Chapter 6
First Hop Redundancy Protocols
NET3011 – 17W

(VRRP & GLBP in a separate slide deck)

FHRP – Basics

- First Hop Redundancy Protocols are the bottom piece in the chain of redundancy from top (network core) to bottom (individual hosts)
- a FHRP allows a set of two (or more) routers (or L3 switches) to cooperatively share the responsibility of acting as a "single" gateway for lower-level hosts by *virtualizing the gateway MAC*
- Hosts do not see any difference and are thus completely *unaware* that the gateway is not a single device
- Like other redundancy schemes, FHRPs provide mechanisms for load-balancing to reduce congestion and get higher utilization from the participating routers (or L3 switches)
- There are 3 protocols in wide use: 2 are Cisco proprietary and 1 is an open standard
- Like almost every other protocol, there are elections, priorities, keep-alives, timers / timeouts, etc. A shift in your perspective?

The Ugly Truth About Default Gateway

- No matter how a host's IP parameters are configured (static configuration or DHCP), once the Default Gateway is set, it typically isn't changed and so it can become a single point of failure.
- If the default gateway fails, the end device is limited to communicating only on the local segment, cut off from the rest of the network.
- Although some TCP/IP stacks allow multiple gateway addresses to be configured, how/when they're used tends to be very implementation dependent and failure detection and switchover is unacceptably slow.

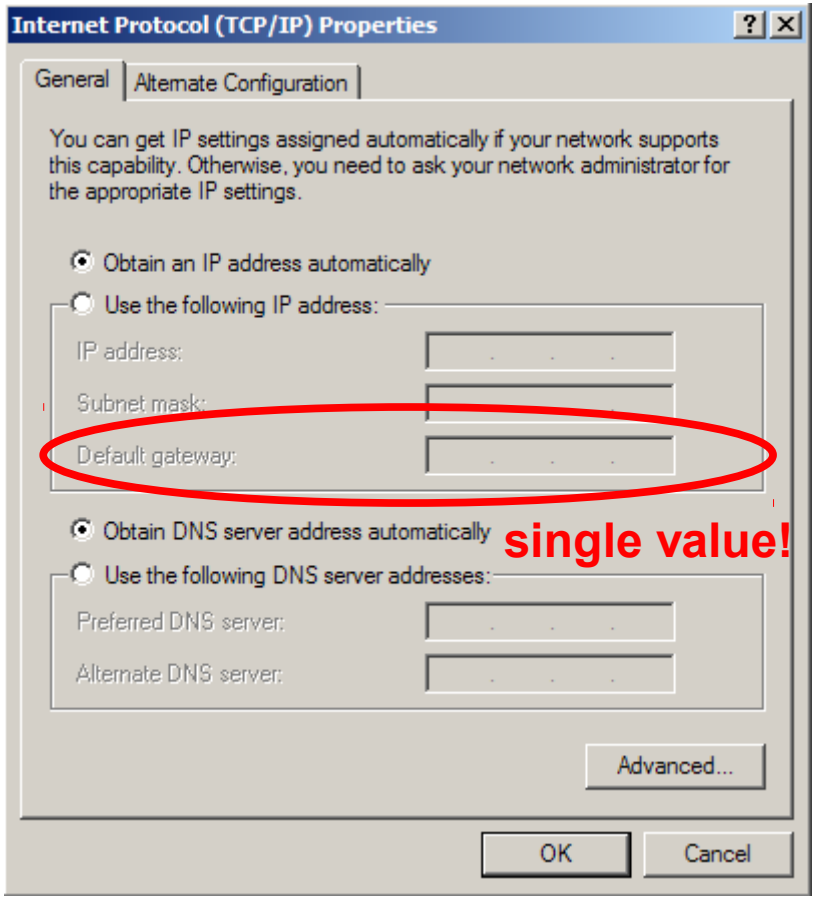

FHRP – Alternative "gateway" options

- Alternative methods for getting off-segment exist … but all of them suffer from one or more significant drawbacks
- None of the alternatives are currently in wide use, nor are they recommended
- We'll look at each one, for comparison purposes, so that the advantages of FHRPs are clear
- The alternatives we'll consider are:
  – Proxy ARP
  – host-based routing protocol
  – ICMP Router Discovery Protocol (IRDP)

FHRP Alternative: Proxy ARP

- Proxy ARP = replying to an ARP for which you do *not* have the IP

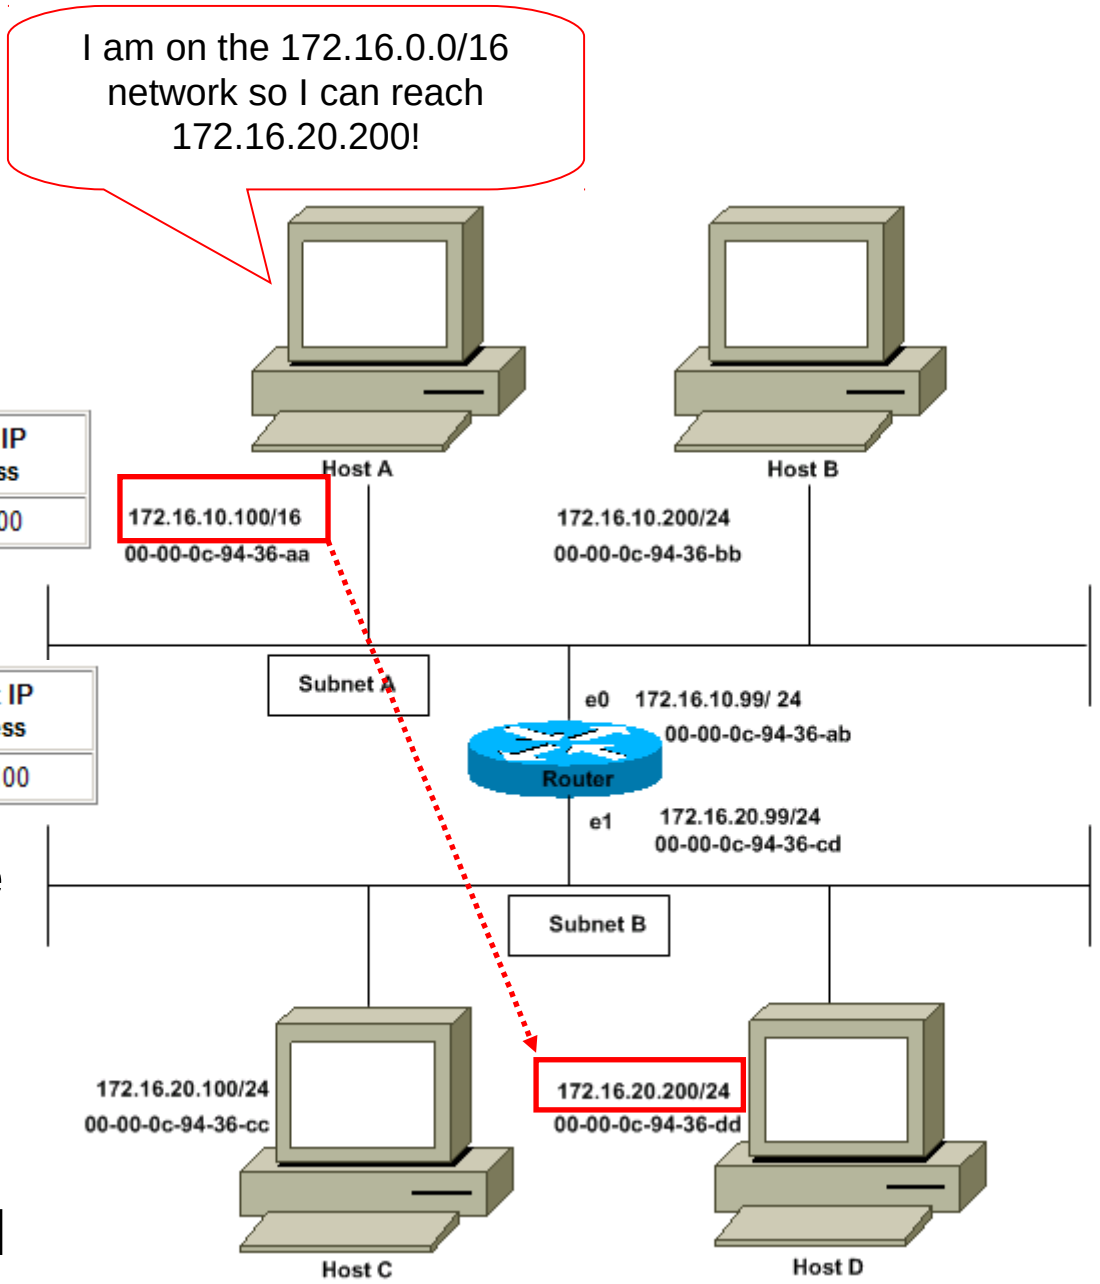

Host A sends ARP request

Sender's MAC Address	Sender's IP Address	Target MAC Address	Target IP Address
00-00-0c-94-36-aa	172.16.10.100	00-00-00-00-00-00	172.16.20.200

Router sends ARP reply

Sender's MAC Address	Sender's IP Address	Target MAC Address	Target IP Address
00-00-0c-94-36-ab	172.16.20.200	00-00-0c-94-36-aa	172.16.10.100

