REDDIG II – Computer Networking Training
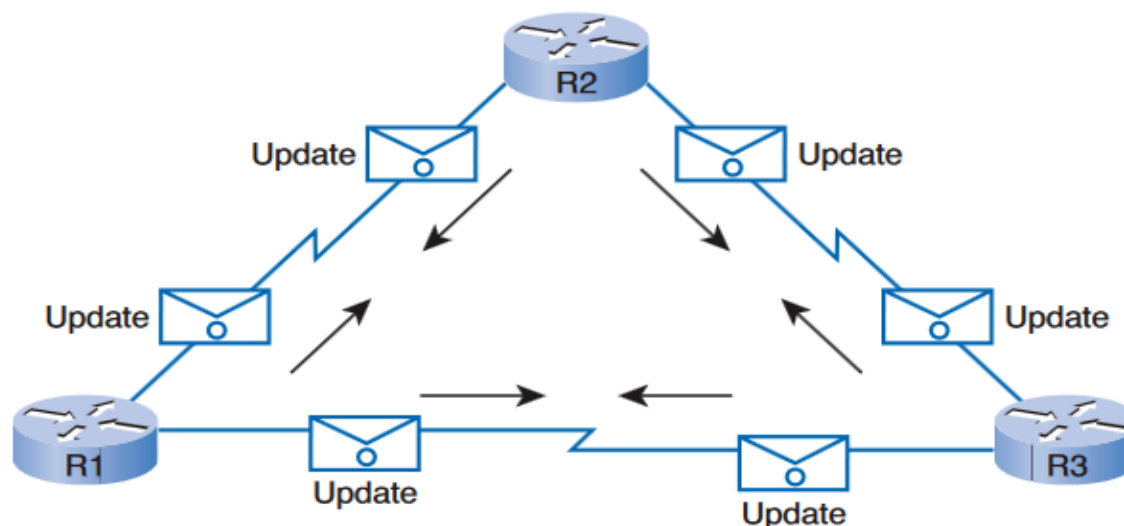
Dynamic Routing

Dynamic Routing

Function(s) of Dynamic Routing Protocols:

- Dynamically share information between routers.
- Automatically update routing table when topology changes.
- Discover remote networks
- Maintaining up-to-date routing information
- Choosing the best path to destination networks
- Ability to find a new best path if the current path is no longer available

Dynamic Routing

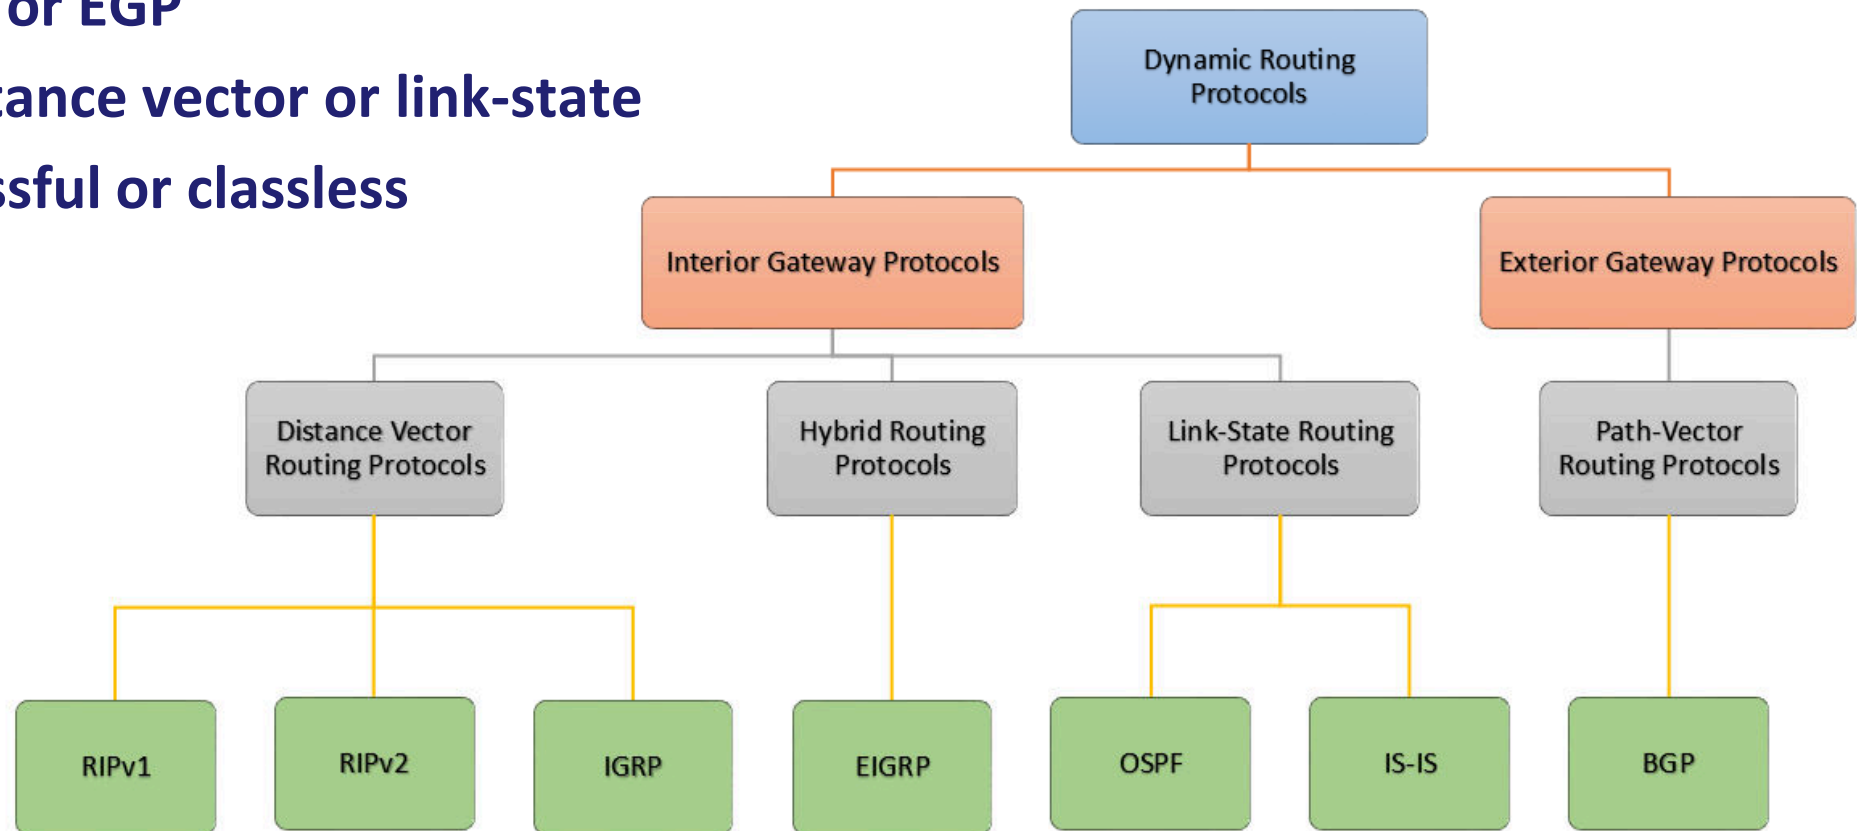**Components of Dynamic Routing Protocols:**

- **Data Structures:**
  - **Tables or databases for their operations, kept in RAM.**

- **Algorithm:**
  - **An algorithm is a finite list of steps used in accomplishing a task.**
  - **Used for processing routing information and for best path determination.**

- **Routing Protocol Messages:**
  - **Discover neighboring routers.**
  - **Exchange, learn and maintain accurate network routing information.**

Dynamic Routing
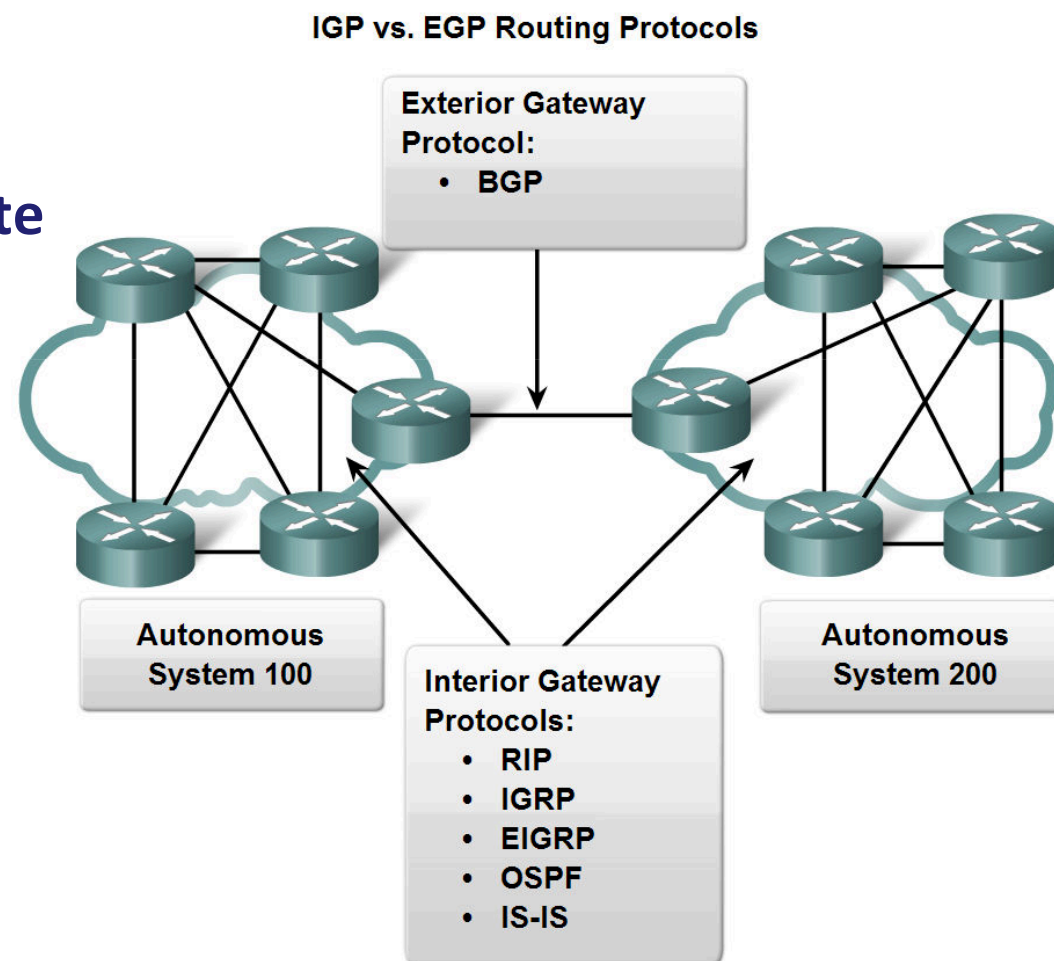
Classifying Dynamic Routing Protocols

• Routing protocols can be classified into different groups according to their characteristics:

  • IGP or EGP

  • Distance vector or link-state

  • Classful or classless

Dynamic Routing

Classifying Dynamic Routing Protocols:

- **Interior Gateway Protocols (IGP)**

  - **Used for routing inside an autonomous system & used to route within the individual networks themselves.**

  - **Examples: RIP, EIGRP, OSPF**

- **Exterior Gateway Protocols (EGP)**

  - **Used for routing between autonomous systems**

  - **Example: BGPv4**

IGP vs. EGP Routing Protocols

Exterior Gateway Protocol:
- BGP

Autonomous System 100

Autonomous System 200

Interior Gateway Protocols:
- RIP
- IGRP
- EIGRP
- OSPF
- IS-IS

Dynamic Routing

Distance Vector & Link State Routing Protocols

- Distance vector
  - East to configure
  - Routes are advertised as vectors of distance & direction.
  - Incomplete view of network topology.
  - Generally, periodic updates.
  - Prone to routing loops
  - Slow convergence
- Link state
  - More difficult to configure
  - Complete view of network topology is created.
  - Updates are not periodic.
  - Less subject to routing loops
 - Fast convergence

Dynamic Routing

Classful and Classless Routing Protocols

- **Classful routing protocols**
  - **Do not send subnet mask information in routing updates**
  - **Example: RIP**

- **Classless routing protocols**
  - **Include the subnet mask with the network address in routing updates.**
  - **Examples: RIP, EIGRP, OSPF**

