Module 4: TCP/IP Dynamic Routing Protocols

>> TCP/IP's distinction is that it can scale to an almost limitless size, whereas IPX and AppleTalk have size and operational limitations that make them undesirable for use in large-scale enterprise and global networks.

>> The flexibility of the TCP/IP protocol suite is a result of IP's connectionless datagram delivery process and the diversity of IP dynamic routing protocols.

Dynamic Routing Protocol Basics

>> To find the best, most efficient route to forward IP network traffic.

>> Also find a secondary route in the event that the best route is lost.

>> Balances the traffic.

>> Easiness to add new routers in the path.

>> IP network routing can be Intranetwork routing and Internetwork routing.

>> Intranetwork routing, or interior routing, exchanges route information between routers within defined routing processes.

>> An intranetwork can have single or multiple routing processes.

>> Protocols that perform such routing are known as interior gateway protocols (IGPs).

>> Eg: Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).

>> Under the OSI model, intranetwork routing is known as **intradomain routing**.

Dynamic Routing Protocol Basics...

>> Internetwork routing has two dimensions:

1) the exchange of routing information between large, centrally administered networks known as autonomous systems (AS). Eg:Exterior Gateway Protocol (EGP) and Border Gateway Protocol (BGP).

2) running multiple distinct routing protocol instances or a collection of different IGP routing protocols to distribute routing information within a network hierarchy.

>> Internetwork routing protocols are also called Exterior routing protocols (ERP).

>> Under the OSI model, internetwork routing is known as **interdomain routing**.

>> Why do we use a Routing Protocol?

* Create host-to-host internetworks and point-to-point links.
* Connectivity between every two site.
* Load balancing
* Sustained networks even on some network or router malfunctionings.

>> *Routing metrics* are used by dynamic routing protocols for;

1) Route diversity: two or more unrelated access points or paths exist for the same network.
2) Route redundancy: two or more access points to the same network exist with equal metrics.
3) Load balancing.

The most common routing metric variables:

1) **Hop count**: It is the number of intermediate systems (routers) between the router and the destination router.
2) **Bandwidth**: This metric reflects the interface's ideal throughput. For example, a serial interface on a Cisco router has a default bandwidth of 1.544Mbps, and Ethernet has a default bandwidth of 10Mbps.
3) **Load**: The load metric varies based on the actual usage (traffic).
4) **Delay**: It is the total time needed to move a packet across the route. The shortest time is the best route.
5) **Reliability**: Reliability estimates the chance of a link failure and can be set by an administrator or established by the given protocol.
6) **Cost**: This metric sets the preference for a given route. The lower the cost, the more preferable the route. The interface's default cost is directly related to its speed.

Network Convergence: The process of bringing all the routing tables of all the routers in the network to a state of consistency. This is done by broadcasting or multicasting the routing information.

