College of DuPage

CCNA3 V3.0

Switching Basics and Intermediate Routing:

Module 1:

Introduction to Classless Routing

3-2004

Overview

- With the phenomenal growth of the Internet and TCP/IP, virtually every enterprise must now implement an IP addressing scheme. The network designer must choose an addressing scheme that allows for growth. Variable-Length Subnet Masking (VLSM) is a technique that allows for the creation of efficient, scalable addressing schemes.
- Over the past two decades, engineers have successfully modified IPv4 so that it can survive the exponential growth of the Internet.
  - VLSM is one of the modifications that has helped to bridge the gap between IPv4 and IPv6.
  - IP version 6 (IPv6), with virtually unlimited address space, is slowly being implemented in select networks and may replace IPv4 as the dominant protocol of the Internet.
- The routing protocol used in a network does much to determine the scalability of the network.
  - Routing Information Protocol (RIP) is still considered suitable for small networks, but is not scalable to large networks because of inherent limitations.
  - To overcome these limitations yet maintain the simplicity of RIP version 1 (RIP v1), RIP version 2 (RIP v2) was developed.

Upon completion of this module, the student will be able to perform tasks related to the following:	
1.1	VLSM
1.2	RIP Version 2

Overview

- Students completing this module should be able to:
  - Define VLSM and briefly describe the reasons for its use
  - Divide a major network into subnets of different sizes using VLSM
  - Define route aggregation and summarization as they relate to VLSM
  - Configure a router using VLSM
  - Identify the key features of RIP v1 and RIP v2
  - Identify the important differences between RIP v1 and RIP v2
  - Configure RIP v2
  - Verify and troubleshoot RIP v2 operation
  - Configure default routes using the **ip route** and **ip default-network** commands

Upon completion of this module, the student will be able to perform tasks related to the following:	
1.1	VLSM
1.2	RIP Version 2

IPv4 Address Classes

Class A	Network	Host		
Octet	1	2	3	4

Class B	Network		Host	
Octet	1	2	3	4

Class C	Network			Host
Octet	1	2	3	4

Class D	Host			
Octet	1	2	3	4

Address Class	First Octet Range	Number of Possible Networks	Number of Hosts per Network
Class A	0 to 127	128 (2 are reserved)	16,777,214
Class B	128 to 191	16,348	65,534
Class C	192 to 223	2,097,152	254

- No medium size host networks
- In the early days of the Internet, IP addresses were allocated to organizations based on request rather than actual need.

IP addressing crisis

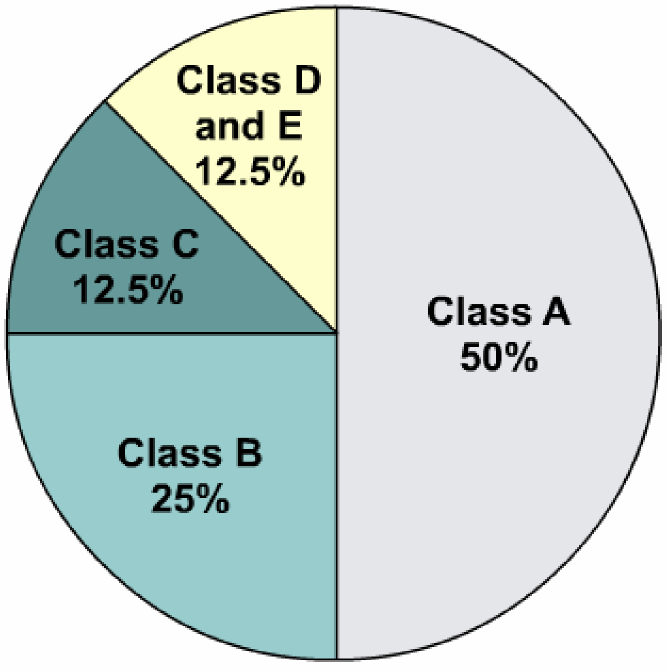

With Class A and B addresses virtually exhausted, Class C addresses (12.5 percent of the total space) are left to assign to new networks.

- Address Depletion
- Internet Routing Table Explosion

IPv4 Addressing

Subnet Mask

- One solution to the IP address shortage was thought to be the subnet mask.
- Formalized in 1985 (RFC 950), the subnet mask breaks a single class A, B or C network in to smaller pieces.

Subnet Example

Given the Class B address 190.52.0.0

Class B	**Network**	**Network**	**Host**	**Host**

Using /24 subnet...	**Network**	**Network**	**Subnet**	**Host**

Internet **routers still "see" this net as 190.52.0.0**

190.52.1.2

190.52.2.2

190.52.3.2

But *internal* routers think all these addresses are on different networks, called subnetworks

Subnet Example

Network address **190.52.0.0 with /16 network mask**

Using Subnets: subnet mask 255.255.255.0 or /24

Network	Network	Subnet	Host
190	52	0	Host
190	52	1	Host
190	52	2	Host
190	52	3	Host
190	52	Etc.	Host
190	52	254	Host
190	52	255	Host

Subnets

255 Subnets

2^8 - 2

Subnet Example

Subnet 0 (all 0's subnet) issue: The address of the subnet, 190.52.0.0/24 is the same address as the major network, 190.52.0.0/16.

Network	Network	Subnet	Host
190	52	0	Host
190	52	1	Host
190	52	Etc.	Host
190	52	254	Host
190	52	255	Host

Subnets

255 Subnets

2^8 - 2

Last subnet (all 1's subnet) issue: The broadcast address for the subnet, 190.52.255.255 is the same as the broadcast address as the major network, 190.52.255.255.

Long Term Solution: IPv6 (coming)

- IPv6, or IPng (IP – the Next Generation) uses a 128-bit address space, yielding

  340,282,366,920,938,463,463,374,607,431,768,211,456

  possible addresses.
- IPv6 has been slow to arrive
- IPv4 revitalized by new features, making IPv6 a luxury, and not a desperately needed fix
- IPv6 requires new software; IT staffs must be retrained
- IPv6 will most likely coexist with IPv4 for years to come.
- Some experts believe IPv4 will remain for more than 10 years.

Short Term Solutions: IPv4 Enhancements

- CIDR (Classless Inter-Domain Routing) – RFCs 1517, 1518, 1519, 1520
- VLSM (Variable Length Subnet Mask) – RFC 1009
- Private Addressing - RFC 1918
- NAT/PAT (Network Address Translation / Port Address Translation) – RFC

What is VLSM and why is it used?

- As IP subnets have grown, administrators have looked for ways to use their address space more efficiently. One technique is called Variable-Length Subnet Masks (VLSM).
  - With VLSM, a network administrator can use a long mask on networks with few hosts, and a short mask on subnets with many hosts.
- VLSM allows an organization to use more than one subnet mask within the same network address space.
  - Implementing VLSM is often referred to as "subnetting a subnet", and can be used to maximize addressing efficiency.
- In order to use VLSM, a network administrator must use a routing protocol that supports it.
  - OSPF, Integrated IS-IS, EIGRP, RIP v2, and static routing.

What is VLSM and why use it?

- The addressing crisis
- Internet Engineering Task Force identified two problems in 1992
- Exhaustion of unassigned IPv4 network address - Class B was on the verge of depletion
- Rapid increase in the size of the internet's routing tables

What is VLSM and why use it?

- Short term extensions to IPv4
- Subnetting 1985
- Variable length subnetting 1987
- Classless Interdomain Routing 1993
- Private IP addresses
- Network Address Translation (NAT)

VLSM is supported by:

- OSPF
- Integrated IS-IS
- EIGRP
- RIP v2
- Static routing

What is VLSM and why use it?

- Ultimate solution: IPv6 128 bit address space
- Allows for: 340,283,366,920,938,463,374,607,431,768,211,456 possibilities

What Is a Variable-Length Subnet Mask?

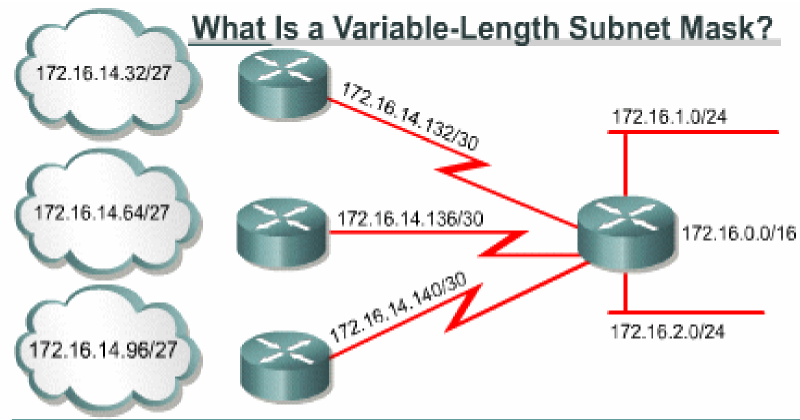

- Subnet 172.16.14.0/24 is divided into smaller subnets
- Subnet with one mask (/27)
- Then further subnet one of the unused /27 subnets into mutiple /30 subnets

What is VLSM and why is it used? (cont.)

- Classful routing protocols (RIP v1 and IGRP) require that a single network use the same subnet mask.
  - Therefore, network 192.168.187.0 must use just one subnet mask such as 255.255.255.0.
- VLSM is simply a feature that allows a single autonomous system to have networks with different subnet masks.
  - If a routing protocol allows VLSM, use a 30-bit subnet mask on network connections, 255.255.255.252, a 24-bit mask for user networks, 255.255.255.0, or even a 22-bit mask, 255.255.252.0, for networks with up to 1000 users.

