated metrics to ease configuration and speed up the decisions about best routing paths for faster convergence.

IGP and EGP Example

The following figure shows three interconnected autonomous systems (domains). Each AS uses an IGP for intra-AS (intra-domain) routing.

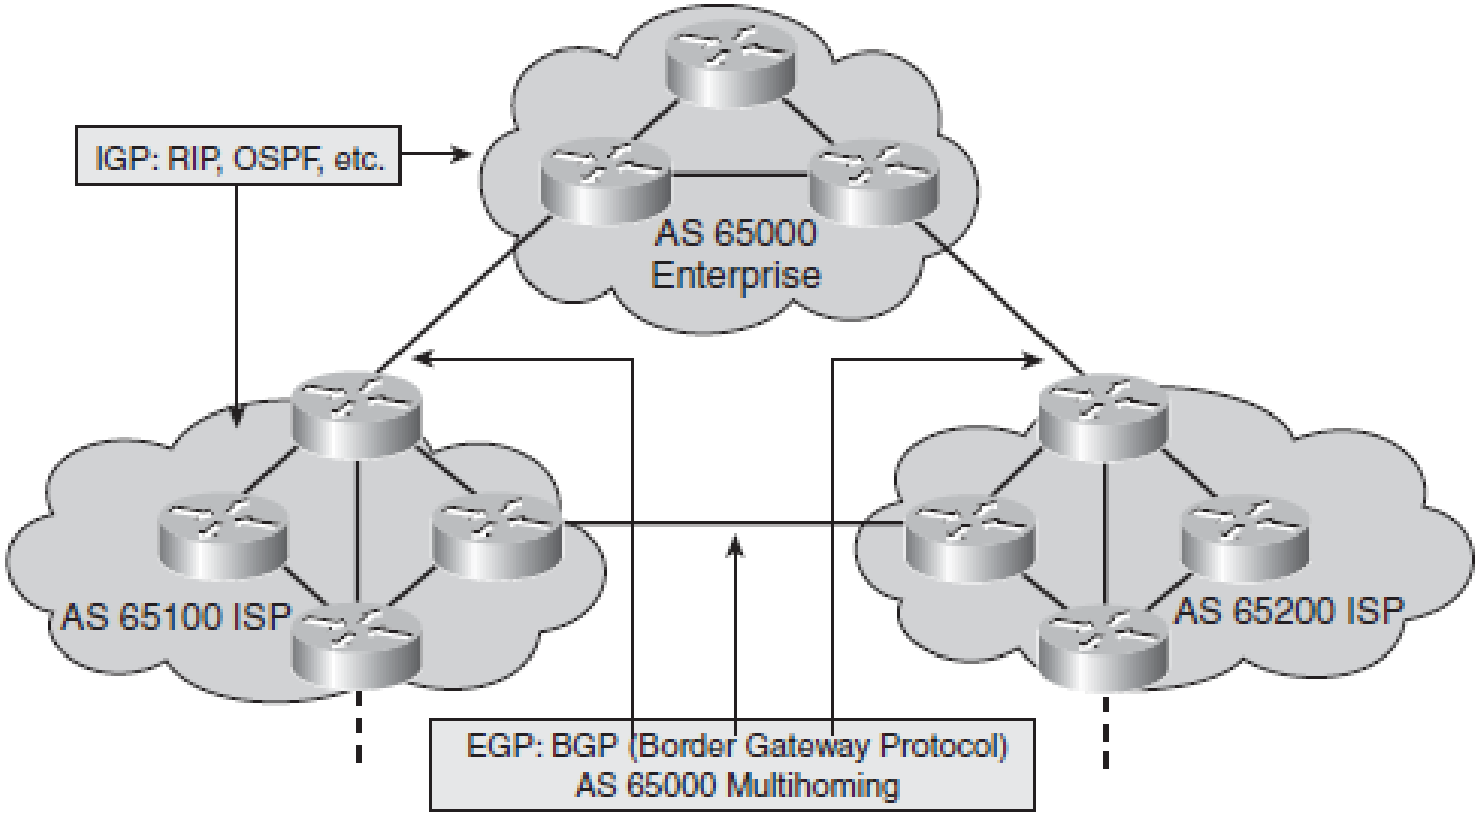

Multihoming is when an AS has more than one connection to the Internet (for redundancy or to increase performance).