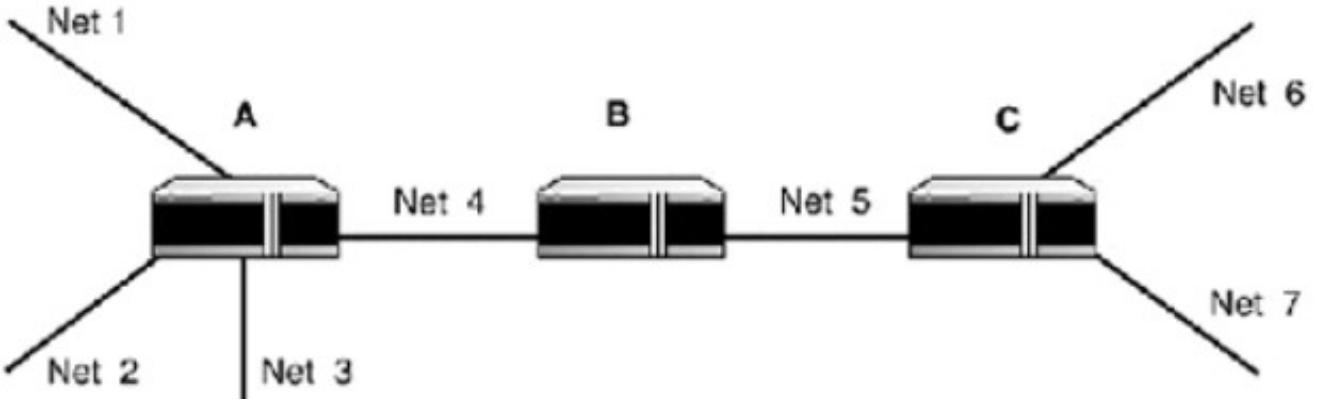

Convergence time: It is the time taken for routers to learn the network topology and/or changes in the network topology (for example, a failed link or the addition of a new network segment). In large networks, it is preferable to use a routing protocol that has a fast convergence time.

TCP/IP Static Routing

>> Better to manage for entry level adinistrators or who does not know deep into the networking.

>> Create a master routing table and copies it io all routers. Done.

>> Strengths of static routing:

1) Ease of use
2) Reliability
3) Control
4) Security
5) Efficiency

>> Weaknesses of Static Routing

- Difficult to manage when the network grows dynamically
- Not practical for large networks
- No adaptive path selection on router failures.

TCP/IP Interior Gateway Protocols

* Routing Information Protocol (RIP), v1 and v2

>> RIP is the base for any dynamic routing protocols.

>> Most of the layer-3 devices (routers, PCs with Unix/Windows NT) supports RIP v1 or v2.

>> RIP uses DV routing.

>> Only one metric to determine the best route: hop count. Hop count ranges from 1 (directly connected) to 16 (unreachable network). Due to low hopcount, RIP is not used in large networks.

>> RIPv1 supports classful addressing only. Uses UDP broadcasts (source and destination port 520)

>> RIPv2 fixes the security holes that existed in RIPv1. Supports CIDR, VLSM, routing authentication and route summarization. RIPv2 uses the multicast address 224.0.0.9 (also to port 520).

>> Extended RIP (ERIP) is also available on some platforms; supports upto 128 hops.

>> DV routing: When booted up, a router creates a local routing table containing its neighbours id and hops to them. It then sends its table to its neighbours. They will update their table using these values. The former router updates its table by receiving the tables from its neighboirs.

* **Routing Information Protocol (RIP): Routing Messages and Convergence**

>> Two types of RIP message updates: when the router first joins the RIP process, and in every 30 seconds.

>> RIP will also send an update in the event of a topology change; this is known as a triggered update.

>> If a router has not received an update message from a particular network after 180 seconds, the route is marked as invalid, but continues to forward the packets through this route until a next alternative path (with a higher metric) is found. If not found and still no update is available from thet network, all other routers are informed using ICMP messages to change the table and flush the unavailable-router-entry.

>> Convergence depends on the number of routers and the topology (router arrangement).

>> To prevent routing-loops and count to infinity problem in DV protocols, RIP employs route poisoning through split horizon and poison reverse methods (both having the same effects).

>> Split horizon: When a router R1 sends an update to another router R2, the update table will not contain any information that is learnt from R2.

>> Poison Reverse: When a router R1 sends an update to another router R2, it sets the hop counts to 16 for all routes learnt from R2.

>> RIP is rarely used today. IPX use RIP for route management.

* Routing Information Protocol (RIP)...

>> RIP v1 and v2 uses the same message format. The major fields are;

Command: request or response.
Version: RIP version 1 or 2
Address Family Identifier: States which protocol is using RIP for routing calculations. ID for IP is 2.
Route Tag: This is a RIPv2-specific field.
IP Address (the vector): This can be an IP address of a host, IP classful network address, IP classful subnet address, or a default route.
Subnet Mask: RIPv2 only. The netmask of the IP address.
Next Hop: RIPv2 only. This is used as a hint; if an alternative same or lower hop gateway exists on the same subnet it will be sent along in the update in this field.
Metric: The distance (hop count) between the router that is sending the update and the destination network.