– host's ARP table gets very large
– very slow detection of failure
– very often proxy ARP masks configuration errors
– coordination with alternative GW is non-existant or ill-defined

FHRP Alternative – Host-based routing

- Every host needing to get off-segment runs a dynamic routing protocol and listens for updates to form it's own routing table
- Many disadvantages for this option:
  - extra complexity and configuration required on *every* host; "explosion" of administrative overhead
  - routing protocol is either simple but *slow* (RIP) or faster but more complicated (OSPF)
  - additional security required in network to prevent clients from hi-jacking routing tables/routes
  - not particularly suitable for small devices (e.g. wireless and IOT)

FHRP Alternative – IRDP

- Operation as per RFC1256:
  - IRDP routers sends multicast advertisements every 7 to 10 *minutes*
  - Default lifetime for advertisements: 30 *minutes*
  - Extremely *slow* by today's standards!
- Advantages
  - no routing protocols
  - no manual config req'd
- Disadvantages:
  - not common; outdated
  - difficult to troubleshoot in larger networks

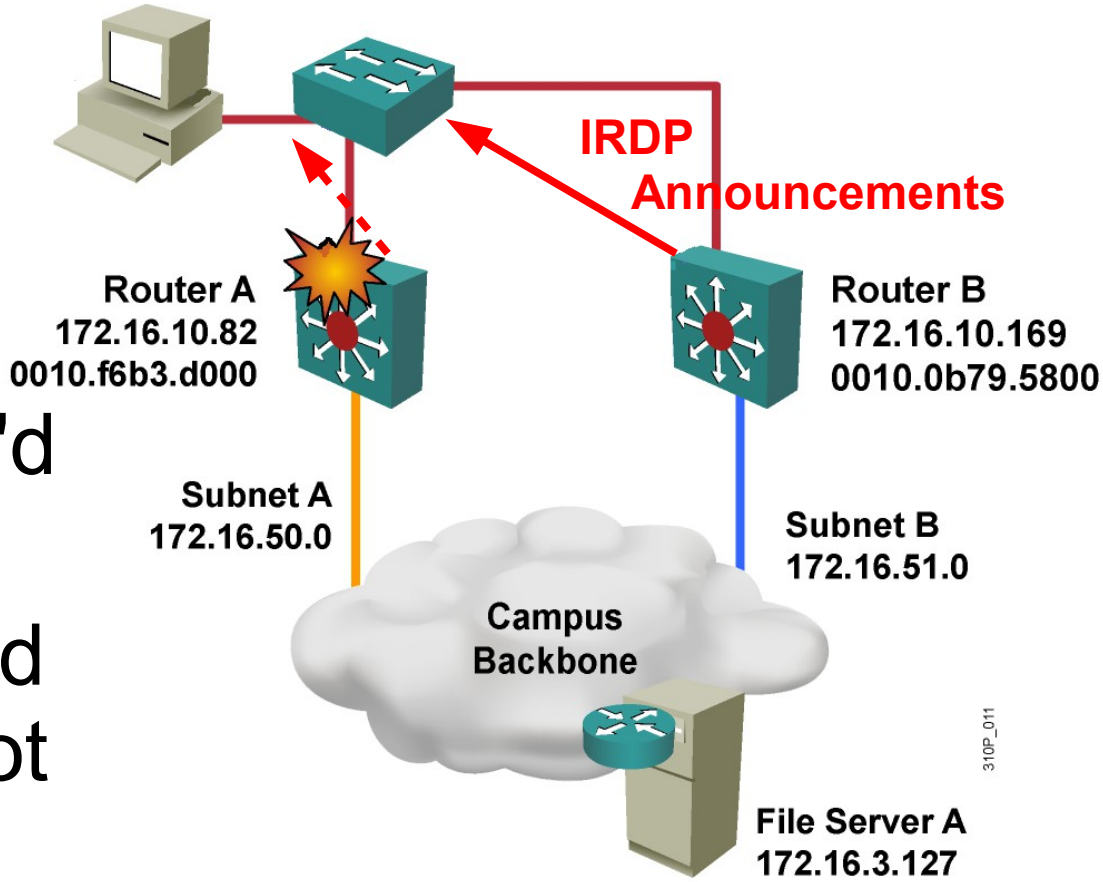

FHRP – Characteristics

- HSRP = Hot Standby Router Protocol
  - Cisco proprietary, but info later provided in RFC2281 (1998)
  - "gateway" MAC and IP are both virtual; 2 active routers
  - load balancing is completely manual
  - supported on most every Cisco router/L3 switch
- VRRP = Virtual Redundant Router Protocol
  - open standard RFC3768 (2004) and RFC5798 (2010)
  - "gateway" MAC is virtual; IP is real or virtual; 2 routers
  - load balancing is completely manual
  - support on Cisco devices varies by model / IOS version
- GLBP = Gateway Load Balancing Protocol
  - Cisco proprietary (~2005)
  - "gateway" MAC and IP are both virtual; max 4 routers
  - load balancing is built-in and automatic
  - support varies by model / IOS version

HSRP – Hot Standby Router Protocol

HSRP – Basics

- 2 routers pro-actively manage the gateway IP & MAC:
  - one is the *Active* router; the other is *Standby* router
  - any other routers are in a passive *Listen* state
- Role of each router determined by the protocol
  - Active router:
    - responds to ARP requests for virtual IP
    - processes frames addressed to virtual MAC
  - Standby router:
    - monitors the Active router
    - takes over immediately after the timeout period
- Both Active and Standby routers transmit Hello msgs at regular intervals to maintain their status (any other routers are silent and don't send Hello's)

HSRP – Steady-state Operation

- Active and Standby *must* have L3 reachability between themselves to maintain correct HSRP state
- HSRPv1 Hello sent to 224.0.0.2 UDP port 1985
- Hello's sent every 3 sec, by default

