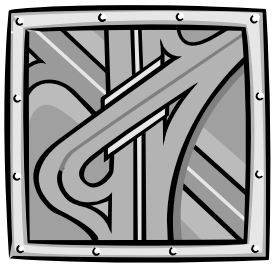

RIP Version 2

The Classless Brother

Why RIPv2

- **Need for subnet information and VLSM**
- **Need for Next Hop addresses for each route entry**
- **Need for external route tags**
- **Need for multicast route updates**
- **RFC 2453**

Because Subnetting and VLSM get more important RIPv2 was created. RIPv2 was introduced in RFC 1388, "RIP Version 2 Carrying Additional Information", January 1993. This RFC was obsolete in 1994 by RFC 1723 and finally RFC 2453 is the final document about RIPv2.

In comparison with RIPv1 the new RIPv2 also support several new features such as, routing domains, route advertisements via EGP – protocols or authentication.

Multicast Updates

- RIPv1 used DA=broadcast
  - Seen by each IP host
  - Slows down other IP stations
- RIPv2 uses DA=224.0.0.9
  - Only RIPv2 routers will receive it

RIPv2 uses the IP-Address 224.0.0.9 to transfer his routing updates. With this advantage only RIPv2 routers see this messages, and will not slow down the different station (RIPv1 and broadcast addresses).

RIPv2 is also an alternative choice to OSPF.

Message Format

Command	Version	Unused or Routing Domain
Address Family Identifier		Route Tag
IP Address		
Subnet Mask		
Next Hop		
Metric		
Address Family Identifier		Route Tag
IP Address		
Subnet Mask		
Next Hop		
Metric		
. . . . .		

Up to 25 route entries

RIPv2 utilizes the unused fields of the RIPv1 message-format. New fields are the “routing tag”, “subnet mask” and the “next hop”.

Version and Routing Domain

- RIPv1 used version "1"
- RIPv2 uses version "2" (*surprise*)
- According RFC the next two bytes are unused
- However, some implementations carry the routing domain here
  - Simply a process number

The routing domain indicates the routing-process for which the routing-update is destined. Now routers can support several domains within the same subnet.

Subnet Mask

- RIPv2 is a classless routing protocol
- For each route a subnet mask is carried
- Discontinuous Subnetting and VLSM is supported

Remember RIP is an classful routing protocol, because RIPv1 does not bind subnet-masks to the routes. So RIPv1 assumes classful addressing.

Next Hop

- **Identifies a better next hop address than implicitly given (SA)**
  - **Only if one exists (better metric)**
  - **0.0.0.0 if the sender is next hop**
- **Especially useful on broadcast multi-access network for peering**
  - **Indirect routing on a broadcast segment would be ...silly.**

With the „next hop“ router announces which networks can be reached over other routers.

Note that the next-hop router must be located in the same subnet as the sender of the routing-update.

Route Tag

- To distinguish between internal routes (learned via RIP) and external routes (learned from other protocols)
- Typically AS number is used
  - Not used by RIPv2 process
  - External routing protocols may use the route tag to exchange information across a RIP domain

Route Tag contains the autonomous system number for EGP and BGP. When the router receive a routing-update with a routing tag unequal zero, the associated path must be
distributed to other routers. In that way interior routers notice the existence of exterior networks (tagging exterior routes).

For example if routes were redistributed from EGP into RIPv2, these routes can be tagged.