Classful vs. Classless Routing
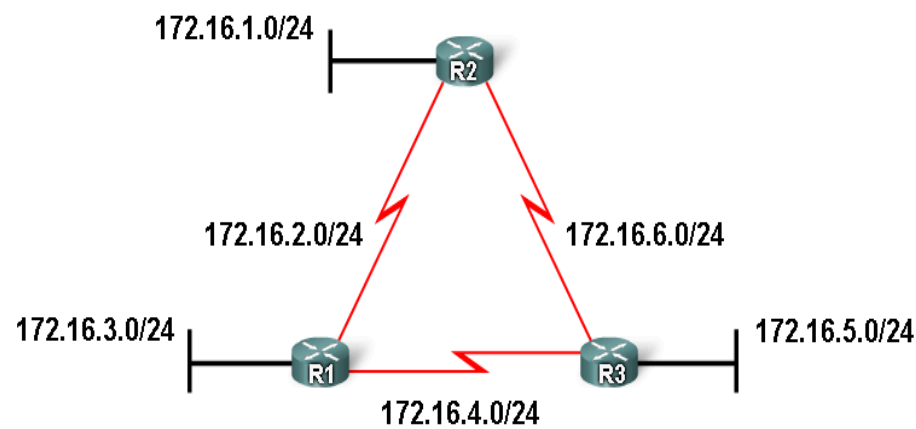
Classful: Subnet mask is the same throughout the topology
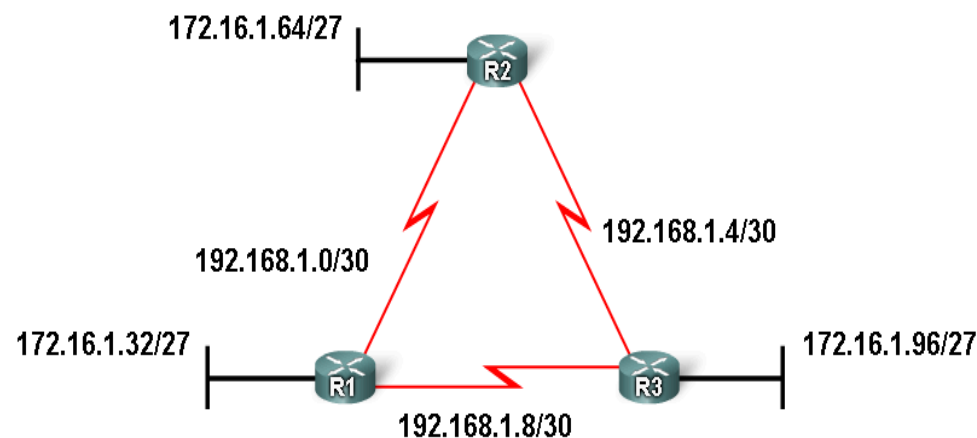
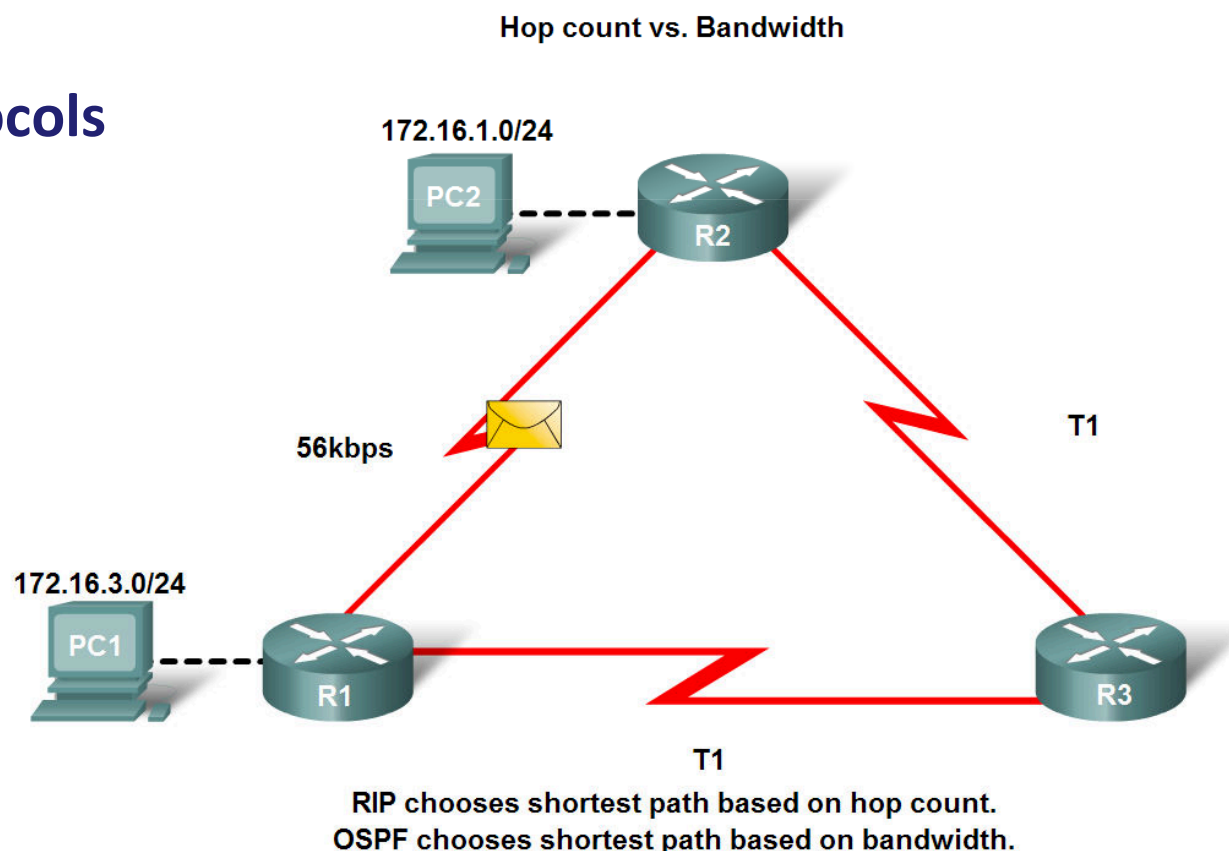Classless: Subnet mask can vary in the topology

Dynamic Routing

Routing Protocols Metrics

- **Metric**
  - A value used by a routing protocol to determine which routes are better than others.
- **Metrics used in IP routing protocols**
  - Bandwidth
  - Cost
  - Delay
  - Hop count
  - Load
  - Reliability

Hop count vs. Bandwidth

172.16.1.0/24

PC2

R2

56kbps

T1

172.16.3.0/24

PC1

R1

R3

T1

RIP chooses shortest path based on hop count.
OSPF chooses shortest path based on bandwidth.

Dynamic Routing

Load Balancing

- Individual routing protocols use metrics to determine the best route to reach remote networks.
- But what happens when two or more routes to the same destination have identical metric values? How will the router decide which path to use for packet forwarding?

  - In this case, the router does not choose only one route. Instead, the router load-balances between these equal-cost paths. The packets are forwarded using all equal-cost paths. Example:

```
R2# show ip route

<output omitted>
R    192.168.6.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
                    [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
```
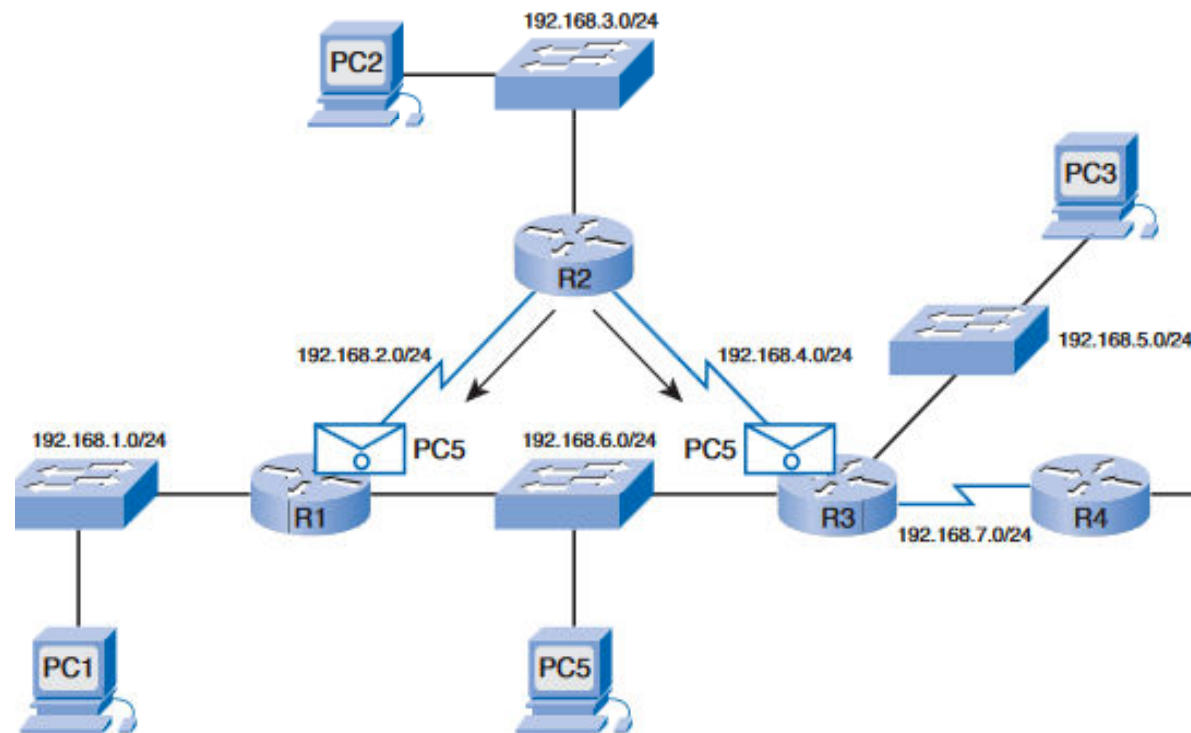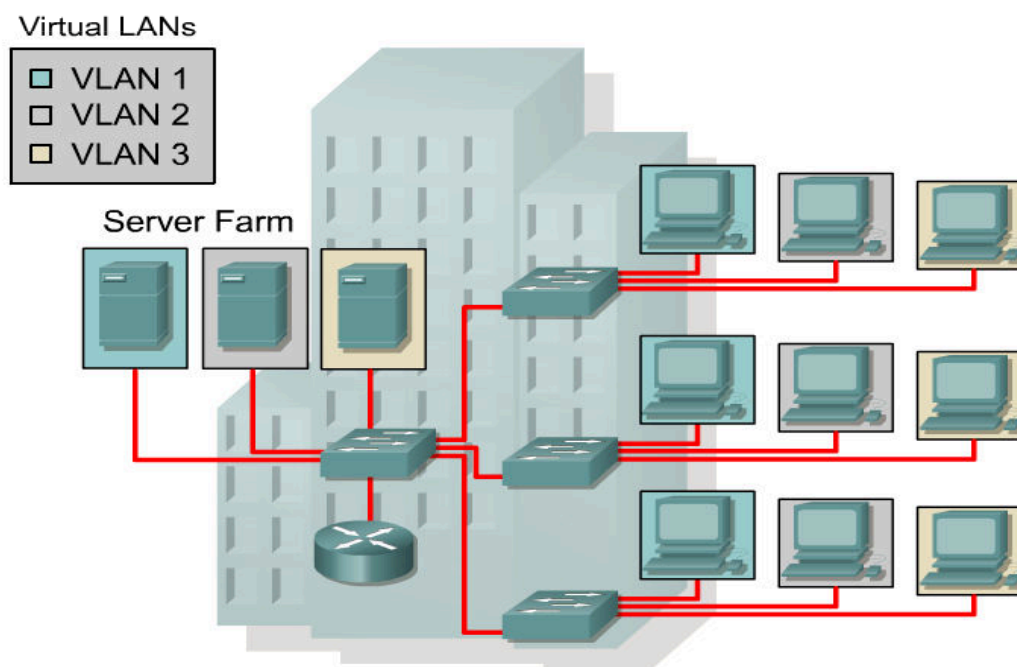
Inter-VLAN Routing

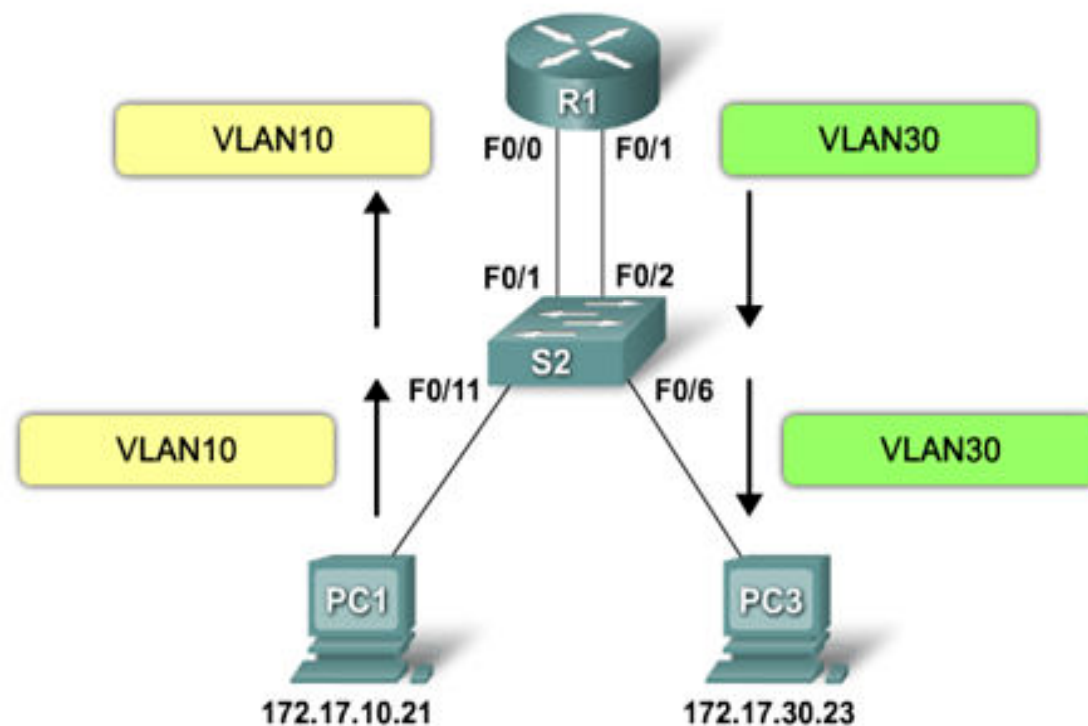# Inter-VLAN Routing

What is Inter-VLAN routing?