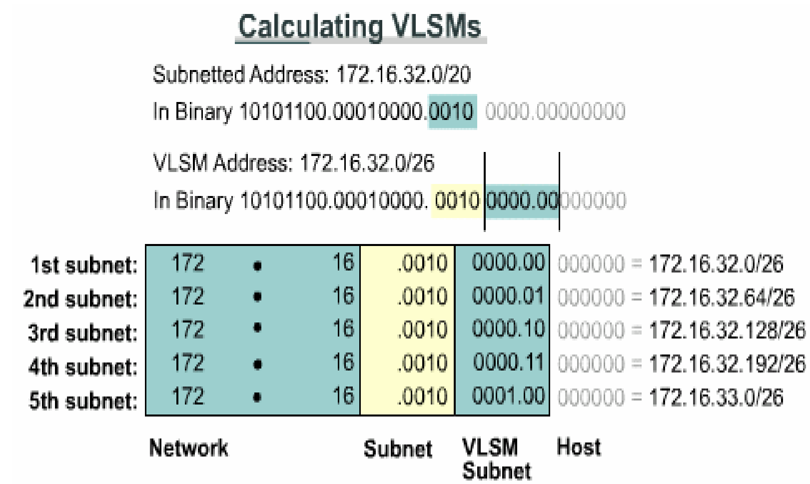

Subnet Masks

255.255.255.252	11111111 11111111 11111111 11111100	30 bits
255.255.255.0	11111111 11111111 11111111 00000000	24 bits
255.255.252.0	11111111 11111111 11111100 00000000	22 bits

A waste of space

- In the past, it has been recommended that the first and last subnet not be used.
  - Use of the first subnet, known as subnet zero, for host addressing was discouraged because of the confusion that can occur when a network and a subnet have the same addresses.
  - The same was true with the use of the last subnet, known as the all-ones subnet. It has always been true that these subnets could be used.
- As networking technologies have evolved, it has become acceptable practice to use the first and last subnets in a subnetted network in conjunction with VLSM.
- On the issue of using subnet zero and the all-ones subnet, RFC 1878 states, "This practice (of excluding all-zeros and all-ones subnets) is obsolete.
  - Modern software will be able to utilize all definable networks."
- <u>However, on certain networks, particularly the ones using legacy software, the use of subnet zero and the all-ones subnet can lead to problems.</u>

A Waste of Space

Subnet Number	Subnet Address	
Subnet 0	192.168.187.0	/27
Subnet 1	192.168.187.32	/27
Subnet 2	192.168.187.64	/27
Subnet 3	192.168.187.96	/27
Subnet 4	192.168.187.128	/27
Subnet 5	192.168.187.160	/27
Subnet 6	192.168.187.192	/27
Subnet 7	192.168.187.224	/27

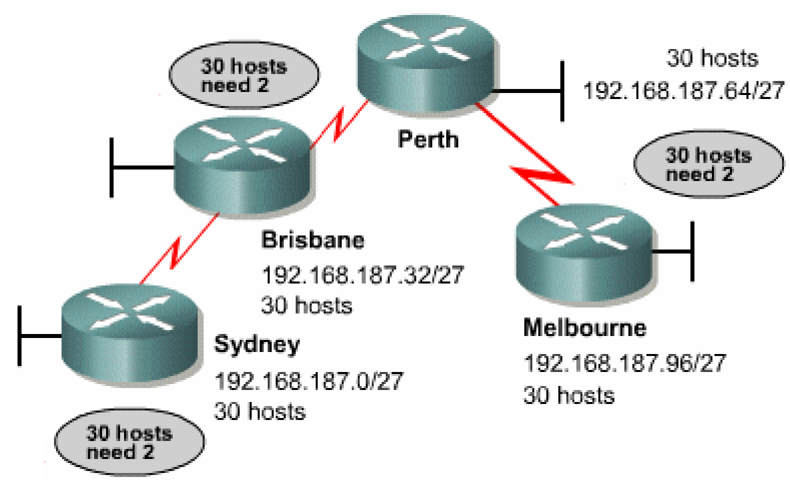

A waste of space - example

- In this network, the network management team has decided to borrow three bits from the host portion of the Class C address that has been selected for this addressing scheme.
- If management decides to use subnet zero, it has eight useable subnets.
- Each may support 30 hosts. If the management decides to use the **no ip subnet-zero** command, it has seven usable subnets with 30 hosts in each subnet. From Cisco IOS version 12.0, remember that Cisco routers use subnet zero by default.
- The team realizes that it has to address the three point-to-point WAN links between Sydney, Brisbane, Perth, and Melbourne.
  - If the team uses the three remaining subnets for the WAN links, it will have used all of the available addresses and have no room for growth.
  - The team will also have wasted the 28 host addresses from each subnet to simply address three point-to-point networks.

A Waste of Space

Subnet Number	Subnet Address	
Subnet 0	192.168.187.0	/27
Subnet 1	192.168.187.32	/27
Subnet 2	192.168.187.64	/27
Subnet 3	192.168.187.96	/27
Subnet 4	192.168.187.128	/27
Subnet 5	192.168.187.160	/27
Subnet 6	192.168.187.192	/27
Subnet 7	192.168.187.224	/27

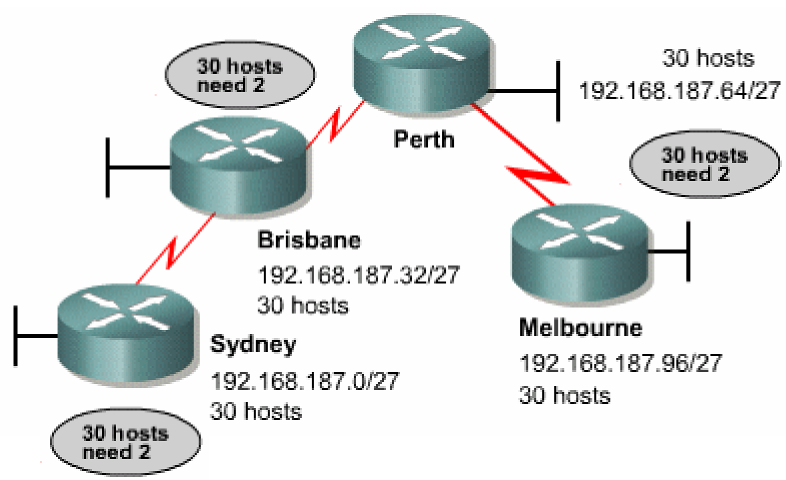

All Zeros and All Ones Subnets

Using the All Ones and All Zeroes Subnet

- There is no command to enable or disable the use of the all-ones subnet, it is enabled by default.

```
Router(config)#ip subnet-zero
```

- The use of the all-ones subnet has always been explicitly allowed and the use of subnet zero is explicitly allowed since Cisco IOS version 12.0.

RFC 1878 states, "This practice (of excluding all-zeros and all-ones subnets) is obsolete! Modern software will be able to utilize all definable networks." Today, the use of subnet zero and the all-ones subnet is generally accepted and most vendors support their use, though, on certain networks, particularly the ones using legacy software, the use of subnet zero and the all-ones subnet can lead to problems.

CCO: Subnet Zero and the All-Ones Subnet http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080093f18.shtml

When to use VLSM?

- It is important to design an addressing scheme that allows for growth and does not involve wasting addresses.
- This section examines how VLSM can be used to prevent waste of addresses on point-to-point links.
- This time the networking team decided to avoid their wasteful use of the /27 mask on the point-to-point links. The team decided to apply VLSM to the addressing problem.

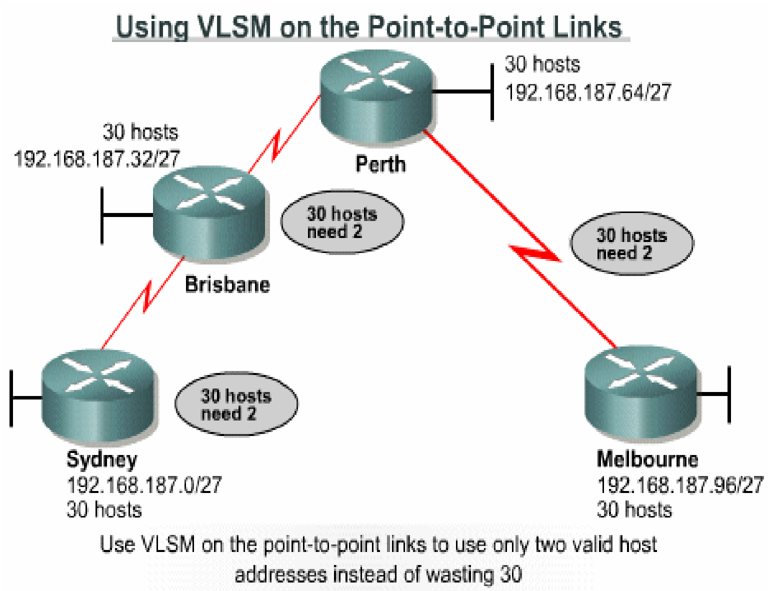

Subnet Number	Subnet Address	
subnet 0	192.168.187.0	/27
subnet 1	192.168.187.32	/27
subnet 2	192.168.187.64	/27
subnet 3	192.168.187.96	/27
subnet 4	192.168.187.128	/27
subnet 5	192.168.187.160	/27
subnet 6	192.168.187.192	/27
subnet 7	192.168.187.224	/27

Subnet Number	Subnet Address	
sub-subnet 0	192.168.187.192	/30
sub-subnet 1	192.168.187.196	/30
sub-subnet 2	192.168.187.200	/30
sub-subnet 3	192.168.187.204	/30
sub-subnet 4	192.168.187.208	/30
sub-subnet 5	192.168.187.212	/30
sub-subnet 6	192.168.187.216	/30
sub-subnet 7	192.168.187.220	/30

When to use VLSM? (cont.)

- To apply VLSM to the addressing problem, the team will break the Class C address into subnets of variable sizes.
  - Large subnets are created for addressing LANs.
  - Very small subnets are created for WAN links and other special cases. A 30-bit mask is used to create subnets with only two valid host addresses.
- In this case this is the best solution for the point-to-point connections.
- In the example, the team has taken one of the last three subnets, subnet 6, and subnetted it again. This time the team uses a 30-bit mask.
- Figures illustrate that after using VLSM, the team has eight ranges of addresses to be used for the point-to-point links.

Subnet Number	Subnet Address	
subnet 0	192.168.187.0	/27
subnet 1	192.168.187.32	/27
subnet 2	192.168.187.64	/27
subnet 3	192.168.187.96	/27
subnet 4	192.168.187.128	/27
subnet 5	192.168.187.160	/27
subnet 6	192.168.187.192	/27
subnet 7	192.168.187.224	/27

Subnet Number	Subnet Address	
sub-subnet 0	192.168.187.192	/30
sub-subnet 1	192.168.187.196	/30
sub-subnet 2	192.168.187.200	/30
sub-subnet 3	192.168.187.204	/30
sub-subnet 4	192.168.187.208	/30
sub-subnet 5	192.168.187.212	/30
sub-subnet 6	192.168.187.216	/30
sub-subnet 7	192.168.187.220	/30

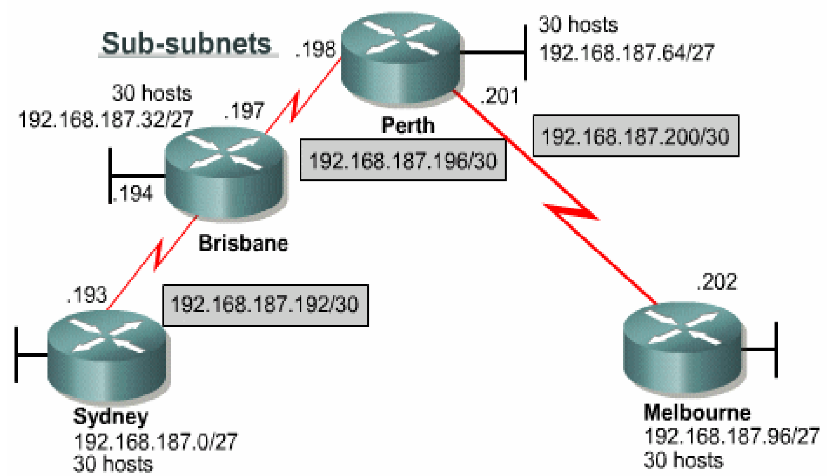

Notice the /27 bit masks for the LANs, and the /30 for the serial links

Calculating subnets with VLSM

- <u>The example contains a Class B address of 172.16.0.0 and two LANs that require at least 250 hosts each.</u>
- If the routers are using a classful routing protocol the WAN link would need to be a subnet of the same Class B network.
- Classful routing protocols such as RIP v1, IGRP, and EGP are not capable of supporting VLSM.
- Without VLSM, the WAN link would have to have the same subnet mask as the LAN segments. A 24-bit mask (255.255.255.0) would support 250 hosts.
- If VLSM were used in this example, a 24-bit mask would still work on the LAN segments for the 250 hosts. A 30-bit mask could be used for the WAN link because only two host addresses are needed.

