First Hop Redundancy Protocols

SWITCH Module 6

Agenda

- Motivation
- Hot Standby Router Protocol
- Virtual Router Redundancy Protocol
- Gateway Load Balancing Protocol

L3 Redundancy

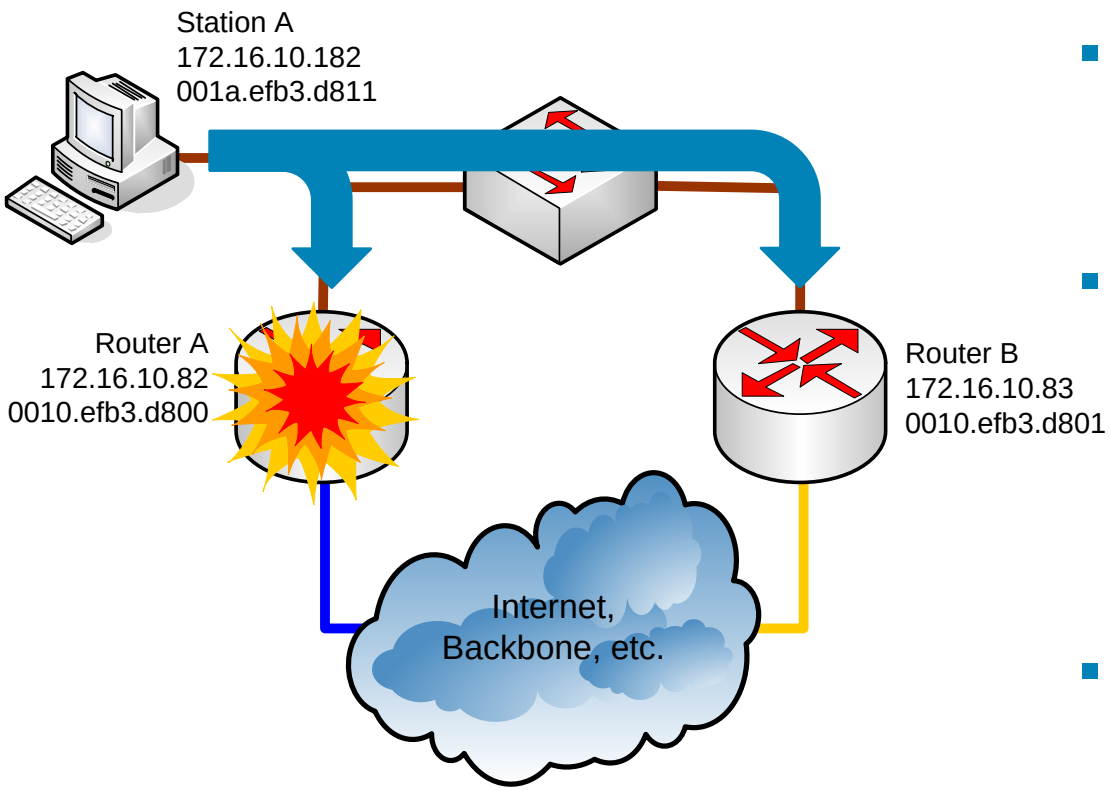

- `IF` the Router *A* fails `THEN` dynamic routing protocol starts to use the Router *B*
- *However, end station does not use routing protocol!*
  - Usually only one IP address of the default gateway is assigned
- "Historical" attempts how to solve redundancy problem
  - **Proxy ARP**
  - **ICMP Router Discovery Protocol**
  - Routing protocol support on the end station
- These attempts
  - Do not scale well
  - A software is usually needed at the end stations