- VLAN is a unique broadcast domain, so, computers on separate VLANs are, by default, not able to communicate
- When a node in one VLAN needs to communicate with a node in another VLAN, a router is necessary to route the traffic between VLANs.
- Without the routing device, inter-VLAN traffic would not be possible.

Inter-VLAN Routing

Traditional Inter VLAN Routing

- **One option is to use a separate link to the router for each VLAN instead of trunk links.**

- **However, this does not scale well.**

- **Be sure hosts and routers have the proper IP addresses, associated with the proper VLANs.**

- **It is common practice to assign VLAN numbers the same as IP addresses when possible.**
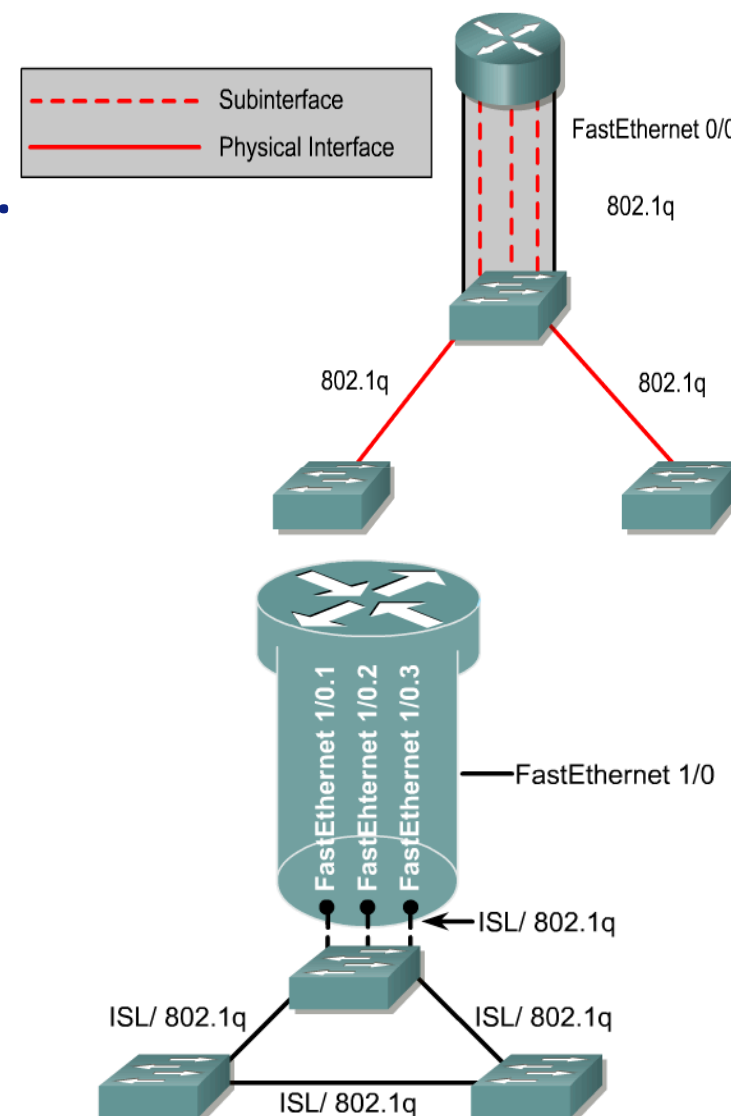
Inter-VLAN Routing

Physical and logical interfaces

- **Subinterfaces on a router can be used to divide a single physical interface into multiple logical interfaces.**

Subinterfaces:

- **Overcomes the hardware limitation of a router.**
- **Subinterfaces are software-based virtual interfaces that are assigned to physical interfaces.**
- **Each subinterface is configured with its own IP address, subnet mask, and unique VLAN assignment.**
- **Connected to a switch trunk link.**
- **Command:**

Rtr(config)#interface fastethernet port.subinterface

Inter-VLAN Routing

Router on a stick Inter VLAN Routing

- Router-on-a-stick is a type of router configuration in which a single physical interface routes traffic between multiple VLANs on a network.

- The router interface is configured to operate as a trunk link and is connected to a switch port configured in trunk mode.

- The router performs the inter-VLAN routing by accepting VLAN tagged traffic on the trunk interface coming from the adjacent switch and internally routing between the VLANs using sub interfaces.

Subinterfaces

F0/0.10: 172.17.10.1/24
F0/0.30: 172.17.30.1/24
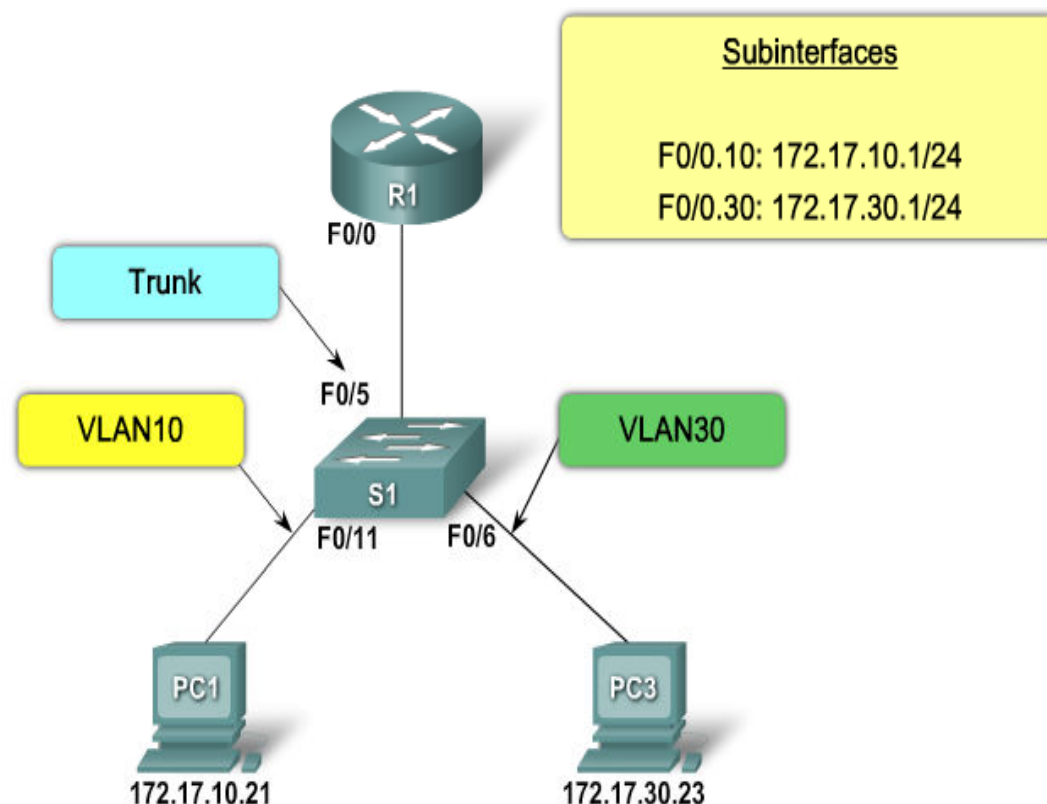
Inter-VLAN Routing

Configuring subinterfaces in a Router

10.10.0.11/16 VLAN 10

VLAN 20 10.20.0.22/16

FastEthernet 0/1

10.1.0.1/16
10.10.0.1/16
10.20.0.1/16
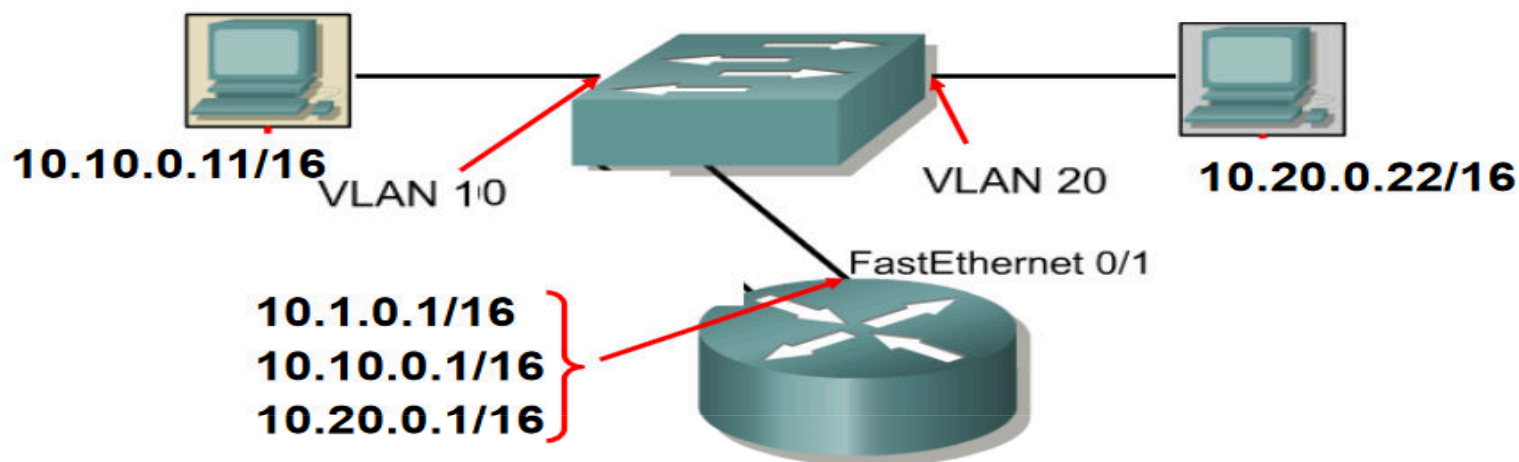
Rtr(config)#interface fastethernet 0/1.10