Next Hop and Route Tag

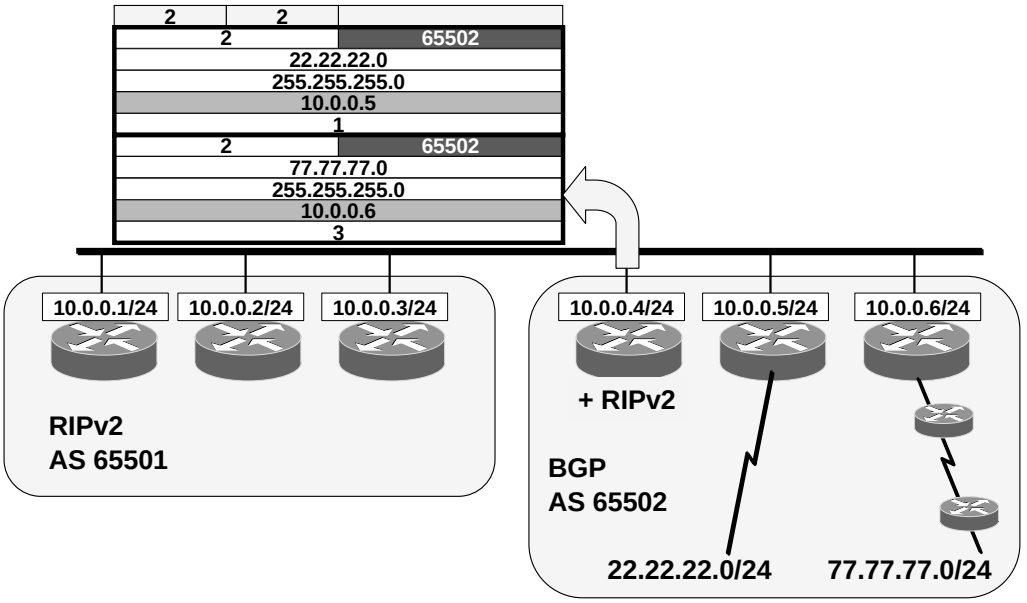

(C) Herbert Haas 2005/03/11 9

In the picture above there are two different autonomous systems on the same LAN. The routers in the first AS use RIPv2 the second AS use BGP. Each entry assigned a AS number (65501/65502). The Left AS could apply policies on these special (external) routes or redistribute them with BGP to some other ASs. Note that only 10.0.0.4 speaks RIPv2, so for efficiency only this one advertises the external routes (22.22.22.0/77.77.77.0) but by indicating the true next hops. This is an important special rule on shared medium (true next hops must be indicated) !

Authentication

- Hackers might send invalid routing updates
- RIPv2 introduces password protection as authentication
- Initially only Authentication Type 2 defined
  - 16 plaintext characters (!)
- RFC 2082 proposes keyed MD-5 authentication (Type 3)
  - Multiple keys can be defined, updates contain a key-id
  - And a unsigned 32 bit sequence number to prevent replay attacks
- Cisco IOS supports MD5 authentication (Type 3, 128 bit hash)

IF a router receives routing updates without valid authentication are ignored by the receiving router, because only trusted router are accepted.

When using MD5 authentication, the first but also the last routing entry space is used for authentication purposes. The MD5 hash is calculated using the routing update plus a password. Thus, authentication and message integrity is provided.

The "Authentication Type" is Keyed Message Digest Algorithm, indicated by the value 3 (1 and 2 indicate "IP Route" and "Password", respectively)

Authentication

Command	Version	Unused or Routing Domain
0xFFFF		Authentication Type
Password		
Password		
Password		
Password		
Address Family Identifier		Route Tag
IP Address		
Subnet Mask		
Next Hop		
Metric		
. . . . .		

Up to 24 route entries

The picture above shows a RIPv2 Message which contains authentication entry's. The password is only a plain text. If the password is under 16 octets, it must be left-justified and padded to the right with nulls.

Key Chain

- Cisco's implementation offers key chains
  - Multiple keys (MD5 or plaintext)
  - Each key is assigned a lifetime (date, time and duration)
- Can be used for migration
  - Key management should rely on Network Time Protocol (NTP)

Several independent routing domains running RIPv2 with different process numbers ("routing domain"). With using key chains this domains can be work together (synchronize) at a special time or date.

RIPv1 Inheritance (1)

- All timers are the same
  - UPDATE
  - INVALID
  - HOLDDOWN
  - FLUSH
- Same convergence protections
  - Split Horizon
  - Poison Reverse
  - Hold Down
  - Maximum Hop Count (also 16 !!!)

RIPv1 uses many timers to regulate its performance. This timers are the same in RIPv2. The routing update timer is set to 30 seconds, with a small random amount of time added whenever the timer is reset. A route is declared invalid without being refreshed by routing updates during 90 seconds. The “holddown” status retains 180 seconds. In this time a router ignore update messages about a special network. After 240 Seconds (Flush timer) a non-refreshed routing table entry will be removed.

RIPv2 also using the same convergence protections such as Split Horizon, Hold Down, etc. Note that the Maximum Hop Count is still **16** to be backwards compatibility.

RIPv1 Inheritance (2)

- **Same UDP port 520**
- **Also maximum 25 routes per update**
  - **Equally 512 Byte payloads**

RIPv2 also inherit the bad consequences of this small routing updates.

What happened if we want to advertise MANY routes with many single updates. There will be a big overhead (IP + UDP + RIP header).