Proxy ARP ①

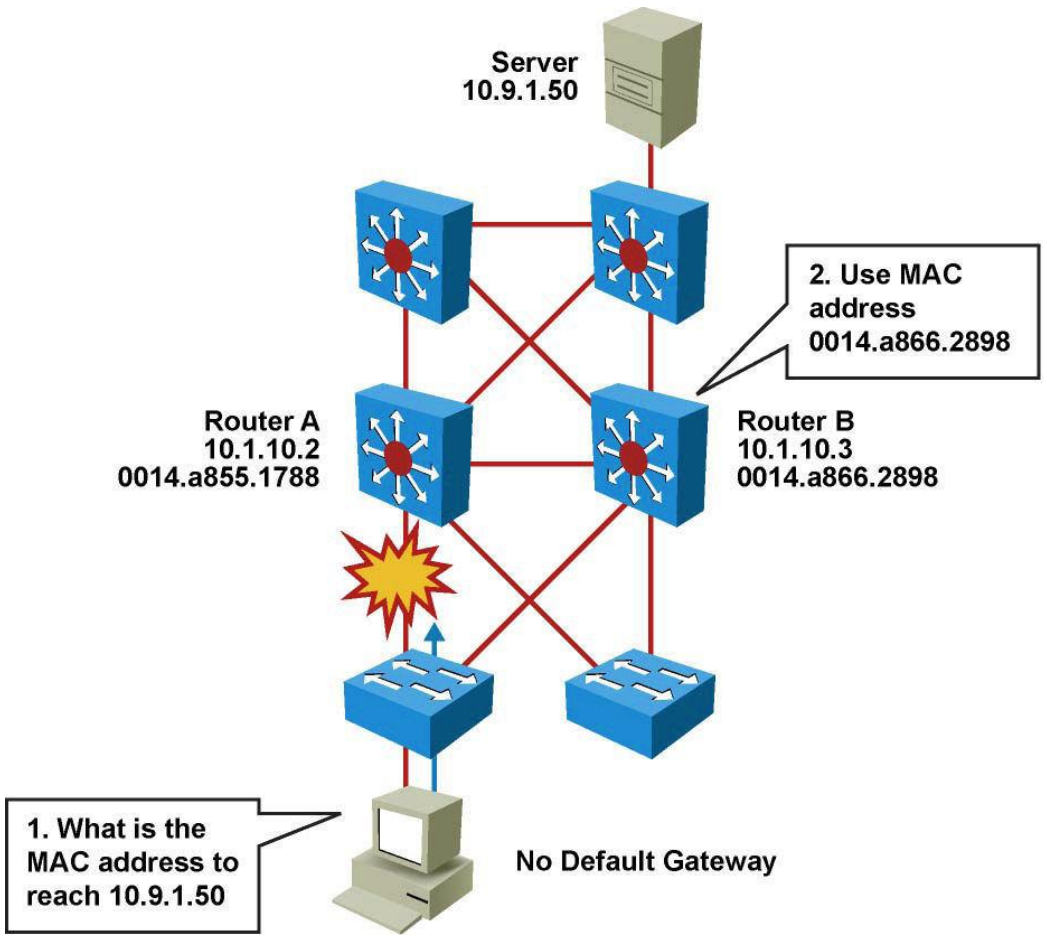

- Enabled by default
- Used before default gateways were supported on IP clients
- End station acts as if destination were on same network segment
- Relatively slow due to reliance on aging out of ARP cache

Proxy ARP ②

- Not used today as a redundancy solution
- *Beware of proxy ARP and default route!*

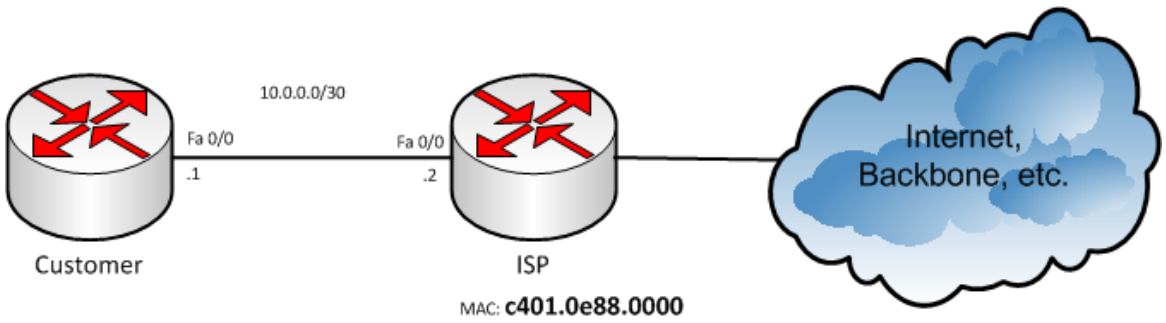


```
Customer(config)# ip route 0.0.0.0 0.0.0.0 FastEthernet 0/0
Customer# ping 1.1.1.1
Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is
2 seconds:
.!!!! Success rate is 80 percent (4/5), round-trip min/avg/max = 8/31/52 ms
Customer# sh arp
Protocol  Address    Age (min)  Hardware Addr       Type      Interface
Internet  1.1.1.1    0          c401.0e88.0000      ARPA      FastEthernet0/0
Internet  10.0.0.2   -          c400.0e88.0000      ARPA      FastEthernet0/0
```