Rtr(config-subif)#description Sales_VLAN 10

Rtr(config-subif)#encapsulation dot1q 10

Rtr(config-subif)#ip address 10.10.0.1 255.255.0.0

Rtr(config)#interface fastethernet 0/1.20

Rtr(config-subif)#description Finance_VLAN 20

Rtr(config-subif)#encapsulation dot1q 20

Rtr(config-subif)#ip address 10.20.0.1 255.255.0.0

VRRP and IP-SLA
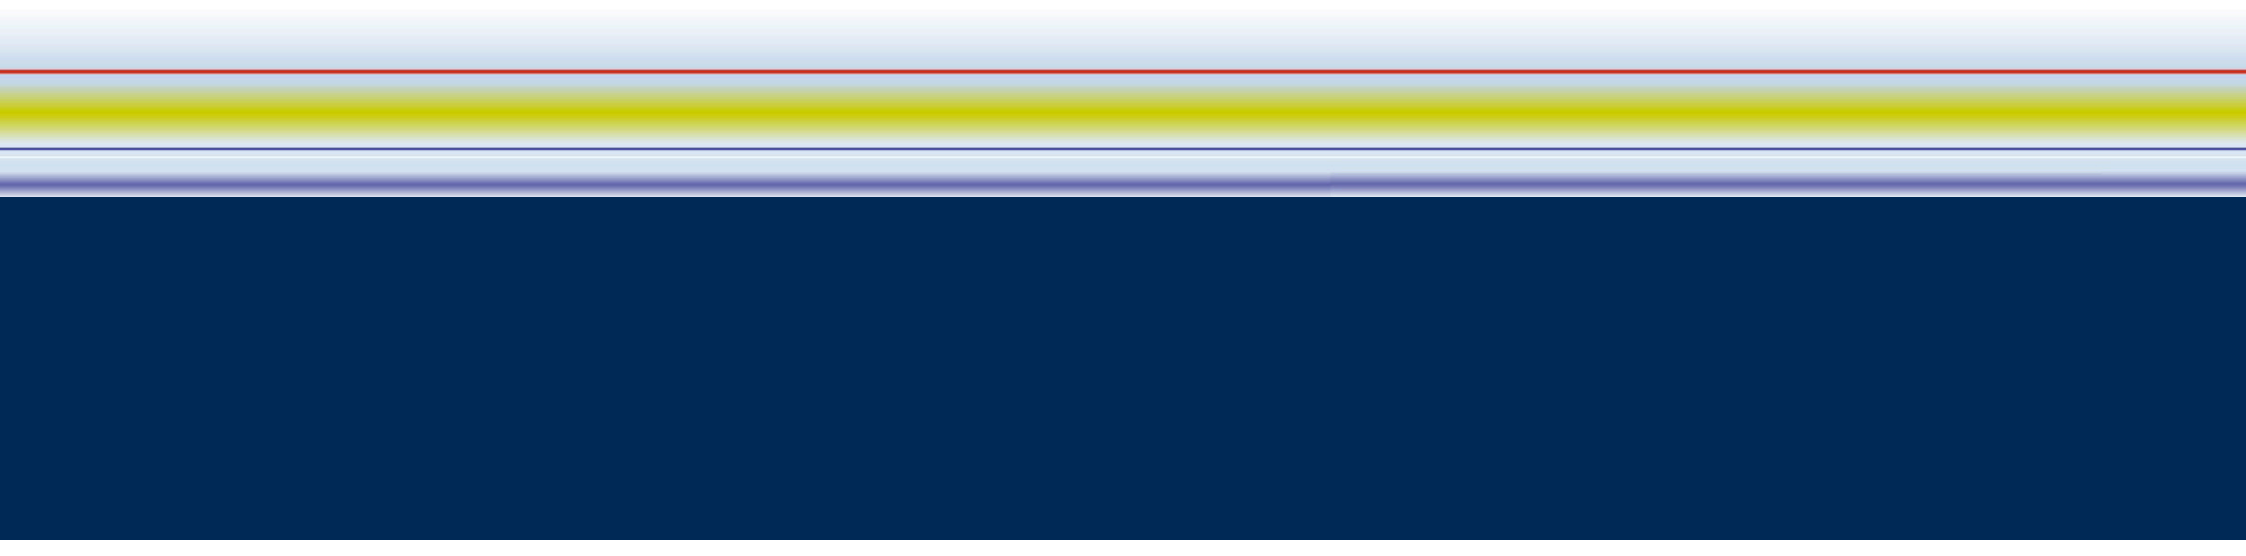
VRRP and IP-SLA

Network Redundancy

- **Networks need redundant links to improve the availability of the network**

- **Failure of a single component might mean an outage that affects at least some part of the network**

- **Single Point of Failure is any one component that if it fails, brings down that part of the network**

- **High Availability: If one of the routers fails the other one will continue to provide services to the segment**
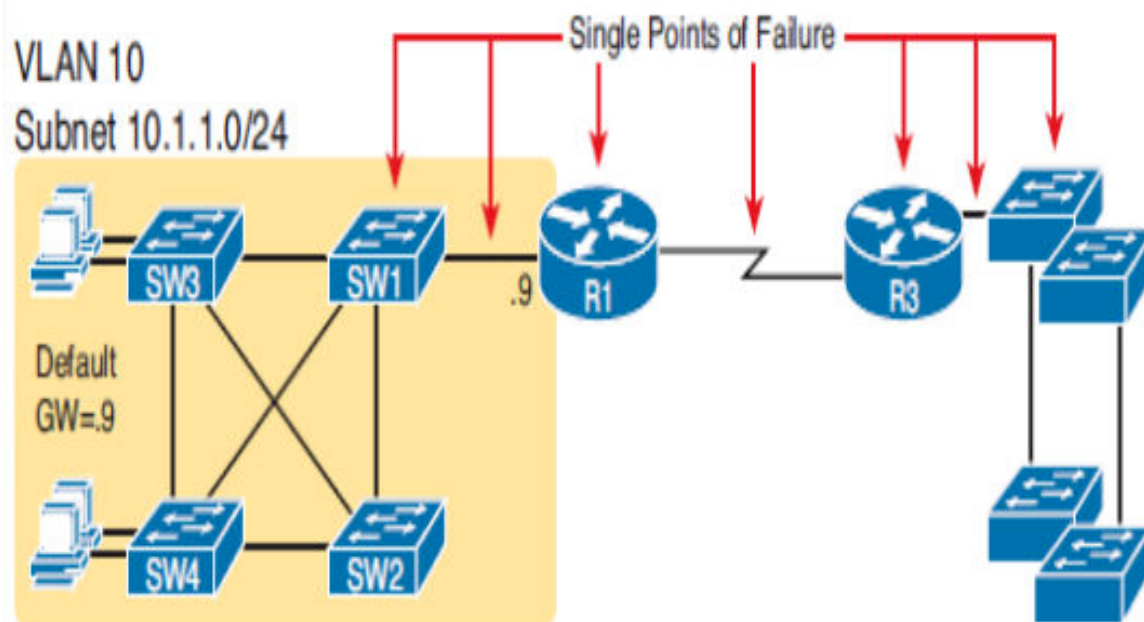
VRRP and IP-SLA

First Hop Redundancy Protocol (FHRP)

- Host configured with default gateway (router) as the IP address of the LAN interface of the router

- If the router fails, so host cannot reach destinations outside the local network

- Using FHRP, two redundant routers appear as single default router

- Routers exchange FHRP protocol messages

- The default routers share a virtual IP address in the subnet, defined by the FHRP

- HSRP, VRRP and GLBP are three FHRP
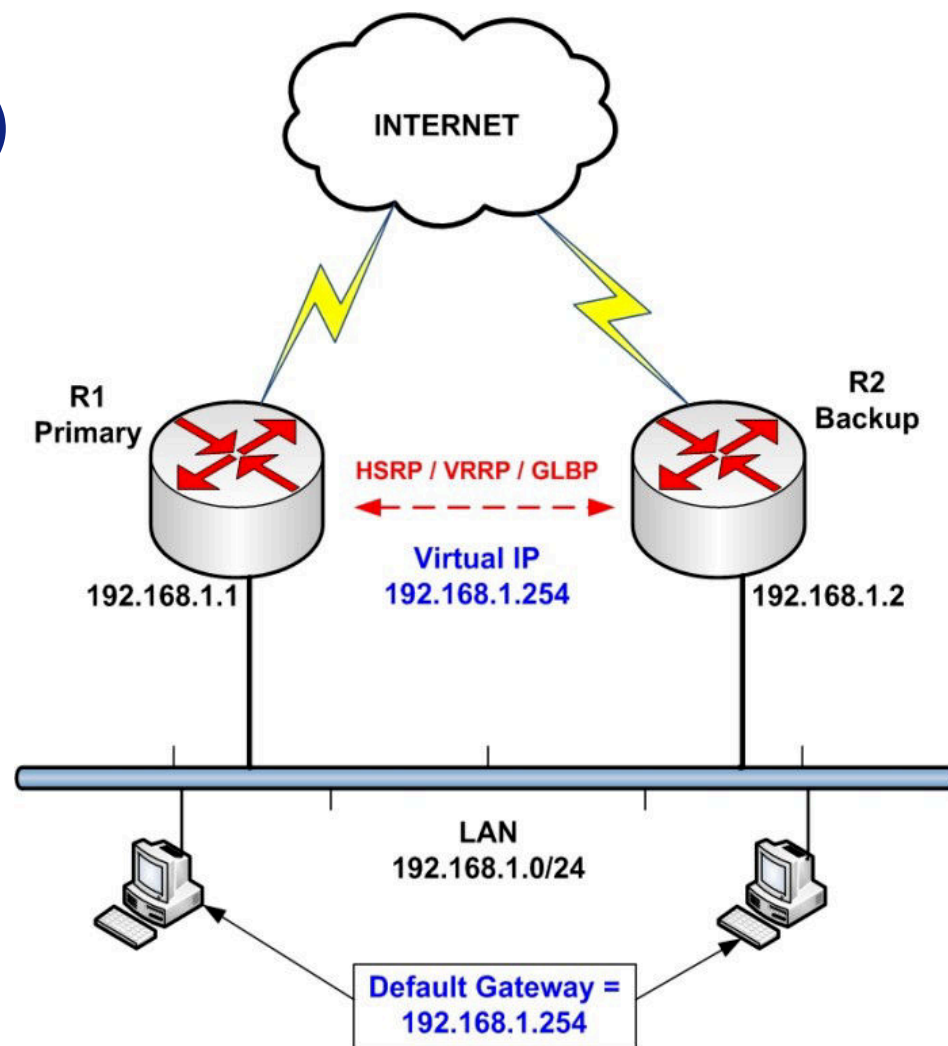
VRRP and IP-SLA

Virtual Router Redundancy Protocol (VRRP)

•VRRP group: It consists of one master router and one or more backup routers.

•Virtual router master: VRRP router that forwards packets.

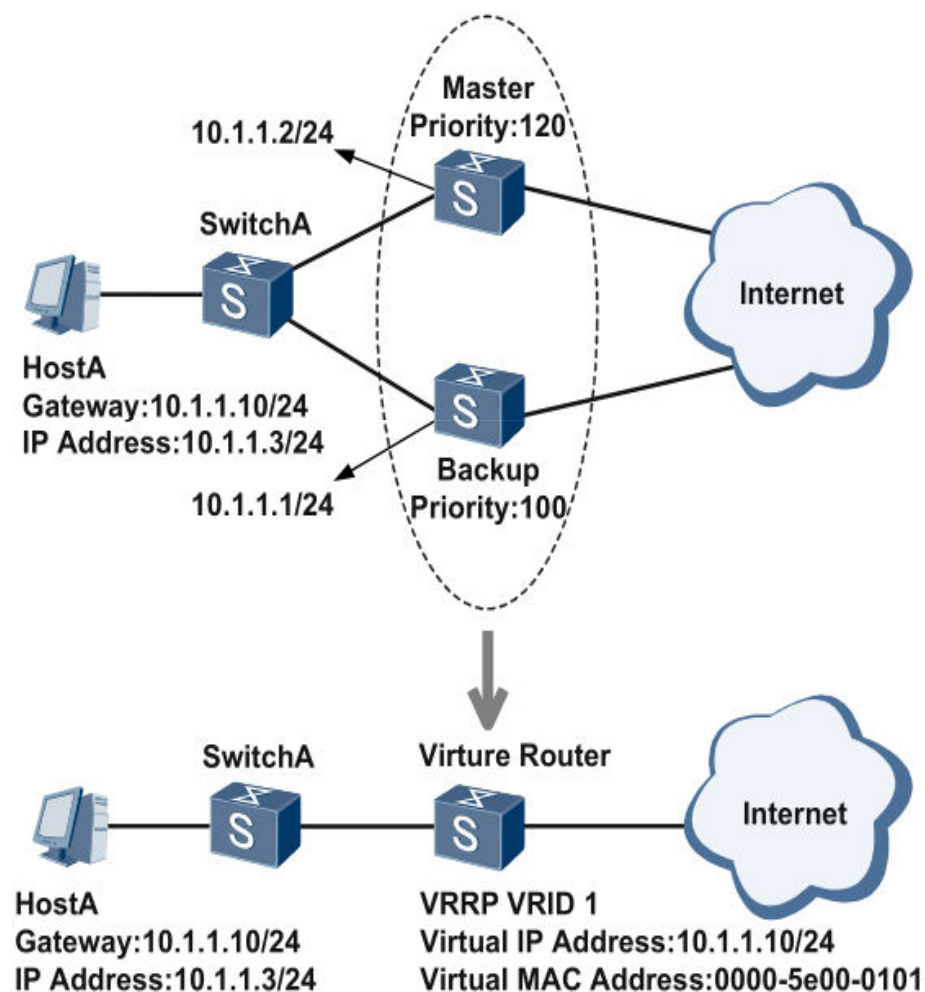•Virtual router backup: a group of VRRP routers that do not forward packets. When the master router is faulty, a backup router preempts to be the new master router.

•VRID: virtual router ID.

•Virtual IP address: IP address of a virtual router. Virtual IP addresses are configurable.

