Chapter 6

First Hop Redundancy Protocols

NET3011 – 17W

VRRP – Basics (1)

- VRRP is the standards based equivalent of HSRP RFC3768 (v2 in 2004) and RFC5798 (v3 in 2010)
- General principle of operation is the same: multiple routers managing a (gateway) IP & MAC

- An opportunity to put your shift in perspective to the test: What questions should we be asking now to understand this next protocol?

VRRP – Basics (2)

Consider the following as a candidate list for learning a protocol. Is it complete or are there items you would add or remove?

1. What is the purpose? What is achieved?
2. Hello
   - Are neighbours configured or discovered?
3. Roles
   - number of participants
   - number & type of elected roles
4. Elections
   - priorities
   - pre-emptive?
   - tie breaker
5. Timers
   - type, number
   - units (msec, sec, or hrs?)
   - reconvergence time?
6. Load balancing
7. States and/or Finite State Machine (FSM)
8. Messages
   - types
   - specific ports and/or addresses?
   - unicast, broadcast, or multi-cast?
   - special or important flags?
9. Authentication
10. Versions
    - most commonly used?
    - backwards compatibility
11. Best commands
    - for viewing overall state
    - for debugging

VRRP – Basics (3)

- Hopefully your questions focus on:
  - multiple roles, elections, priorities, timers, preempt, MAC addresses, groups √, tracking √, etc.
- The concept and operation of groups, for load balancing, is identical to HSRP
- The concept and operation of tracking is identical
  … except that VRRP only tracks objects
  … except that a tracked object can be an interface!

VRRP – Steady-state Operation

- **Master** and **Backup** *must* have L3 reachability between themselves to maintain correct VRRP state
- VRRP Hello sent to 224.0.0.18 protocol 112 (~~TCP~~, ~~UDP~~)
- Hello's sent every **1 sec**, by default