Distance Vector Versus Link-State Versus Hybrid Protocols

There are two main types of routing protocols:

- **Distance vector protocol:** In a distance vector protocol, routing decisions are made on a hop by-hop basis. Each router relies on its neighbor routers to make the correct routing decisions. The router passes only the results of this decision (its routing table) to its neighbors.

Distance vector protocols are typically slower to converge and do not scale well; however, they are easy to implement and maintain.

Examples of distance vector protocols include RIPv1, RIPv2, and Interior Gateway Routing Protocol (IGRP).

Distance Vector Versus Link-State Versus Hybrid Protocols

- **Link-state protocol:** Each router floods information about itself (its link states) either to all other routers in the network or to a part of the network (area). Each router makes its own routing decision based on all received information and using the shortest path first (SPF) algorithm which calculates the shortest path to any destination. Link-state protocols are fast to converge, have less routing traffic overhead, and scale well. However, because of their complexity, link-state protocols are more difficult to implement and maintain. The IP link-state protocols are OSPF and Integrated IS-IS.

Distance Vector Versus Link-State Versus Hybrid Protocols

A third type of protocol also exists: the **Hybrid interior gateway protocol**, which is the Cisco EIGRP. EIGRP has characteristics of both distance vector and link-state protocols; it combines distance vector behavior with some link-state characteristics and some proprietary features.

EIGRP is a fast-converging and scalable routing protocol.

Distance Vector Versus Link-State Versus Hybrid Protocols

Note: In the name link-state, link refers to the interface, and state refers to the link's characteristics, such as whether it is up or down.

Note: A network is converged when routing tables on all routers in the network are synchronized and contain a route to all destination networks.

Note: Routers running link-state and hybrid protocols use multicast packets to communicate with each other.

Distance Vector Versus Link-State Versus Hybrid Protocols

** When a network is using a distance vector routing protocol, all the routers periodically send their routing tables, or a portion of their tables, to only their neighboring routers.

** In contrast, when a network is using a link-state routing protocol, each of the routers sends the state of its own interfaces (its links) to all other routers, or to all routers in a part of the network known as an area, only when there is a change.

** Routers running a hybrid protocol send changed information only when there is a change (similar to link-state protocols), but only to neighboring routers (similar to distance vector protocols).

Distance Vector Example

A distance vector router's understanding of the network is based on its neighbor's perspective of the topology; consequently, **the distance vector approach is sometimes referred to as routing by rumor**. Routers running traditional distance vector protocols periodically send their complete routing tables to all connected neighbors.

The following figure shows a sample network that runs a distance vector protocol. In this network, the routing updates are periodic and include the entire routing table.