•Virtual MAC address: MAC address that is generated by the virtual router based on the virtual router ID. A virtual router has one virtual MAC address and is in the format of 00-00-5E-00-01-{VRID}

VRRP - Selecting the Master

- **Priority: priority of a VRRP router. The virtual router selects the master and backup routers based on the priority.**
- **Preemption mode: If the priority of a virtual router backup is higher than the priority of the current virtual router master, the virtual router backup automatically becomes the virtual router master.**
- **The Hellos in VRRP are sent to 224.0.0.18 multicast address.**
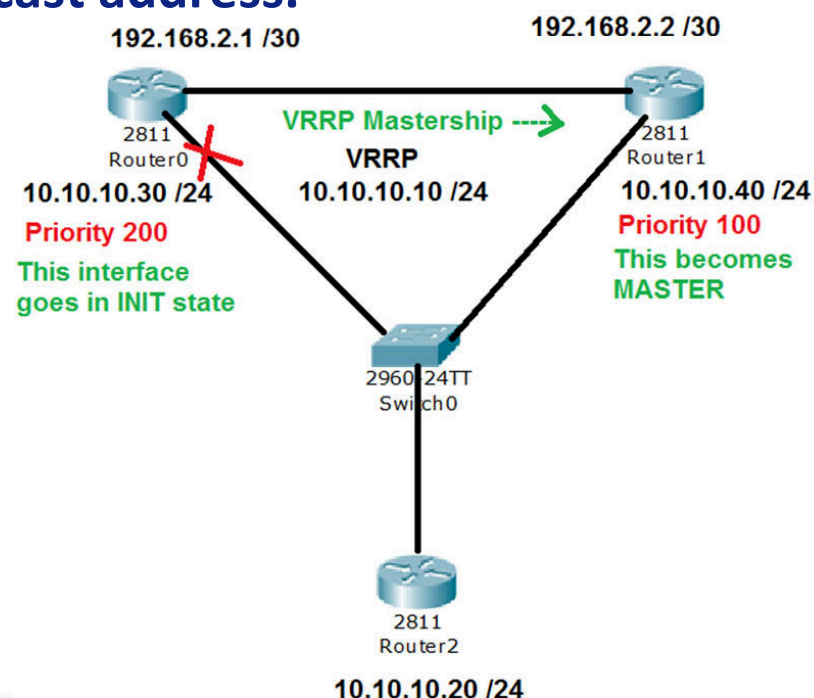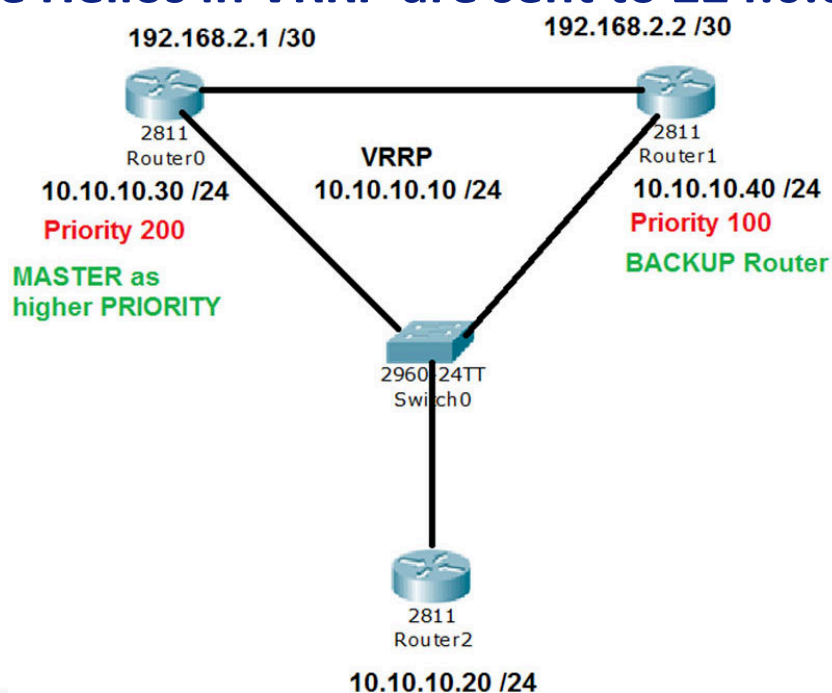
VRRP and IP-SLA

Cisco IP-SLA

- **Cisco IOS IP Service Level Agreements (SLAs) uses active traffic monitoring for measuring network performance.**
- **Cisco IOS IP SLAs send simulated data across the network and measure performance between network locations. One-way travel times and packet loss are gathered.**
- **IP SLAs collects a unique subset of these performance metrics:**

  - **Delay (both round-trip and one-way)**

  - **Jitter (directional)**

  - **Packet loss (directional)**

  - **Packet sequencing (packet ordering)**

  - **Path (per hop)**

  - **Connectivity (directional)**

  - **Server or website download time**

VRRP and IP-SLA

IP SLA ICMP Echo

- The ICMP Echo IP SLA Operation measures end-to-end response time between a Cisco router and any device with an IP Address.
- The response time is computed by measuring the time taken between sending an ICMP Echo request and receiving an Echo reply.
- Parameters:

```
Router(config-ip-sla)#

icmp-echo {destination-ip-address | destination-hostname} [source-
    ip {ip-address | hostname} | source-interface interface-name]
```

Parameter	Description
destination-ip-address | destination-hostname	Destination IPv4 or IPv6 address or hostname.
source-ip {ip-address | hostname}	(Optional) Specifies the source IPv4 or IPv6 address or hostname. When a source IP address or hostname is not specified, IP SLAs chooses the IP address nearest to the destination.
source-interface interface-name	(Optional) Specifies the source interface for the operation.

VRRP and IP-SLA

Configuring IP SLA ICMP Echo

➢ **Step 1. Configure IP SLA probe**

➢ **Step 2. Activate probe**

➢ **Step 3. Configure tracking object**

➢ **Step 4. Configure action on tracking object**

• **The first step is to use the command ip sla monitor followed by a number to enter in IP SLA configuration mode. The number identifies the SLA test.**

• **Example: The IP SLA test is done by sending an ipIcmpEcho message to the IP address 10.1.1.1 from the local interface Fa0/1 every 10 seconds Fa0/1 every 10 seconds**


```
SwitchB(config)# ip sla monitor 11
SwitchB(config-sla)# type echo protocol ipIcmpEcho 10.1.1.1 source-
int fa0/1
SwitchB(config-sla)# frequency 10
SwitchB(config-sla)# exit
SwitchB(config)# ip sla monitor schedule 11 life forever start-time
now
SwitchB(config)# track 1 ip sla 11 reachability
```

VRRP and IP-SLA

REDDIG – VRRP Configuration with IP SLA tracks

```
SBCT-CISCO-VSAT-1-A-V15#sh vrrp brief
Interface              Grp Pri Time  Own Pre State   Master addr      Group addr
Gi0/0                  30  0   4000          Y  Backup  10.100.30.102   10.100.30.254

SBCT-CISCO-VSAT-1-B-V15#sh vrrp brief
Interface              Grp Pri Time  Own Pre State   Master addr      Group addr
Gi0/0                  30  50  3804          Y  Master  10.100.30.102   10.100.30.254
Gi0/1                  120 50  3804          Y  Master  10.120.30.102   10.120.30.254
Gi0/1.100              100 50  3804          Y  Master  10.0.80.252     10.0.80.254
Gi0/1.101              101 50  3804          Y  Master  10.0.81.252     10.0.81.254
Gi0/1.102              102 50  3804          Y  Master  10.0.82.252     10.0.82.254
|
SBCT-CISCO-VSAT-1-B-V15#sh vrrp interface g0/0
GigabitEthernet0/0 - Group 30
    State is Master
    Virtual IP address is 10.100.30.254
    Virtual MAC address is 0000.5e00.011e
    Advertisement interval is 1.000 sec
    Preemption enabled
    Priority is 50
       Track object 1 state Up decrement 100
    Authentication text, string "SBCTVSAT"
    Master Router is 10.100.30.102 (local), priority is 50
    Master Advertisement interval is 1.000 sec
    Master Down interval is 3.804 sec
```

VRRP MAC address: 00-00-5E-00-01-{VRID} (in hex in internet standard bit-order)
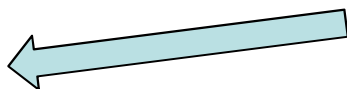
OSPF Routing Protocol

OSPF Routing Protocol

- **OSPF (Open Shortest Path First) is a standardized Link-State routing protocol, designed to scale efficiently to support larger networks.**
  RIP, IGRP and EIGRP are distance-vector (routing by rumor) routing protocols, susceptible to routing loops, split-horizon, and other issues.

- **OSPF has fast convergence**
  RIP and IGRP hold-down timers can cause slow convergence.

- **OSPF supports VLSM and CIDR**
  RIPv1 and IGRP do not

  - **Cisco's OSPF metric is based on bandwidt**
  RIP is based on hop count

- **OSPF only sends out changes when they occur.**
  RIP sends entire routing table every 30 seconds, IGRP every 90 seconds

  - **OSPF also uses the concept of areas to implement hierarchical routing**

OSPF Routing Protocol

OSPF characteristics

- OSPF employs a hierarchical network design using <u>Areas</u>.

- OSPF will form <u>neighbor</u> relationships with <u>adjacent</u> routers in the same Area.

- Instead of advertising the distance to connected networks, OSPF <u>advertises</u> the <u>status</u> of directly connected links using <u>Link-State Advertisements (</u>LSAs).

- OSPF sends updates (LSAs) when there is a change to one of its links, and will <u>only send the change in the update</u>. LSAs are additionally refreshed every 30 minutes.

- OSPF traffic is <u>multicast</u> either to address <u>224.0.0.5</u> (all OSPF routers) or 224.0.0.6 (all Designated Routers).

- OSPF uses the <u>Dijkstra</u> Shortest Path First (SPF) algorithm to determine the shortest path.

- OSPF is a <u>classless</u> protocol, and thus supports <u>VLSMs</u>.

- OSPF uses <u>cost</u> as its <u>metric</u>. OSPF routes have an <u>administrative distance is 110</u>.

OSPF Routing Protocol

OSPF Algorithm

OSPF routers build & maintain link-state database containing LSA received from other routers.

- **Information found in database is utilized upon execution of Dijkstra SPF algorithm**
- **SPF algorithm used to create SPF tree**
- **SPF tree used to populate routing table**
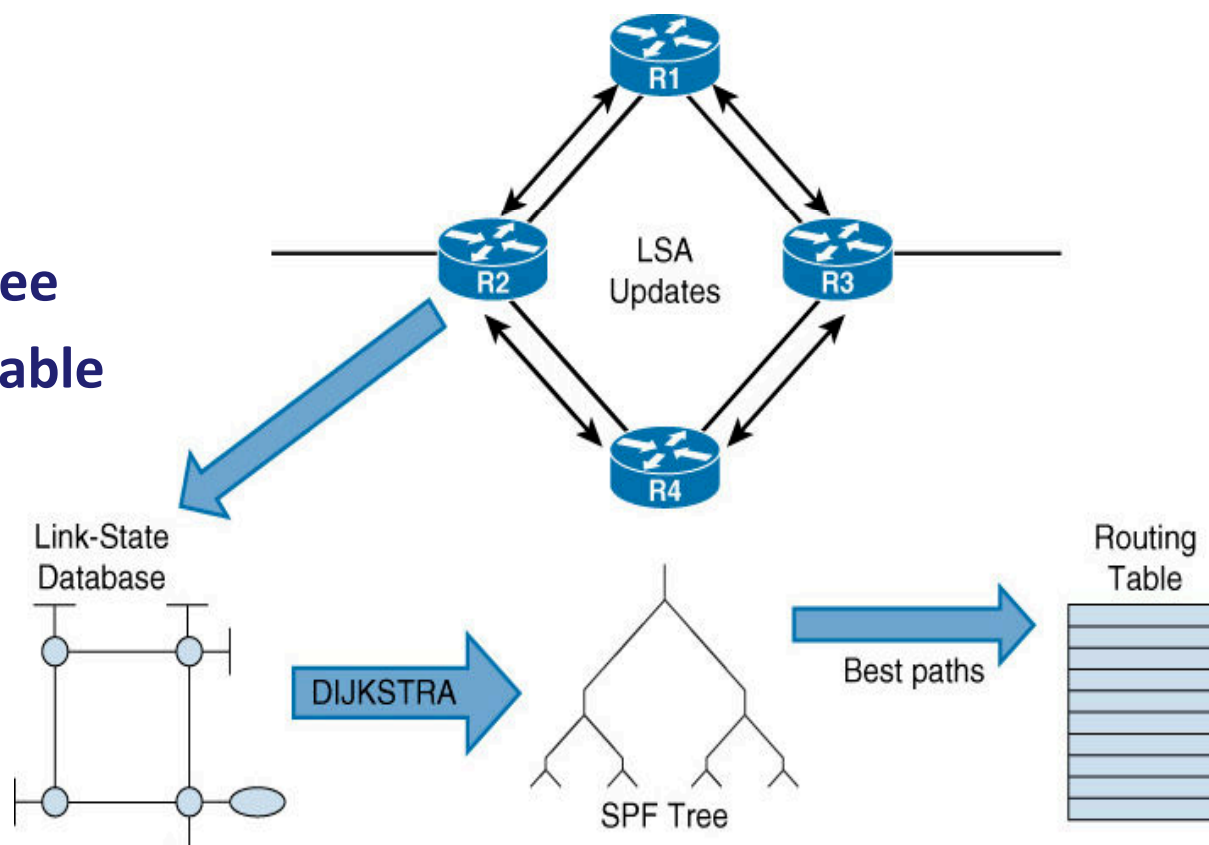
OSPF Routing Protocol

OSPF Tables

The OSPF process builds and maintains three separate tables:

- A neighbor table – contains a list of all neighboring routers.