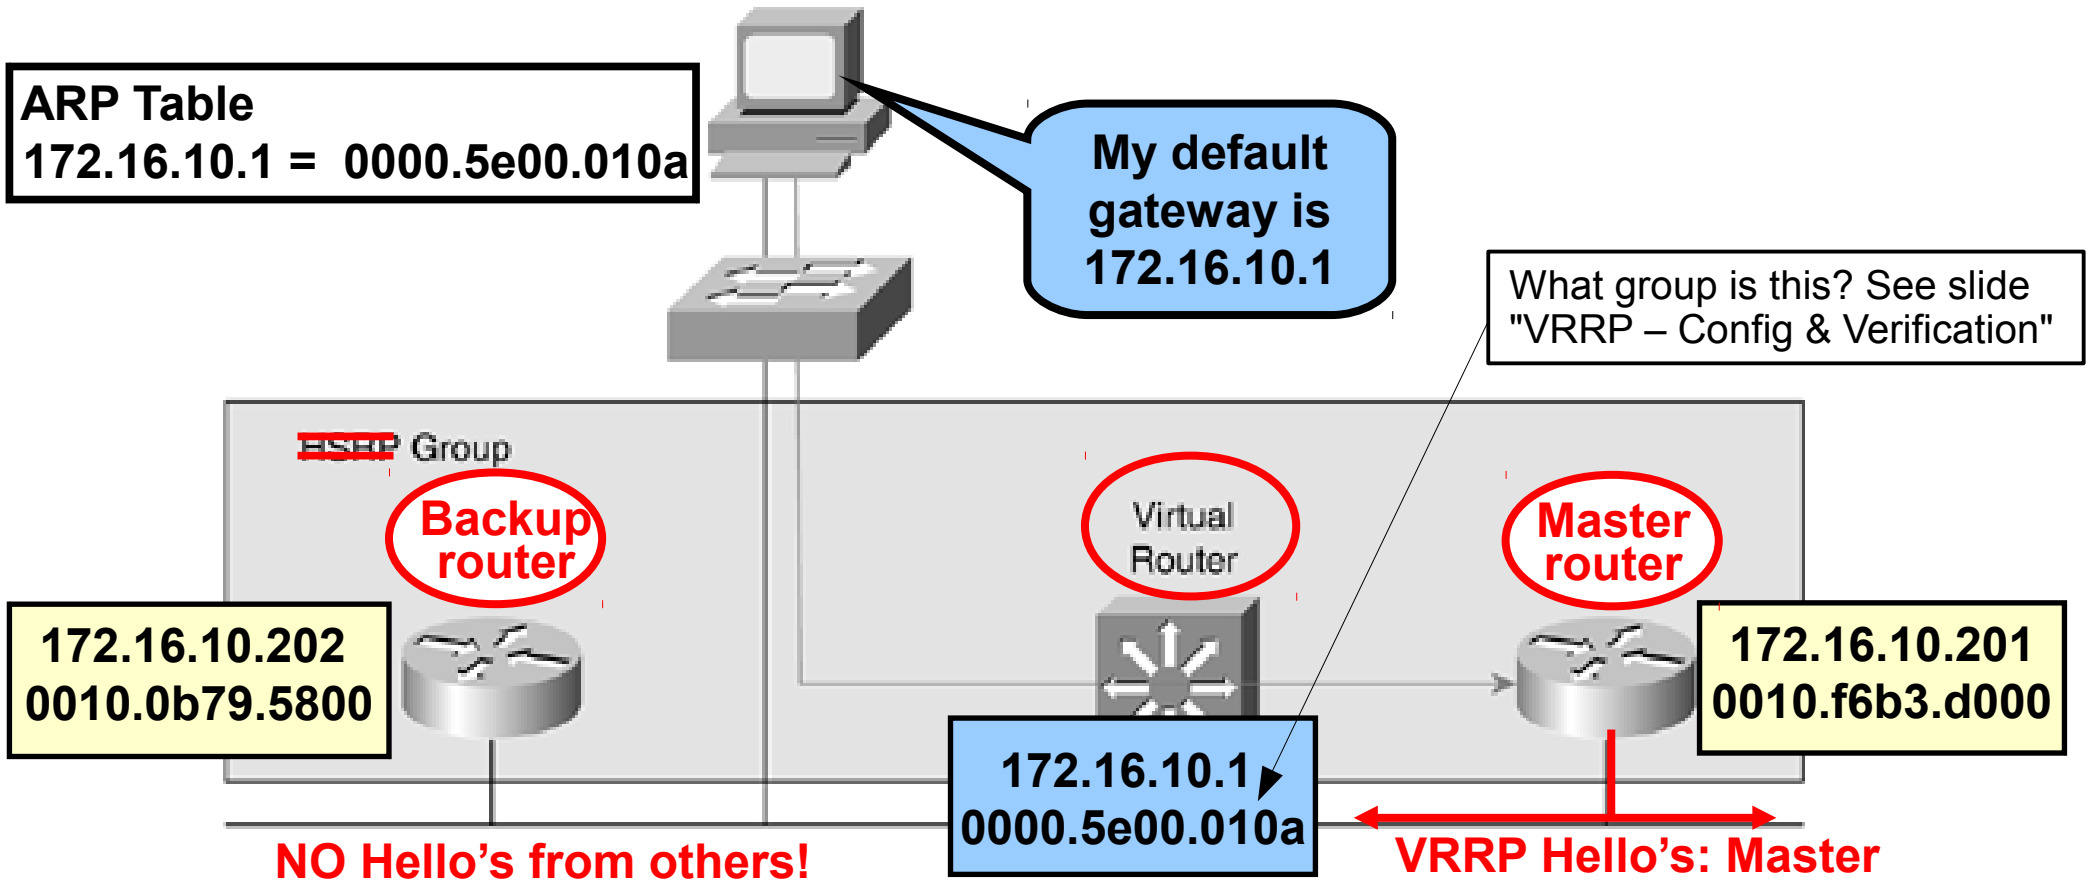

VRRP – Roles

- 1 router pro-actively manages the gateway IP & MAC
  - One *Master* router plus any number of *Backup*s
- Role of each router determined by the protocol
  - Master router:
    - responds to ARP requests for (virtual) IP
    - processes frames addressed to virtual MAC
  - Backup routers:
    - monitor the Master router
    - have a rapid election after the timeout period
- ONLY Master router transmits Hello messages
  - all Backups are silent and don't send Hello's

VRRP – Elections

- *When* an election occurs, highest priority wins
  - priority range is 0-255, default = 100, but configurable range = 1-254
  - tie-breaker is highest IP address on the link
- **`preempt`** is enabled by default, so routers with a better priority always force an election
- Other routers in group simply monitor Hello msgs, but don't play any other role until an election occurs

VRRP – Priorities and IP (virtual/real)

- Priorities are the same as HRSP with 2 exceptions
  - When a VRRP group's managed IP is used as the real IP for a VRRP router's interface it is automatically given a priority of 255

```
L3(config)#interface fa 0/1
L3(config-if)#ip address 10.0.0.1 255.255.255.0
L3(config-if)#vrrp 1 ip 10.0.0.1
```

  - When a Master undergoes an orderly shutdown, it advertises a priority of 0 to force an election to instate a new Master without holdtime delay

VRRP – Timers

- Hello time, or Advertisement interval, is 1 sec by default
- Fail-over time is 3 Advertisement intervals … plus a little bit extra
- The extra bit of time is called skew time:
  - skew time = ((256 – priority) / 256) secs (if you work it through, that's 4-1000msec)
  - higher priority => lower skew time => quicker
  - this guards against race conditions and allows the highest priority Backup to assume the Master role more quickly (uncontested?)

VRRP – Load Balancing

Do you see the mixing of priority methods in the config below?

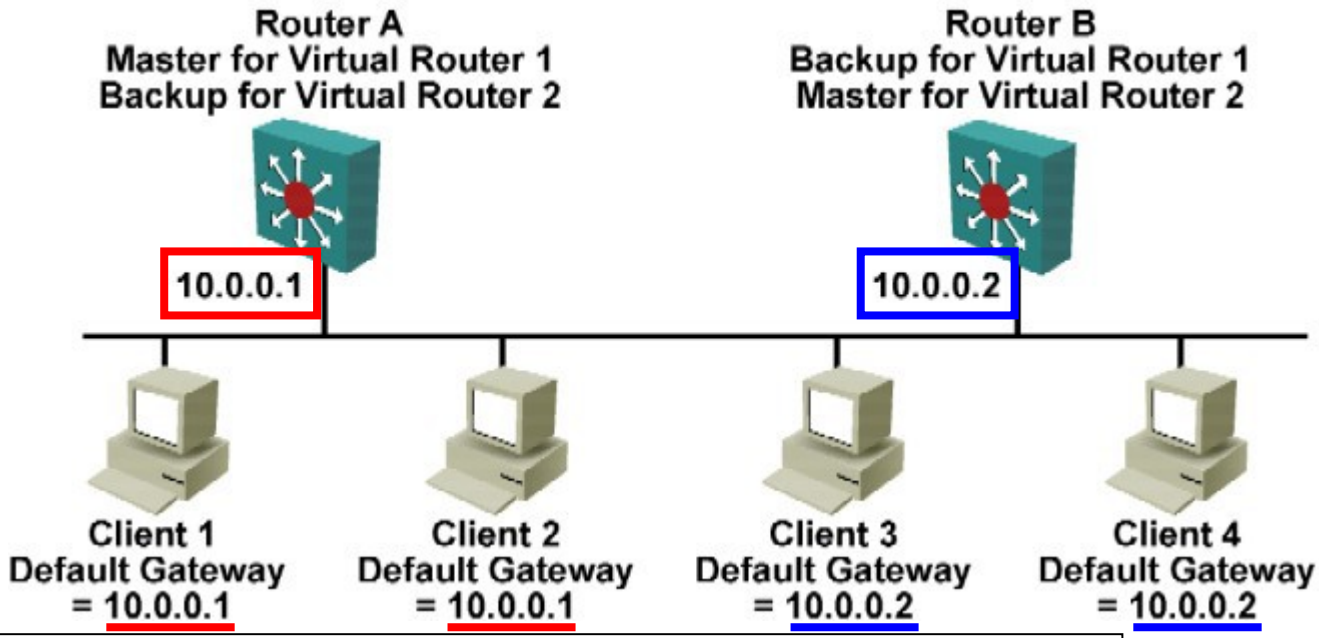


```
RouterA(config)#interface fa 0/1
RouterA(config-if)#ip address 10.0.0.1 255.255.255.0
RouterA(config-if)#vrrp 1 ip 10.0.0.1
RouterA(config-if)#vrrp 2 ip 10.0.0.2
RouterA(config-if)#vrrp 2 priority 110
```

```
RouterB(config)#interface fa 0/1
RouterB(config-if)#ip address 10.0.0.2 255.255.255.0
RouterB(config-if)#vrrp 1 ip 10.0.0.1
RouterB(config-if)#vrrp 1 priority 110
RouterB(config-if)#vrrp 2 ip 10.0.0.2
```

VRRP – Config and Verification (1)

- All commands are identical except use "vrrp" – as with HSRP, MAC addrs are calculated
- Extra: timers can be learned from Master

```
L3(config-if)# vrrp {grp} timers learn
```

```
L3# show vrrp
 Ethernet1/0 - Group 1
  State is Master
  Virtual IP address is 12.0.0.2
  Virtual MAC address is 0000.5e00.0101
  Advertisement interval is 1.000 sec
  Preemption is enabled
    min delay is 0.000 sec
  Priority is 105
    Track object 1 state Down decrement 15
  Master Router is 12.0.0.1 (local), priority is 105
  Master Advertisement interval is 1.000 sec
  Master Down interval is 3.531 sec
```

Where did this value come from?