RIPv1 Compatibility

- **RIPv1 Compatibility Mode**
  - **RIPv2 router uses broadcast addresses**
  - **RIPv1 routers will ignore header extensions**
  - **RIPv2 performs route summarization on address class boundaries**
    - **Disable: `(config-router)# no auto-summary`**
- **RIPv1 Mode**
  - **RIPv2 sends RIPv1 messages**
- **RIPv2 Mode**
  - **Send genuine RIPv2 messages**

RIPv2 is totally backwards compatible with existing RIP implementations. There is also an compatibility switch, which allows to chance between three different settings:

1. RIP-1 Modus. Only RIP-1 packets are sent
2. RIP-1 compatibility Modus. RIP-2 packets are broadcast
3. RIP-2 Modus. RIP-2 packets are multicast.

The recommended default for this switch is RIP-1 compatibility.

Classful Routing

- **routing protocols like RIP, IGRP cannot carry subnetmask information in routing updates**
- **this has several consequences**
  - if a given class A, B or C address is subnetted the subnetmask must be constant in the whole area
    - no variable length subnet mask (VLSM) can be used
  - if a routing update is sent to an interface with an network number different to the subnetted network
    - only the major class A, B or C network number will be announced
    - route summarization will be performed on class boundaries
    - hence a subnetted area must be contiguous
  - classful routing