Command to see: #sh ip ospf neighboor

- A topology table – contains a list of all possible routes to all known networks within an area.

Command to see: sh ip ospf topology

- A routing table – contains the best route for each known network.

Command to see: #sh ip route ospf

OSPF Routing Protocol

OSPF Metric

- **OSPF determines the best (or shortest) path to a destination network using a cost metric, which is based on the bandwidth of interfaces. The total cost of a route is the sum of all outgoing interface costs. Lowest cost is preferred.**

- Cost is based on bandwidth of an interface
  - Cost is calculated using the formula
        10^8 / bandwidth
- Reference bandwidth:
  - defaults to 100Mbps
  - can be modified using auto-cost reference-bandwidth command

  - Fast Ethernet Cost = 1

OSPF Accumulates Cost
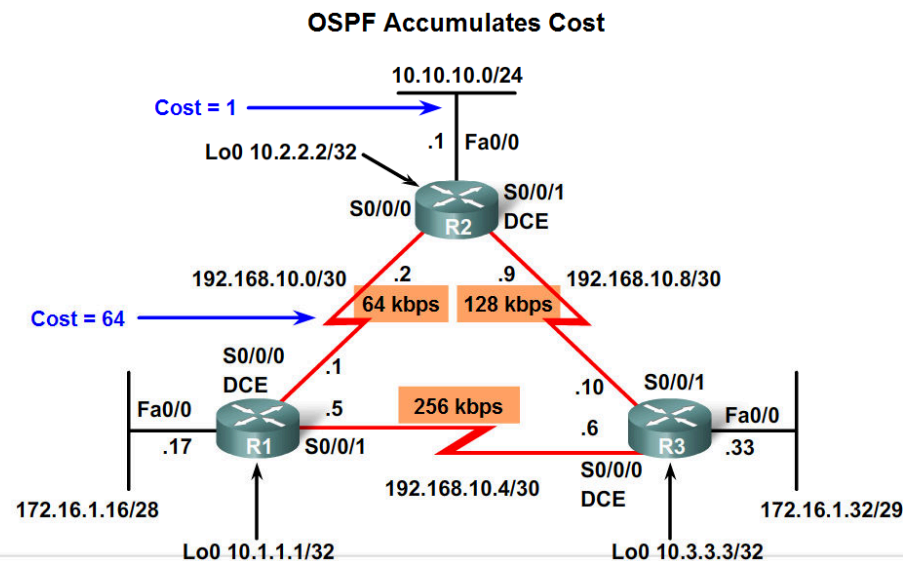

```
R1#show ip route
Codes: <some code output omitted>
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

<route ouput omitted.
O        10.10.10.0/24 [110/65] via 192.168.10.2, 14:27:57, Serial0/0/0
```

Accumulated Cost = 65

OSPF Routing Protocol

OSPF Network Types

OSPF interfaces automatically recognize three types of networks
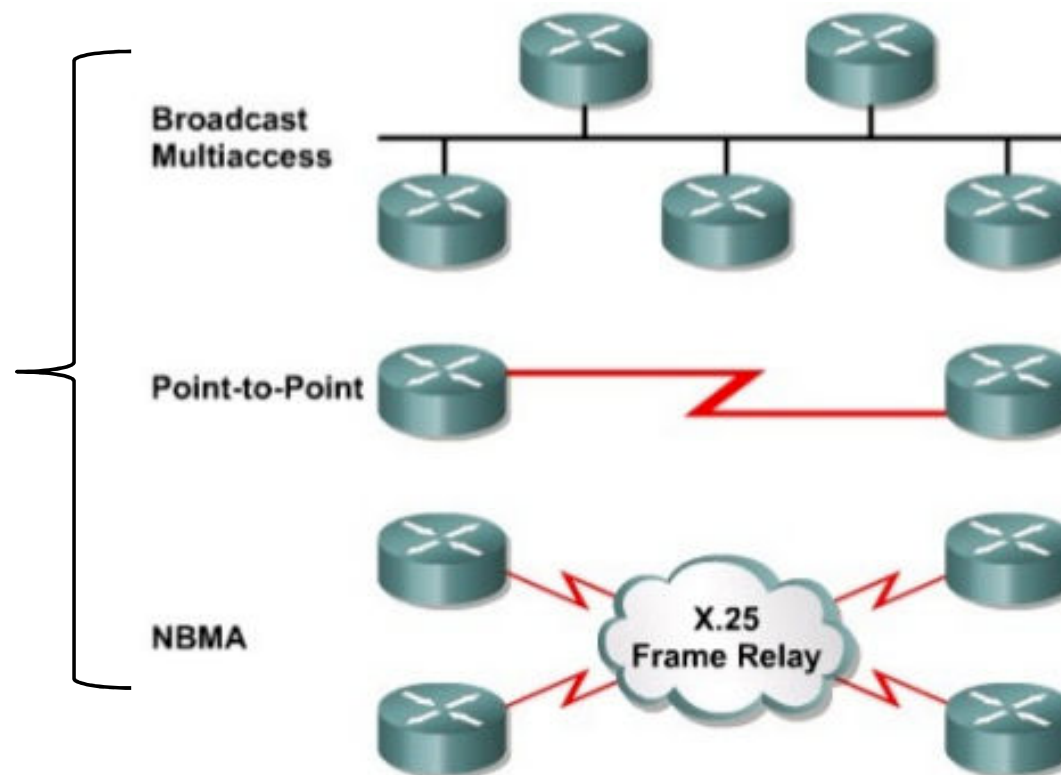
Network Type	Characteristics	DR election?
Broadcast multiaccess	Ethernet, Token Ring, or FDDI	Yes
Nonbroadcast multiaccess	Frame Relay, X.25, SMDS	Yes
Point-to-point	PPP, HDLC	No
Point-to-multipoint	Configured by an administrator	No

OSPF Routing Protocol

OSPF Neighbors

• **OSPF forms neighbor relationships, called adjacencies, with other routers in the same area by exchanging Hello packets to multicast address 224.0.0.5. Only after an adjacency is formed can routers share routing information.**

• **By default, Hello packets are sent out OSPF-enabled interfaces every 10 seconds.**

• **OSPF routers will only become neighbors if the following parameters within a Hello packet are identical on each router:**
  - **Area ID**
  - **Area Type (stub, NSSA, etc.)**
  - **Prefix**
  - **Subnet Mask**
  - **Hello Interval**
  - **Dead Interval**
  - **Network Type (broadcast, point-to-point, etc.)**
  - **Authentication**
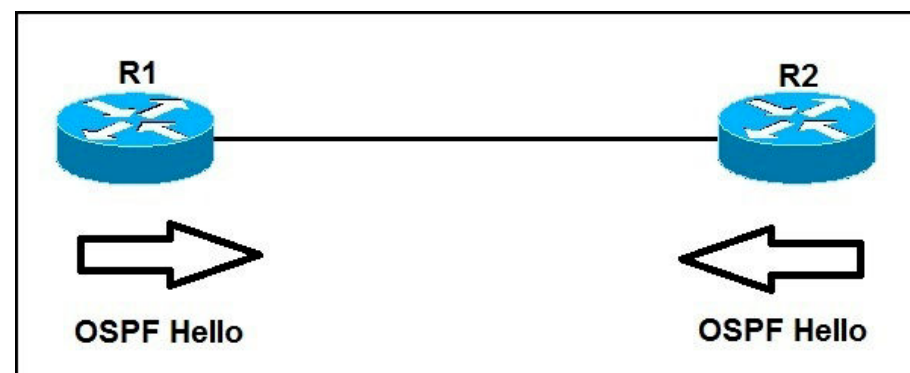
OSPF Routing Protocol

OSPF Neighboor Table and Router ID

- A neighbor table is constructed from the OSPF Hello packets, which includes the following information:
  - The Router ID of each neighboring router
  - The current "state" of each neighboring router
  - The interface directly connecting to each neighbor
  - The IP address of the remote interface of each neighbor

The Router ID can be determined in one of three ways:
  - The Router ID can be manually specified.
  - If not manually specified, the highest IP address configured on any Loopback interface on the router will become the Router ID.
  - If no loopback interface exists, the highest IP address configured on any Physical interface will become the Router ID

OSPF Routing Protocol

OSPF in Multiaccess Networks

- **2 challenges presented by multiaccess networks:**
  - **Multiple adjacencies**
  - **Extensive LSA flooding**

Number of Adjacencies Grows Exponentially
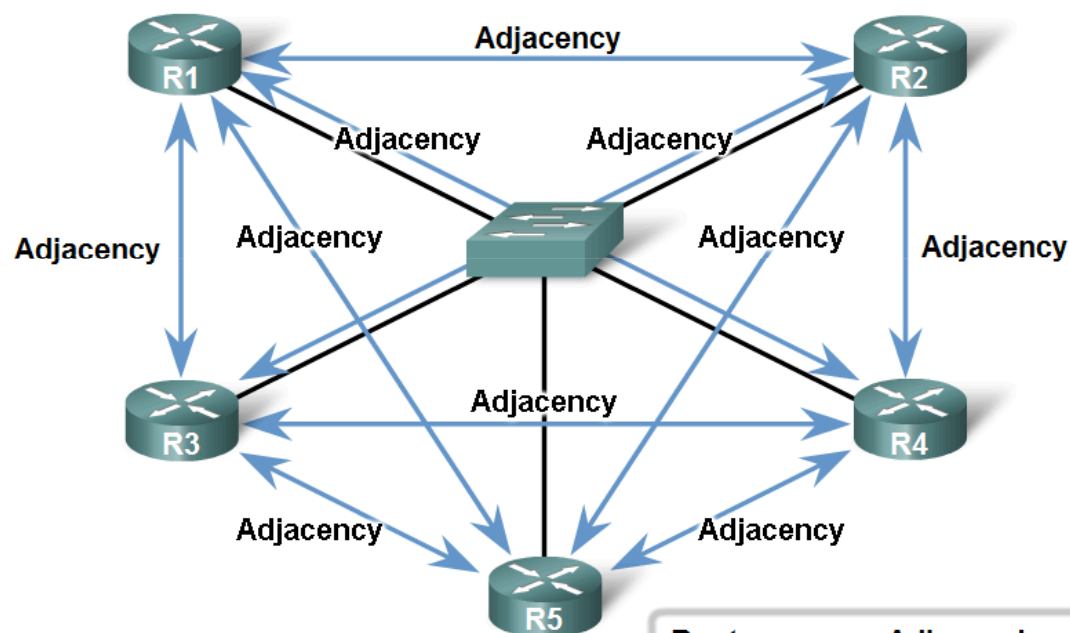
Number of Adjacencies = n(n-1)/2
n = number of routers
Example: 5 routers (5 - 1)/2 = 10 adjacencies

Routers	Adjacencies
n	n(n-1)/2
5	10
10	45
20	190
100	4,950

OSPF Routing Protocol

OSPF in Multiaccess Networks

- **Solution to LSA flooding issue is the use of:**
  - Designated router (DR)
  - Backup designated router (BDR)

- **DR & BDR selection:**
  - Routers are elected to send & receive LSA

- **Sending & Receiving LSA:**
  - DRothers send LSAs via multicast 224.0.0.6 to DR & BDR
  - DR forward LSA via multicast address 224.0.0.5 to all other routers

Adjacencies are formed with DR and BDR only.
LSAs are sent to the DR. BDR listens.