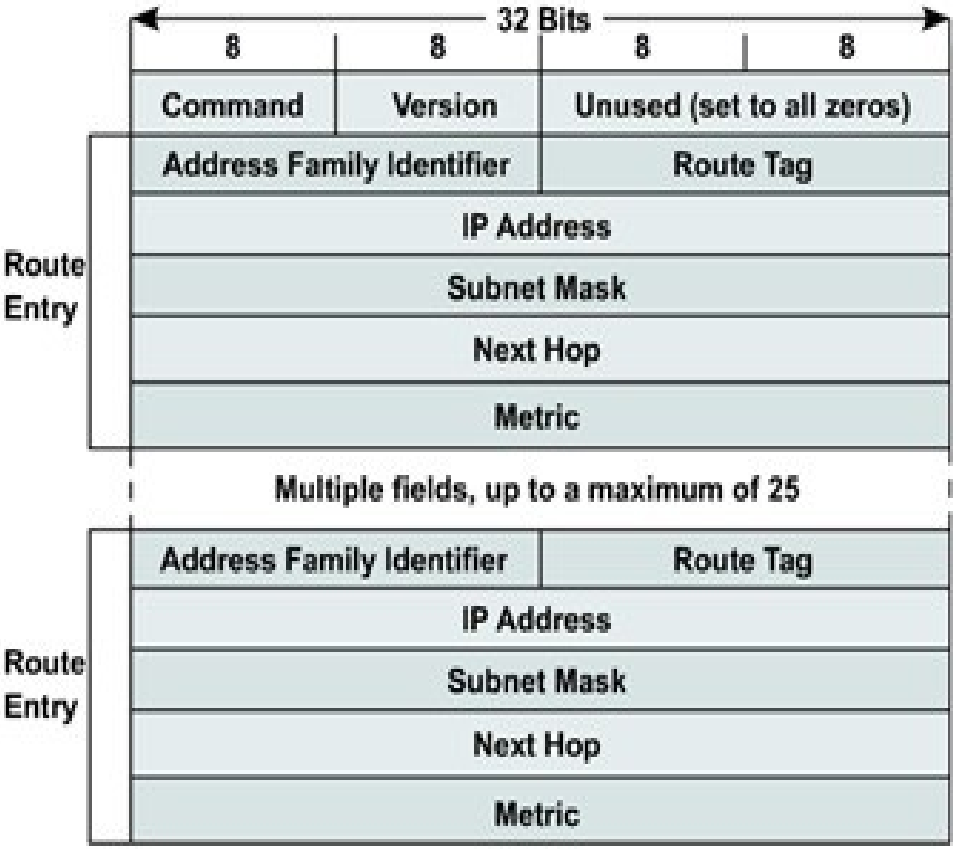

IGRP and EIGRP

>> Interior Gateway Routing Protocol (IGRP) and Enhanced Interior Gateway Routing Protocol (EIGRP) are Cisco Systems' proprietary protocols, which means they are only supported on Cisco routers.

>> IGRP was created as a more robust alternative to RIPv1 and became more poupular.

>> a DV protocol.

>> implements split horizon to perform route poisoning.

>> no support for VLSM (like RIPv1).

++ IGRP's advantages over RIPv1 were;

1) Support for multiple route metrics: Where RIPv1 uses the hop count metric to determine a route's desirability, IGRP supports four metrics (bandwidth, delay, load, and reliability) all of which are tunable by the administrator.
2) Support for larger network diameters: IGRP can support a network of up to 255 hops in size versus RIP's 15 and ERIP's 128.
3) Support for AS and multiple routing process domains: IGRP has support for external routes; RIP can only operate as a single routing domain.

- - The disadvantage of IGRP is that it has a 90 second update interval.

IGRP and EIGRP...

>> EIGRP is a hybrid protocol that uses the diffusing update algorithm (DUAL) than DV routing. Hence no count to infinity problems.