Classful Routing

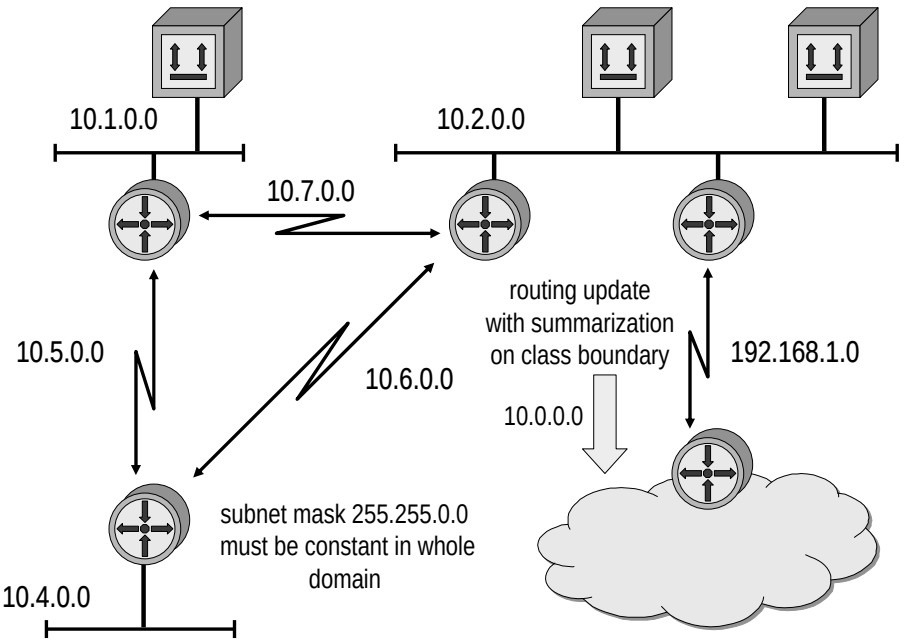

Discontiguous Subnetting Classful

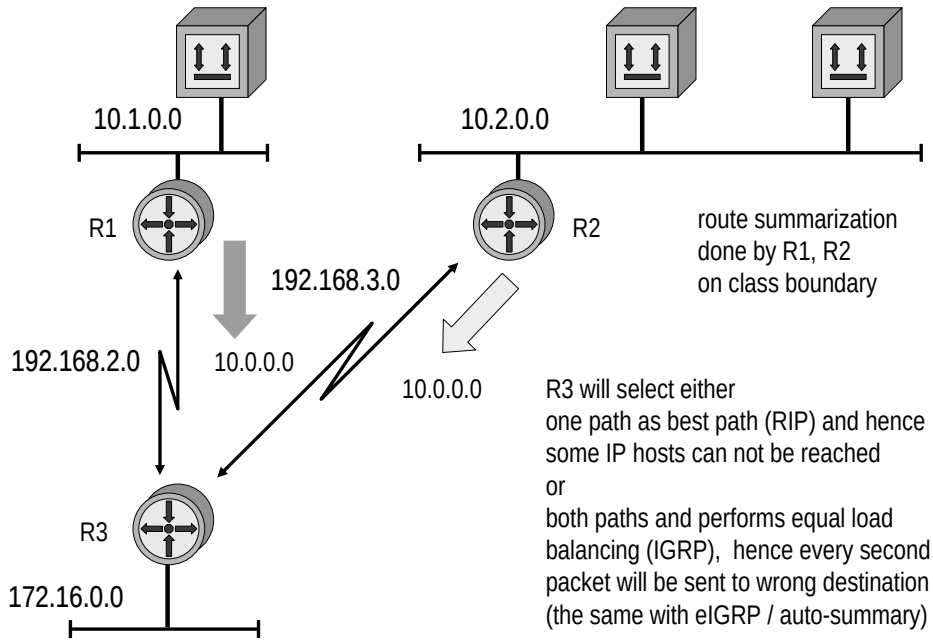

route summarization
done by R1, R2
on class boundary

R3 will select either
one path as best path (RIP) and hence
some IP hosts can not be reached
or
both paths and performs equal load
balancing (IGRP), hence every second
packet will be sent to wrong destination
(the same with eIGRP / auto-summary)

Routing Table Lookup (Classful)

- **assumption:**
  - IP datagram with a given IP address is received by a classful router
- **IP address is interpreted as class A, B or C**
  - the major net is determined
- **next a lookup in the routing table for the major net is performed**
  - if there is no entry the IP datagram will be discarded
- **if there is a match the IP address is compared to every known subnet of this major network**
  - if there is no such subnet the IP datagram will be discarded

Routing Table Lookup (Classful) cont.

- **hence a problem may arise with default routing**
  - if the major network is known by the router, but the subnet does not exist, the IP datagram will be discarded even if a default route exists
- **therefore**
  - subnetted area must be contiguous
  - all subnets of a given major net must be reachable using only paths with these subnet-IDs
- **remark:**
  - Cisco´s configuration command ***ip classless*** will change such an behavior in case of default routing to the behavior of classless routing even if classful routing is used

Classful route match (1/2)

- **1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)**
- **2) If the classful network number is listed in the routing table:**
  - If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route
  - If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)

Example:

IP Packet
DA = 10.35.72.26
SA = …

Routing Table:

```
10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16   via 172.17.7.19
10.31.0.0/16   via 172.17.8.31
10.34.0.0/16   via 172.18.1.254
10.35.0.0/16   via 192.186.176.254
0.0.0.0/0      via 172.19.41.254
```

Classful route match (2/2)

- **1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)**
- **2) If the classful network number is listed in the routing table:**
  - If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route
  - If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)

Example:

IP Packet
DA = 10.35.72.26
SA = …

Routing Table:

```
10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16   via 172.17.7.19
10.31.0.0/16   via 172.17.8.31
10.34.0.0/16   via 172.18.1.254
0.0.0.0/0      via 172.19.41.254
```

DISCARD THE PACKET (!)

Classless Routing

- **routing protocols like RIPv2, OSPF, eIGRP can carry subnet mask information in routing updates**
- **this has several advantages**
  - variable length subnet mask (VLSM) can be used
    - subnetting of a given address can be done according to the number of hosts required on a certain subnet
    - more efficient use of address space ⇨ sub-subnetting
  - route summarization can be performed on any address boundary and not only on class boundaries
    - a routing update contains prefix (relevant part of IP address) and length (number of ones used in subnetmask)
    - supernetting
      - actual subnetmask is smaller than natural subnetmask of given class