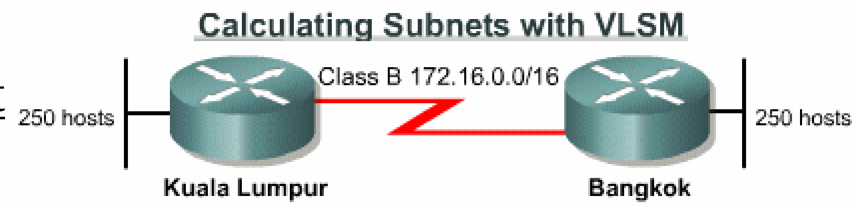

Each LAN must support over 250 hosts. The Class B network 172.16.0.0/16 can be subnetted with a 24-bit mask (255.255.255.0) to create large enough subnets for each LAN.

Subnetted Class B as 255.255.255.0

#	ID	Range	Broadcast
0	172.16.0.0	172.16.0.1 - 172.16.0.254	172.16.0.255
1	172.16.1.0	172.16.1.1 - 172.16.1.254	172.16.1.255
2	172.16.2.0	172.16.2.1 - 172.16.2.254	172.16.2.255
3	172.16.3.0	172.16.3.1 - 172.16.3.254	172.16.3.255
4	172.16.4.0	172.16.4.1 - 172.16.4.254	172.16.4.255
5	172.16.5.0	172.16.5.1 - 172.16.5.254	172.16.5.255
6	172.16.6.0	172.16.6.1 - 172.16.6.254	172.16.6.255
7	172.16.7.0	172.16.7.1 - 172.16.7.254	172.16.7.255
8	172.16.8.0	172.16.8.1 - 172.16.8.254	172.16.8.255
9	172.16.9.0	172.16.9.1 - 172.16.9.254	172.16.9.255
10	172.16.10.0	172.16.10.1 - 172.16.10.254	172.16.10.255
11	172.16.11.0	172.16.11.1 - 172.16.11.254	172.16.11.255
12	172.16.12.0	172.16.12.1 - 172.16.12.254	172.16.12.255
13	172.16.13.0	172.16.13.1 - 172.16.13.254	172.16.13.255
14	172.16.14.0	172.16.14.1 - 172.16.14.254	172.16.14.255
15	172.16.15.0	172.16.15.1 - 172.16.15.254	172.16.15.255

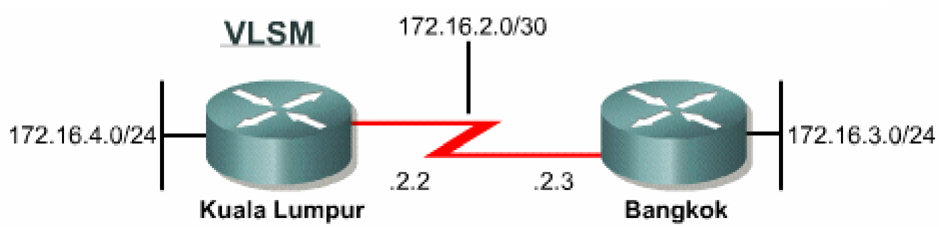

The /30 means that fewer addresses are lost.

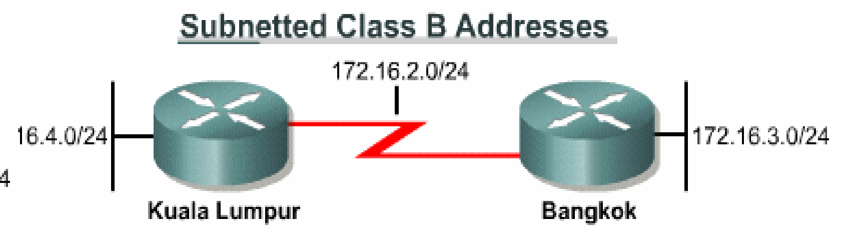

Each Link can support over 250 hosts, but the WAN link only needs two, one for each router interface. Therefore, 252 addresses would be wasted.

Calculating subnets with VLSM (cont.)

- Use this procedure to further subnet:
  Step 1 Write 172.16.32.0 in binary form.
  Step 2 Draw a vertical line between the 20th and 21st bits.
  Step 3 Draw a vertical line between the 26th and 27th bits.
  Step 4 Calculate the 64 subnet addresses using the bits between the two vertical lines, from lowest to highest in value.
- In Figure the subnet addresses used are subdividing the 172.16.32.0/20 subnet into multiple /26 subnets.
- In this example, 172.16.33.0/26 is further subnetted with a prefix of /30. This provides four more subnet bits and therefore 16 (2^4) subnets for the WANs.
  - The subnet address 172.16.32.0/20 there are over 4000 ($2^{12} – 2 = 4094$) addresses.
  - By subnetting 172.16.32.0/20 to 172.16.32.0/26, there is a gain of 64 (2^6) subnets, each of which could support 62 ($2^6 – 2$) hosts.

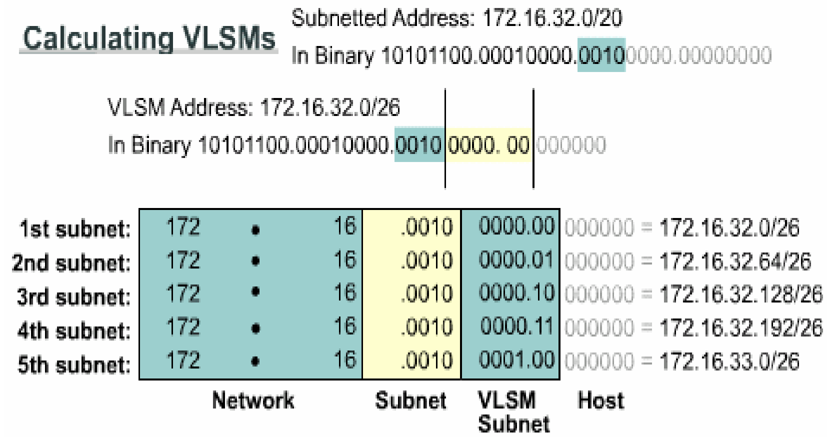

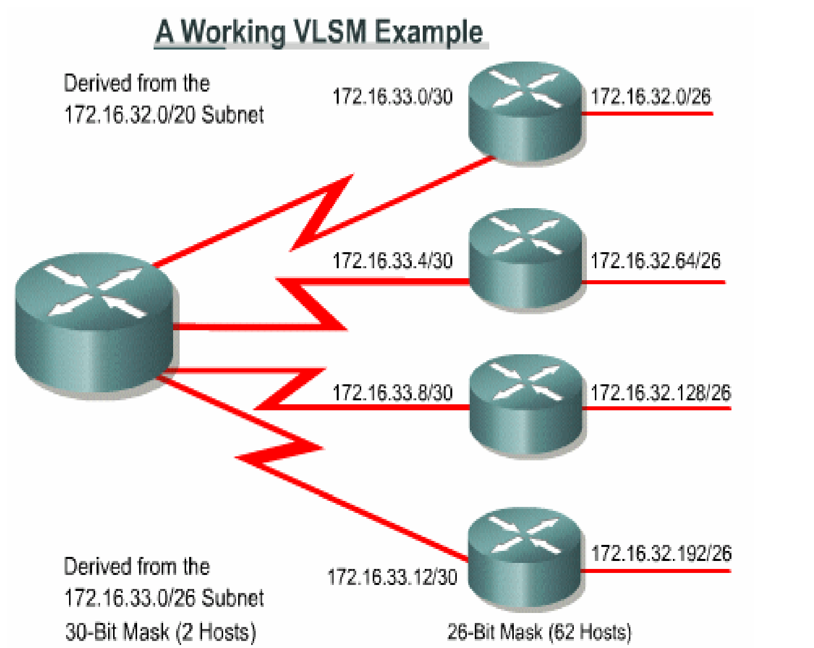

Calculating subnets with VLSM (cont.)

Sub-subnetting 172.16.33.0/26

•It is important to remember that only unused subnets can be further subnetted. If any address from a subnet is used, that subnet cannot be further subnetted. In the example, four subnet numbers are used on the LANs. Another unused subnet, 172.16.33.0/26, is further subnetted for use on the WANs.

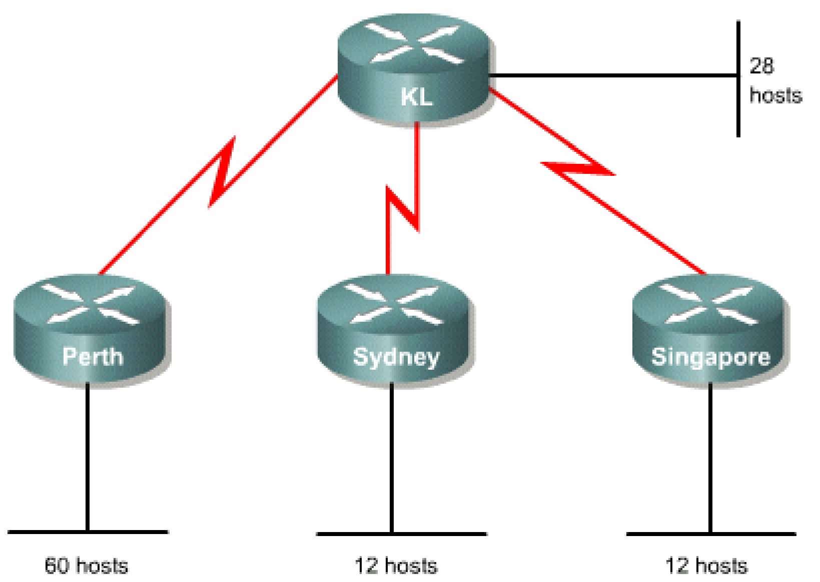

next page for answer

Calculating subnets with VLSM (cont.)