>> Route computations are "shared" across routers.

>> Routing tables are still exchanged between routers (such as RIP and IGRP), but updates are only sent when changes in the network topology occur.

>> Routing metrics parameters are same as IGRP.

>> Supports VLSM, route redistribution, and multiprotocol routing (IPX and AppleTalk).

++ Fast convergence, low bandwidth consumption.

- - Only supported by Cisco routers.

OSPF (Open Shortest Path First)

>> A link state (LS) protocol.

>> Exchange link state information about the network in which they are directly connected.

>> Each router constructs a map of the entire network topology from its perspective.

>> Advantages:

1) Better reliability: OSPF routers construct their own routing table describing network reachability from their perspective form within the network. All network routes are stored in the topological or link state database, where they can be retrieved quickly in the event that a routing change must be made.
2) Fast convergence: After the network has converged and all the routers have constructed their own routing tables and network maps, updates are sent out only when changes in the network topology occur. When changes occur, they are flooded.
3) Unlimited network size: Operate in both large and small-scale networks.
4) VLSM support
5) Type of Service routing: Supports Layer 4 Quality of Service routing.
6) Low bandwidth usage: OSPF uses multicast instead of local network broadcasts.
7) Dynamic load balancing and route selection

>> Needs more processing power and memory.

OSPF...

>> OSPF implementation elements:

1) The OSPF area hierarchy
2) Router designations
3) OSPF network types
4) Inter/intra routing

1) The OSPF area hierarchy

>> Since OSPF needs more memory to calculate shortest path, low end routers will fail for large network computations. To solve this, OSPF uses a hierarchical network segmentation structure called **areas**.

>> Every OSPF network has at least one area, known as a **root** or **backbone** area (called area 0 or area 0.0.0.0).

>> All other areas are directly or virtually connected to the root area.

>> Each area performs route computation separately and the root area exchanges these inter-area routing information.

OSPF...

2) Router Designations

>> There are four different designations for routers in an OSPF network.

Backbone routers: Maintain complete routing information for all the areas (or domains) that are connected to the backbone.

Area border routers: Connect two or more areas (inter-area routing). These routers usually sit on the network backbone and connect the branch areas to the root area.

Internal routers: Involved only in routing their internal area (intra-area routing) and only maintain information about the area in which they operate.

Autonomous system (AS) boundary routers: These routers have interfaces on the network that are outside the internetwork routing domain.