L3 Redundancy Using Virtual Router

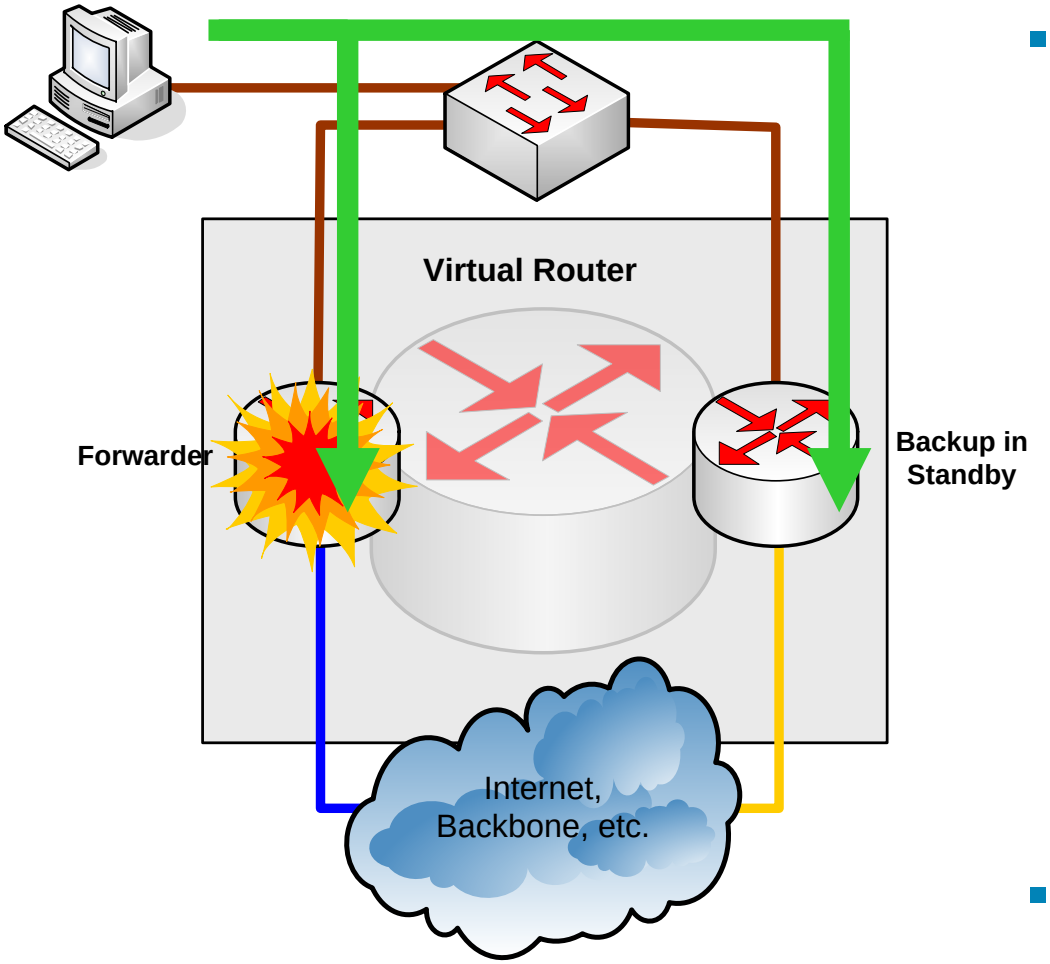

- Routers can create a virtual router
  - Virtual router has own virtual IP and MAC addresses
  - Virtual IP is used as a default gateway
  - A real router serves as forwarder for the virtual IP address
  - `IF` the forwarder fails `THEN` another router starts to be the forwarder
  - **First Hop Redundancy Protocols (FHRP)**
  - Only one FHRP protocol could be run on Cisco device
- Everything is transparent for the end station – virtual IP and MAC stay same

Hot Standby Routing Protocol

Hot Standby Router Protocol (HSRP)

- HSRP - Cisco proprietary protocol
  - Cisco Document ID: 10583, „Understanding and Troubleshooting HSRP Problems in Catalyst Switch Networks“
- Two HSRP versions exist:
  - **HSRPv1** (RFC 2281)
    - Packets are sent to 224.0.0.2, UDP port 1985
    - Group numbers are restricted to the range from 0 to 255
  - **HSRPv2**
    - Packets are sent to 224.0.0.102, UDP port 1985
      - 224.0.0.2 can conflict with Cisco Group Management Protocol
    - Group numbers range from 0 to 4095
    - Supports for millisecond timer values
    - Supports IPv6 gateway
  - Default is version 1

Device Roles

- **Active router**
  - One router is elected within an HSRP group
  - Physically forwards packets sent to the MAC address of the virtual router
- **Standby router**
  - Backup active router (similar as DR and BDR in OSPF)
  - When the active router fails, standby router then assumes the role of the active router
  - One standby router is elected within and HSRP group
- **Other routers**
  - Other routers within an HSRP group
  - Remain in the initial state
  - `IF` both the active and standby routers fail `THEN` other routers in the group contend for the active and standby roles
- **Virtual router**
  - Virtual router with virtual IP and MAC address pair

Active and Standby Router

- Active router forwards packets sent to the virtual MAC / IP address
  - Elected according a priority (range from 0 to 255, default 100)
  - `IF` priorities are same `THEN` router with higher IP address wins election
  - Virtual IP address must be set in configuration
  - Virtual MAC address depends on the HSRP group identifier
    - **HSRPv1**: `0000.0C07.ACXX`
    - **HSRPv2**: `0000.0C9F.FXXX`
- The Standby router is elected similar as the Active router (the second best in priority / higher IP)
- Active and Standby routers exchange Hello packet periodically
  - Hello packet informs other routers in a HSRP group that the Active/Standby router is still operational
  - Other routers in a HSRP group DO NOT send Hello packets

HSRP States

- **Init / Disabled**
  - This is the starting state and indicates that HSRP is not running
- **Learn**
  - The router has not determined the virtual IP address
  - The router has not yet seen an authenticated Hello message from the active router
- **Listen (10 sec)**
  - Router listens for hello messages from Active/Standby router
  - Knows the virtual IP address, but the router is neither the active nor the standby router
- **Speak (10 sec)**
  - Router sends and receives periodic hello messages
  - Actively participates in the election of the active or standby router
- **Standby**
  - The router is a candidate to become the next active router
  - Sends periodic Hello messages.
- **Active**
  - Sends periodic Hello messages
  - Forwards packets that are sent to the group virtual MAC/IP address