Here are my LSAs
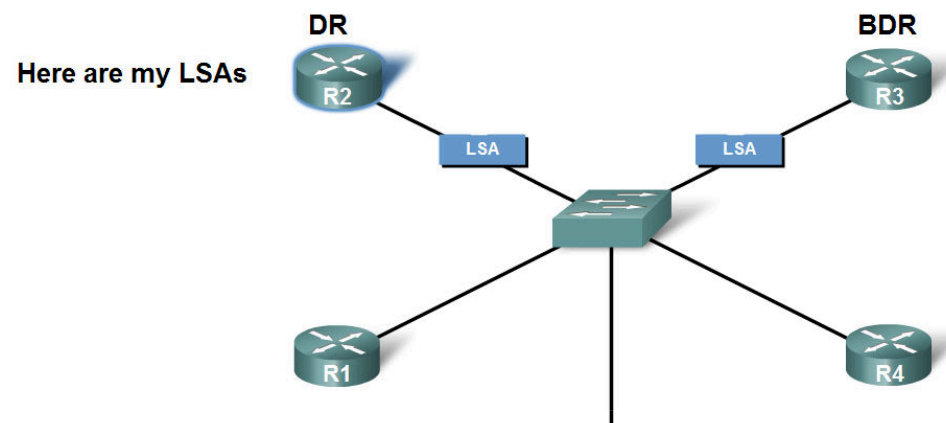
DR sends out any LSAs to all other routers.

Here are 10.1.1.1's LSAs
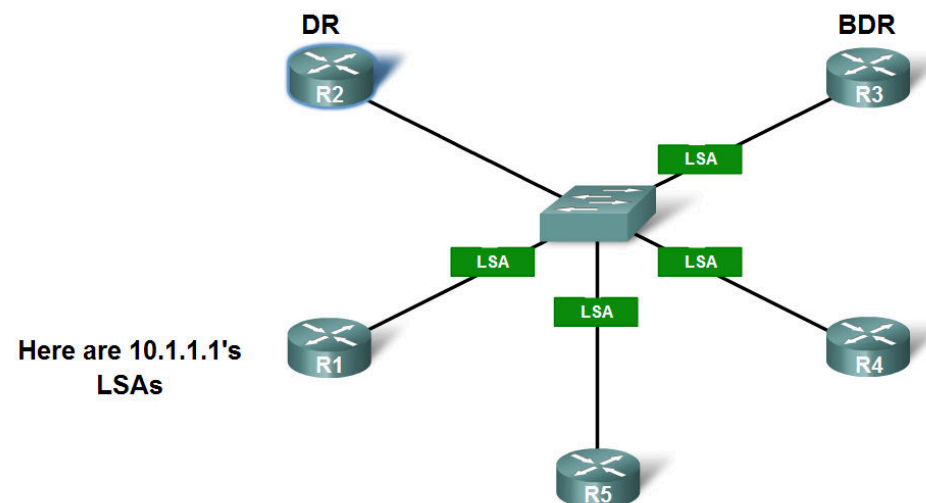
Electing the DR and BDR

- **Router with the highest Router ID is elected the DR, next is BDR.**

- **But like other elections, this one can be rigged. The router's priority field can be set to either ensure that It becomes the DR or prevent it from being the DR.**

Rtr(config-if)# ip ospf priority <0-255>
      Higher priority becomes DR/BDR
      Default = 1
      0 = Ineligible to become DR/BDR

The router can be assigned a priority between 0 and 255, with 0 preventing this router from becoming the DR (or BDR) and 255 ensuring at least a tie. (The highest Router ID would break the tie.)

Next highest router ID → Lo0 192.168.31.22/32

Highest router ID

RouterB — BDR

Fa0/0 192.168.1.2./24

Lo0 192.168.31.33/32

Lo0 192.168.31.11/32

DROther — RouterA

RouterC — DR

Fa0/0 192.168.1.1./24

Fa0/0 192.168.1.3./24
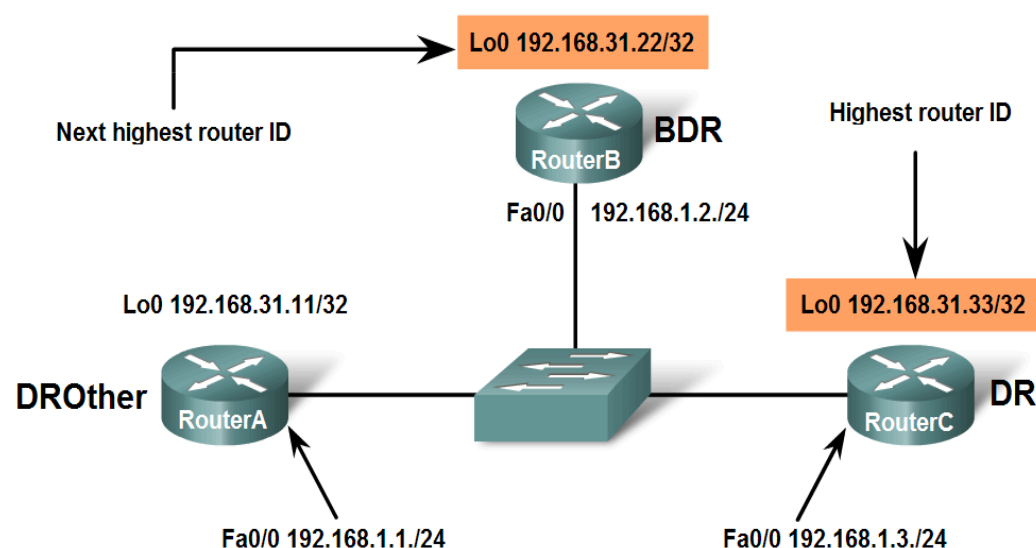
OSPF Routing Protocol

OSPF Neighbor States

- **Down – indicates that no Hellos have been heard from the neighboring router.**
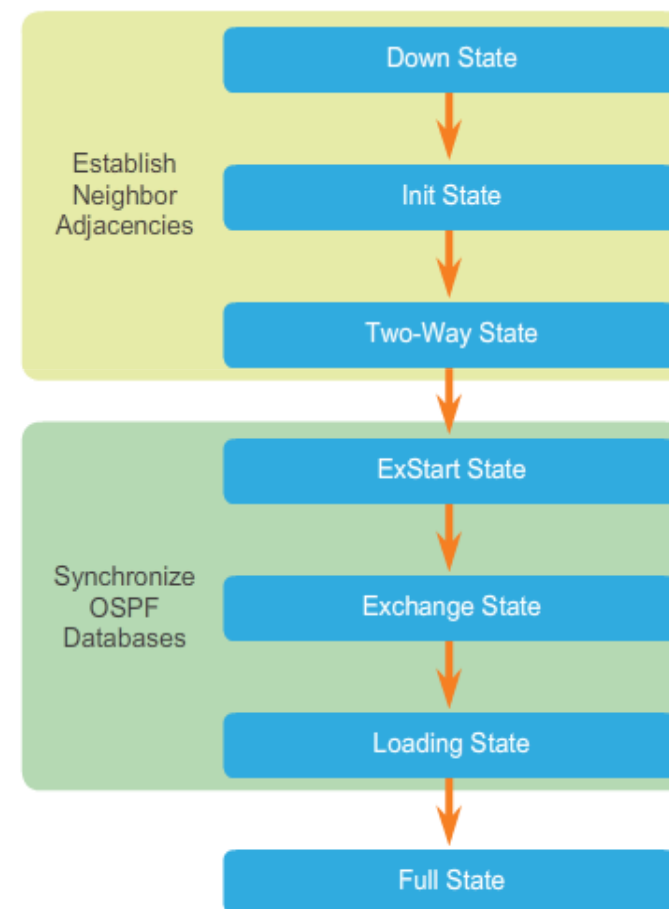
- **Init – indicates a Hello packet has been heard from the neighbor, but two-way communication has not yet been initialized.**

- **2-Way – indicates that bidirectional communication has been established.**

- **ExStart – indicates that the routers are preparing to share link state information.**

- **Exchange – indicates that the routers are exchanging Database Descriptors (DBDs).**

- **Loading – indicates the routers are finally exchanging Link State. Essentially, routers are sharing their topology tables with each other.**

- **Full – indicates that the routers are fully synchronized.**

Establish Neighbor Adjacencies
- Down State
- Init State
- Two-Way State

Synchronize OSPF Databases
- ExStart State
- Exchange State
- Loading State
- Full State

OSPF Routing Protocol

OSPF Areas

- **Group of contiguous nodes/networks**
- **Per area topology database**
  - **Invisible outside the area**
  - **Reduces routing traffic**

- **Area 0 is required for OSPF to function, and is considered the "Backbone"**

- **Backbone Area is contiguous**
  - **All others areas must connect to the backbone**