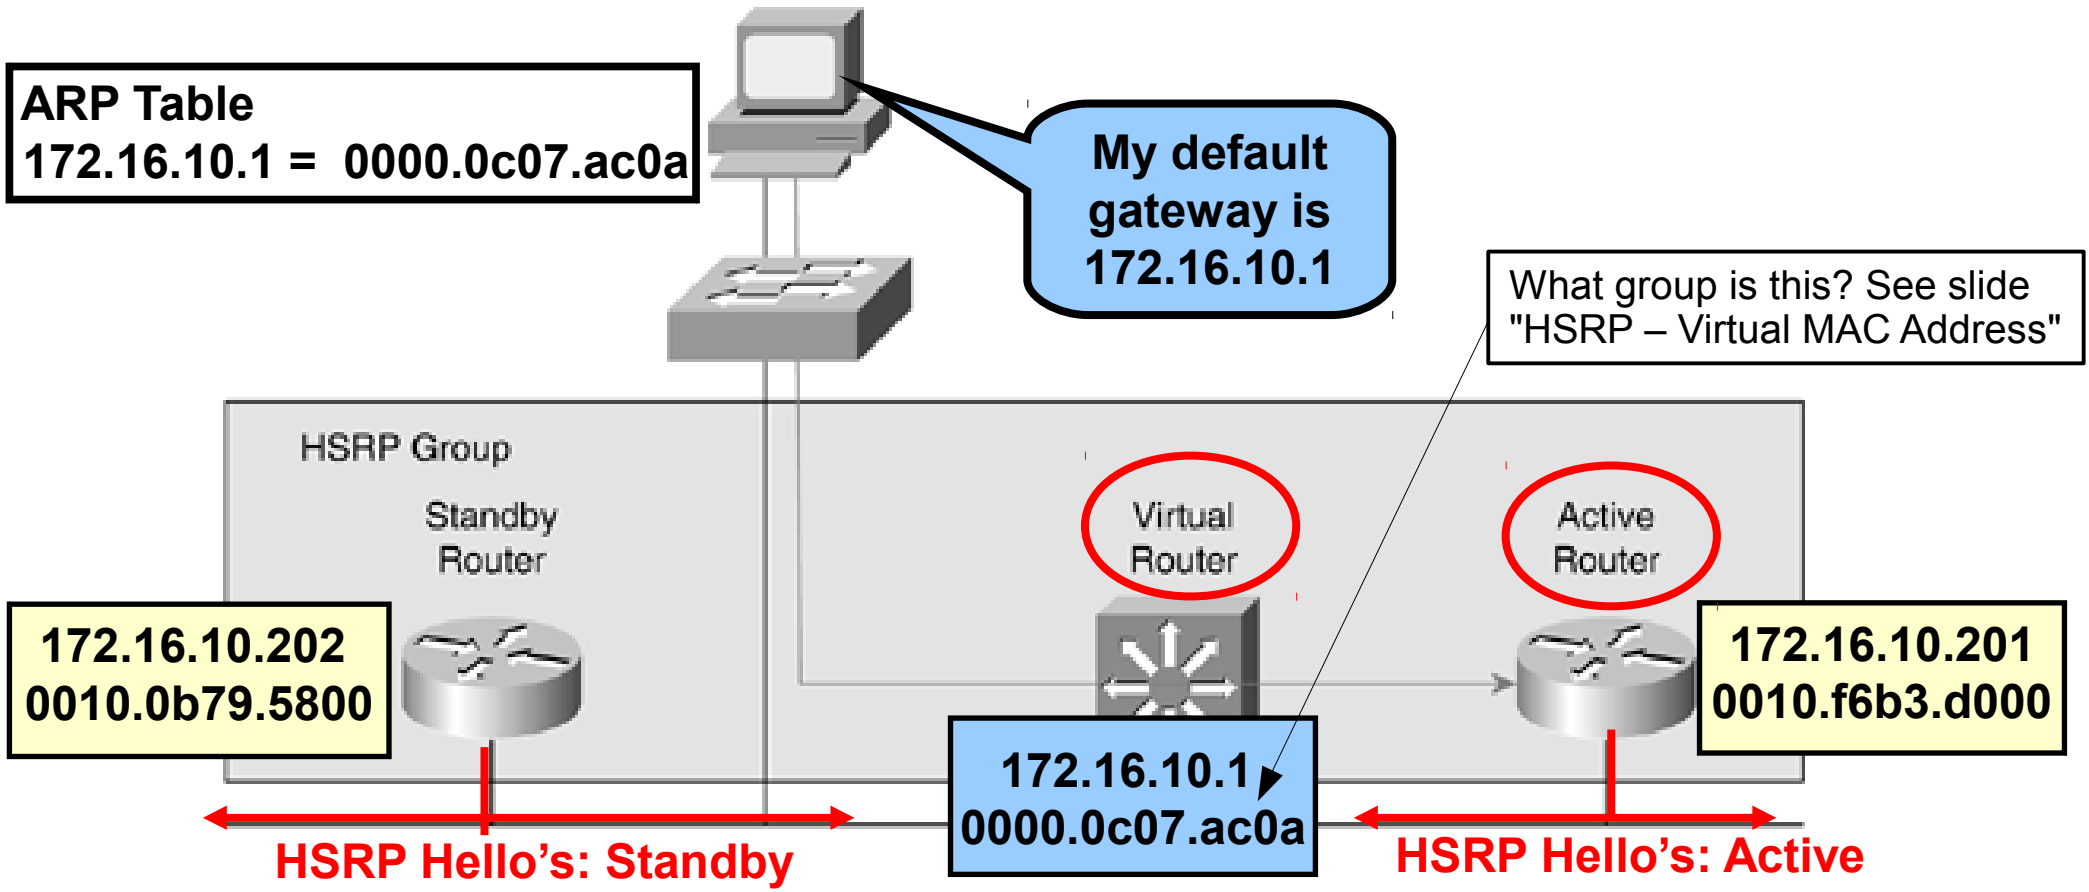

HSRP – Elections

- *When* an election occurs, highest priority wins
  – values 0-255, default = 100
  – tie-breaker is highest IP address on the link
- By default, can *not* force an election, even with a better priority, so Active router *can't* be overthrown
  – first router that is powered-on could (always) win!
  – generally *not* desirable behaviour (troubleshooting!)
  – handled by configuring **`preempt`** option
- Standby router is 2nd highest priority (tie: highest IP)
  – Standby role is always preemptive
- Other routers in group simply monitor Hello msgs, but don't play any other role until an election occurs

HSRP – Election times

- Elections occur when other routers stop receiving Hello's, i.e. holdtime expires (default: 10 sec)
- End-hosts may experience packet loss during holdtime (depends on nature of failure)
- If only the Active router fails, Standby takes over immediately upon expiry of holdtime
  – other routers in group (if any) elect new Standby
- New Active router remains in the role permanently, unless **`preempt`** is configured
- If both Active and Standby fail, other routers in group (if any) compete for both roles

HSRP – Election State Diagram

(No preemption configured)

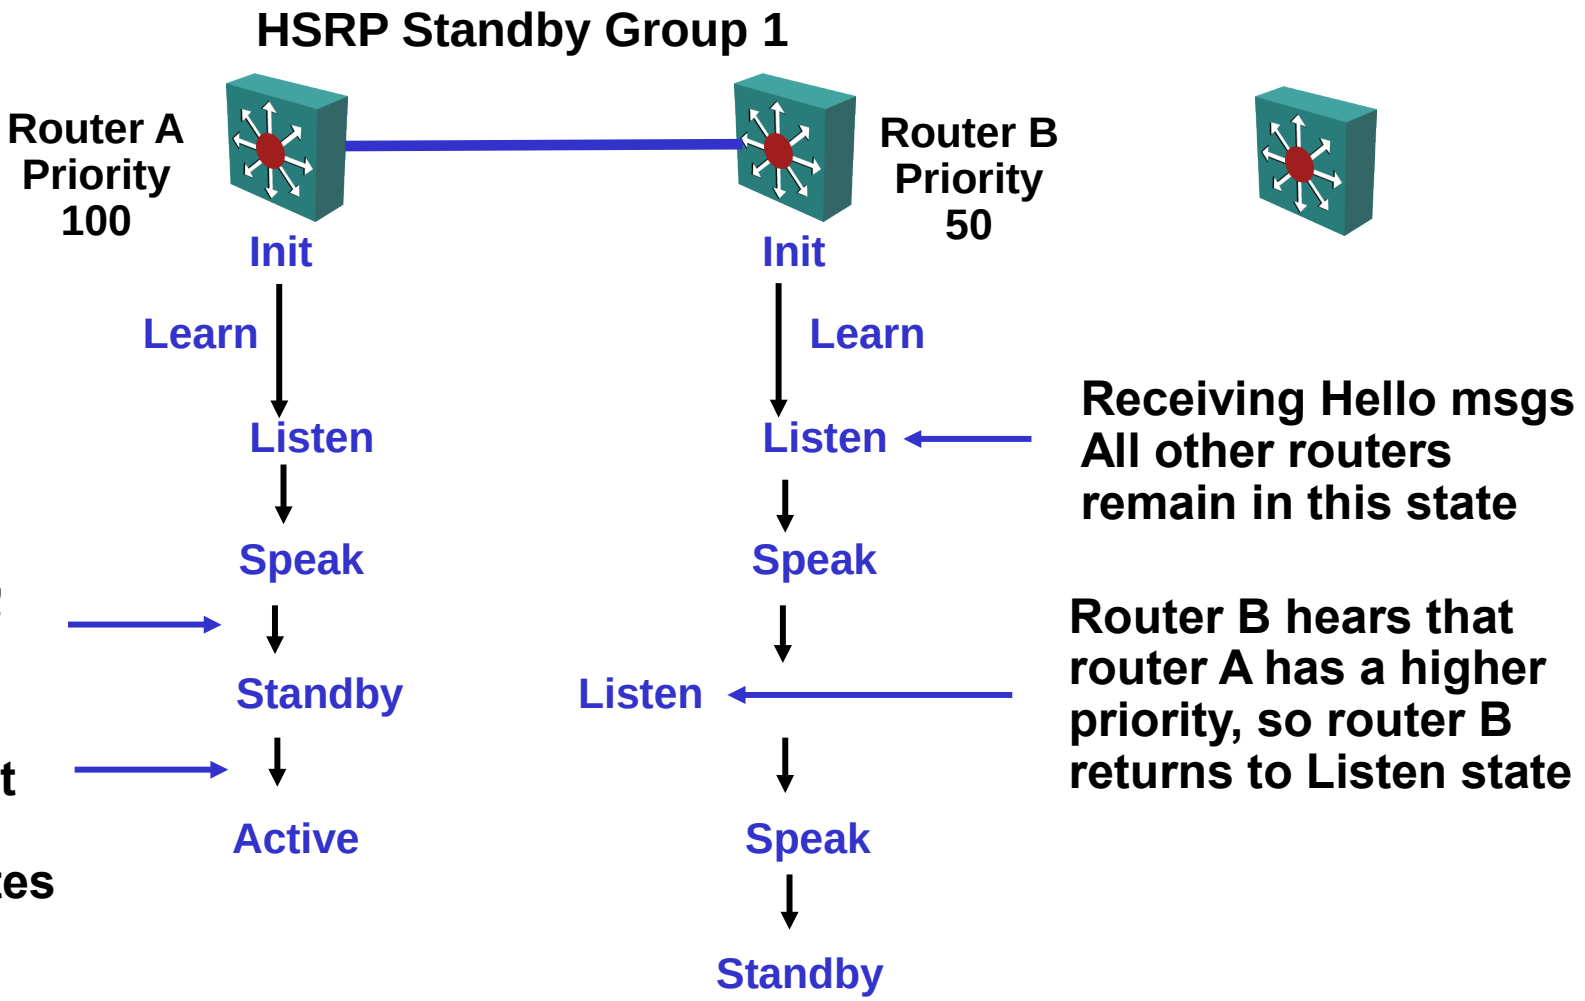

(see next slide for definition of states)

HSRP – Election State Definitions

At any time, a router configured with HSRP is in one of the following states:

- Initialize– Beginning state for all HSRP routers. Not yet ready or able to participate in HSRP, typically because the associated interface is not up. Groups with an interface that is down or groups without a specified interface IP address appear in the Init state.
- Learn – Router doesn't know the virtual IP address (wasn't configured!) and has not yet seen an authenticated hello message from the active router. In this state, the router still waits to hear from the active router.
- Listen – Router is receiving hello messages
- Speak – Router is sending and receiving hello messages
- Standby– Router is prepared to take over if active router fails
- Active – Router is processing & forwarding packets