Example: State Transition

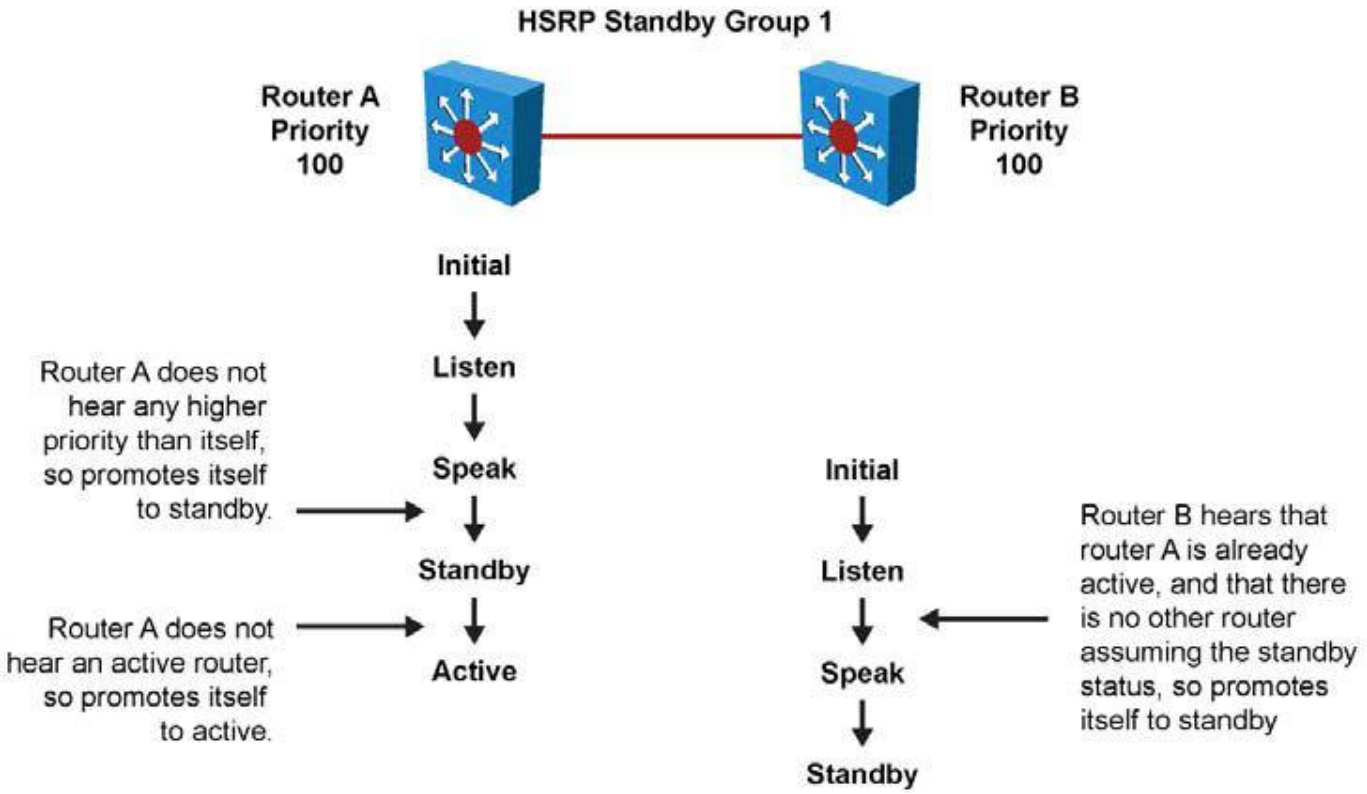

- Router *A* starts
  - As it is the first router for standby Group 1 in the subnet, it transits through the listen and speak states and then becomes the active router
- Router *B* starts after Router *A*
  - While *B* is in listen state, *A* is already assuming the standby and then the active role
  - As there is already an existing active router, *B* assumes the standby role

Timers

Timer	Description
Hellotime	It contains the approximate period between the Hello messages that the router sends. The time is given in seconds. Recommended value is 3 sec.
Holdtime	Time, in seconds, before the active or standby router is declared to be down. This is an integer from 1 through 255. The default is 10 seconds.
Active timer	Every router in a HSRP group monitor the active router. The active timer is started anytime an authenticated Hello message is seen from the active router. It is set to expire in the Holdtime field, seen in the Hello message.
Standby timer	The Standby timer is used to monitor the standby router. The Standby timer is started anytime an authenticated Hello message is seen from the standby router. It is set to expire in the Holdtime field seen in the Hello message.