VRRP – Study shortcut

Compared to HSRPv1, terminology and functional differences …

1. Master router serves (possibly) virtual IP [vs. Active router]
   all others are in backup state [1 standby + listeners]
2. Group numbers 1-255 [0-255; though v2 allows 0-4095]
3. priorities 1-254, with 0 & 255 treated specially [0-255]
4. MAC ID: 0000.5e00.01xx [0000.0c07.acxx]
5. Hellos every 1 sec by default [3 sec default]
6. Hellos sent by Master only [Hellos sent by active & standby, both]
7. Master can publish timer values [individually set timers must agree]
8. Pre-emptive behaviour is the default, but can be disabled
   [pre-emptive behaviour must be explicitly configured]
9. Object tracking supported [interface & object tracking are available]
10. Multicast to 224.0.0.18; protocol 112 [224.0.0.2; 17[UDP], port 1985]

GLBP

GLBP – Basics

- Gateway Load Balancing Protocol manages IP & MAC, with additional benefit of built-in load balancing
- 4 routers (max) of the group pro-actively share the work:
  - only one is the *Active Virtual Gateway* (AVG) router, as well as being an *Active Virtual Forwarder* (AVF)
  - 3 more routers (max) act only as *Active Virtual Forwarders*
- Role of each router determined by the protocol
  - AVG:
    - assigns a virtual MAC to itself and each AVF
    - responsible for / responds to ARP requests for virtual IP, giving one of the AVF virtual MAC addresses
    - processes frames addressed to *it's* virtual MAC
  - AVF:
    - accepts and processes frames sent to *their* virtual MAC
    - monitors AVG router to take over after timeout period