Classless Routing

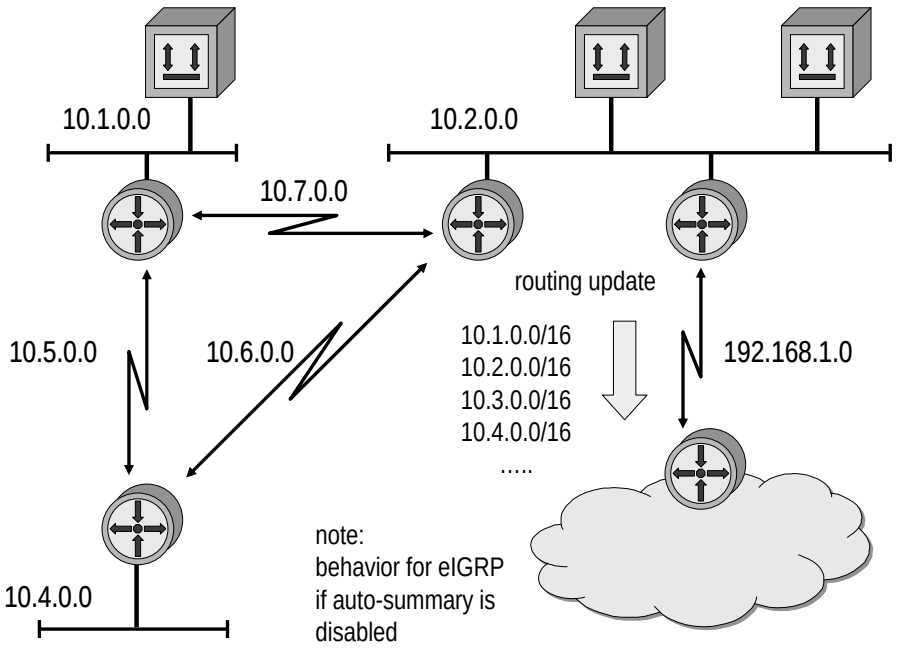

Discontiguous Subnetting Classless

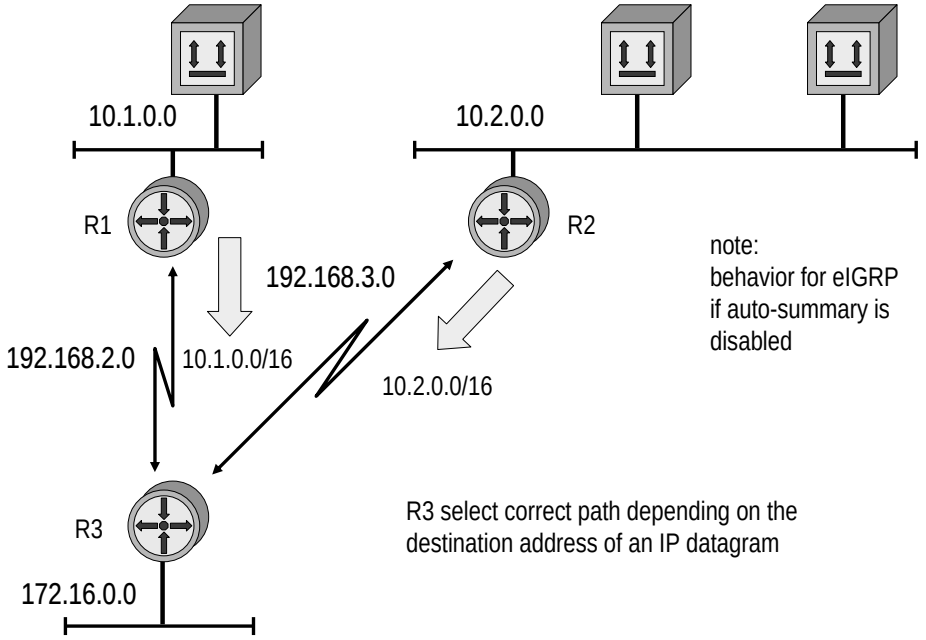

Routing Table Lookup (Classless)

- **assumption:**
  - IP datagram with a given IP address is received by a classless router
- **IP address is not interpreted as class A, B or C**
- **a lookup in the routing table for the best match for this IP address is performed**
  - IP prefixes of the routing table are compared with the given IP address bit by bit from left to right
  - IP datagram is passed on to the network which matches best
  - “Longest Match Routing Rule”
  - result: IP addresses with any kind of subnetting can be used independent from the underlying network topology

Classless routing: Longest match

- **The subnet mask of each route entry tells how many bits must be compared with the IP-packet's destination address**
- **The router takes the route with the longest match**

Example:

IP Packet
DA = 10.35.72.26
SA = …

Routing Table:

```
10.0.0.0/8      via 172.16.1.1
10.22.0.0/16    via 172.17.7.19
10.31.0.0/16    via 172.17.8.31
10.34.0.0/16    via 172.18.1.254
10.35.0.0/16    via 192.186.176.254
10.35.64.0/19   via 192.186.177.254
10.35.192.0/19  via 172.19.54.1
```

VLSM Example (1)