HSRPv1 Message Format

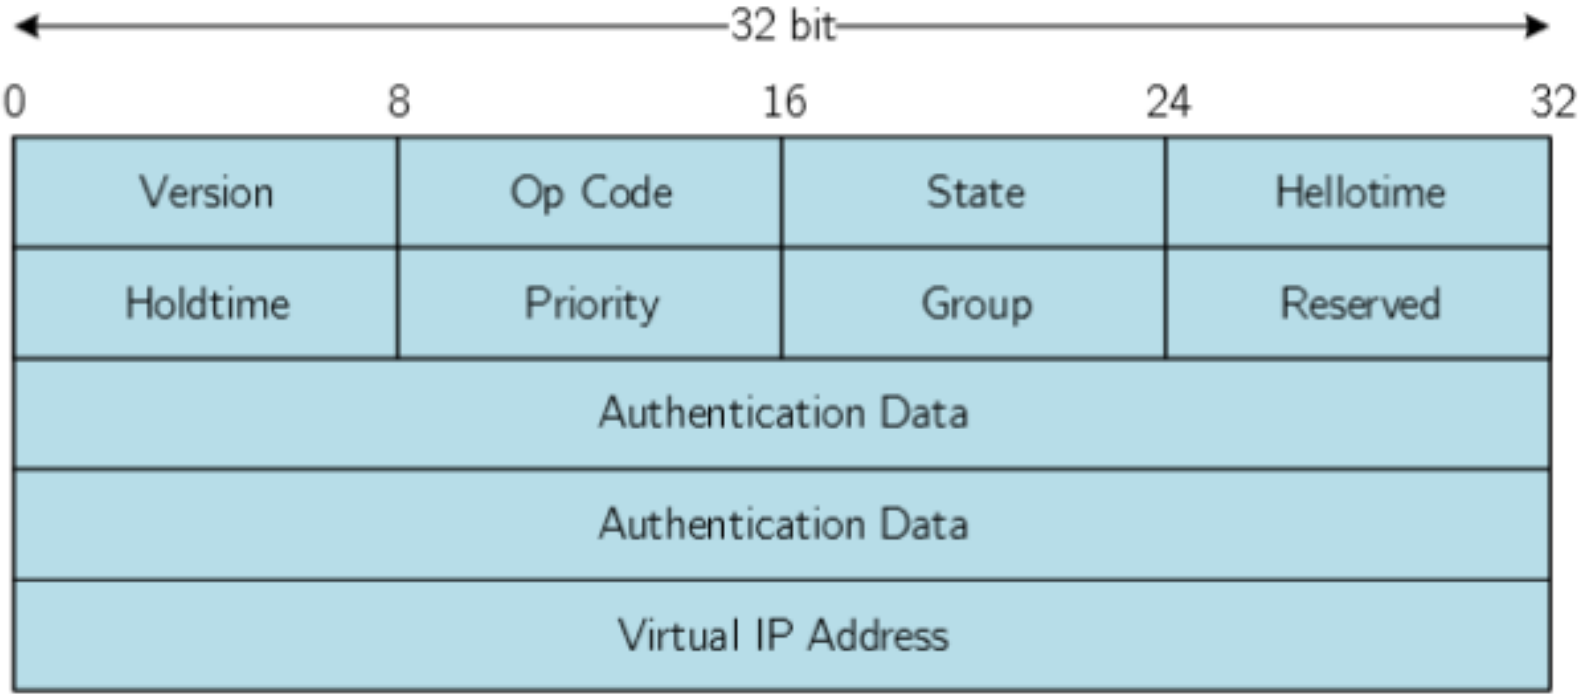

- **Op Code**
  - Hello
  - Coup
  - Resign
- **State**
  - Initial
  - Learn
  - Listen
  - Speak
  - Standby
  - Active