Area 2

Area 3

Area 0
Backbone Area

Area 4

Area 1
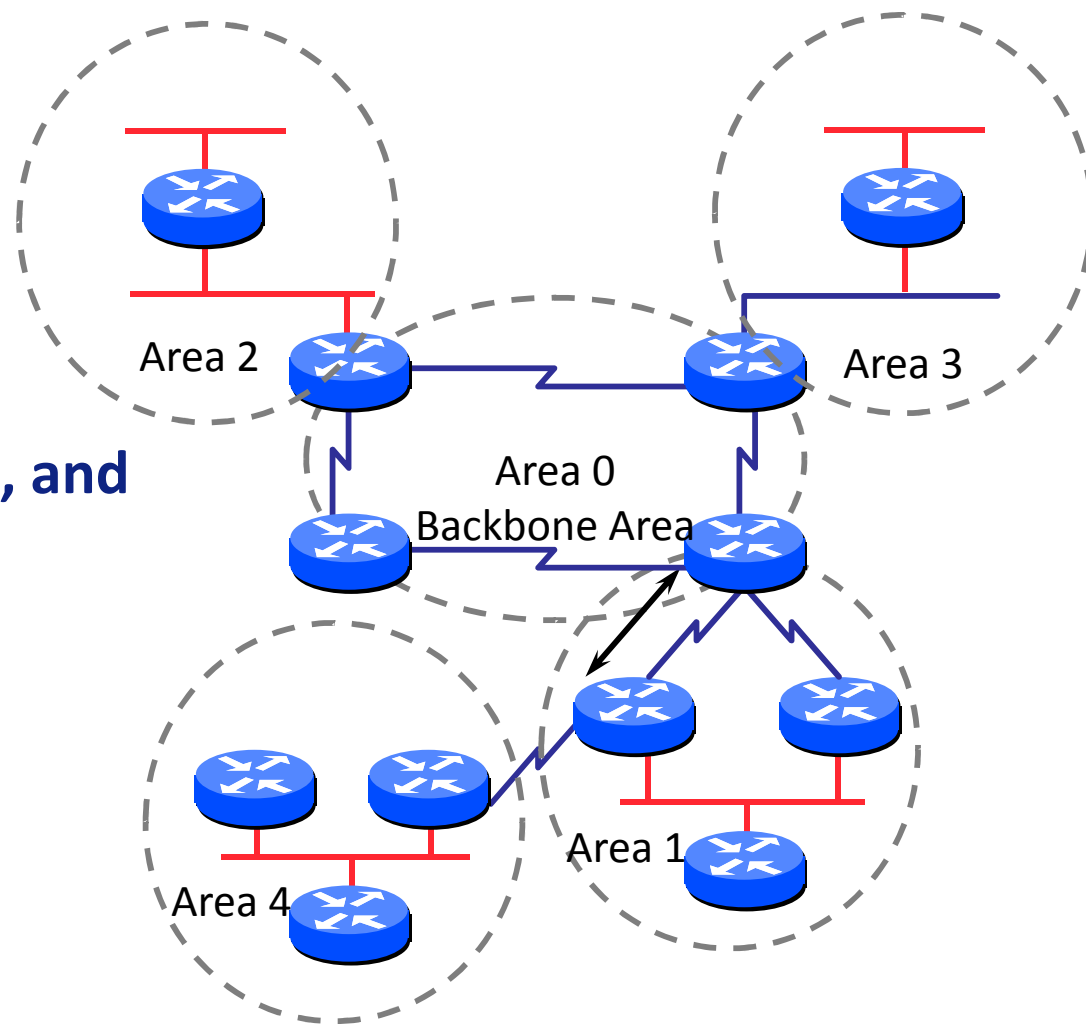
OSPF Routing Protocol

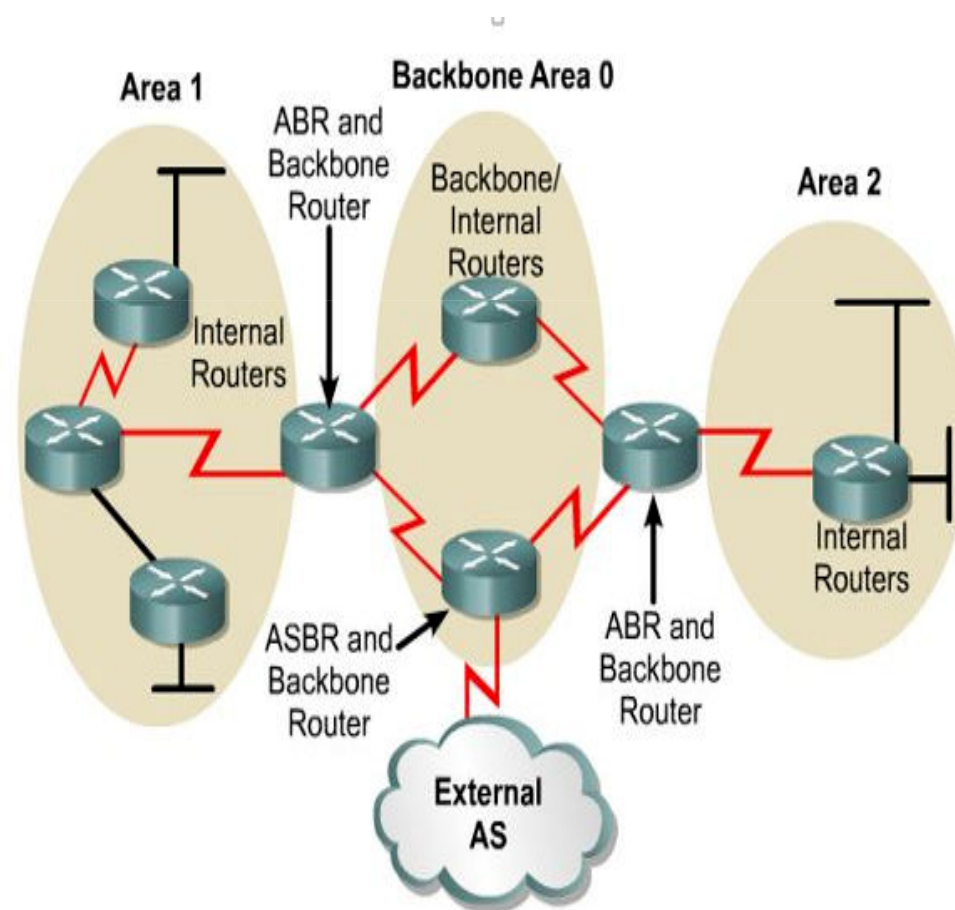## OSPF Router Types

The four separate OSPF router types are as follows:

- **Internal Routers** – all router interfaces belong to only one Area.

- **Area Border Routers (ABRs)** – contains interfaces in at least two separate areas

- **Backbone Routers** – contain at least one interface in Area 0

- **Autonomous System Border Routers (ASBRs)** – contain a connection to a separate Autonomous System

OSPF Routing Protocol

LSA TYPES

• **OSPF routers keep track of the status of links within their respective areas. OSPF routers forward link-state advertisements (LSAs) to ensure the topology database is consistent on each router within an area.**

Several LSA types exist:

• **Router LSA (Type 1) –** Contains a list of all links local to the router, and the status and "cost" of those links.
• **Network LSA (Type 2) –** Generated by all Designated Routers in OSPF, and contains a list of all routers attached to the Designated Router.
• **Network Summary LSA (Type 3) –** Generated by all ABRs in OSPF, and contains a list of all destination networks within an area. Type 3 LSAs are sent between areas to allow inter-area communication to occur.
• **ASBR Summary LSA (Type 4) –** Generated by ABRs in OSPF, and contains a route to any ASBRs in the OSPF system. Type 4 LSAs are sent from an ABR into its local area, so that internal routers know how to exit the Autonomous System.

OSPF Routing Protocol

Configuring Basic OSPF

- **Enabling OSPF**

Rtr(config)# router ospf *process-id*

- **process-id: 1 - 65,535**

- **Cisco feature, which allows you to run multiple, different OSPF routing processes on the same router.**

- **Process-id is locally significant, and does <u>not</u> have to be the same number on other routers (they don't care).**

- **This is different than the process-id used for IGRP and EIGRP which must be the same on all routers sharing routing information.**

OSPF Routing Protocol

Configuring Basic OSPF

- **Network Command**

Rtr(config)# router ospf *process-id*

Rtr(config-router)#network *address wildcard-mask* area *area-id*

• Tells OSPF which interfaces to enable OSPF on (send and receive updates), matching the address and wildcard mask. Also, tells OSPF to include this network in its routing updates.

• Wildcard is necessary because OSPF supports CIDR and VLSM

• Most of the time you can just use an inverse-mask (like access-lists) as the network wildcard mask.

Special Cases:

➢ Rtr(config-router)#network 0.0.0.0 255.255.255.255 area 0

  - Matches all interfaces on this router, <u>not</u> recommended

➢ Rtr(config-router)#network 10.5.1.2 0.0.0.0 area 0

  - Matches only the interface 10.5.1.2 and not any other 10.5.1.n interfaces.

OSPF Routing Protocol

Configuring Basic OSPF

• OSPF Passive-Interface

• It is possible to control which router interfaces will participate in the OSPF process. OSPF will no longer form neighbor relationships out of a "passive" interface
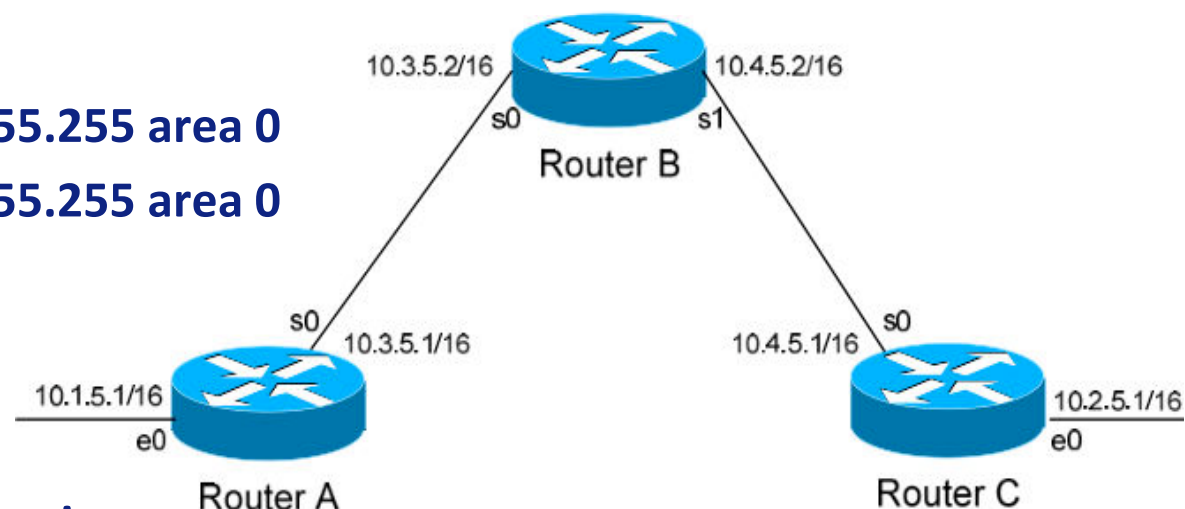
Example:

Router(config)# router ospf 1

Router(config-router)# network 10.4.0.0 0.0.255.255 area 0

Router(config-router)# network 10.2.0.0 0.0.255.255 area 0

Router(config-router)# passive-interface s0

It is possible to configure all interfaces to be passive:

RouterC(config-router)# passive-interface default

OSPF Routing Protocol

Configuring Basic OSPF

- OSPF Authentication
- OSPF supports authentication to secure routing updates
- Two forms of OSPF authentication exist, using either clear-text or an MD5 hash
- To configure MD5-hashed authentication, the first step is also to enable authentication for the area under the OSPF process:

Router(config)# router ospf 1
Router(config-router)# area 0 authentication message-digest

- Next, the hashed authentication key must be configured on the interface:

Router(config)# interface s0
Router(config-router)# ip ospf message-digest-key 10 md5 MYKEY

OSPF Routing Protocol

Configuring Basic OSPF

- **Inter-Area OSPF Summarization**
- **OSPF is a classless routing protocol, thus all of the listed networks would be advertised individually.**
- **This increases the size of the topology databases and routing tables on routers in the domain. Advertising only a summary route for inter-area communication can reduce the load on router CPUs**
- **Example: All of the networks in Area 1 can be summarized as 10.1.0.0/21**
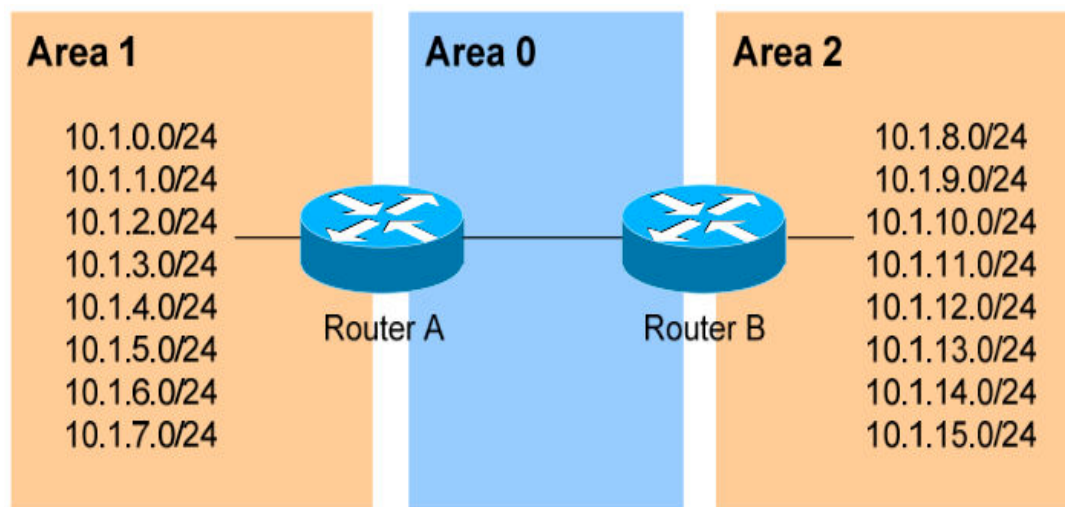
RouterA(config)# router ospf 1
RouterA(config-router)# network 10.1.0.0 0.0.7.255 area 1
RouterA(config-router)# area 1 range 10.1.0.0 255.255.248.0

The area 1 range command creates a summary route for those networks, which is then advertised into Area 0.

OSPF Routing Protocol

Verifying OSPF Operation

Commands:

- **show ip protocols**—Displays IP routing protocol parameters about timers, filters, metrics, networks, and other information for the entire router.

- **show ip route ospf**—Displays the OSPF routes known to the router. This command is one of the most useful in determining connectivity between the local router and the rest of the internetwork.

- **show ip ospf interface**—Verifies that interfaces are configured in the intended areas. In addition, this command displays the timer intervals (including the hello interval) and shows the neighbor adjacencies.

- **show ip ospf**—Displays the OSPF router ID, OSPF timers, the number of times the SPF algorithm has been executed, and LSA information.

- **show ip ospf neighbor**—Displays a list of neighbors, including their OSPF router ID, their OSPF priority, their neighbor adjacency state (for example, init, exstart, or full), and the dead timer.

- **show ip ospf database**— Command that lists a summary of the LSAs in the database

OSPF Routing Protocol

Troubleshooting OSPF

- Ensure you're using a wildcard mask and not a subnet mask when configuring OSPF
- Check IP addressing and subnet masks are correctly configured on interfaces
- Routers running OSPF must have the same hello and dead timers to form a adjacency
- If using OSPF authentication must be match on other routers
- You may have accidentally configured passive-interface on a interface running OSPF – this will stop advertisements
- debug ip ospf adj – shows elections for DR and BDR
- debug ip ospf packet – shows ospf hello packets being received by the router
- debug ip ospf hello – shows more in depth information relating to hello packets including being sent and received by the router

REDDIG2 - IP Routing Scheme

REDDIG IP ROUTING SCHEME

OSPF area 0

SITE A

10.100.XX.90

10.100.XX.32
ABR

Default Static Route
Metric 100

Cisco Level-3

SKYWAN-A   SKYWAN-B

OSPF
area 10.100.YY.0

SITE B

SKYWAN-A   SKYWAN-B

OSPF
area 10.100.ZZ.0

SITE C

SKYWAN-A

SKYWAN-B

10.100.XX.31
ABR

OSPF
area 10.100.XX.0

Cisco 1A

VRRP
10.100.XX.254

Cisco 1B

Cisco GBB 1

10.100.XX.101
DROTHER

10.100.XX.102
BDR

10.100.XX.121
DR

Default Static Route
Metric 100

Default Static Route
Metric 1

Cisco 2A

VRRP
10.100.XX.253

Cisco 2B

Cisco GBB 2

10.100.XX.103

10.100.XX.104

10.100.XX.122