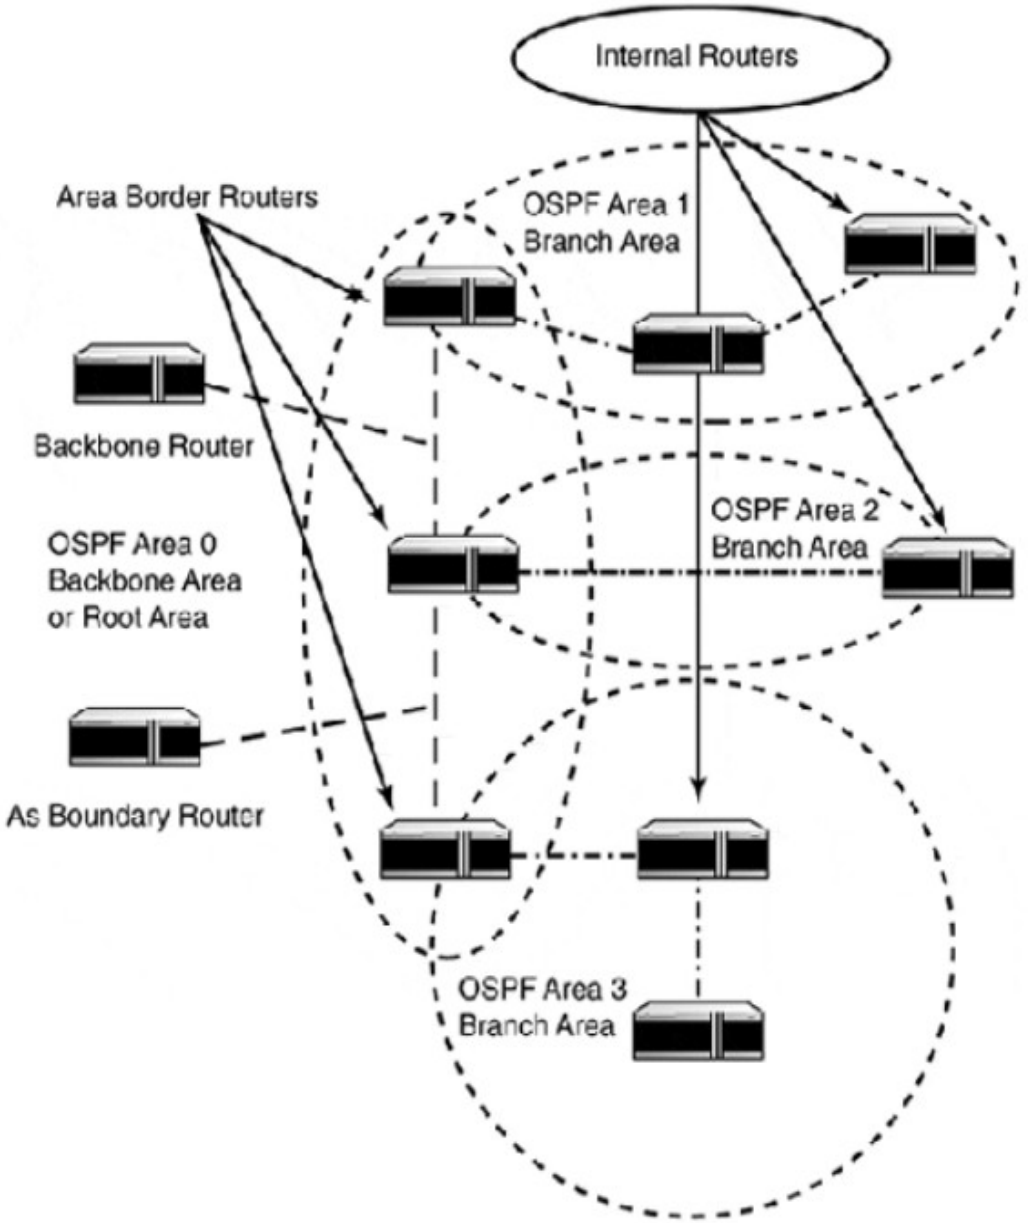

OSPF...

3) OSPF Network Types

>> There are three types of networks that OSPF recognizes:

Broadcast networks: These networks allow multiple routers to be connected to the same physical network segment. The network must also support broadcast and multicast capabilities. OSPF uses multicast addresses 224.0.0.5 and 224.0.0.6 to send and receive messages. All LAN media types and some WAN media types support these capabilities.

Point-to-point networks: used toconnect a pair of routers. A serial line connection is a common example of a point-to-point network.

Nonbroadcast multi-access networks: These networks allow multiple routers to be connected to the same network segment, but do not support the broadcast and multicast capabilities needed to send OSPF messages. In these cases, messages are sent using unicast.

TCP/IP Exterior Gateway Protocols

>> Exterior routing protocols (ERPs) are used primarily in Internet backbone networks that connect internetworks.

>> Exterior routing is the routing between Autonomous Systems (AS).

An autonomous system (AS) is a network or a collection of networks that are all managed and supervised by a single entity or organization.

>> There are three types of AS:

1) Stub: Stub AS reach other networks through one gateway. Routing for stub AS is commonly achieved with static routes. Stub ASs are not common now-a-days.
2) Multihomed nontransit: Large private enterprise networks, with multiple Internet access points, commonly operate as AS. they do not want to have traffic other than their own traversing their network backbone.
3) Multihomed transit: A multihomed transit network allows traffic belonging to other networks to traverse across its network in order to reach its destination.