GLBP – Load Balancing Concept

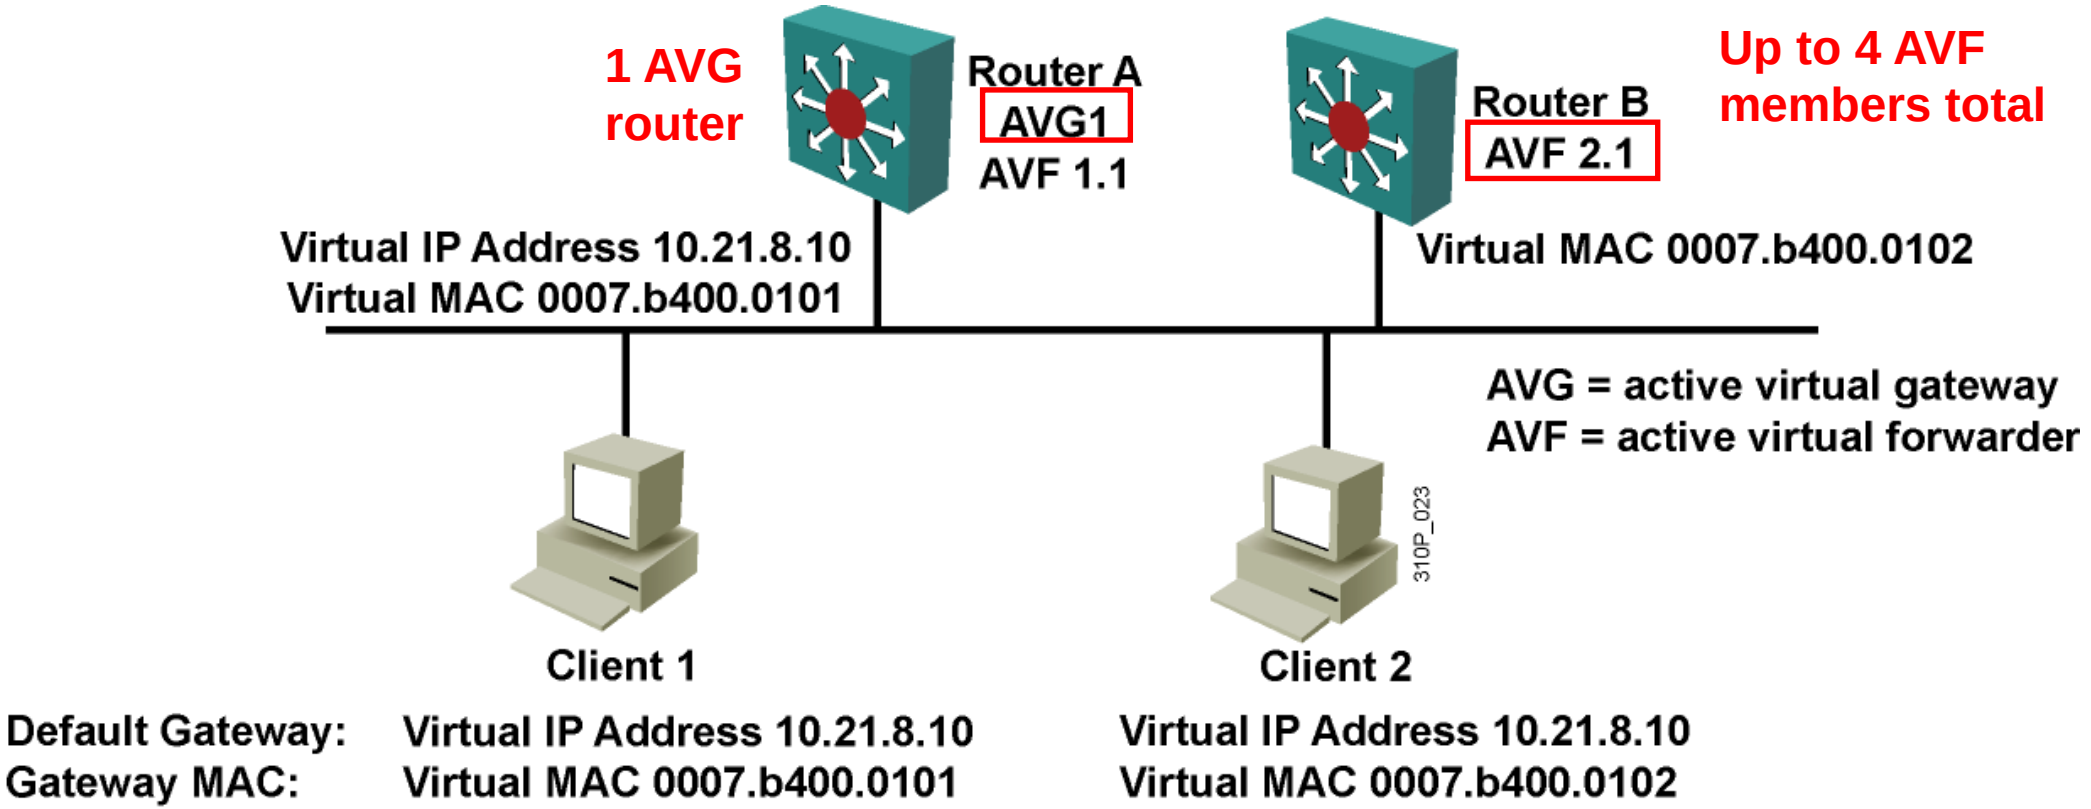

- The AVG automatically manages the virtual MAC addr assignments
  - calculated MAC address format is: 0007.b4xx.xxyy
  - xx.xx = (16 bits) six 0 bits, followed by 10 bit GLBP group number
  - yy = is the virtual forwarder number

GLBP – Load Balancing Operation (1)