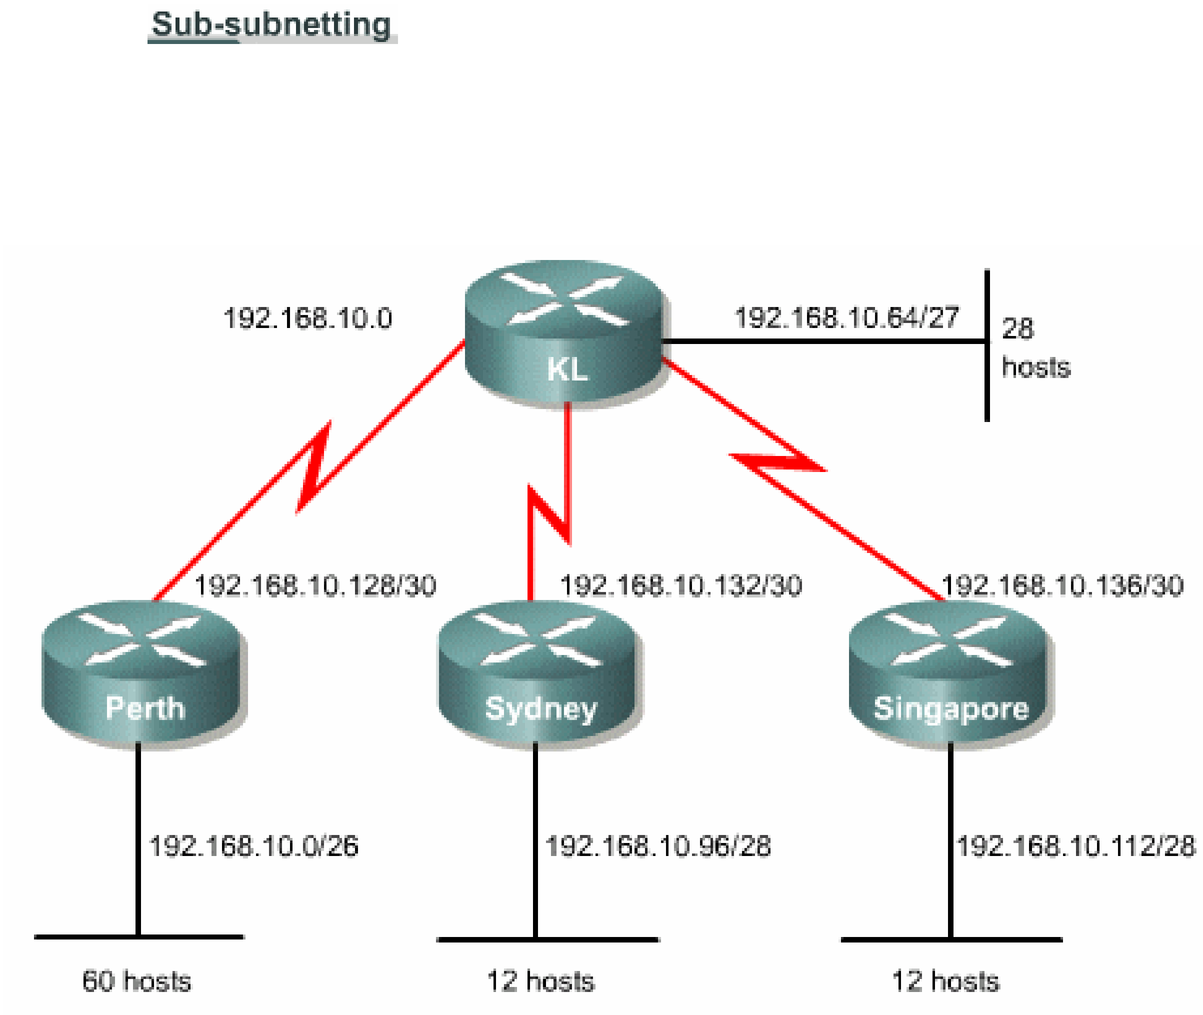

Route aggregation with VLSM

- When using VLSM, try to keep the subnetwork numbers grouped together in the network to allow for aggregation.
  - This means keeping networks like 172.16.14.0 and 172.16.15.0 near one another so that the routers need only carry a route for 172.16.14.0/23.
- The use of Classless InterDomain Routing (CIDR) and VLSM not only prevents address waste, but also promotes route aggregation, or summarization.

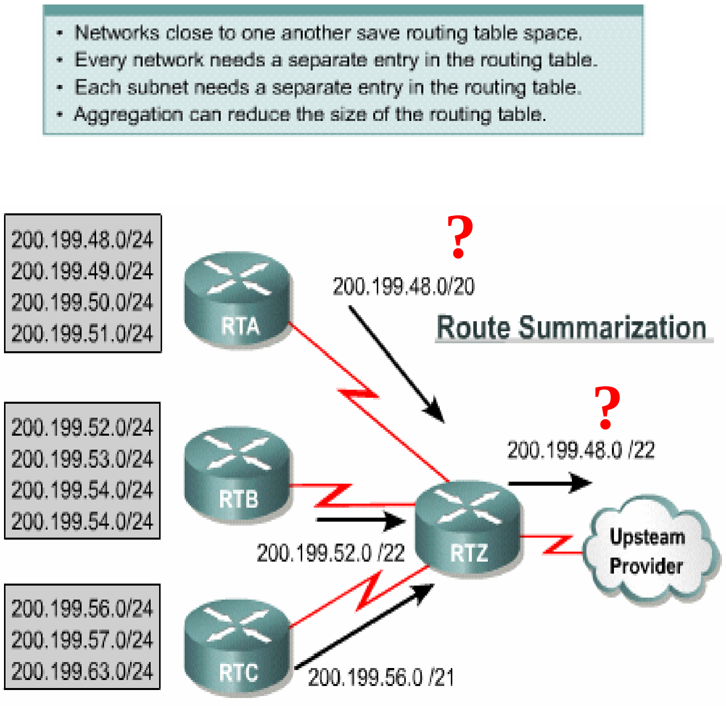

Route summarization reduces routing table size by aggregating routes to multiple networks into one supernet.

See next page for explanation?

- VLSM vs. CIDR
  - **VLSM is similar to CIDR**
    - Both recursively divide networks into small sub networks
- **Difference**
  - **VLSM:** The recursion is performed on the address space previously assigned to an organization and is invisible to the global Internet.
  - **CIDR:** CIDR permits the recursive allocation of an address blocked by an Internet Registry to a high-level ISP, to a mid level to a lower level ISP and finally to a private organization's network.

CIDR (Classless Inter-Domain Routing)

- By 1992, members of the IETF were having serious concerns about the exponential growth of the Internet and the scalability of Internet routing tables.
- The IETF was also concerned with the eventual exhaustion of 32-bit IPv4 address space.
- Projections were that this problem would reach its critical state by 1994 or 1995.
- IETF's response was the concept of Supernetting or CIDR, "cider".
- To CIDR-compliant routers, address class is meaningless.
  - The network portion of the address is determined by the network subnet mask or prefix-length (/8, /19, etc.)
  - The first octet (first two bits) of the network address (or network-prefix) is NOT used to determine the network and host portion of the network address.
- CIDR helped reduced the Internet routing table explosion with supernetting and reallocation of IPv4 address space.

Route aggregation with VLSM (cont.)

- VLSM allows for the summarization of routes and increases flexibly by basing the summarization entirely on the higher-order bits shared on the left, even if the networks are not contiguous.
- The 1st figure shows that the addresses, or routes, share each bit up to and including the 20th bit.
  - The 21st bit is not the same for all the routes. Therefore the prefix for the summary route will be 20 bits long.
- The 2nd figure shows that the addresses, or routes, share each bit up to and including the 21st bit.
  - The 22nd bit is not the same for all the routes. Therefore the prefix for the summary route will be 21 bits long.

Addresses	First Octet	Second Octet	Third Octet	Fourth Octet
192.168.98.0	11000000	10101000	0110.0010	00000000
192.168.99.0	11000000	10101000	0110 0011	00000000
192.168.100.0	11000000	10101000	0110 0100	00000000
192.168.101.0	11000000	10101000	0110 0101	00000000
192.168.102.0	11000000	10101000	0110 0110	00000000
192.168.105.0	11000000	10101000	0110 1001	00000000

Summary route is 192.168.96.0/20

192.168.96.0	11000000	10101000	0110 0000	00000000

Addresses	First Octet	Second Octet	Third Octet	Fourth Octet
172.16.0.0	10101100	00010000	00000 000	00000000
172.16.2.0	10101100	00010000	00000 010	00000000
172.16.3.128	10101100	00010000	00000 011	10000000
172.16.4.0	10101100	00010000	00000 100	00000000
172.16.4.128	10101100	00010000	00000 100	10000000

Answer:

172.16.0.0/21	10101100	00010000	00000 000	00000000

Route aggregation with VLSM

- Figure illustrates how route summarization reduces the burden on upstream routers.
  - These subnetworks are summarized at various points, until the entire network is advertised as a single aggregate route, 200.199.48.0/20.
- Route summarization, or supernetting, is only possible if the routers run a classless routing protocol,
  - such as OSPF or EIGRP.
  - Classless routing protocols carry a prefix that consists of 32-bit IP address and bit mask in the routing updates.
- Remember the following rules:

1. A router must know in detail the subnet numbers attached to it.
2. A router does not need to tell other routers about each individual subnet if the router can send one aggregate route.
3. A router using aggregate routes would have fewer entries in its routing table.

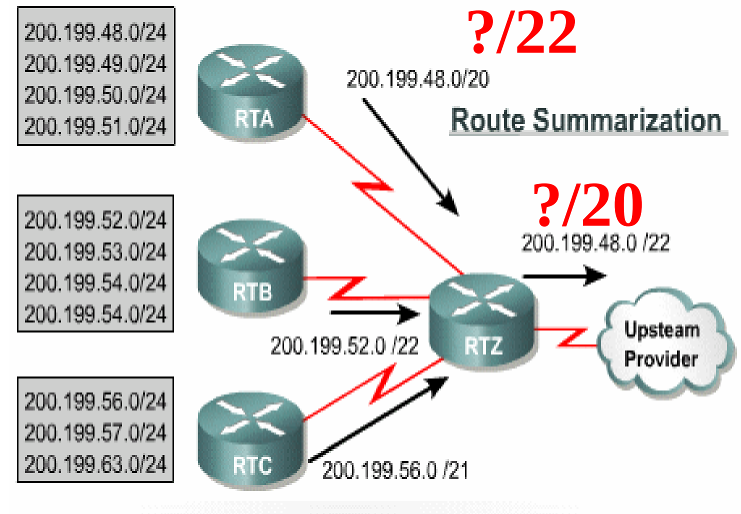

Route summarization reduces routing table size by aggregating routes to multiple networks into one supernet.

In Figure, the summary route that eventually reaches the provider contains a 20-bit prefix common to all of the addresses in the organization, 200.199.48.0/20 or 11001000.11000111.0011.