IGP and ERP Differences

>> Context—IGPs and ERPs both operate in different routing contexts. IGP protocols are interested in building and exchanging routing information. ERPs are interested in network reachability information.

>> Connectivity—IGP networks are connected physically to one another. Routing tables are built on direct connectivity and local adjacency, using the hop-to-hop routing paradigm common to IP. ERPs exchange network reachability information. Information is exchanged by designated peers. Routes are based on AS paths that need to be traversed in order to reach the destination network.

>> Choice—IGP protocols support route metrics that enable the router to choose the best route path. ERPs do not support routing metrics. ERPs use policy routing to manage traffic behavior.

BGP (Border Gateway Protocol) Version 4

>> Most commonly used ERP.

>> Works by exchanging network reachability information with other BGP systems.

>> Based on the vector distance algorithm. Routing information is sent using the AS number as the vector and the gateway address as distance.

>> Need not be directly connected, but need reachability across the network.

>> Uses a TCP unicast (port 179), peer-to-peer-based message exchange.

>> All BGP messages use a 19-byte header.

>> There are four types of BGP messages:

1) OPEN— Contains generic information needed to establish peer-to-peer BGP communication. The fields are version, AS number, hold time, BGP identifier, and options.
2) UPDATE— Used to send routing update messages. Routing information, also known as network layer reachability information (NLRI), is sent as tuples: netmask and network address.
3) NOTIFICATION— Used to communicate errors in the peer-to-peer exchange.
4) KEEPALIVE— Sent to keep the peer-to-peer session open.

>> Unlike IGP, BGP messages contain full path of the traversal.

Configuring IP Routing Protocols on Cisco Routers

Choosing the Right Protocol

>> Selection of routing (static/dynamic) depends on the following factors;

1) Kind of network (ISP backbone, Single gateway LAN/WAN, multipoint LAN/WAN, etc)
2) Network Diameter (how many routers)
3) CIDR/VLSM support requirement
4) Need of redundant paths between network segments
5) Types of routing equipments
6) Performance factor (convergence time, etc)

Route Selection

>> Routers use reachability information from all 'available' routers to create routing table.

>> Each information source is assigned an administrative distance, by Cisco.

>> IOS administrative distances are given here.

Protocol	Distance
Connected interface	0
Static route	1
EIGRP summary route	5
BGP (external)	20
EIGRP (internal)	90
IGRP	100
OSPF	110
RIP	120
EIGRP (external)	170
BGP (internal)	200
Unknown	255

Displaying General Routing Information

>> Display commands:

<show ip route>	IP routing table
<show ip route connected>	conneted routes
<show arp>	ARP table
<show ip protocols>	IP routing protocol process parameters and statistics

>> Control commands:

<clear ip route *>	Delete all routes
<clear ip route [network] [mask]>	Destination network route to delete
<clear arp-cache>	Clear the entire ARP cache

Displaying IP Network Information

>> <show ip route> will show the following information.

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, GigabitEthernet0/0
L        10.0.0.1/32 is directly connected, GigabitEthernet0/0
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.16.0.0/16 is directly connected, GigabitEthernet0/1
L        172.16.0.2/32 is directly connected, GigabitEthernet0/1
S    192.168.0.0/24 [1/0] via 10.0.0.2
                    [1/0] via 172.16.0.1
```

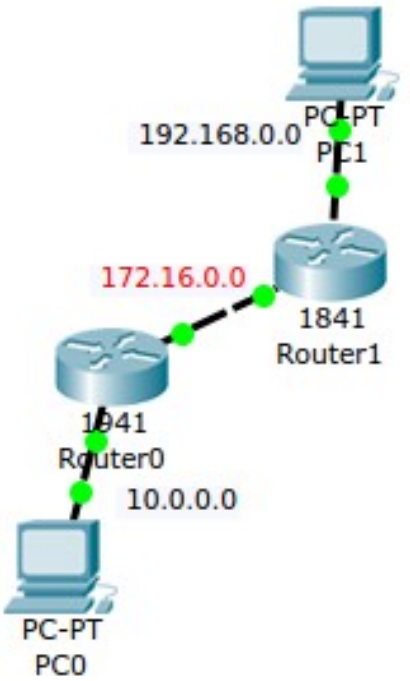

Viewing from left to right of the information;

- The source of the route (type).
- The network and netmask (displayed using bitcount).
- The route's administrative distance and routing metric.
- The network's next hop gateway.
- The gateway of last resort, if one has been set.

Displaying IP Network Information

>> <show ip route connected> displays all the network routing information about all the active router interfaces.