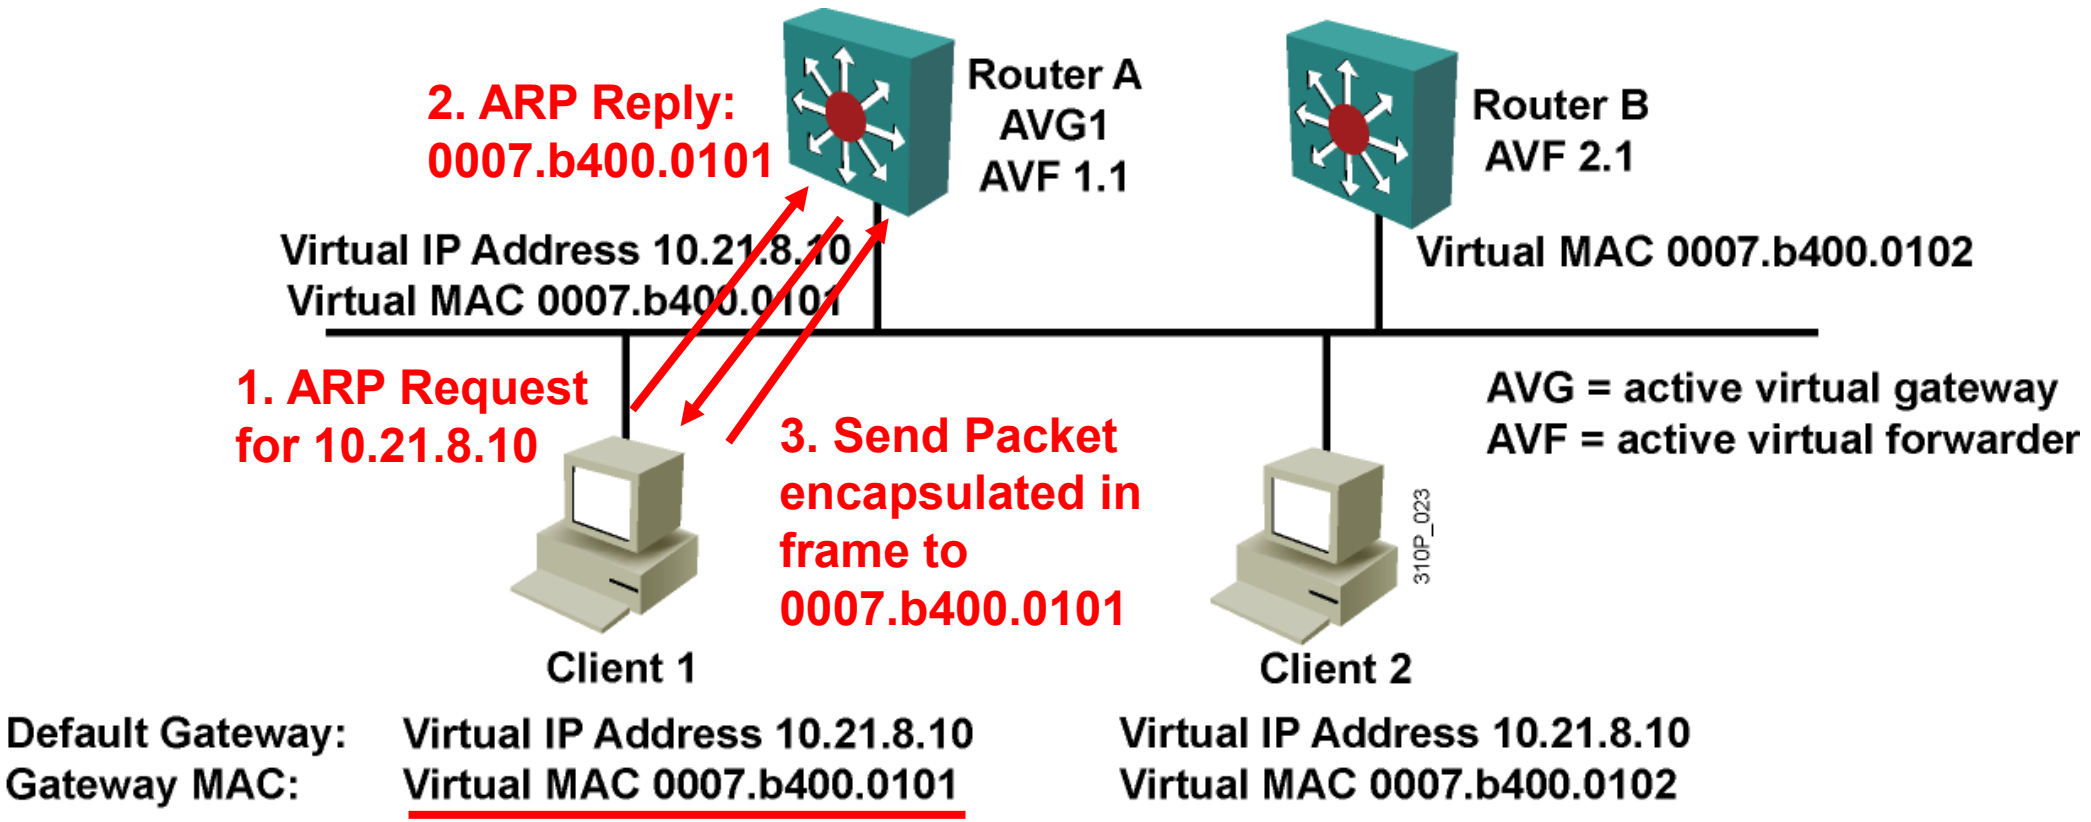

Ethernet Header			IP Header			
Destination MAC Add.	Source MAC Address	Type	Source IP Address	Destination IP Address	Rest of IP Hdr	Data
0007.b400.0101	000C.0417.91C1		10.21.8.101	8.8.8.8	…	

GLBP – Load Balancing Operation (2)

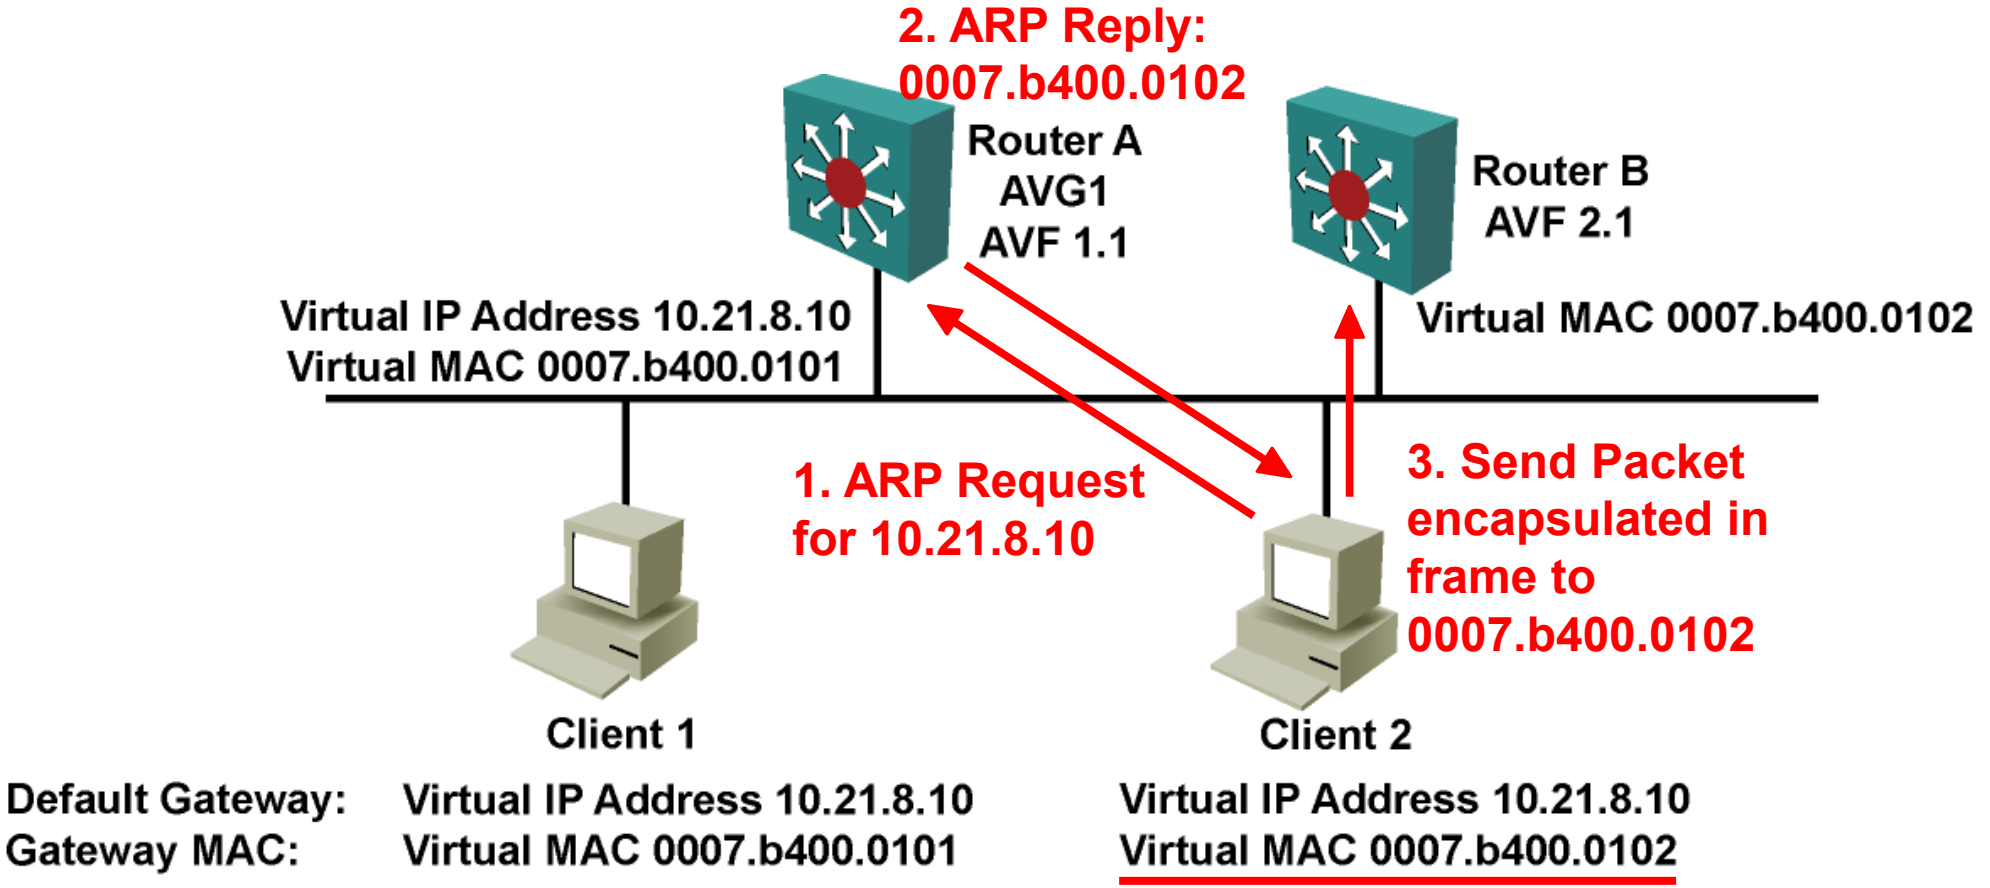

Ethernet Header			IP Header			
Destination MAC Add.	Source MAC Address	Type	Source IP Address	Destination IP Address	Rest of IP Hdr	Data
0007.b400.0102	000C.0417.91C2		10.21.8.102	8.8.8.8	…	

GLBP – Did you notice?