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/12-4/fhp-12-4-book/fhp-hsrp.html#GUID-A9D4F853-5515-4386-A5EB-282576C3FA39

HSRP – Full State Diagram

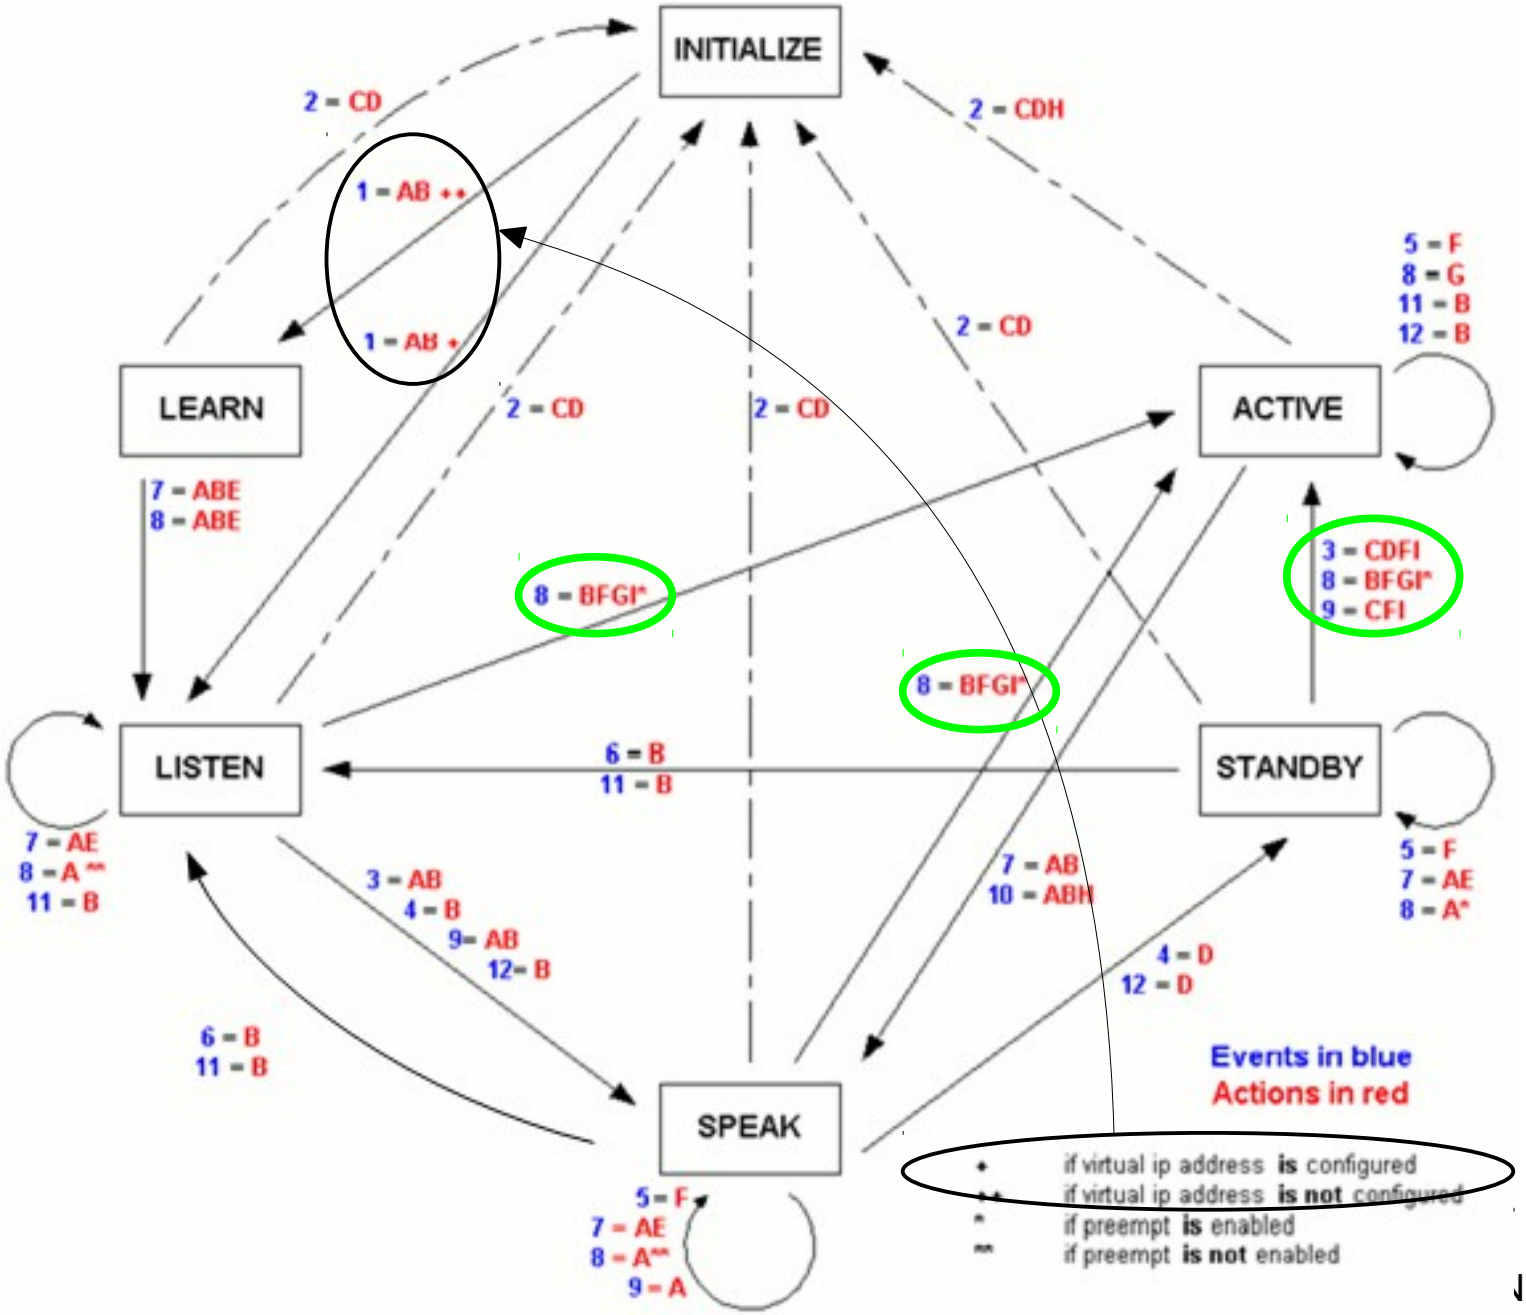

See if you can find parallels with the STP startup sequence

(see next page for source URL)

HSRP – Full State Diagram: Events

https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/10583-62.html

1	HSRP is configured on an enabled interface
2	HSRP is disabled on an interface or the interface is disabled
3	Active timer expiry: the active timer is set to the hold time when the last hello message is seen from the active router
4	Standby timer expiry: the standby timer is set to the hold time when the last hello message is seen from the standby router
5	Hello timer expiry: the periodic timer for the send of hello messages is expired.
6	Receipt of a higher priority Hello msg from router in speak state
7	Receipt of a higher priority Hello message from the active router
8	Receipt of a lower priority Hello message from the active router
9	Receipt of a Resign message from the active router
10	Receipt of a Coup message from a higher priority router
11	Receipt of a higher priority Hello message from the standby router
12	Receipt of a lower priority Hello message from the standby router

HSRP – Full State Diagram: Actions

https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/10583-62.html

A	Start active timer: If this action occurrs as the result of the receipt of an authenticated hello message from the active router, the active timer is set to the hold time field in the hello message. Otherwise, the active timer is set to the current hold time value that is in use by this router. The active timer then starts.
B	Start standby timer: If this action occurs as the result of the receipt of an authenticated hello message from the standby router, the standby timer is set to the hold time field in the hello message. Otherwise, the standby timer is set to the current hold time value that is in use by this router. The standby timer then starts.
C	Stop active timer: The active timer stops.
D	Stop standby timer: The standby timer stops.
E	Learn parameters: This action is taken when an authenticated message is received from the active router. If the virtual IP address for this group is not manually configured, the virtual IP address can be learned from the message. The router can learn hello time and hold time values from the message.
F	Send hello message: router sends a hello message with its current state, hello time, and hold time
G	Send coup message: The router sends a coup message in order to inform the active router that there is a higher-priority router available.
H	Send resign message: The router sends a resign message in order to allow another router to become the active router.
I	Send gratuitous ARP message: The router broadcasts an ARP response packet that advertises the group virtual IP and MAC addresses. The packet is sent with the virtual MAC address as the source MAC address in the link layer header, as well as within the ARP packet.

… Moral of the story:

Everything is discoverable and knowable,
if you're willing to dig into it.