HSRPv2 Message Format

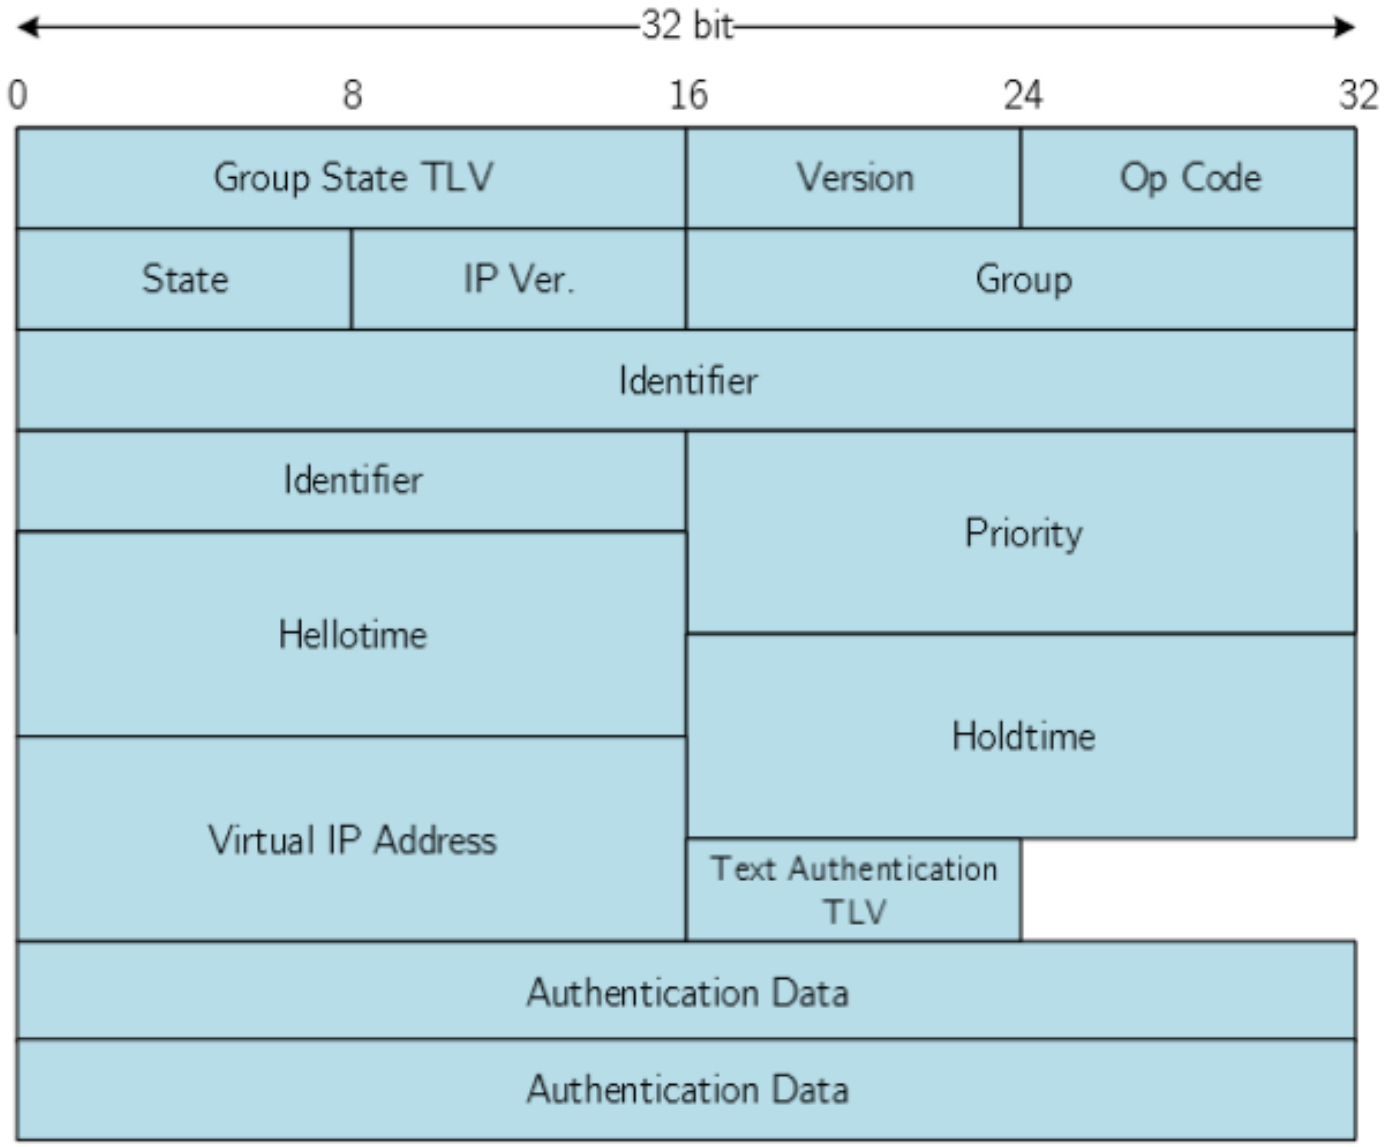

Message Types

- Message **Hello**
  - Sent by Active and Standby routers
- Message **Coup**
  - Sent when a router wishes to become the active router
  - Used together with preemption capability
- Message **Resign**
  - Resign messages are sent when a router no longer wishes to be the active router
- Preemption capability
  - `IF` a router has higher priority than the active router and preemption is configured `THEN` it may take over as the active router using a Coup message
  - Disabled by default – whenever the Active router fails standby router take over as the active router only

Virtual IP Address

- HSRP group creates a virtual router with a virtual IP and MAC address
  - Every member of a HSRP group is configured with the same virtual IP address
- Virtual IP address MUST BE from the IP address space of a HSRP enabled interface
  - AND MUST NOT be same as a real IP address of a HSRP group member
  - Best practices
    - Virtual IP is the lowest, real routers have the highest IP
    - Virtual IP is the highest, real routers have the lowest IP

Basic Configuration

- Configure HSRP on the interface:

```
Router(config-if)# standby [ group-number ] ip virtual-IP
```

- All members of the group must have the same virtual IP address and group number (default group is 0)
- Disabling HSRP with all relevant commands:

```
Router(config-if)# no standby group-number
```

- To set the HSRP priority value of a router in range of 0 and 255, enter this command:

```
Router(config-if)# standby group-number priority prio
```

Configuring Preemption

- IF routers are not preemptive THEN a router that boots up significantly faster than the others in the standby group becomes the active router, regardless of the configured priority of the others
- The former active router can be configured to resume the forwarding router role by preempting a router with a lower priority:

```
Router(config-if)# standby group-number preempt
```

- Preemption capability could be delayed:

```
Router(config-if)#
  standby group-number preempt delay minimum SECONDS
  standby group-number preempt delay reload SECONDS
```

Configuring Timers and Version

- By default Hellotime is 3s, Holdtime 10s
- Hold should be at least 3× higher than Hello
- Timers SHOULD be consistent within a HSRP group

```
Router(config-if)#
   standby group-number timers Hello Holddown
```

- Timers can be set in milliseconds for HSRPv2

```
Router(config-if)#
   standby version 2
   standby group-number timers msec Hello msec Holddown
```

- Different versions CAN NOT be present on same router

Simple Example

- Routers A and B are configured with priorities of 110 and 90, respectively
- The **`preempt`** keyword ensures that Router A will be the HSRP active router as long its interface is active

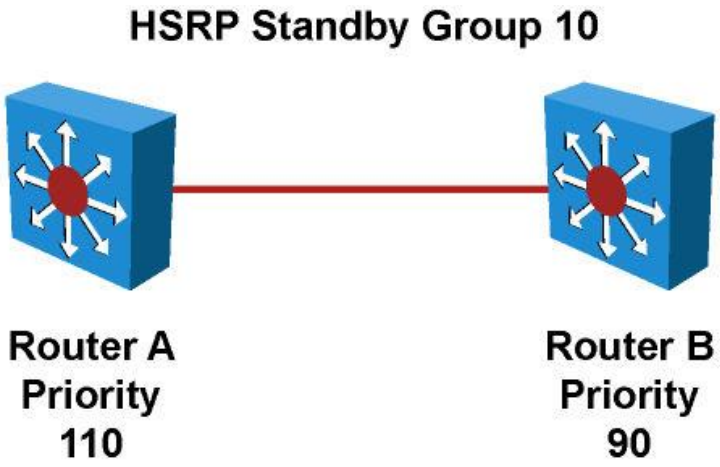


```
RouterA(config)# interface vlan 10
RouterA(config-if)# ip address 10.1.1.2 255.255.255.0
RouterA(config-if)# standby version 2
RouterA(config-if)# standby 10 ip 10.1.1.1
RouterA(config-if)# standby 10 priority 110
RouterA(config-if)# standby 10 preempt
```