Route aggregation with VLSM

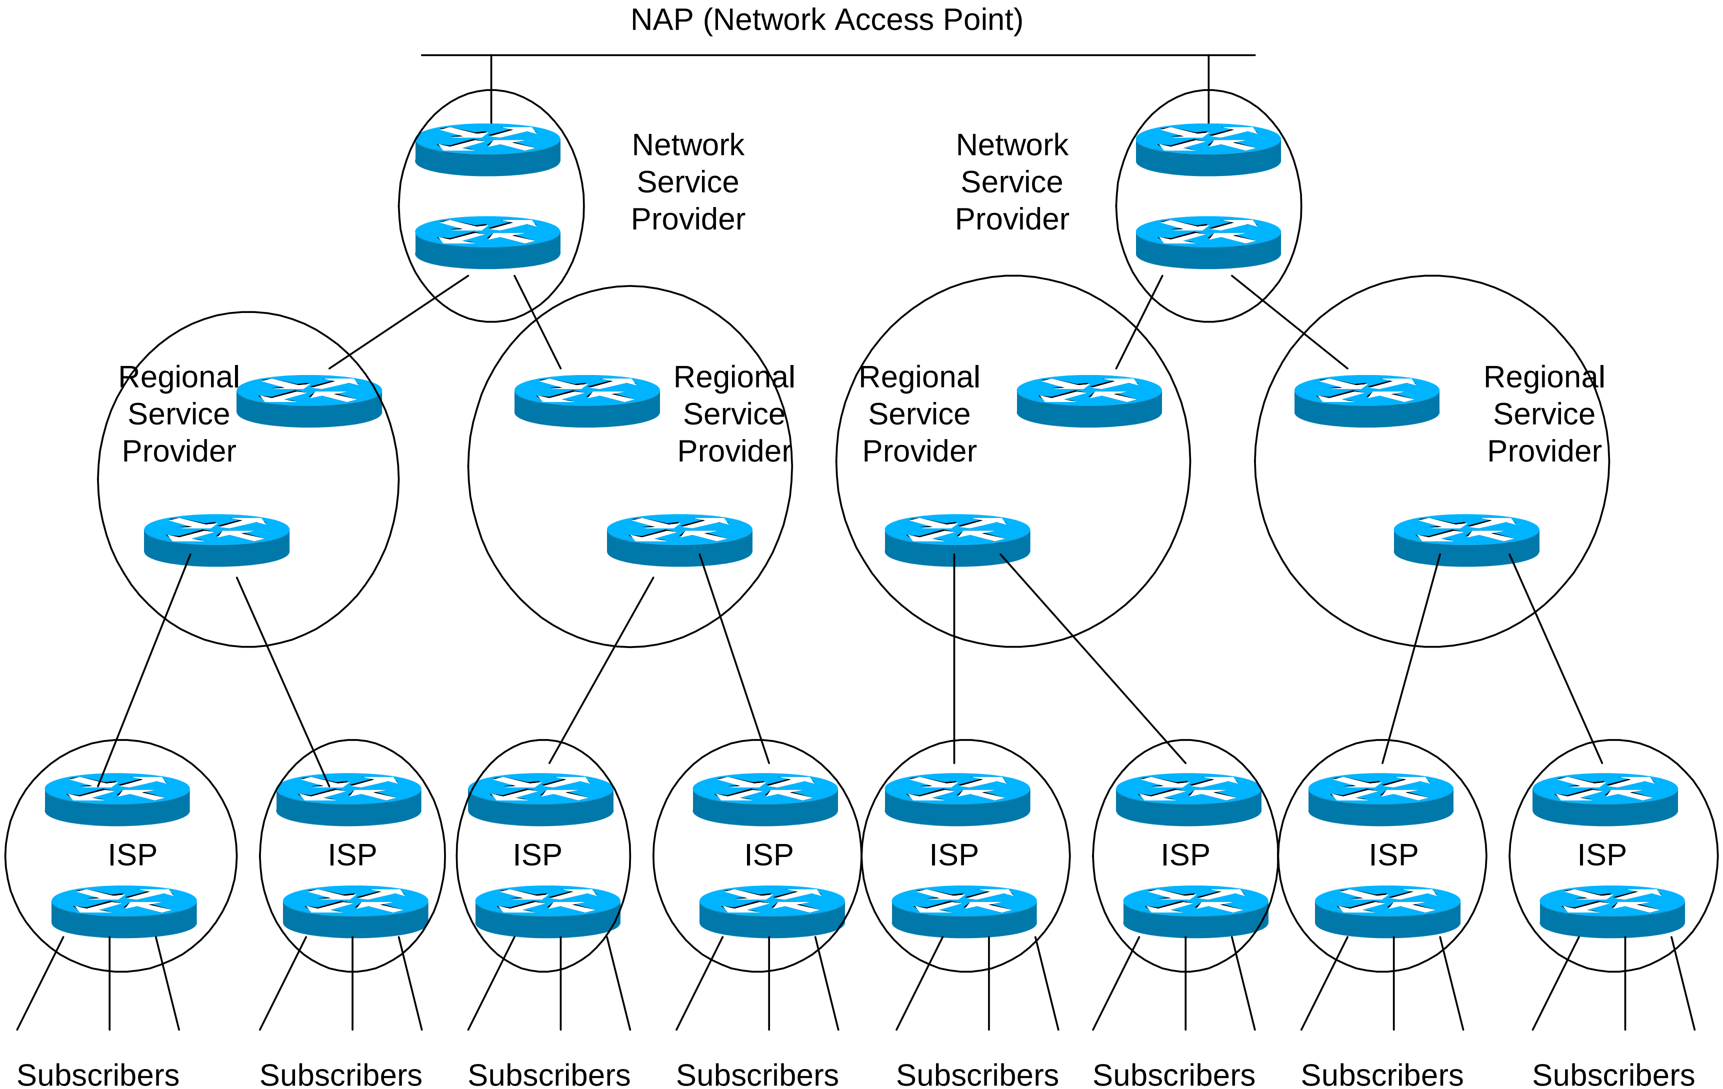

VLSM and the Routing Table

Routing Table without VLSM

```
RouterX#show ip route
     207.21.24.0/27 is subnetted, 4 subnets
C       207.21.24.192 is directly connected, Serial0
C       207.21.24.196 is directly connected, Serial1
C       207.21.24.200 is directly connected, Serial2
C       207.21.24.204 is directly connected, FastEthernet0
```

Routing Table with VLSM

```
RouterX#show ip route
     207.21.24.0/24 is variably subnetted, 4 subnets, 2 masks
C       207.21.24.192 /30 is directly connected, Serial0
C       207.21.24.196 /30 is directly connected, Serial1
C       207.21.24.200 /30 is directly connected, Serial2
C       207.21.24.96  /27 is directly connected, FastEthernet0
```

Configuring VLSM

- In this example address: 192.168.10.0
  - The Perth router has to support 60 hosts.
    - $2^6 = 64 - 2 = 62$, so the division was 192.168.10.0/26.
  - The Sydney and Singapore routers have to support 12 hosts each.
    - $2^4 = 16 - 2 = 14$, so the division is 192.168.10.96/28 for Sydney and 192.168.10.112/28 for Singapore.
  - The KL router requires 28 hosts.
    - $2^5 = 32 - 2 = 30$, so the division here is 192.168.10.64/27.
- The following are point-to-point:
  - **Perth to KL 192.168.10.128/30**
  - **Sydney to KL 192.168.10.132/30**
  - **Singapore to KL 192.168.10.136/30**
- The example for Singapore to KL is configured as follows:
  - Singapore(config)#**interface serial 0** Singapore(config-if)#**ip address 192.168.10.137 255.255.255.252**
  - KualaLumpur(config)#**interface serial 1** KualaLumpur(config-if)#**ip address 192.168.10.138 255.255.255.252**

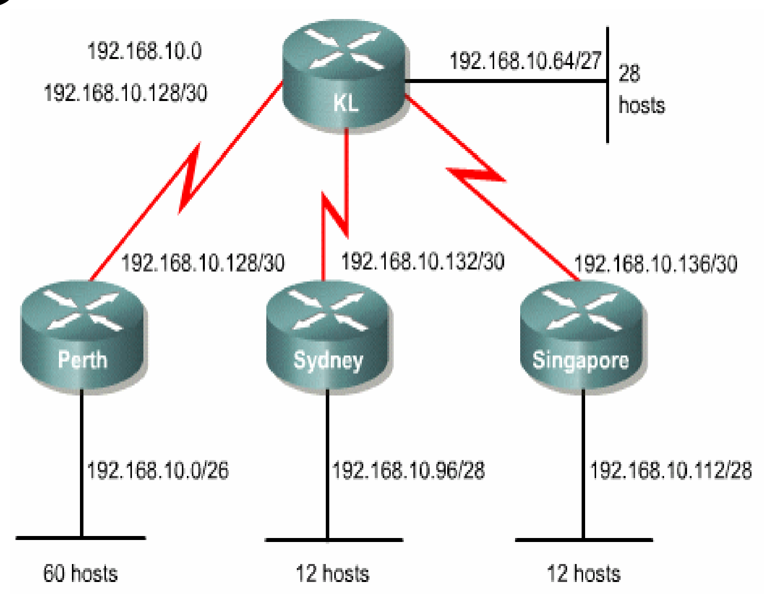

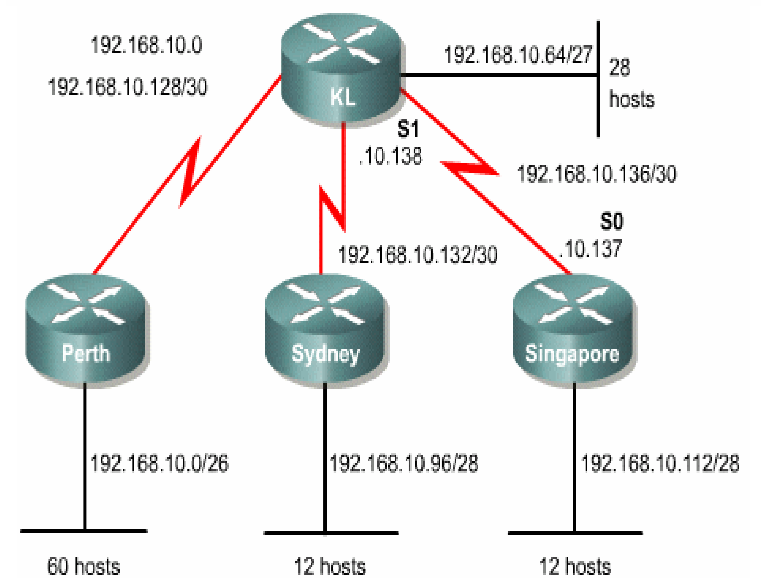

RIP history

- The Internet is a collection of autonomous systems (AS). Each AS is generally administered by a single entity. Each AS will have its own routing technology.
- RIP was designed to work as an IGP in a moderate-sized AS.
  - RIP v1 is a distance vector protocol that broadcasts its entire routing table to each neighbor router at predetermined intervals.
  - The default interval is 30 seconds.
  - RIP uses hop count as a metric, with 15 as the maximum number of hops.