Distance Vector Example

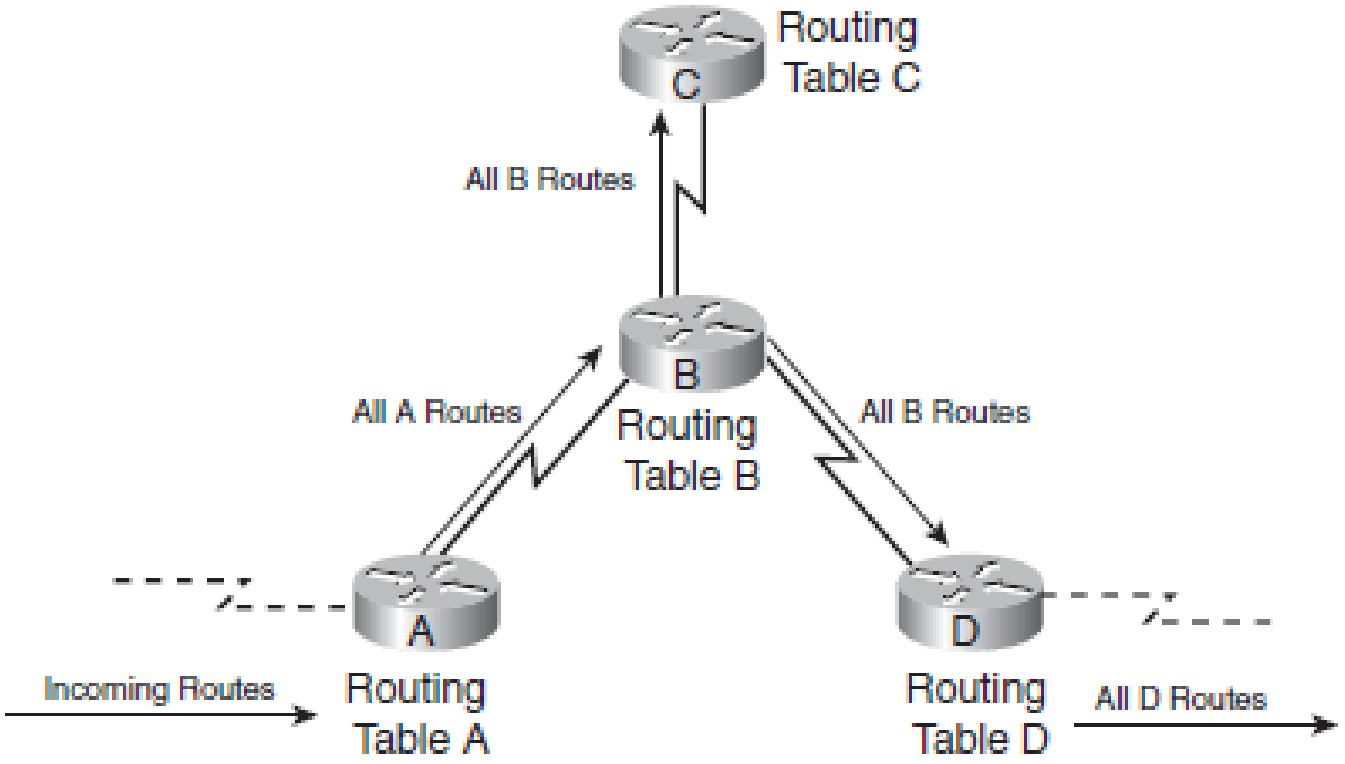

Link-State Example

Both OSPF and Integrated IS-IS use the Hello protocol for establishing neighbor relationships.

Those relationships are stored in a neighbor table (also called an adjacencies database). Each router learns a complete network topology from information shared through these neighbor relationships. That topology is stored in the router's link-state database (LSDB), also called the topology table or topology database. Each router uses this topology and the SPF algorithm to create a shortest-path tree for all reachable destinations. Each router selects the best routes from its SPF tree and places them in its routing table (also called the forwarding database).

Link-State Example

The following figure shows a network that uses a link-state protocol. Triggered updates, which include data on the state of only links that have changed, are sent in this network.