- Can you think of a protocol that *doesn't* use timers and/or events to transition between states? (BGP FSM, STP, etc)

HSRP – Configuration: Virtual IP

- HSRP can be activated with just a single line:
  `L3(config-if)# standby {grp-num} ip {virtual-ip}`
  (you should probably add **`preempt`** too)
- What are potential **implications for DHCPv4**??

```
DLS1#
interface fa0/1                    !Configure on routed ports
  no switchport
  ip address     172.16.10.2 255.255.255.0
  standby 10 ip 172.16.10.1
  standby 10 preempt

interface vlan 20                    ! Configure on SVI's too
  ip address    172.16.20.2 255.255.255.0
  standby 2 ip 172.16.20.1
  standby 2 preempt

interface vlan 30                    ! Configure on SVI's too
  ip address    172.16.30.2 255.255.255.0
  standby 3 ip 172.16.30.1
  standby 3 preempt
```

HSRP – Configuration: Group ID

- Another characteristic of HSRP:
  every HSRP instance has a group number
- Why? One reason: groups allow us to obtain load-balancing
  (see "HSRP – Load Balancing" slides)
- Numbering is per-interface, not per-platform:
  – (correlate this with MPLS label spaces from NET3012!)
  – so numbers are locally significant to a broadcast domain
  – so each bcast domain does NOT need a unique group #
  … but all group members must be in the same bcast domain
- Details on group numbering:
  – an 8 bit number, so ranges from 0 to 255 (default = 0)
  – similar to VLANs, some switches may only support a limited quantity of groups simultaneously (think Lotto 649)

HSRP – Virtual MAC Address

- MAC addresses are ***not*** configured, but predetermined according to group number (and HSRP version!)
- HSRPv1 virtual MAC address: 00:00:0c:07:ac:XX
  - 00:00:0c = an OUI used by many Cisco protocols
  - 07:ac = designates HSRP ver 1
  - XX = 8 bit group number
- HSRPv2 virtual MAC address: 00:00:0c:9f:fX:XX
  - 9f:f = designates HSRP ver 2
  - XXX = 12 bit group number (able to match VLAN #)
- HSRPv2 IPv6 MAC address: 00:05:73:a0:0X:XX
  - 00:05:73 a Cisco OUI, a0:0 for HSRPv2 IPv6
- Now, go back to "HSRP – Steady-state Operation" !

HSRP – Configuration: Priority

- During an *election* , the router with the highest priority takes the Active role, 2nd highest takes Standby; tie-breaker is highest IP address on segment
- Priority is an 8 bit value (range: 0 to 255); default = 100
- Without preempt, first router to boot becomes Active

```
DLS1#
interface fa0/1
  no switchport
  ip address     172.16.10.2 255.255.255.0
  standby 10 ip 172.16.10.1
  standby 10 priority 110
  standby 10 preempt   !No point to priority w/o preempt

interface vlan 20
  ip address    172.16.20.2 255.255.255.0
  standby 2 ip 172.16.20.1
  standby 2 priority 220
  standby 2 preempt   !No point to priority w/o preempt
```

HSRP – Configuration: Timers

- The two keep-alive timers are hello and holdtime:
  `# standby {grp} timers [msec]{hello} [msec]{hold}`
  - both are 8 bits
  - measured in secs unless preceded by **msec**
  - defaults are: hello = 3sec; hold time = 10 sec
- for defaults: `# no standby {grp} timers`

```
DLS1#
interface fa0/1
  no switchport
  ip address    172.16.10.2 255.255.255.0
  standby 10 ip 172.16.10.1
  standby 10 priority 110
  standby 10 preempt
  standby 10 timers  2  7
```

HSRP – Configuration: Preempt Delay

- N.B: HSRP is L3 so may depend on routing and other protocols to converge; a delay may be beneficial and/or necessary:

```
#standby {id} preempt [delay [minimum {secs}] [reload {secs}]]
```

  - default values: force an election immediately
  - minimum: wait (0-3600 secs) before overthrowing an active router with a lower priority
  - reload: wait (0-3600 secs) after router has been reloaded or restarted before attempting to overthrow an active router with a lower priority
  - timer starts when router can take on active role i.e. HSRP configured, interface comes up

HSRP – Configuration: Authentication

- Current best practices say do *not* use authentication – if it fails, two routers are acting as same GW!
- If using it, use MD5 and *not* plain text (unsafe!)

```
DLS1#
interface fa0/1
  no switchport
  ip add 172.16.10.2 255.255.255.0
  standby 10 ip 172.16.10.1
  standby 10 preempt
  standby 10 authentication md5 key-string SlightlySafer

interface vlan 20
  ip add 172.16.20.2 255.255.255.0
  standby 2 ip 172.16.20.1
  standby 2 preempt
  standby 2 authentication NotSecretAtAll
```

Back to FHRP generics for a moment

- The interface and object tracking introduced in the next few slides is generic for all 3 FHRPs

FHRP – Tracking: Scenarios (1)

- Sometimes loss of *upstream*, or external connectivity should cause a change of FHRP GW