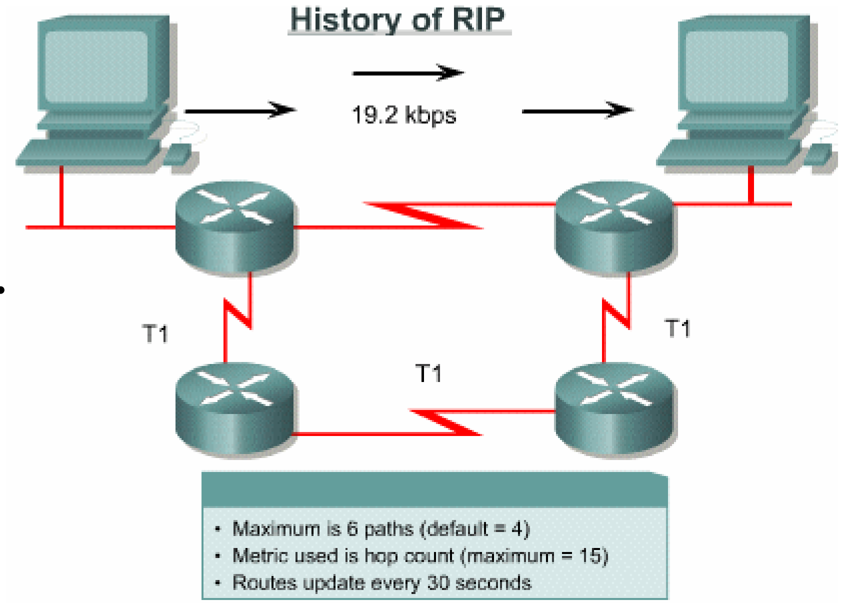

RIP v1 Configuration

```
Sydney(config)#router rip
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
```

RIP history (cont.)

- RIP v1 is considered an interior gateway protocol that is classful.
- If the router receives information about a network, and the receiving interface belongs to the same network but is on a different subnet, the router applies the one subnet mask that is configured on the receiving interface:
  - For Class A addresses, the default classful mask is 255.0.0.0.
  - For Class B addresses, the default classful mask is 255.255.0.0.
  - For Class C addresses, the default classful mask is 255.255.255.0.

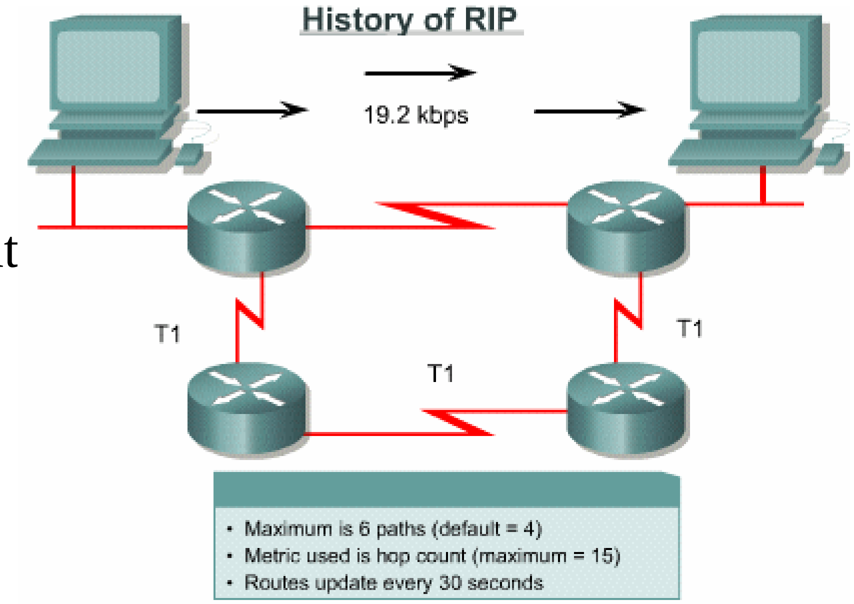

RIP v1 Configuration

```
Sydney(config)#router rip
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
```

RIP history (cont.)

- RIP v1 is a popular routing protocol because virtually all IP routers support it.
- RIP v1 is capable of load balancing over as many as six equal-cost paths, with four paths as the default.
- RIP v1 has the following limitations:
  - It does not send subnet mask information in its updates.
  - It sends updates as broadcasts on 255.255.255.255.
  - It does not support authentication.
  - It is not able to support VLSM or classless interdomain routing (CIDR).

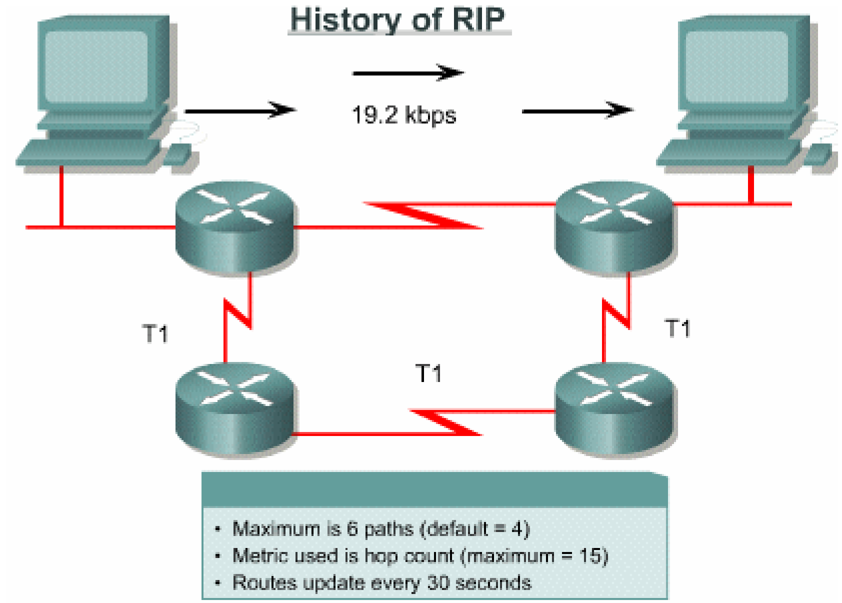

RIP v1 Configuration

```
Sydney(config)#router rip
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
Sydney(config-router)#network network-number
```

RIP v1

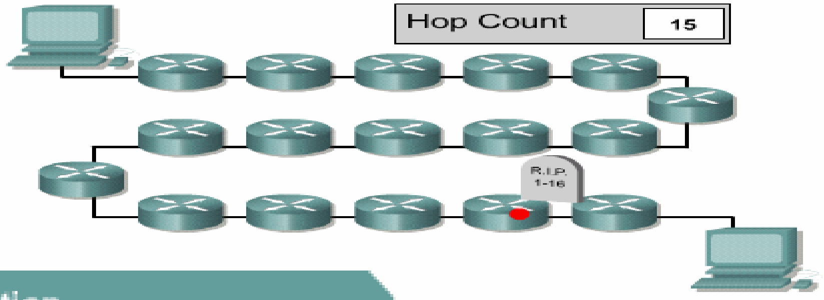

Behavior of RIP v1	Explanation
Directly connected subnets are already known to the router	These routes are advertised to neighboring routers
Routing updates are broadcast	All neighboring routers learn via single broadcast
Routers listen for updates	Helps routers to learn new routes
A metric describes each route in the update	Describes how good route works. If there are many routes, the lowest metric route is used
Topology information in routing updates	At a minimum, this includes metric information
Periodic updates are expected from neighboring routers	Failure to receive updates in timely manner results in removal of routes previously learned from the neighbor
Routes learned from neighboring routers are presumed to be from that router	
A failed route is advertised for a time with a metric that implies an "infinite" distance	RIP v1 uses 16 for infinite distance as RIP v1 maximum valid hop count is 15

RIP v2 features

- RIP v2 is an improved version of RIP v1 and shares the following features:
  - It is a distance vector protocol that uses a hop count metric.
  - It uses holddown timers to prevent routing loops – default is 180 seconds.
  - It uses split horizon to prevent routing loops.
  - It uses 16 hops as a metric for infinite distance.
- RIP v2 provides prefix routing, which allows it to send out subnet mask information with the route update.
  - Therefore, RIP v2 supports the use of classless routing in which different subnets within the same network can use different subnet masks, as in VLSM.
- RIP v2 provides for authentication in its updates.
  - A set of keys can be used on an interface as an authentication check. The choice can be either clear text or Message-Digest 5 (MD5) encryption.
- RIP v2 multicasts routing updates using the Class D address 224.0.0.9, which provides for better efficiency.

Features of RIP v2

Feature	Description
Transmits subnet mask with route	Enables VLSM by passing the mask along with each route so that the subnet is exactly defined.
Provides authentication	Both clear text and/or MD5
Includes a next-hop route IP address in its routing update	A router can advertise a route and direct any listeners to a different router on the same subnet (if the other router has a better route).
Uses external route tags	RIP can pass information about routes learned from an external source and redistributed into RIP. This is used to separate RIP routes from externally learned routes.
Provides multicast routing updates	Instead of sending updates to 255.255.255.255, the destination IP address is 224.0.0.9. This reduces the amount of processing required on non-RIP speaking hosts on a common subnet.

Configuring authentication (EXTRA)

The steps for setting up RIP v2 authentication are as follows:

1. Define a key chain with a name.
2. Define the key or keys on the keychain.
3. Enable authentication on an interface and specify the key chain to be used.
4. Specify whether the interface will use clear text or MD5 authentication.
5. Optionally configure key management.

```
Router(config)#key chain Romeo
Router(config-keychain)#key 1
Router(config-keychain-key)#key-string Juliet
```

The password must be the same on both routers (`Juliet`), but the name of the key (`Romeo`) can be different.

```
Router(config)#interface fastethernet 0/0
Router(config-if)#ip rip authentication key-chain Romeo
Router(config-if)#ip rip authentication mode md5
```

- If the command **`ip rip authentication mode md5`** is not added, the interface will use the default clear text authentication. Although clear text authentication may be necessary to communicate with some RIP v2 implementations, for security concerns use the more secure MD5 authentication whenever possible.

Comparing RIP v1 and v2

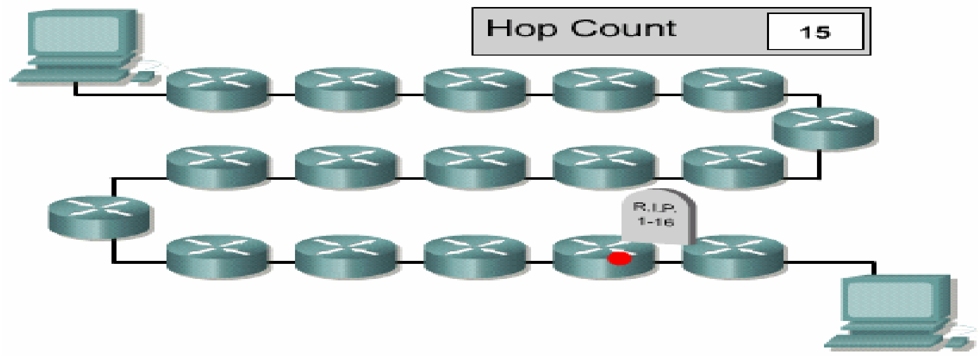

RIP v1 versus RIP v2

RIP v1	RIP v2
Easy to configure	Easy to configure
Only supports classful routing protocol	Supports use of classless routing
No subnet information with the routing update	Sends subnet mask information with the routing updates
Does not support prefix routing - all the devices in the same network must use the same subnet mask.	Supports prefix routing - different subnets within the same network can have different subnet masks (VLSM)
No authentication in updates	Provides for authentication in its updates
Broadcasts over 255.255.255.255	Multicasts routing updates over the Class D address 224.0.0.9 - makes it more efficient

Configuring RIP v2

- To enable a dynamic routing protocol, the following tasks must be completed:
  - Select a routing protocol, such as RIP v2.
  - Assign the IP network numbers without specifying the subnet values.
  - Assign the network or subnet addresses and the appropriate subnet mask to the interfaces.
- RIP v2 is a dynamic routing protocol that is configured by naming the routing protocol RIP Version 2, and then assigning IP network numbers without specifying subnet values.