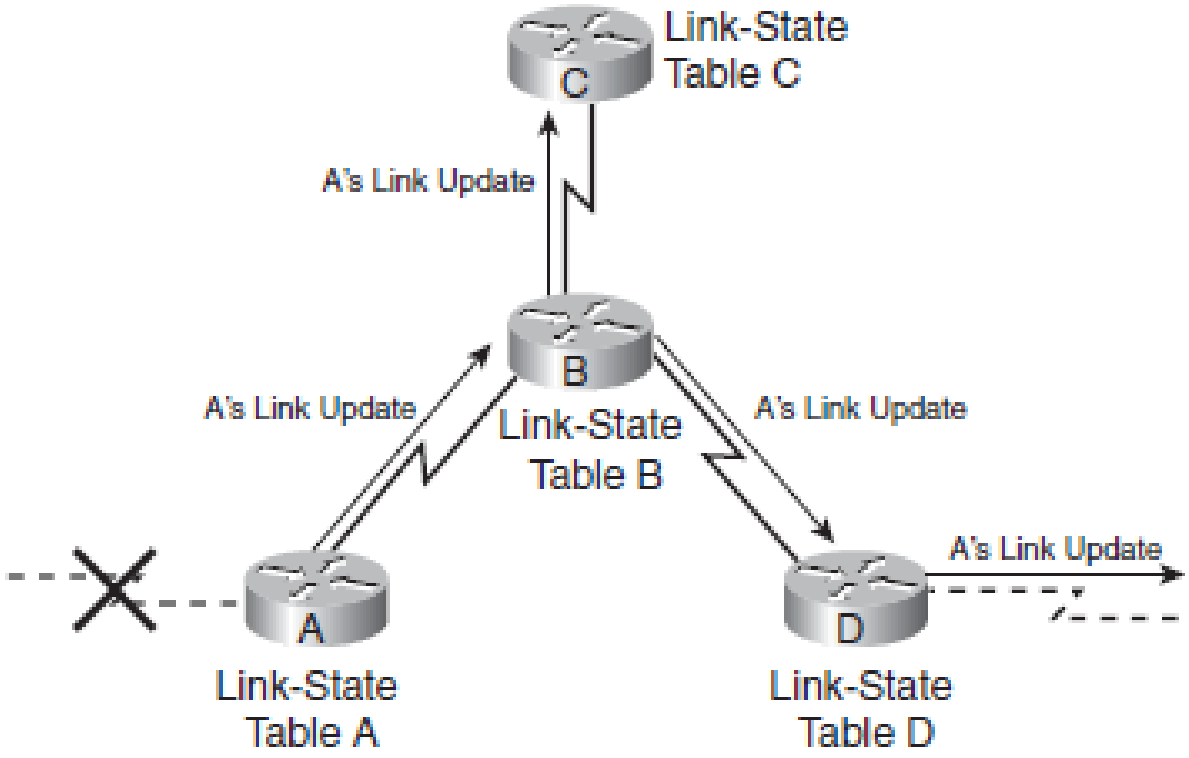

Link-State Example

In link-state protocols, the information about connected links (including the subnets on those links) on all routers is flooded throughout the network or to a specific area of the network. Therefore, all routers in the network have detailed knowledge of the entire network. In contrast, routers running a distance vector routing protocol receive knowledge about only the best routes from their neighbors.

After the initial exchange of all link states and on reaching the full (converged) state of operation, almost no periodic updates are sent through the network. (In OSPF, periodic updates are sent every 30 minutes for each specific route, but not at the same time for all routes, reducing the routing traffic volume.) Triggered updates are flooded through the network only when a change in a link state occurs (the link goes down, comes up, or link parameters that affect routing—such as bandwidth—are changed). Only periodic hello messages are sent between neighbors to maintain and verify neighbor relationships.

Routing Protocol Metrics

A **metric** is a value (such as path length) that routing protocols use to measure paths to a destination.

Different routing protocols base their metric on different measurements, including hop count, interface speed, or more-complex metrics.

Most routing protocols maintain databases containing all the networks that the routing protocol recognizes and all the paths to each network. If a routing protocol recognizes more than one way to reach a network, it compares the metric for each different path and chooses the path with the lowest metric. If multiple paths have the same metric, a maximum of 16 can be installed in the routing table, and the router can perform load balancing among them. EIGRP can also perform load balancing between unequal-cost paths.

Routing Protocol Metrics

The following figure shows network 172.16.1.0, which is connected to Router A. The parameters for route metric calculation are forwarded in routing protocol updates.