The show standby brief Command

```
Sumperk# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0.1     1   150  P Active   local           172.16.10.2     172.16.10.110
Fa0/0.2     2   100  P Standby  172.16.20.2     local           172.16.20.120
```

```
Jesenik# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0.1     1   100  P Standby  172.16.10.1     local           172.16.10.110
Fa0/0.2     2   150  P Active   local           172.16.20.1     172.16.20.120
```

Interface/Object Tracking

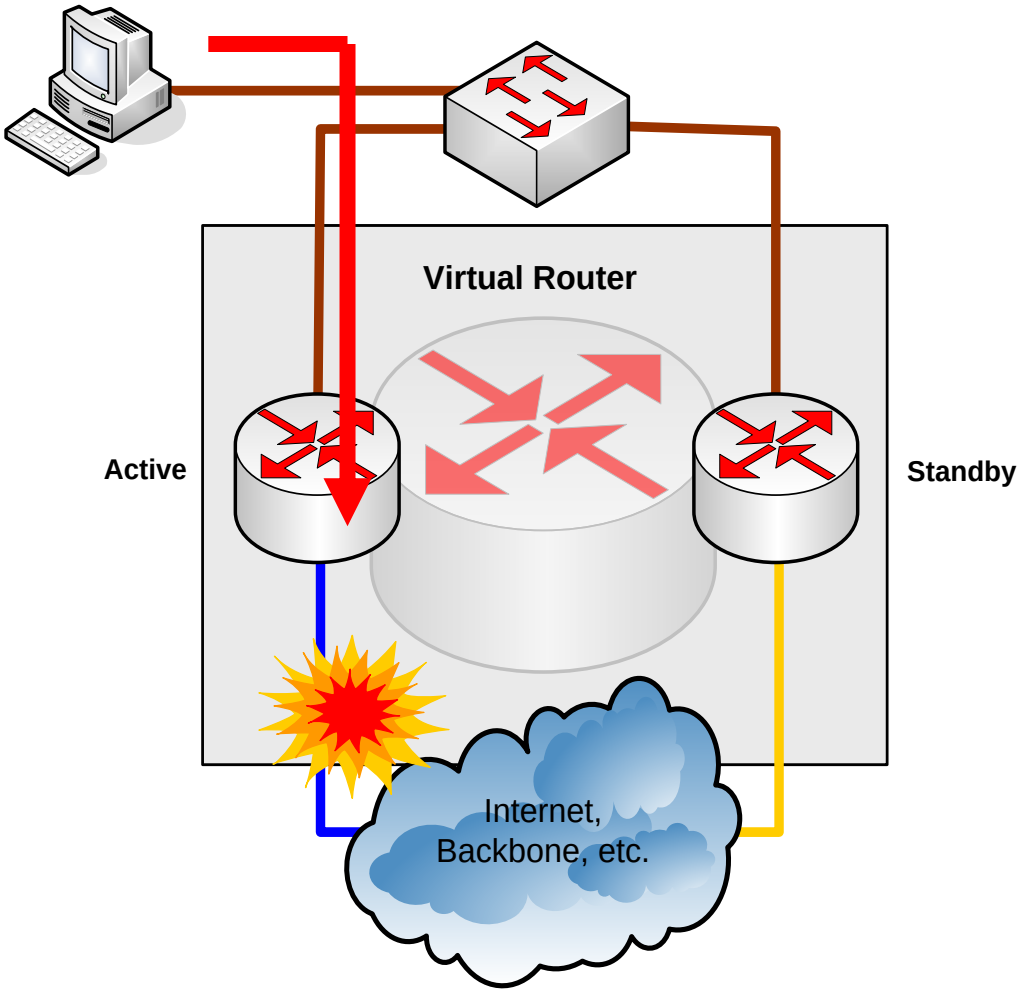

- *What if the link connecting active router to Internet fails?*
  - HSRP interfaces use limited ICMP redirect support
- Active router should renounce its role if the router cannot serve as a default gateway
- **Interface/Object tracking**
  - `IF` a monitored interface (object) fails `THEN` HSRP priority is decreased

Configure Interface Tracking

- Configure interface tracking:

```
Router(config-if)#
  standby [group-number] track IFACE [penalty]
```

Variable	Description
group-number	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
IFACE	Indicates the interface type and number that will be tracked.
Penalty	(Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. The default value is 10.

Configure Object Tracking

- Instead of particular interface it uses more general track object:

```
Switch(config)#
  track object-id interface IFACE {line-protocol|ip-routing}
Switch(config)# interface ...
Switch(config-if)#
  standby group-number track object-id [decrement penalty|shutdown]
```

- *E.g.:*

```
DLS1(config)# track 100 interface Port-channel 1 line-protocol
DLS1(config-track)#exit
DLS1(config)# int vlan 20
DLS1(config-if)# standby 1 track 100 ?
  decrement  Priority decrement
  shutdown   Shutdown group
  <cr>

DLS1(config-if)# standby 1 track 100 decrement 60
```