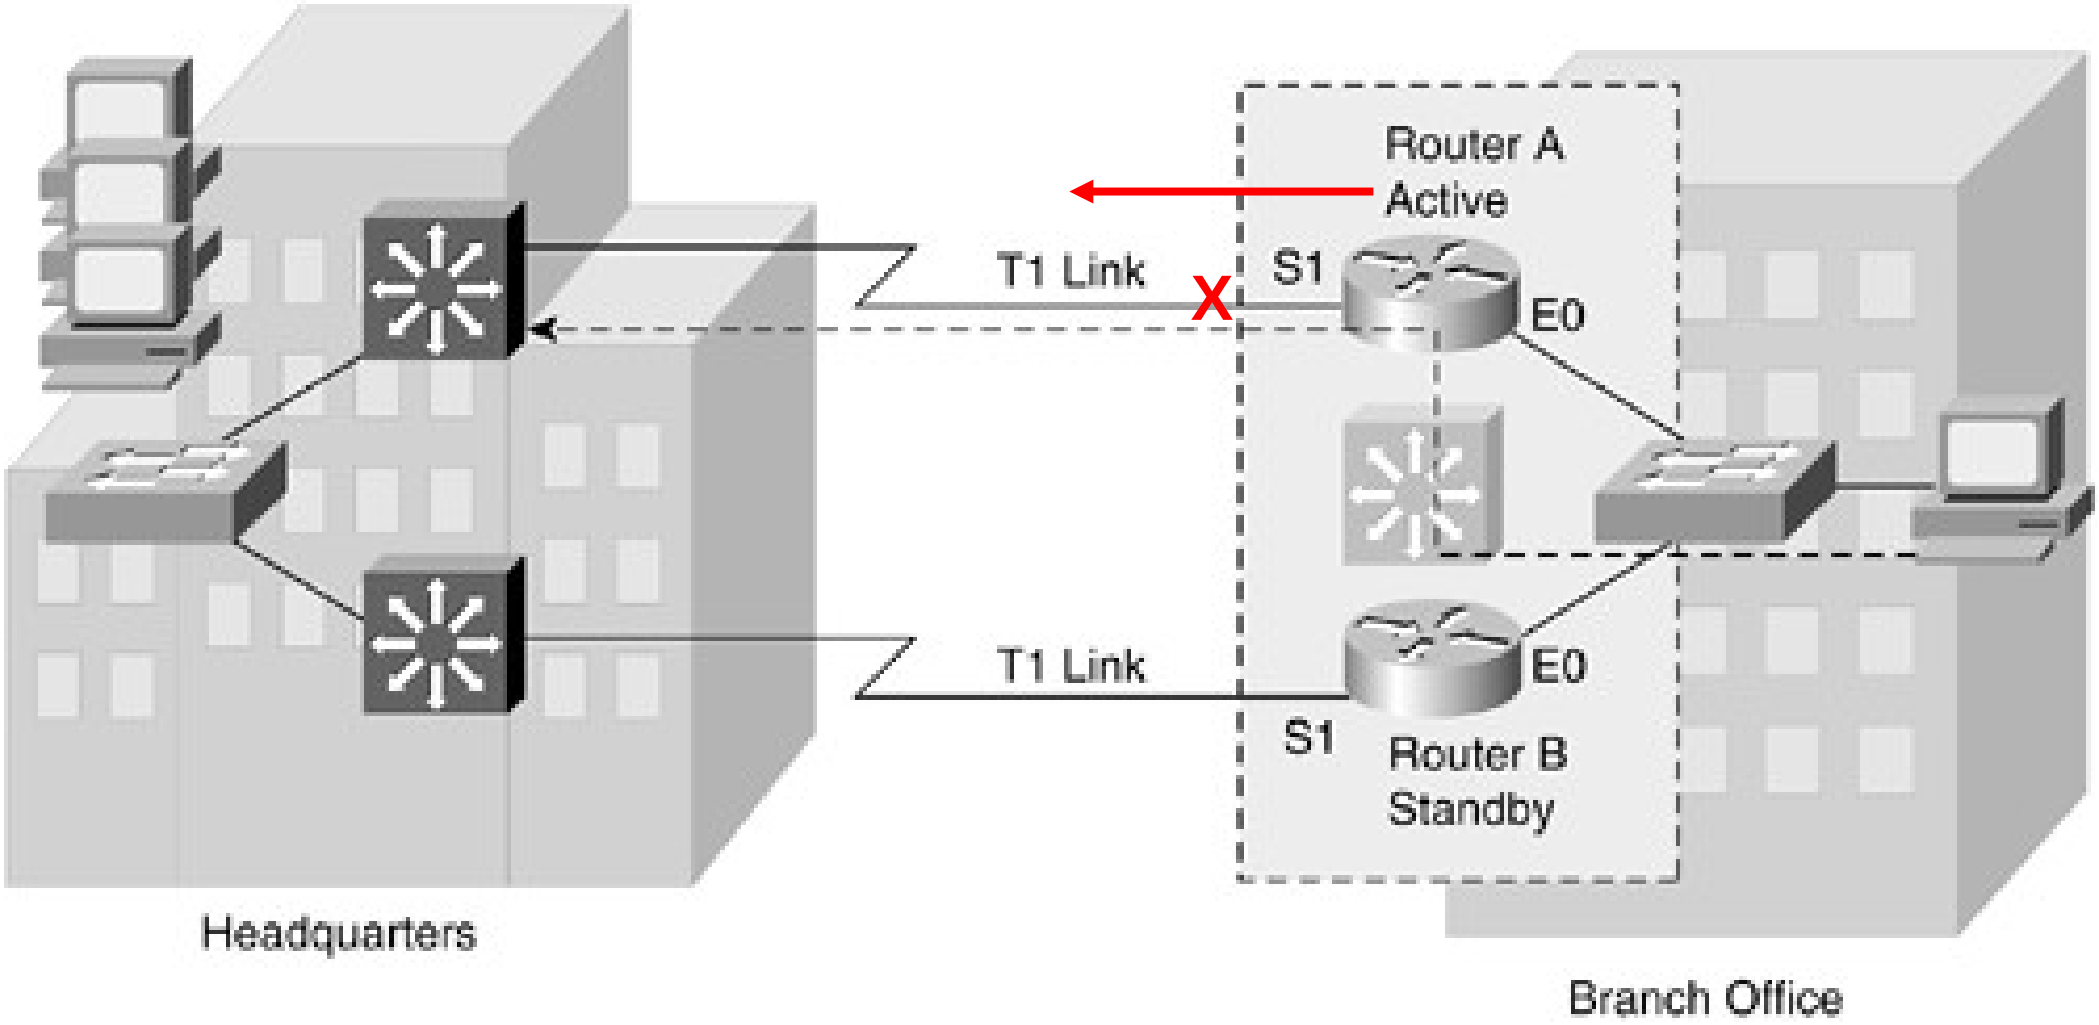

FHRP – Tracking: Scenarios (2)