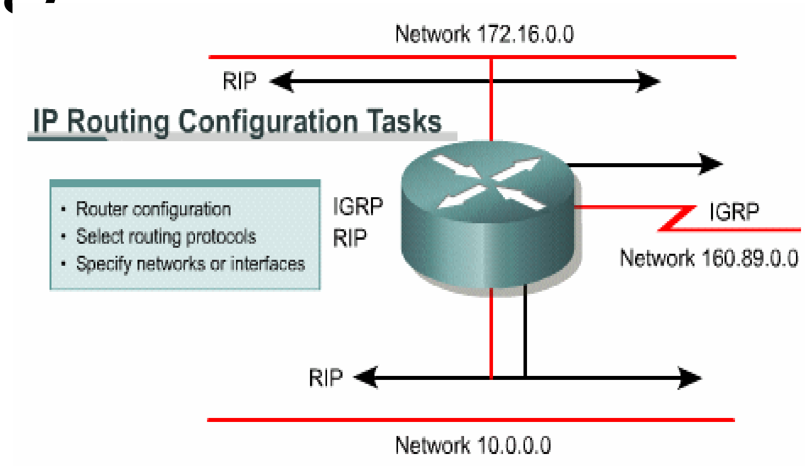

Dynamic Routing Configuration

```
Router(config)#router protocol [keyword]
```

• Defines an IP routing protcol

```
Router(config-router)#network network-number
```

• Mandatory configuration command for each IP routing process
• Identifies the physically connected network that routing updates are forwarded to

Configuring RIP v2 (cont.)

- The **router** command starts the routing process. The **network** command causes the implementation of the following three functions:
  - The routing updates are multicast out an interface.
  - The routing updates are processed if they enter that same interface.
  - The subnet that is directly connected to that interface is advertised.
- The **network** command is required because it allows the routing process to determine which interfaces will participate in the sending and receiving of routing updates.
- The **network** command starts up the routing protocol on all interfaces that the router has in the specified network. The **network** command also allows the router to advertise that network.

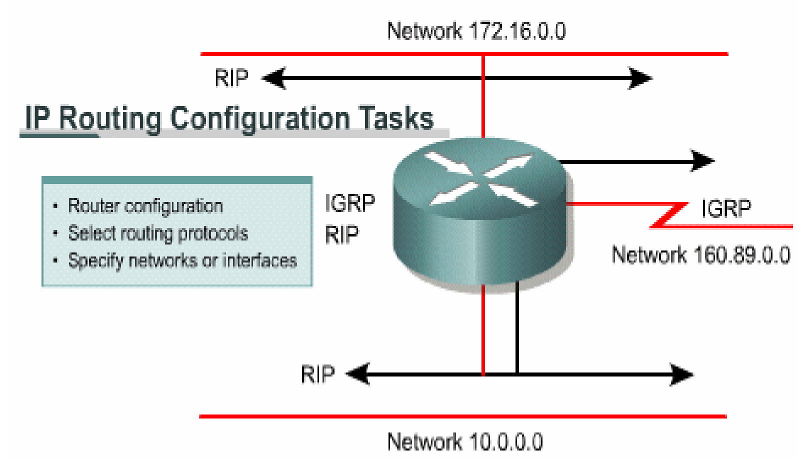

Dynamic Routing Configuration

```
Router(config)#router protocol [keyword]
```

• Defines an IP routing protcol

```
Router(config-router)#network network-number
```

• Mandatory configuration command for each IP routing process
• Identifies the physically connected network that routing updates are forwarded to

Configuring RIP v2 (cont.)

- The **router rip version 2** command specifies RIP v2 as the routing protocol, while the **network** command identifies a participating attached network.
- In this example, the configuration of Router A includes the following:
  - **router rip version 2** – Selects RIP v2 as the routing protocol.
    - **network 172.16.0.0** – Specifies a directly connected network.
    - **network 10.0.0.0** – Specifies a directly connected network.
- Routers B and C have similar RIP configurations but with different network numbers specified.

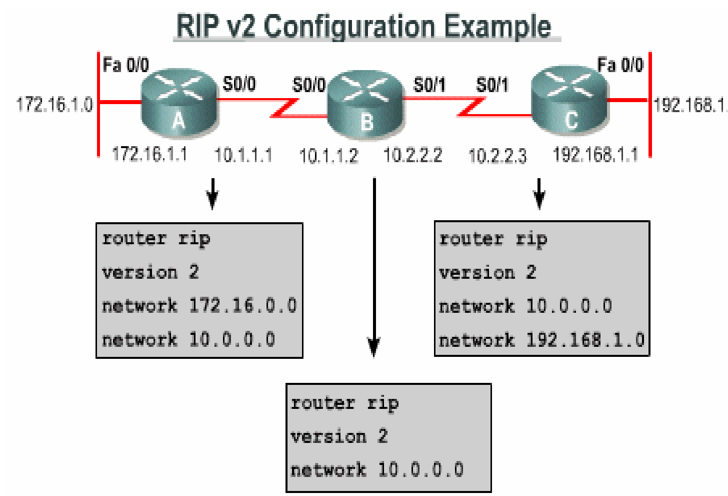

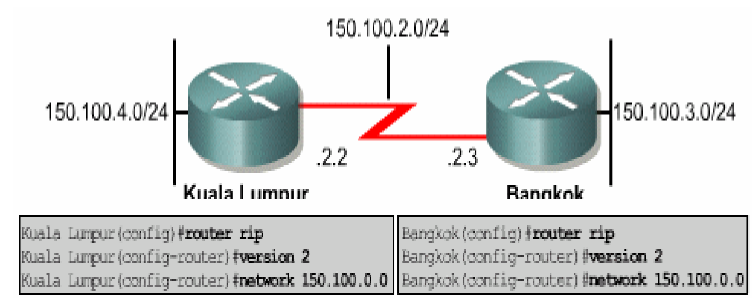

Discontiguous subnets and classless routing

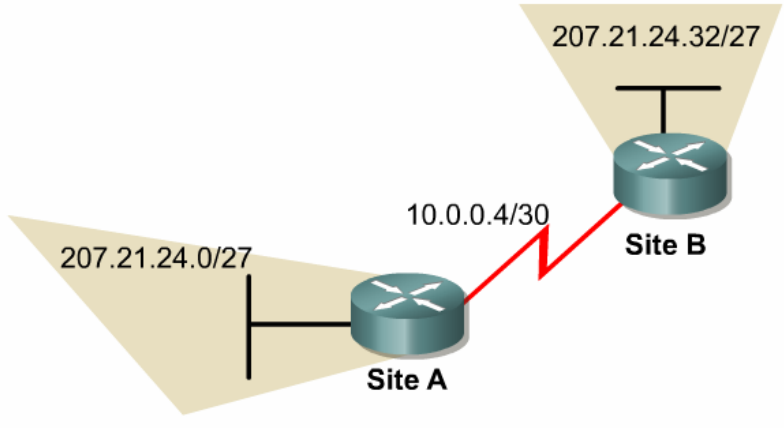

Two subnets from the same major network 207.21.24.0, are separated by a different major network 10.0.0.0, creating discontiguous subnets.

```
router rip
  version 2
  no auto-summary
```

- RIP v1 always uses automatic summarization.
- The default behavior of RIP v2 is to summarize at network boundaries the same as RIP v1.

Verifying RIP v2

- The **show ip protocols** and **show ip route** commands display information about routing protocols and the routing table.
- The **show ip protocols** command displays values about routing protocols and routing protocol timer information associated with the router.
  - In the example, the router is configured with RIP and sends updated routing table information every 30 seconds. <u>This interval is configurable</u>.
  - If a router running RIP does not receive an update from another router for 180 seconds or more, the first router marks the routes served by the non-updating router as being invalid.
  - If there is still no update after 240 seconds the router removes the routing table entries.
  - The distance default of 120 refers to the administrative distance for a RIP route.

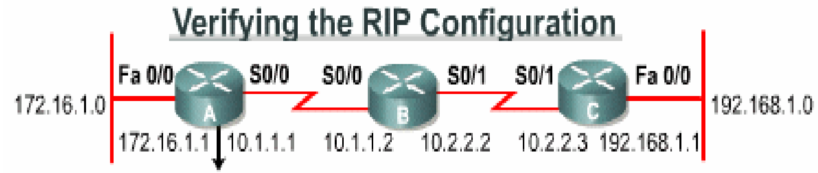


```
RouterA#show ip protocols
Routing Protocol is "rip"
  Sending updates ever 30 seconds, next due in 12 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter lists for all interfaces is
  Incoming update filter lists for all interfaces is
  Redistributing: rip
  Default version control: send version 1, receive any version
     Interface        send   Recv  Triggered  RIP  Keychain
     Ethernet         1      1 2
     Serial2          1      1 2
Routing for Networks:
     10.0.0.0
     172.16.0.0
Routing Information Sources:
     Gateway          Distance        Last Update
     (this router)    120               0:2:12:15
     10.1.1.2         120               0:1:09:01
Distance: (default is 120)
```

Verifying RIP v2 (cont.)

- The **show ip interface brief** command can also be used to list a summary of the information and status of an interface.
- The **show ip route** command displays the contents of the IP routing table.
  - The routing table contains entries for all known networks and subnetworks, and contains a code that indicates how that information was learned.
- If entries are missing, routing information is not being exchanged.
  - Use the **show running-config** or **show ip protocols** privileged EXEC commands on the router to check for a possible misconfigured routing protocol.

Displaying the IP Routing Table


```
RouterA#show ip route
Codes:C - connected, S - static, I - IGRP, 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
      N1 - OSPF external type 1, E2 - OSPF external type 2, * -
candidate
default
       U - Per-user static route, 0 = CCR
       T - Traffic engineered route

Gateway of last resort is not set
   172.16.0.0/24 is subnetted, 1 subnets
C      172.16.1.0 is directly connected, Ethernet0
   10.0.0.0/24 is subnetted, 2 subnets
R      10.2.2.0 (120/1) via 10.1.1.2, 00:00:07, Serial 0/0
C      10.1.1.0 is directly connected, Serial 0/0
R    192.168.1.0/24 (120/2) via 10.1.1.2, 00:00:07, Serial 0/0
```

Troubleshooting RIP v2

- Use the **debug ip rip** command to display RIP routing updates as they are sent and received. The **no debug all** or **undebug all** commands will turn off all debugging.
  - The example shows that the router being debugged has received updates from one router at source address 10.1.1.2.
  - The router at source address 10.1.1.2 sent information about two destinations in the routing table update.
  - The router being debugged also sent updates, in both cases to broadcast address 255.255.255.255 as the destination.
  - The number in parentheses is the source address encapsulated into the IP header.