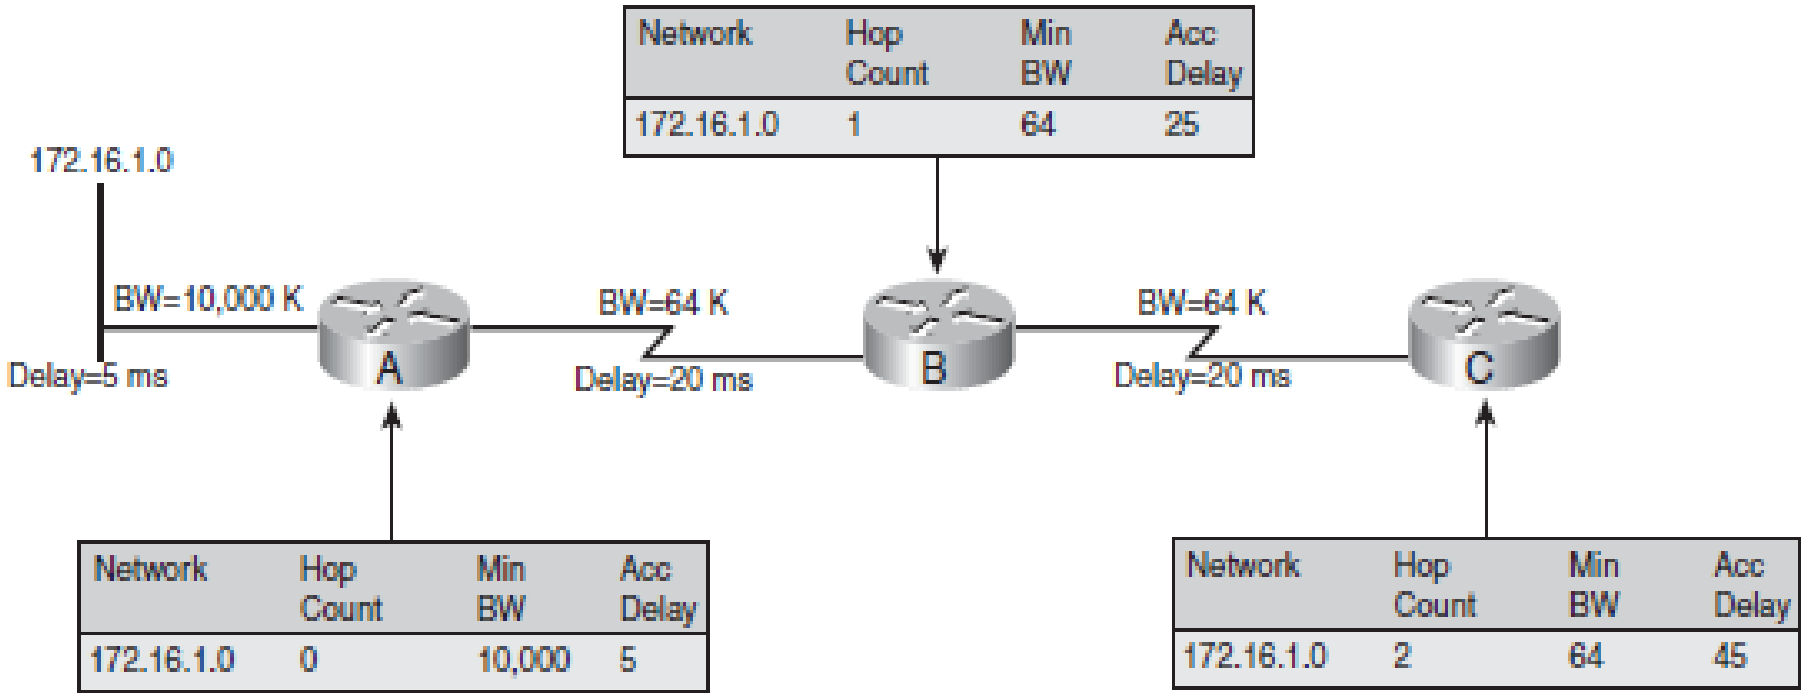

Routing Protocol Metrics

In this case, the EIGRP method of route metric parameters is used, and the minimum bandwidth and cumulative delay influence best path selection (the path with the highest minimum bandwidth and lowest delay is preferred). The previous figure shows the following steps:

Step 1: Router A, which is the originator of the route 172.16.1.0, sends the initial metric values to Router B.

Step 2: Router B takes into account the parameters of its link toward Router A, adjusts the parameters (bandwidth, delay, hop count) appropriately, calculates its metric toward the 172.16.1.0 network, and sends the routing update to Router C.

Step 3: Router C adjusts the parameters again and calculates its metric toward the destination network 172.16.1.0 from those parameters.

Routing Protocol Convergence

Whenever a change occurs in a network's topology, all the routers in that network must learn the new topology. This process is both collaborative and independent; the routers share information with each other, but they must calculate the impact of the topology change independently. Because they must mutually develop an independent agreement on the new topology, they are said to converge on this consensus.

The quicker the convergence, the more optimal the routing protocol is said to be.

Network convergence must occur whenever a new routing protocol starts and whenever a change takes place in the network. It occurs in both new networks and those that are already operational.

Convergence is also important when changes occur in the network.

Flat Versus Hierarchical Routing Protocols

Flat routing protocols propagate all routing information throughout the network, whereas hierarchical routing protocols divide large networks into smaller areas.