- Fault could be intra-network

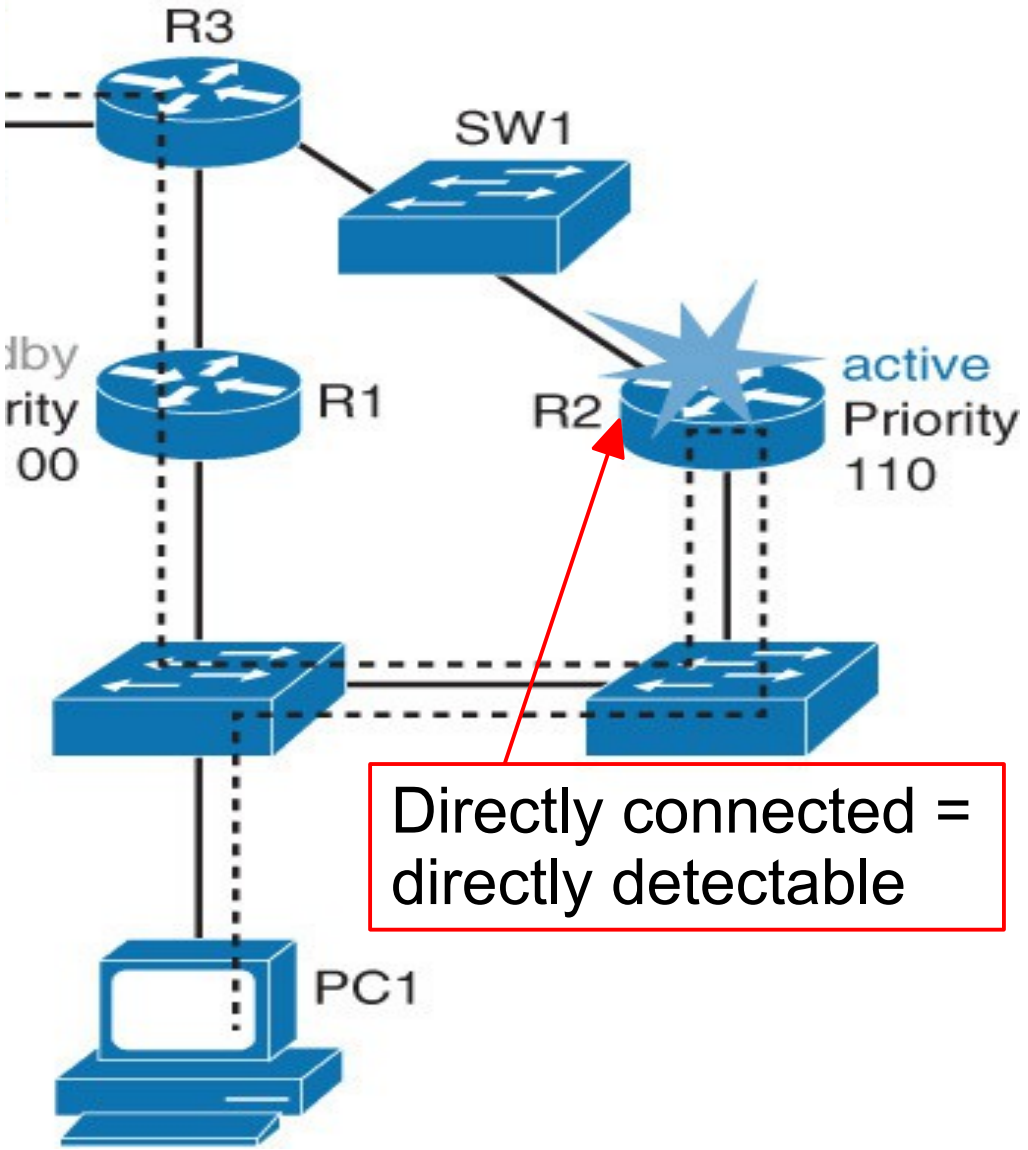

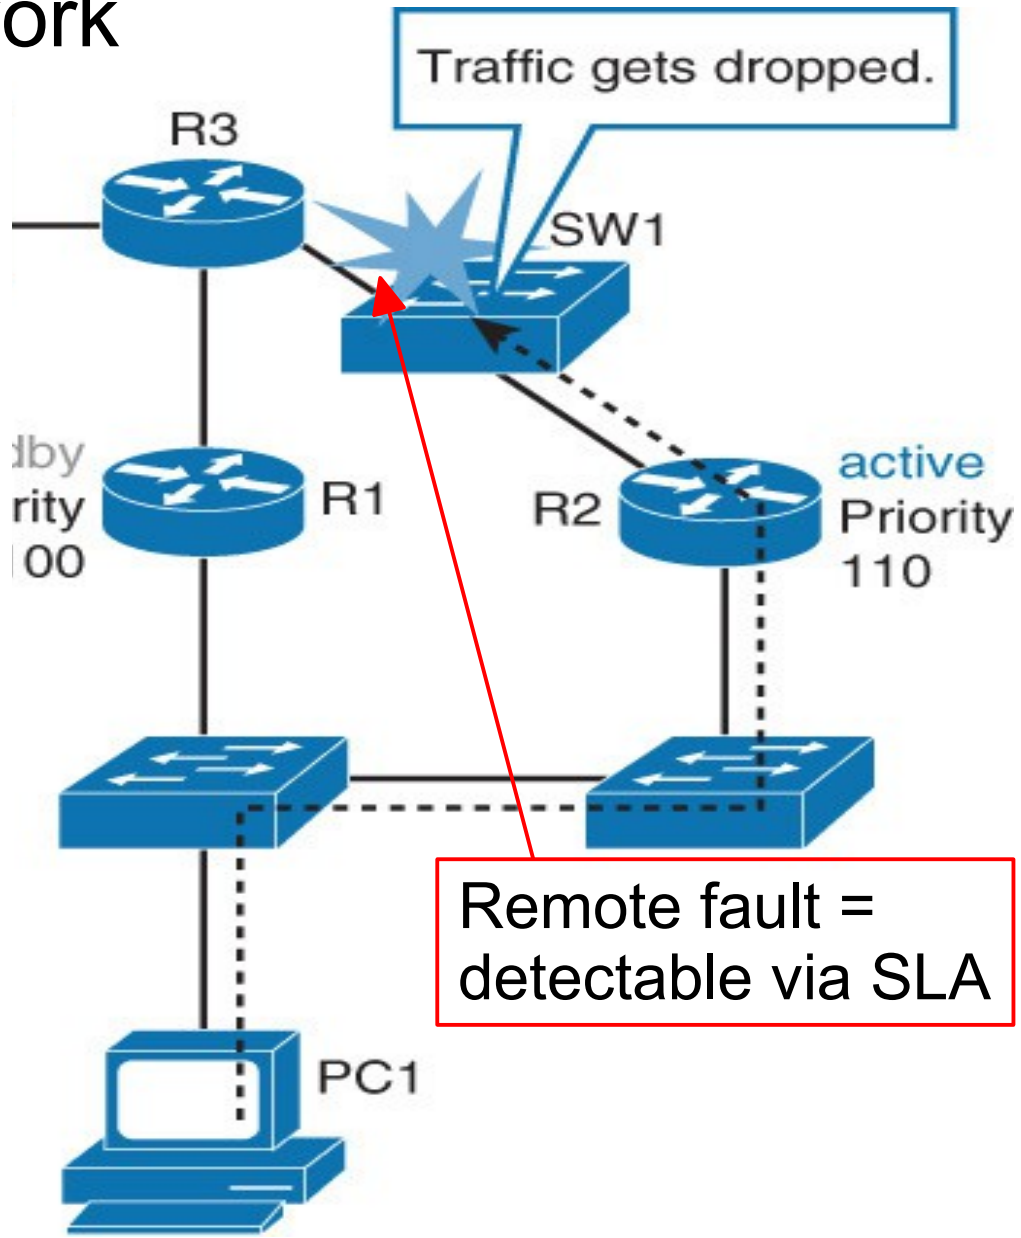

FHRP – Tracking: Basics

- There are situations when the active gateway needs to change, even though FHRP routers are still exchanging Hello's
- FHRPs may support either or both types of tracking:
  – direct detection = interface tracking
  – indirect detection = object tracking:
  for example IP SLA results (IP SLA covered in NET3008?)
- Tracking creates *automatic priority adjustments* for the group based on availability of interfaces or objects
  – as long as conditions defined by the object are fulfilled, the group priority remains the same
  – when verification defined by the object fails, group priority is decremented by configured value
- You can potentially track 100's of objects simultaneously

HSRP – Tracking Characteristics

- The number & type of objects varies with IOS
- Tracked objects are defined in global config mode

```
L3(config)# track 1 ?
interface   Select an interface to track
ip          IP protocol
list        Group objects in a list
rtr         Response Time Reporter (RTR) entry
```

```
L3(config)# ip sla 10
L3(config-sla)# icmp-echo 10.9.9.1
L3(config-sla)# exit
L3(config)# ip sla schedule 10 start now life forever
L3(config)# track 1 rtr 10 state
```

HSRP – Configuration: Tracking (1)

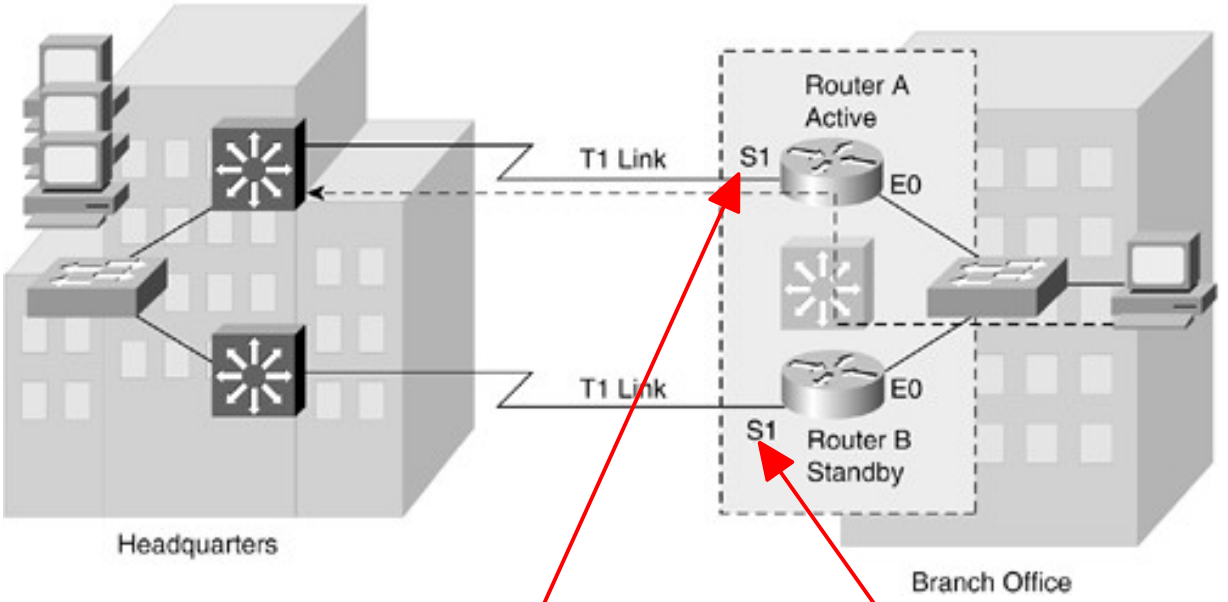

- In HSRP, default decrement value is 10

```
Router A
interface E0
 ip address 171.16.6.5 /24
 standby 1 priority 110
 standby 1 preempt
 standby 1 ip 171.16.6.100
 standby 1 track S1 20
```

```
Router B
interface E0
 ip address 171.16.6.6 /24
 standby 1 priority 100
 standby 1 preempt
 standby 1 ip 172.16.6.100
 standby 1 track S1 20
```

HSRP – Verification (1)

- Principal show command is: **`show standby [...]`**
- Debug commands are:
  `debug standby terse`
  `debug standby [errors] [events] [packets]`

```
L3# show standby brief
                    P indicates configured to preempt.
                    |
I/f    Grp Pri   P State    Active      Standby     Virtual IP
fa0/1 10  110   P Active   local       10.1.10.3   10.1.10.1
Vl20   2   220   P Standby 10.1.20.3   local       10.1.20.1

L3#show standby neighbor fa0/1
HSRP neighbors on fastEthernet0/1
10.1.10.3
Active groups: 10
No standby groups
```

HSRP – Verification (2)

```
Switch# show standby

fastEthernet0/1 - Group 10
State is Active
Virtual IP address is 10.1.10.1
Active virtual MAC address is 0000.0c07.ac0a
Local virtual MAC address is 0000.0c07.ac0a (v1 default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 1.248 secs
Preemption enabled
Active router is local
Standby router is 10.1.10.3, priority 90 (exp in 10.096s)
Priority 110 (configured 110)
Group name is "hsrp-Fa0/0-10" (default)

[...output omitted...]
```

HSRP – Verification: Tracking (1)

Before failure

```
RtrA# show standby
Ethernet0 - Group 1
Local state is Active, priority 110, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:01.028
Hot standby IP address is 171.16.6.100 configured
Active router is local
Standby router is 171.16.6.6 expires in 00:00:08
Tracking interface states for 1 interface, 1 up:
Up Serial1
```

Before failure

```
RtrB# show standby
Ethernet0 - Group 1
Local state is Standby, priority 100, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.772
Hot standby IP address is 171.16.6.100
Active router is 171.16.6.5 expires in 00:00:09
Standby router is local
Standby virtual mac address is 0000.0c07.ac01
Tracking interface states for 1 interface, 1 up:
Up Serial1
```

HSRP – Verification: Tracking (2)

After failure

```
RtrA# show standby
Ethernet0 - Group 1
Local state is Standby, priority 90, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:01.028
Hot standby IP address is 171.16.6.100 configured
Active router is 171.16.6.6 expires in 00:00:08
Standby router is local
Tracking interface states for 1 interface, 0 up:
Down Serial1
```

After failure

```
RtrB# show standby
Ethernet0 - Group 1
Local state is Active, priority 100, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.772
Hot standby IP address is 171.16.6.100
Active router is local
Standby router is 171.16.6.5 expires in 00:00:09
Standby virtual mac address is 0000.0c07.ac01
Tracking interface states for 1 interface, 1 up:
Up Serial1
```

HSRP – Load Balancing (1)

- Within a single subnet, it's possible to have multiple gateways (… but each host uses only 1)