The debug ip rip Command Description

Command	Explanation
debug ip rip	displays RIP routing updates as sent and received
no debug all	turns off debugging

The debug ip rip Command

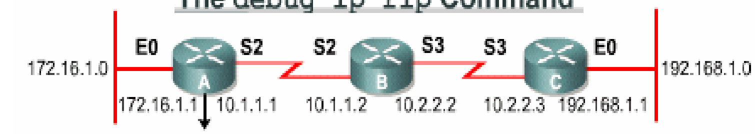


```
RouterA#debug ip rip
RIP protocol debugging is on
RouteA#
00:06:24:RIP:received v1 update from 10.1.1.2 on Serial2
00:06:24:     10.2.2.0 in 1 hops
00:06:24:     192.168.1.0 metric 3
00:06:33:RIP sending v1 update to 255.255.255.255 via Ethernet0
              (172.16.1.1)
00:06:34:     network 10.0.0.0, metric 2
00:06:34:RIP sending v1 update to 255.255.255.255 via Serial2
              (10.1.1.1)
00:06:34:     network 172.16.0.0, metric 1
```

The debug ip rip Outputs and Meanings

Output	Possible significance
RIP: broadcasting general request on Ethernet0	At startup Transition of interface at startup User manually clearing interface
RIP: bad version 128 from 160.89.80.43	Malformed packet from the transmitter
RIP: received v2 update from 150.100.2.3 on Serial0	Shows Version 2 RIP is receiving
RIP: sending v1 update to 255.255.255 via Serial0 (150.100.2.2)	Shows Version 1 RIP is operating
RIP: ignored v1 packet from 150.100.2.2 (illegal version)	Shows that the router is not going to deal with a RIP v1 packet
RIP: sending v2 update to 224.0.0.9 via FastEthernet0 (150.100.3.1)	Shows RIP version 2 is sending
RIP: build update entries 150.100.2.0/24 via 0.0.0.0 metric 1, tag	Shows use of default route and tag

Default routes

- By default, routers learn paths to destinations three different ways:
  - **Static routes** – The system administrator manually defines the static routes as the next hop to a destination. Static routes are useful for security and traffic reduction, as no other route is known.
  - **Default routes** – The system administrator also manually defines default routes as the path to take when there is no known route to the destination. Default routes keep routing tables shorter. When an entry for a destination network does not exist in a routing table, the packet is sent to the default network.
  - **Dynamic routes** – Dynamic routing means that the router learns of paths to destinations by receiving periodic updates from other routers.
- In Figure, the ~~default~~ route is indicated by the following command:
  - Router(config)#**ip route 172.16.1.0 255.255.255.0 172.16.2.1**

The ip route Command

Command	Description
Router(config)#**ip route 172.16.1.0 255.255.255.0 172.16.2.1**	IP network number or subnet number defined as the default
255.255.255.0	Subnet mask indicates that 8 bits of subnetting are in effect
172.16.2.1	ip address of next-hop router in the path to the destination

The ip default-network Command

Command	Description
Router(config)#**ip default-network 192.168.20.0**	IP network number or subnet number defined as the default

Configuring Default Route

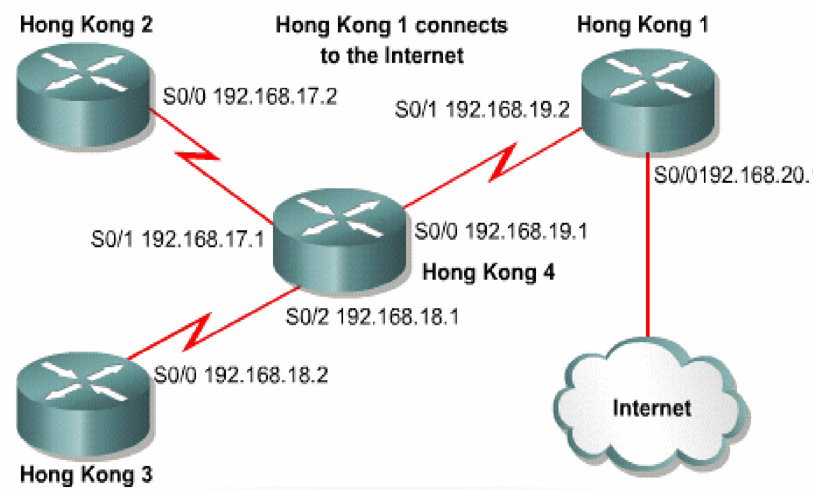

Configure Hong Kong 2, Hong Kong 3, and Hong Kong 4 using **ip default-network 192.168.20.0**

Default routes (cont.)

- The **ip default-network** command establishes a default route in networks using dynamic routing protocols:
  - Router(config)#**ip default-network 192.168.20.0**
- One example is a router that connects to the Internet. This is called the default route for the router. All the packets that are not defined in the routing table will go to the nominated interface of the default router.
- In Figure, Hong Kong 2 and Hong Kong 3 would use Hong Kong 4 as the default gateway. Hong Kong 4 would use interface 192.168.19.2 as its default gateway. Hong Kong 1 would route packets to the Internet for all internal hosts. To allow Hong Kong 1 to route these packets it is necessary to configure a default route as:
  - HongKong1(config)#**ip route 0.0.0.0 0.0.0.0 192.168.20.~~1~~**
- The zeros represent any destination network with any mask. Default routes are referred to as quad zero routes.

The ip route Command

Command	Description
Router(config)#**ip route 172.16.1.0 255.255.255.0 172.16.2.1**	IP network number or subnet number defined as the default
255.255.255.0	Subnet mask indicates that 8 bits of subnetting are in effect
172.16.2.1	ip address of next-hop router in the path to the destination

The ip default-network Command

Command	Description
Router(config)#**ip default-network 192.168.20.0**	IP network number or subnet number defined as the default

Configuring Default Route

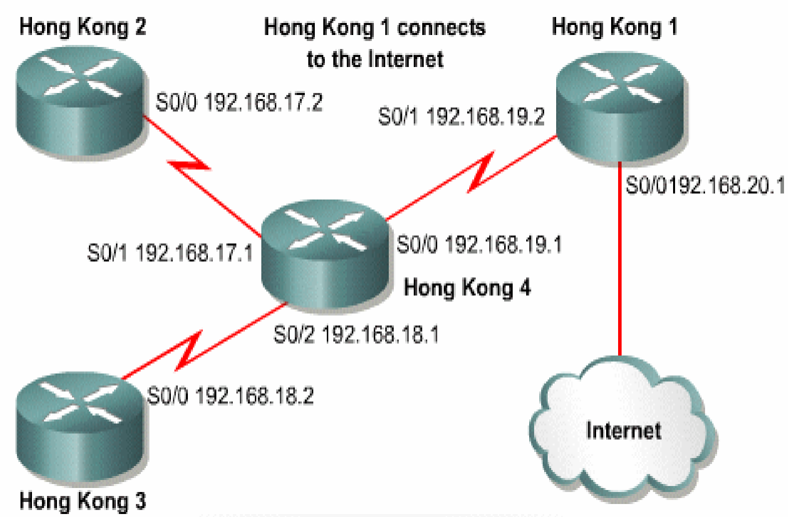

Configure Hong Kong 2, Hong Kong 3, and Hong Kong 4 using `ip default-network 192.168.20.0`

Adding a default Routes to

```
ISP
router rip
   redistribute static
   network 10.0.0.0
   network 192.168.4.0
   version 2
   no auto-summary
   default-information originate

ip route 207.0.0.0 255.0.0.0 null0
ip route 0.0.0.0 0.0.0.0 10.0.0.2
   etherenet0
```

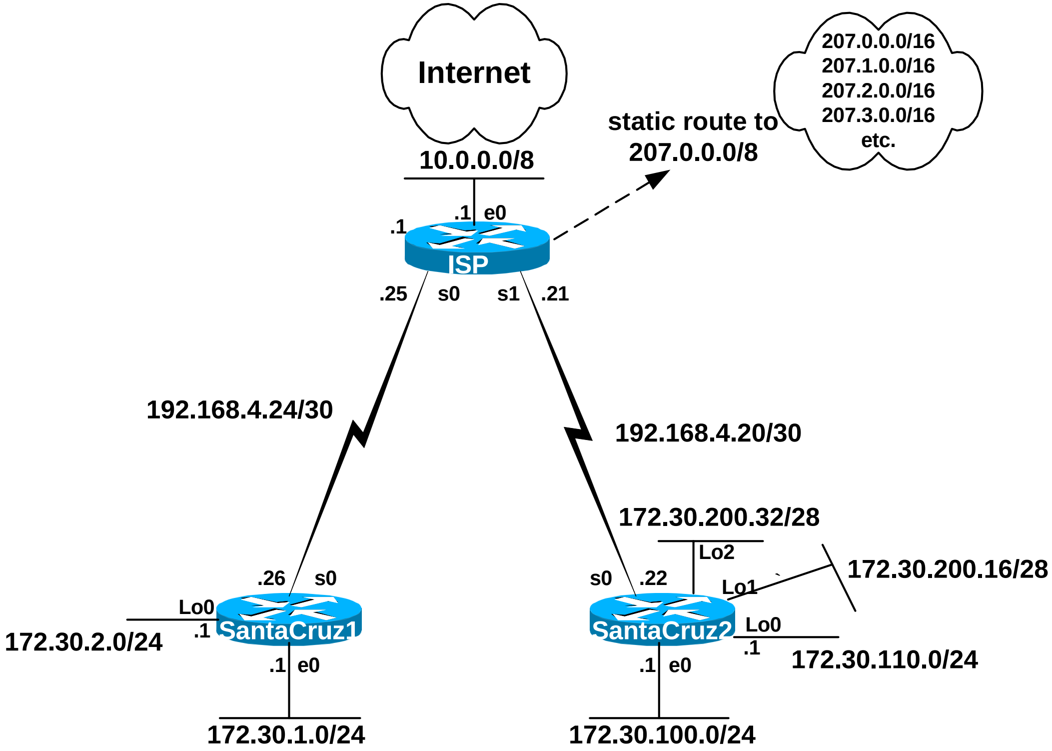

Labs

- **Lab:**
  - **1.1.4 Calculating VLSM Subnets**
  - **1.2.3 Review of Basic Router Configuration with RIP**
  - **1.2.4 Converting RIP v1 to RIP v2**
  - **1.2.5 Verifying RIP v2 Configuration**
  - **1.2.6 Troubleshooting RIP v2 using Debug**
- **e-Lab:**
  - **1.2.3 Review of Basic Router Configuration including RIP**
  - **1.2.4 Converting RIP v1 to RIP v2**
  - **1.2.6 RIP v2 using Debug**