```
Router#show ip route connected
    C   10.0.0.0/8  is directly connected, GigabitEthernet0/0
    C   172.16.0.0/16  is directly connected, GigabitEthernet0/1
```

>> To display the ARP table, the <show arp> command is used:

```
Router#sh arp
Protocol  Address      Age (min)  Hardware Addr   Type   Interface
Internet  10.0.0.1       -   0001.64D3.CA01  ARPA   GigabitEthernet0/0
Internet  172.16.0.2     -   0001.64D3.CA02  ARPA   GigabitEthernet0/1
```

Managing Static Routing

>> Static routing is ideal for small (single gateway) networks and provides needed gateway and route redirection services.

>> Used for route advertisement in enterprise networks because it is stable and removes the possibility of IP traffic being misdirected by some runaway, misconfigured, or broken dynamic routing process.

>> Static routing tables created in one router can be distributed to other routers using <copy> and tftp. This has an advantage of keeping a copy on a server other than router.

>> <copy running-config tftp> and <copy tftp running-config> can be used to store the running- configuration to tftp server and to load the configuration to another router.

>> Classful addressing and classless addressing can be used in static routing.

>> ICMP packets are allowed by default on router interfaces; but this can make potential hazards if the routes are visible to all (such as when using the command *tracert*). Thus the ICMP service can be disabled by the global command *<no ip routing>*. Enabled by *<ip routing>*.

>> *<ip forward-protocol udp [port]>* can be used to enable IP broadcast on a particular port.

>> <ip mtu [bytes]> command is used on a particular 'interface' to set the MTU (Maximum Transfer Unit: 68 - 1500) limit of a packet. The default and maximum MTU for ethernet is 1500.

>> <bandwidth [kilobytes]> also can be set within the maximum allowed value; it is a parameter for route metrics.

Route Control and Redistribution

>> Different parts of a network may use different dynamic routing protocols. Redistribution helps to announce each other's routing information.

>> In static routing, instead of adding static routes on every router on the network, a single router can redistribute a collection of static routes.

>> Bad redistribution may results in routing loops and the difficulty in translation of different routing protocol metrics and distances.

>> To avoid ill-effects of redistribution:

1) If you have only a few networks to redistribute, use static routes or use a single protocol.
2) If possible, avoid redistributing between classless and classful protocols since classful addressing doesnot use subnet masks in their messages.
3) When using redistribution with multigateway networks, it is essential that the gateway routers use metrics and distances that favor one router over the other (ie, select a 'preferable' route).

>> Basic route redistribution is enabled with the routing protocol configuration subcommand <redistribute [source] [process id] [metrics]>

Eg: If a network segment uses RIP, redistribution can be enabled by,

Router(config)#router rip

Router(config-router)#redistribute rip metric 3

Route Control and Redistribution...

>> The metric denotes the routing parameter of the protocol (for RIP, the metric is hop count).

>> If OSPF is used, the *subnet* command also is to be used; this command is specific to OSPF.

>> To control the flow of full information in redistribution, IOS provides Route Filtering.

>> Route Filtering gives strict control over the route announcements exchanged between routers.

>> Route Filters suppress unwanted routing information from being redistributed, entered in or advertised out of the routing table.

>> Since DV routers announces their routing tables as such, the filters affects the local routing tables and all the routers which receives these tables.

>> Since LS routers keeps their routing tables with them and announces link-state information only, the filters affect the local routing tables only.