- With HRSP and VRRP, load balancing occurs by having *multiple* gateways (i.e. multiple virtual IPs)
- With GLBP, all hosts see a *single gateway* with a common IP address

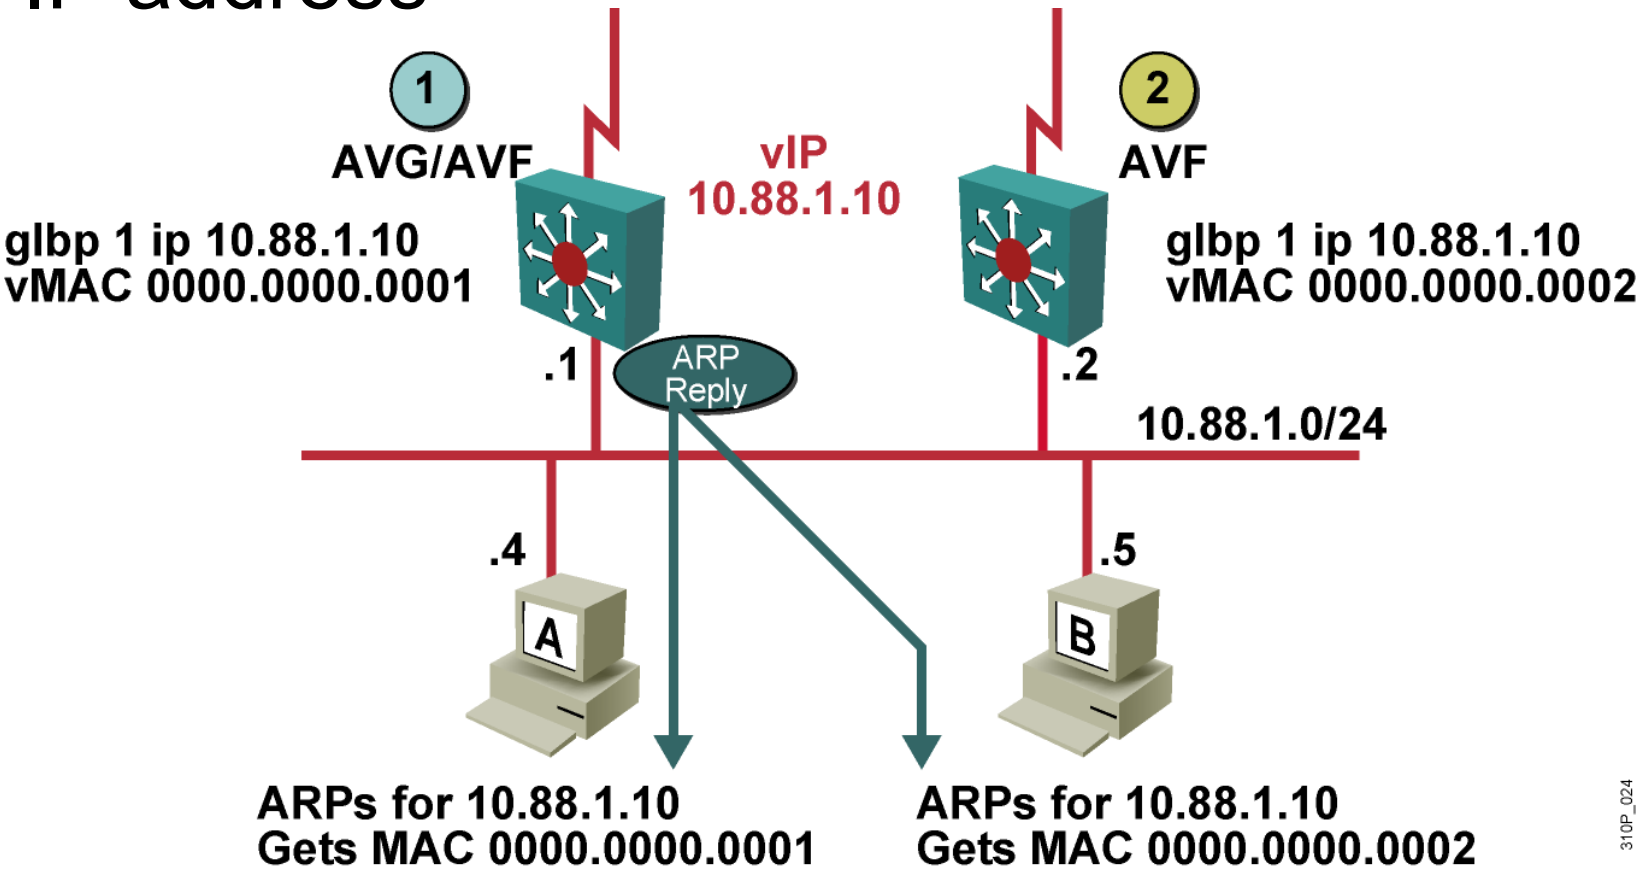

GLBP – Load Balancing Algorithms (1)

GLBP supports the following modes for load balancing:

- Round-robin load-balancing (default):
  - AVG gives out each virtual MAC address in turn for ARP replies to the virtual IP address
- Weighted load-balancing:
  - amount of load directed to a router depends on the weighting value advertised by that AVF
- Host-dependent load-balancing:
  - hosts are guaranteed to receive the same virtual MAC address as long as that MAC address is participating in the GLBP group

Perhaps the reason MS changed ARP cache refresh scheme to unicast as of Windows Vista/7?