- **First step 6 bit subnetting of 172.16.0.0**
  - 172.16.0.0 with 255.255.252.0 (172.16.0.0 / 22)
  - subnetworks:
    - 172.16.0.0
    - 172.16.4.0
    - 172.16.8.0
    - 172.16.12.0
    - 172.16.16.0

      …….
    - 172.16.248.0
    - 172.16.252.0
  - 64 subnetworks each of them capable of addressing 1022 IP systems

VLSM Example (2)

- **next step sub-subnetting**
  - basic subnet 172.16.4.0 255.255.252.0 (172.16.4.0 / 22)
  - sub-subnetworks with mask 255.255.255.252 (/ 30):
    - 172.16.4.0 / 30
    - 172.16.4.4 / 30
      - 172.16.4.4 net-ID
      - 172.16.4.5 first IP host of subnet 172.16.4.4
      - 172.16.4.6 last IP host of subnet 172.16.4.4
      - 172.16.4.7 directed broadcast of subnet 172.16.4.4
    - 172.16.4.8 / 30
    - 172.16.4.12 / 30
    - …….
    - 172.16.4.252 / 30
  - 64 sub-subnetworks each of them capable of addressing 2 IP systems

VLSM Example (3)

- **next step sub-subnetting**
  - basic subnet 172.16.8.0 255.255.252.0 (172.16.8.0 / 22)
  - sub-subnetworks with mask 255.255.255.0 (/ 24):
    - 172.16.8.0 / 24
    - 172.16.9.0 / 24
      - 172.16.9.0 net-ID
      - 172.16.9.1 first IP host of subnet 172.16.9.0
      - ----------
      - 172.16.9.254 last IP host of subnet 172.16.9.0
      - 172.16.9.255 directed broadcast of subnet 172.16.9.0
    - 172.16.10.0 / 24
    - 172.16.11.0 / 24
  - 4 sub-subnetworks each of them capable of addressing 254 IP systems

VLSM Example (4)

- **no sub-subnetting for basic subnet 172.16.12.0**
  - 172.16.12.0 with 255.255.252.0 (172.16.12.0 / 22)
    - 172.16.12.0 net-ID
    - 172.16.12.1 first IP host of subnet 172.16.12.0
    - ----------
    - 172.16.15.254 last IP host of subnet 172.16.12.0
    - 172.16.15.255 directed broadcast of subnet 172.16.12.0
  - one subnetwork capable of addressing 1022 IP systems

VLSM Classless

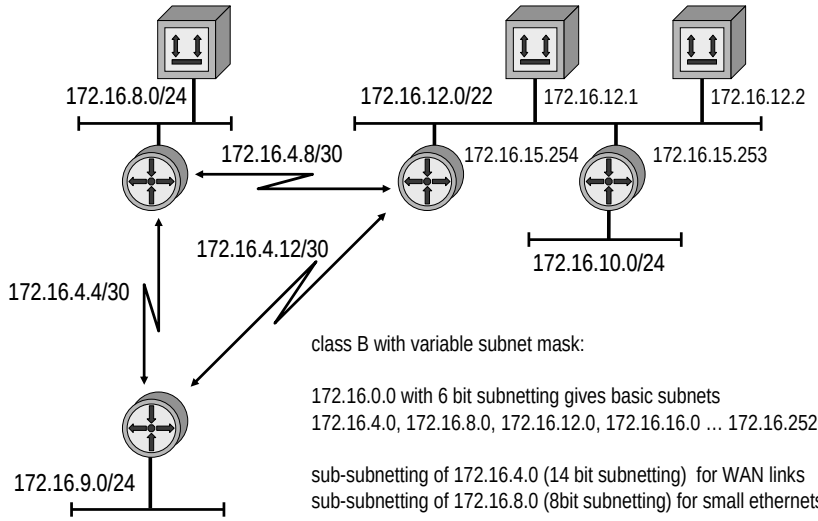

class B with variable subnet mask:

172.16.0.0 with 6 bit subnetting gives basic subnets
172.16.4.0, 172.16.8.0, 172.16.12.0, 172.16.16.0 ... 172.16.252.0

sub-subnetting of 172.16.4.0 (14 bit subnetting) for WAN links
sub-subnetting of 172.16.8.0 (8bit subnetting) for small ethernets

Summary

- Most important: RIPv2 is classless
  - Subnet masks are carried for each route
- Multicasts and next hop field increase performance
- But still not powerful enough for large networks

Quiz

- What is a routing domain?
- Why is "infinity" still 16?