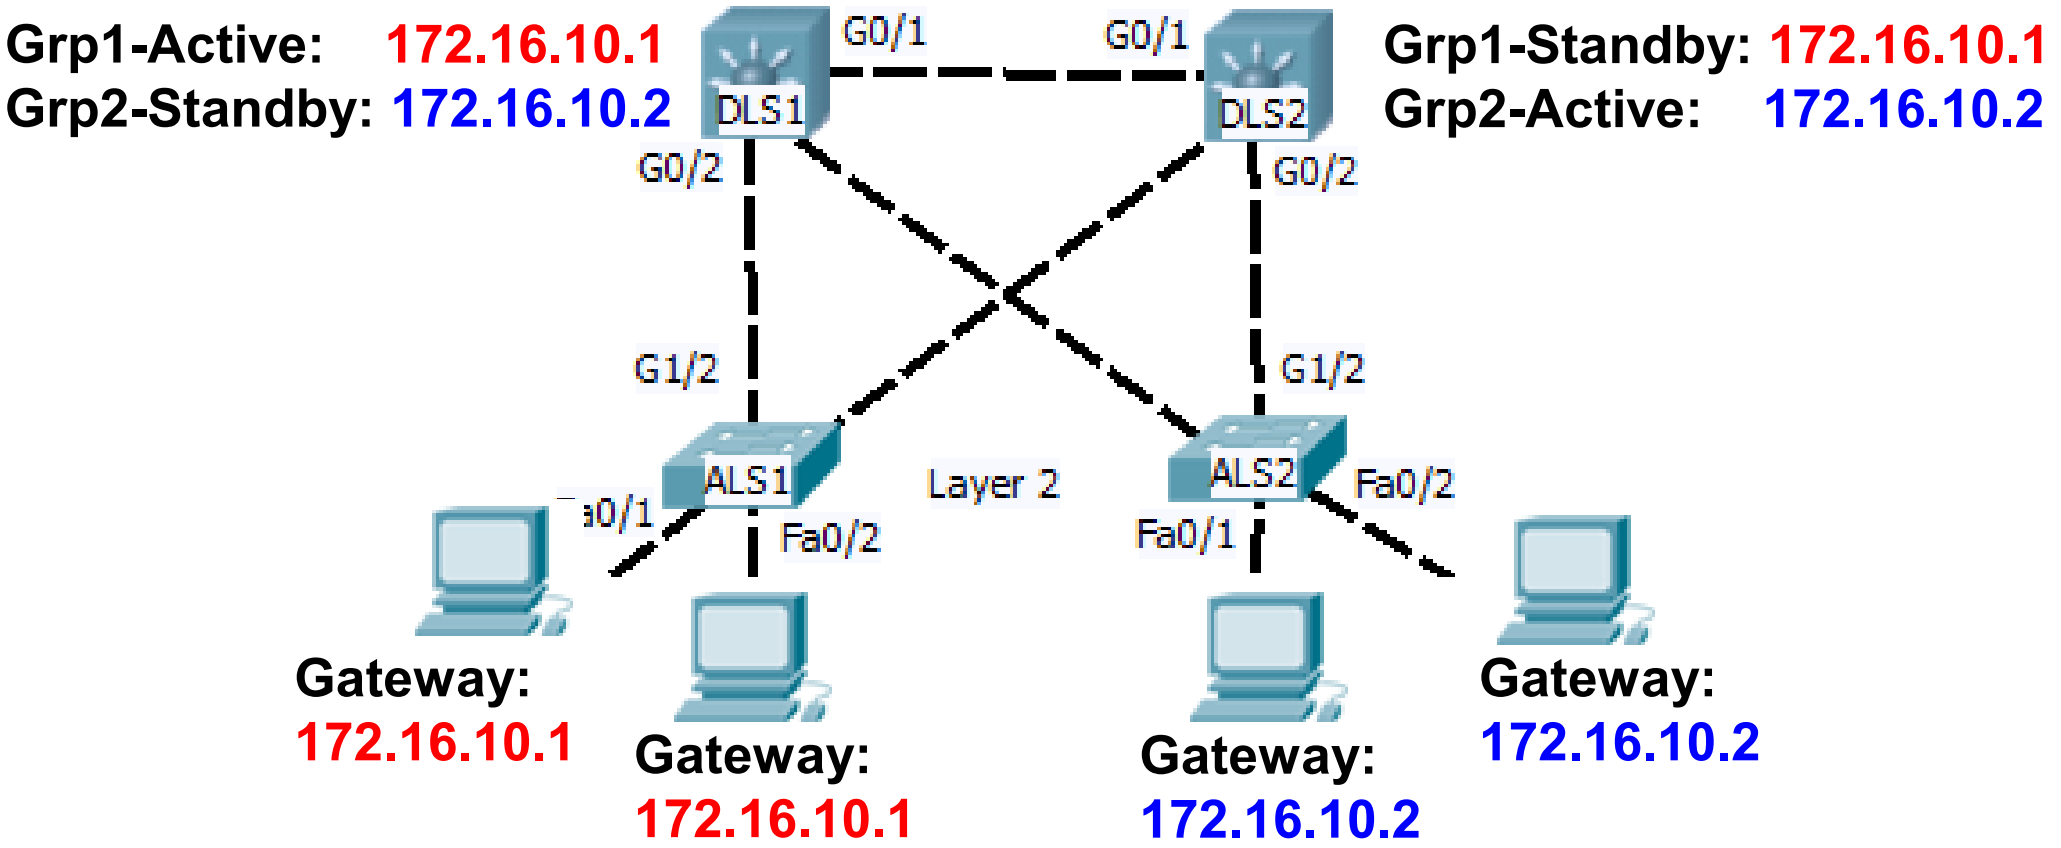

- a router can be in Active role for one group while simultaneously being in Standby role for another

HSRP – Load Balancing (2)

- Notice the priorities and virtual IPs, and that all balancing is done manually.

Do these actually achieve anything?

```
DLS1                                    DLS2
interface vlan 10                       interface vlan 10
  ip add 172.16.10.254 …                  ip add 172.16.10.253 …
  ----------------------------------------------------------------------
  standby 1 priority 200                  standby 1 priority 100
  standby 1 ip 172.16.10.1                standby 1 ip 172.16.10.1
  standby 1 preempt                       standby 1 preempt
  ----------------------------------------------------------------------
  standby 2 priority 100                  standby 2 priority 200
  standby 2 ip 172.16.10.2                standby 2 ip 172.16.10.2
  standby 2 preempt                       standby 2 preempt
```

- Although there can be up to 255 standby groups on any interface (SVI), every additional group represents additional loading on the control plane (CPU) which may affect performance

HSRP – Aligning L2 & L3 Topologies

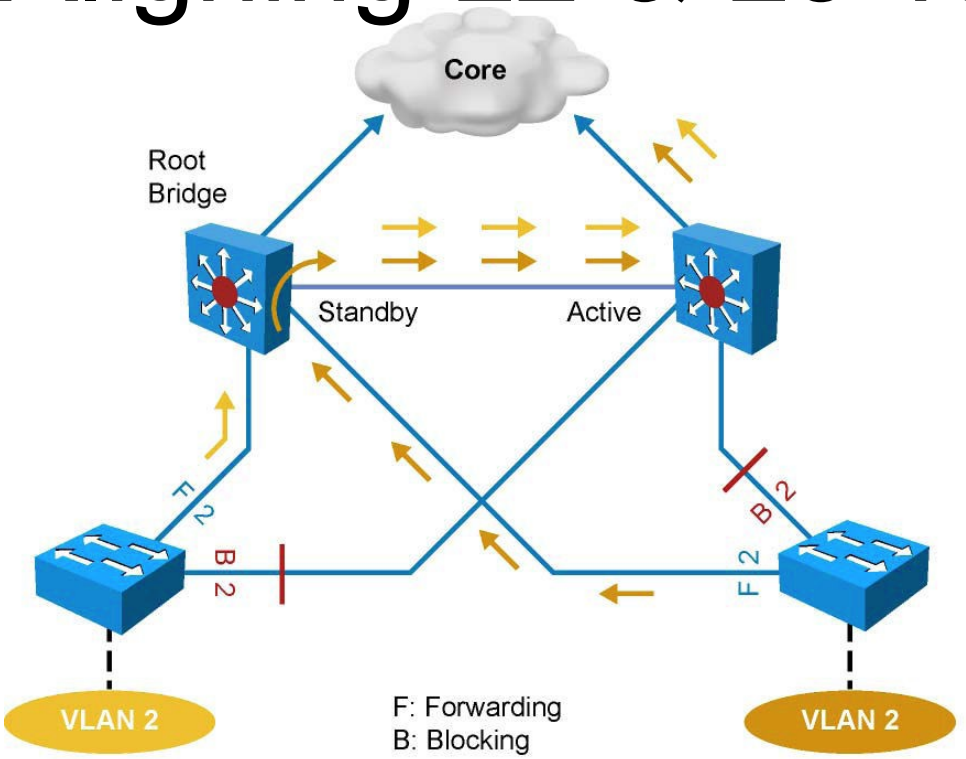

- In a redundant spanning-tree topology, some links are blocked. The spanning-tree topology has no awareness about the HSRP configuration. There is no automatic relationship between the HSRP active router election process and the Spanning Tree Root Bridge election.
- When configuring both spanning tree and HSRP (or any other first hop redundancy protocol), you must make sure that the active router is the same as the root bridge for the corresponding VLAN. When the root bridge is different from the HSRP active router, a suboptimal path can result, as illustrated.

HSRP – Versions

- HSRP ver 2 was introduced to allow 12 bit group numbers which allows group numbers to exactly match SVI VLAN numbers (group numbers may be re-used but may cause [human] confusion)
- HSRPv1:
  - is the default version
  - allows 8-bit group numbers up to 255
  - multicasts HSRP Hello messages to UDP 224.0.0.2 ("all routers")
  - responds to ARPs with virtual MAC addresses: 0000.0C07.ACXX
- HSRPv2:
  - is configurable per interface: **`standby {grp} version 2`**
  - allows 12-bit group numbers up to 4095
  - is configured the same in all other respects as ver 1
  - multicasts to 224.0.0.102, still to UDP port 1985
  - responds to ARPs with virtual MAC addresses 0000.0C9F.FXXX
  - is incompatible with HSRPv1 (different packet format), so *must* ensure the same version is configured on all routers in an HSRP group (otherwise hello messages are not understood)

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)