GLBP – Load Balancing Algorithms (2)

Processing in the event of a failure is a little more involved than with HSRP

- If an AVG fails, the standby AVG must:
  – continue handling it's assigned virtual MAC
  – assume responsibility for handling the original AVG's assigned virtual MAC
  – assume responsibility for handling ARP requests as per it's configured load-balancing algorithm
  (… you did remember to configure all group members identically, didn't you?? If not, painful troubleshooting!)
- if either an AVF or the AVG fail:
  – "the load-balancing ratio is adjusted among the remaining AVFs so that the resources are used in the most efficient way" (FLG p. 295)

GLBP – Weighting (1)

- Weighting is related but not identical to priority; mentioned but not described in textbook

Consider the scenario:

– SW4 is forwarding

– Initial weight is 110

– SW4 tracks both Fa0/23 and Fa0/24 interfaces

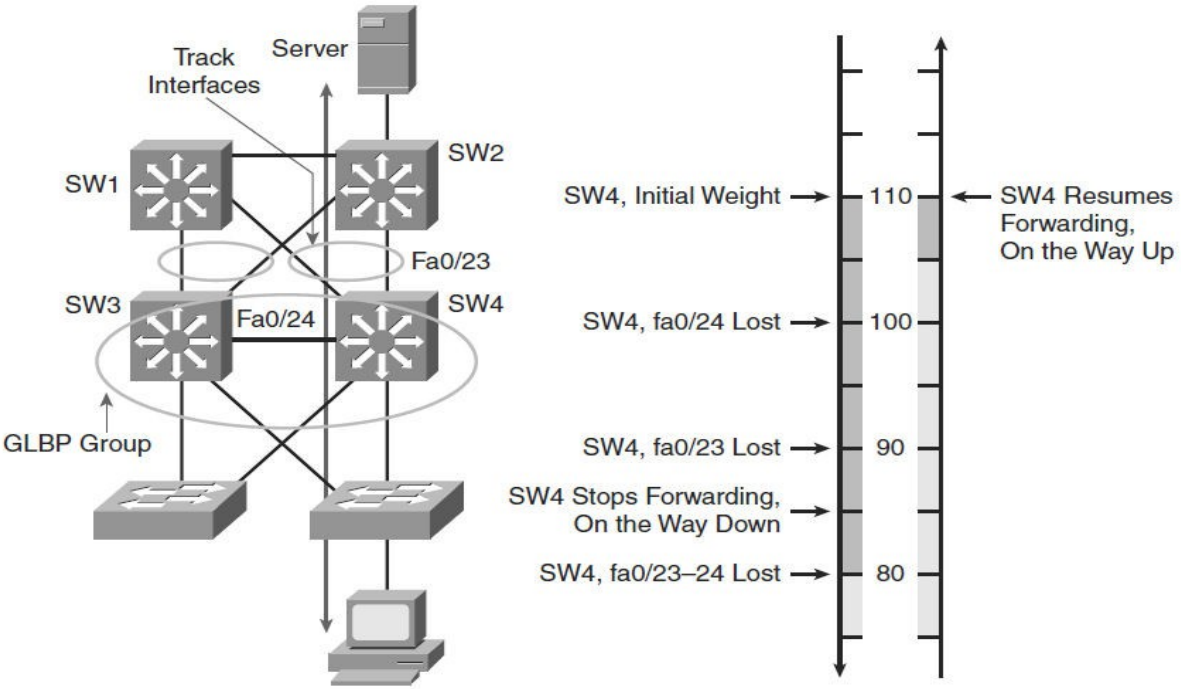

– Fa0/23 is the active interface. It is tracked with a decrement of 20. If it goes down, SW4 is decremented 20 points, bringing it's weight down to 90

– Fa0/24 is a backup interface. It is tracked with a decrement of 10. If it goes down, SW4 is decremented 10 points, bringing it's weight down to 100

– Losing both Fa0/23 and Fa0/24 brings SW4's weight down to 80

GLBP – Weighting (2)

Weighting specifies the cut-off / cut-in points for forwarding

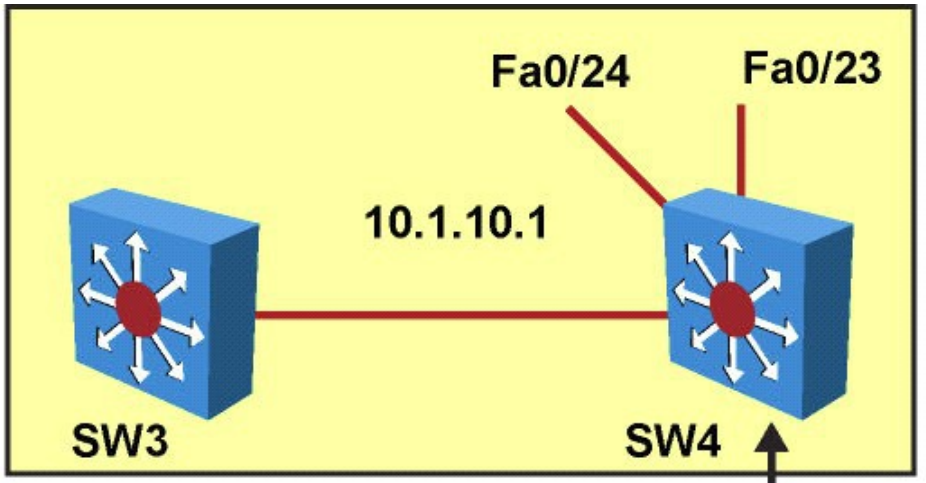


```
SW4(config)# track 90 interface fa0/24 line-protocol
SW4(config)# track 91 interface fa0/23 line-protocol
SW4(config)# interface vlan10
SW4(config-if)# ip address 10.1.10.2 255.255.255.0
SW4(config-if)# glbp 1 ip 10.1.10.1
SW4(config-if)# glbp 1 weighting 110 lower 85 upper 105
SW4(config-if)# glbp 1 timers msec 200 msec 700
SW4(config-if)# glbp 1 preempt delay minimum 300
SW4(config-if)# glbp 1 authentication md5 keystring xyz123
SW4(config-if)# glbp 1 weighting track 90 decrement 10
SW4(config-if)# glbp 1 weighting track 91 decrement 20
```