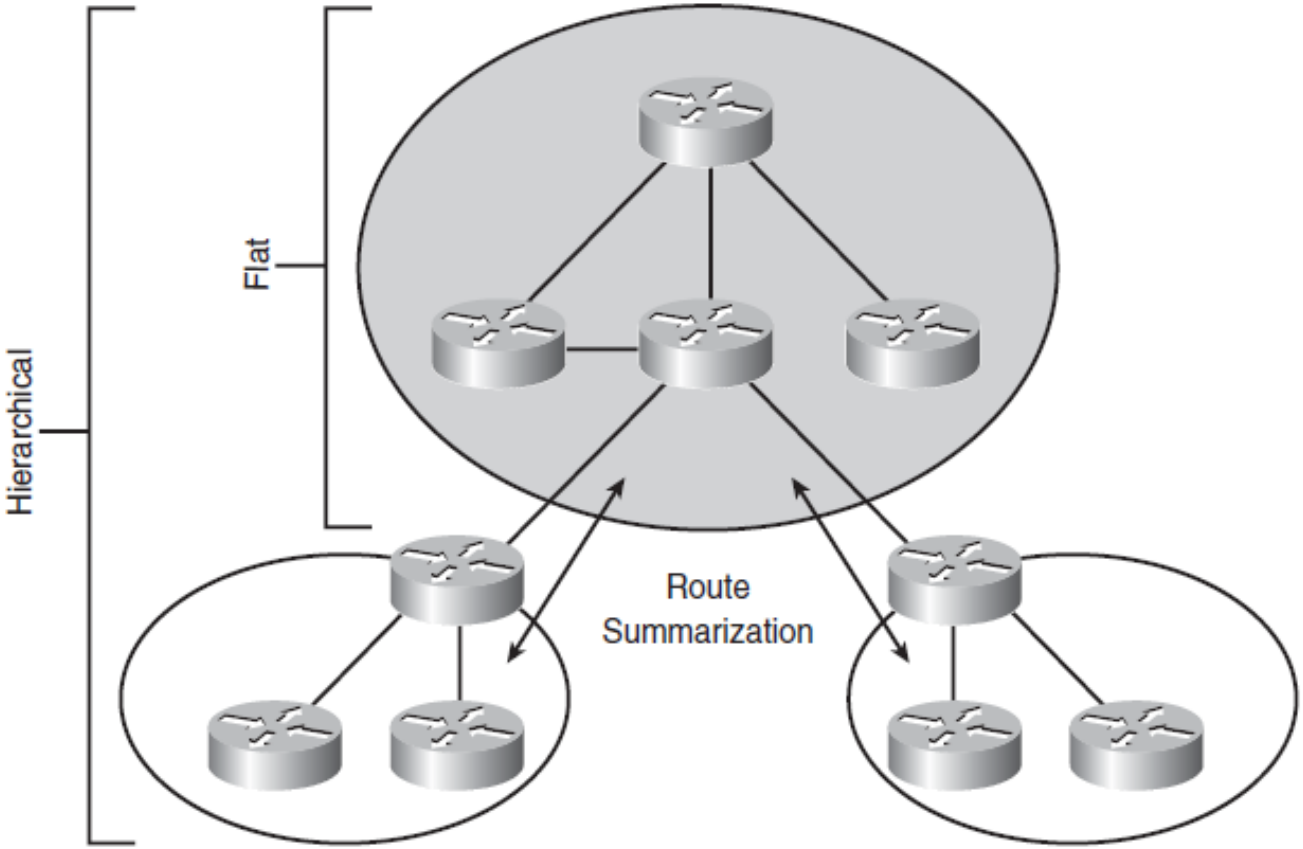

Flat Routing Protocols

Flat routing protocols have no means of limiting route propagation in a major network (within a Class A, B, or C network) environment. These protocols are typically classful distance vector protocols.

Classful means that routing updates do not include subnet masks and that the protocol performs automatic route summarization on major network (class) boundaries. Summarization cannot be done within a major network. These protocols support only fixed-length subnet masking (FLSM); they do not support VLSM.

Distance vector protocols periodically send entire routing tables to neighbors. Distance vector protocols do not scale well, because in a large network, they produce significant volumes of routing information that consume too many network resources (CPU, bandwidth, memory). These resources should be available to the routed traffic (application data and user traffic) instead.

Two examples of flat routing protocols are RIPv1 and RIPv2. Note, however, that RIPv2 is a classless protocol.

Hierarchical Routing Protocols

To solve the problems associated with flat routing protocols, additional features are implemented in hierarchical routing protocols to support large networks—for example, some support an area based design.

Hierarchical routing protocols are typically classless link-state protocols. Classless means that routing updates include subnet masks in their routing updates; therefore, the routing protocol supports VLSM.

Hierarchical Routing Protocols

Hierarchy is part of the implementation of link-state protocols with the concept of backbone and non backbone areas. With link-state protocols such as OSPF and IS-IS, large networks are divided into multiple areas.

Route summarization can be performed manually in hierarchical protocols and is required in most cases. With the help of route summarization, smaller routing updates propagate among areas, resulting in higher scalability. Instabilities in one part of the network are isolated, and convergence is greatly improved. Summarization can be performed on an arbitrary bit boundary within an IP address. Note, however, that OSPF supports summarization on only specific routers called area border routers and autonomous system boundary routers.

Although it is a classless hybrid protocol, EIGRP is considered a flat routing protocol because it is not area-based. Because EIGRP also supports manual summarization, EIGRP can be used in a hierarchical network design by dividing the network into areas. A hierarchical design is not necessary in EIGRP, but one is recommended for large networks.

Thank you