Sets range of priorities for which the interface will act as an AVF

GLBP – Basic Configuration

- Basic config is identical to HSRP / VRRP

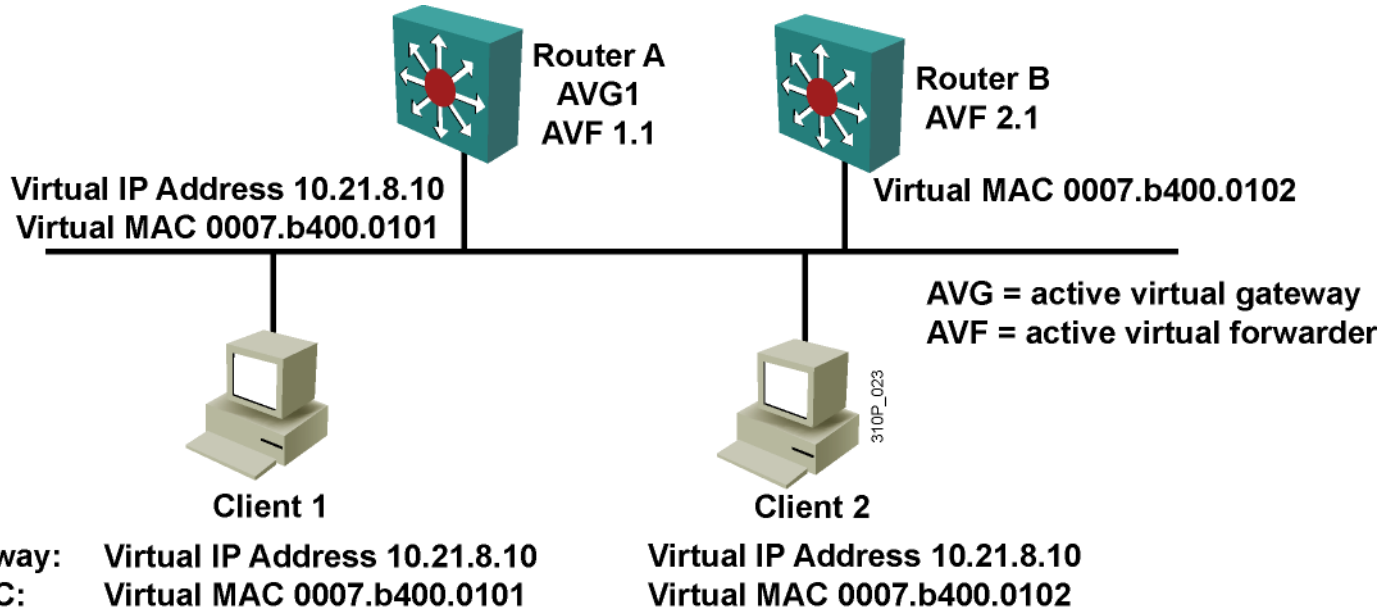


```
RtrA(config)#interface vlan 21
RtrA(config-if)#ip address 10.21.8.1 255.255.255.0
RtrA(config-if)#glbp 21 ip 10.21.8.10
RtrA(config-if)#glbp 21 priority 200
```

```
RtrB(config)#interface fa 0/1
RtrB(config-if)#ip address 10.21.8.2 255.255.255.0
RtrB(config-if)#glbp 21 ip 10.21.8.10
RtrB(config-if)#glbp 21 priority 110
```

GLBP – What's the same?

Many features are similar or identical to HSRP:

- AVG role is based on highest priority (then IP)
- 1 *Active* Virtual Gateway, 1 AVF+standby AVG, 2 others can be AVF (… and candidate AVG), any extras are passive candidate AVF
- same states for AVG as HSRP: disabled, init, listen, speak, standby, active
  AVF states: disabled, init, listen, active
- Hello's sent by AVG and standby AVG
- preempt disabled for AVGs (default), or after hold time
- preempt (default) for AVFs after 30 sec delay (default)
- timers: Hello = 3 sec, hold time = 10 sec
- multicasts to 224.0.0.102 (=HSRPv2!) UDP:3222
- authentication: plain text (yuck!) and MD5
- tracking of objects only

GLBP – Pros and Cons

- Good: Automatic selection and simultaneous use of multiple forwarding routers
  - *no* administrative requirement to configure multiple groups or manage multiple default gateways for hosts (only 1 DHCP server config!)
- Weak: more groups (0-1023) than HSRP but not enough to match all 4096 VLANs (only VTPv2 VLANs)
- Bad: Possible extra complexity & difficulty for trouble-shooting since never quite sure what path is taken
- Bad: May be difficult or impossible to create matching L2 and L3 topologies, resulting in congested or under-utilized links

FHRP – Chart

- You'll be happy to know that we're not the only ones working through these protocols!

EINCOP NetWax Lab

Difference between HSRP, VRRP and GLBP Protocols

http://www.slideshare.net/Netwaxlab/33-difference-between-hsrp-vrrp-and-glbp-protocols

Difference between HSRP, VRRP and GLBP Protocols

Protocol Features	HSRP	VRRP	GLBP
Scope	Cisco Proprietary	IEEE standard	Cisco Proprietary
Standard	RFC 2281	RFC 3768	none
OSI Layer	Layer 3	Layer 3	Layer 2
Load Balancing	No	No	Yes
Multicast Group IP Address	224.0.0.2 in v1 224.0.0.102 in v2	224.0.0.18	224.0.0.102
Transport Port Number	UDP 1985	UDP 112	UDP 3222
Timers	Hello 3 sec	Advertisement 1 sec	Hello 3 sec
	Hold 10 sec	Master down Time= 3xAdvertisement Time+Skew Time(Skew Time)(256-Priority)/256	Hold 10 sec
Election	Active Router: 1. Highest Priority 2. Highest IP (tie breaker)	Master Router: (*) 1. Highest Priority 2. Highest IP (tie breaker)	Active Virtual Gateway: 1. Highest Priority 2. Highest IP (tie breaker)
Router Role	1. One Active Router, One Standby Router 2. One or more Listening Routers	1. One Active Router 2. One or More Backup Routers	1. One AVG (Active Virtual Gateway) 2. Up to 4 AVF Routers on the group (Active Virtual Forwarder) passing traffic. 3. Up to 1024 virtual Routers (GLBP groups) per physical interface.
Preempt	If Active Router(Highest Priority) is down and up again, Preempt should be configured to become a Active Router again	By default Preempt is ON in VRRP, If Active Router is down and up again, It will automatically become a Master Router	If Active Router(Highest Priority) is down and up again, Preempt should be configured to become a Active Router again.
Group Virtual Mac Address	0000.0c07.acxx	0000.5e00.01xx	0007.b4xx.xxxx
IPv6 Support	Yes	No	Yes

Reminder

- *LOTS* of details do not appear in these slides
- You are responsible for *reading* the textbook to gain the knowledge (memorization) and understanding (apply the knowledge)