Authentication

- In case of FHRP authentication DOES NOT imply increased security
  - IF different passwords are used on two routers THEN both routers became the Active router which leads to vIP/vMAC conflict
  - Problem with virtual IP address configured as a real IP
- HSRP supports plain-text, MD5 (+key-chain authentication)
  - The plaintext authentication string is a max. 8 characters long (default keyword is 0x63 0x69 0x73 0x63 0x6F 0x00 0x00 0x00 a.k.a. “cisco”)

```
Switch(config-if)# standby group-number authentication string
```

```
Switch(config-if)#
   standby group-number authentication md5 key-string string
! Or variant benefiting existing key-chain
   standby group-number authentication md5 key-chain chain-name
```

Remarks

- It is important to prevent hosts from discovering the real IP/MAC address
  - `IF` a host knows and uses the real IP/MAC address of a router and router later fails `THEN` packets from the host will be lost
- Whenever HSRP is enabled on an interface it modifies behavior for ARP and ICMP protocols
  - The Active router replies with the MAC address of the virtual router if an ARP request is received from a host that is not on the local LAN
  - Other routers have Proxy ARP disabled
  - ICMP redirects are supported with some limitations

Load Balancing

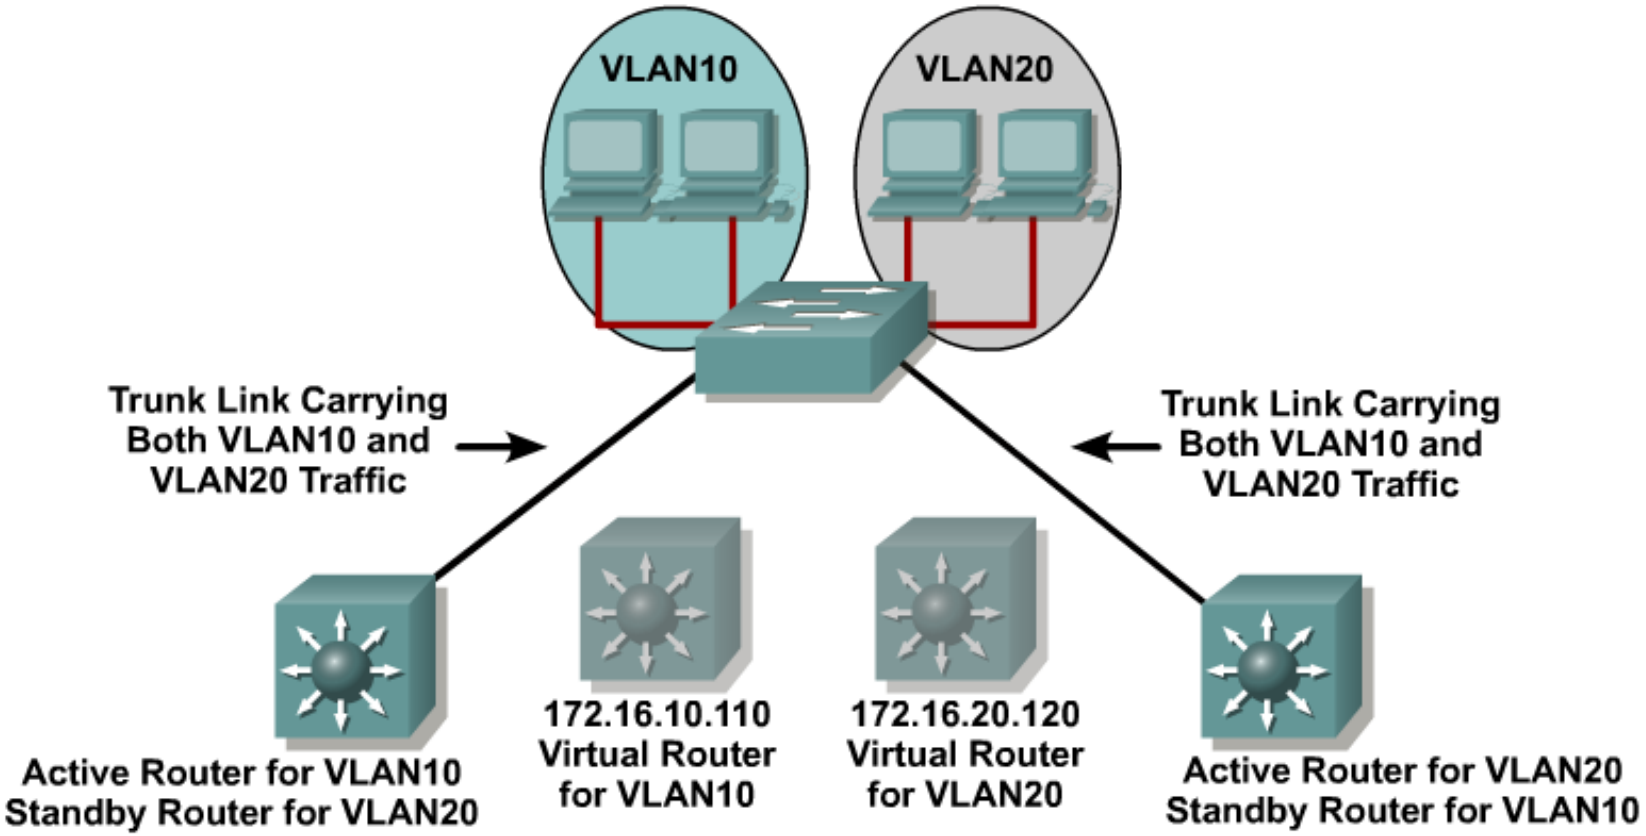

To load balance routers and links:

- Per VLAN, configure the HSRP active router and the spanning tree root to be the same multilayer switch.

Example: Load Balancing

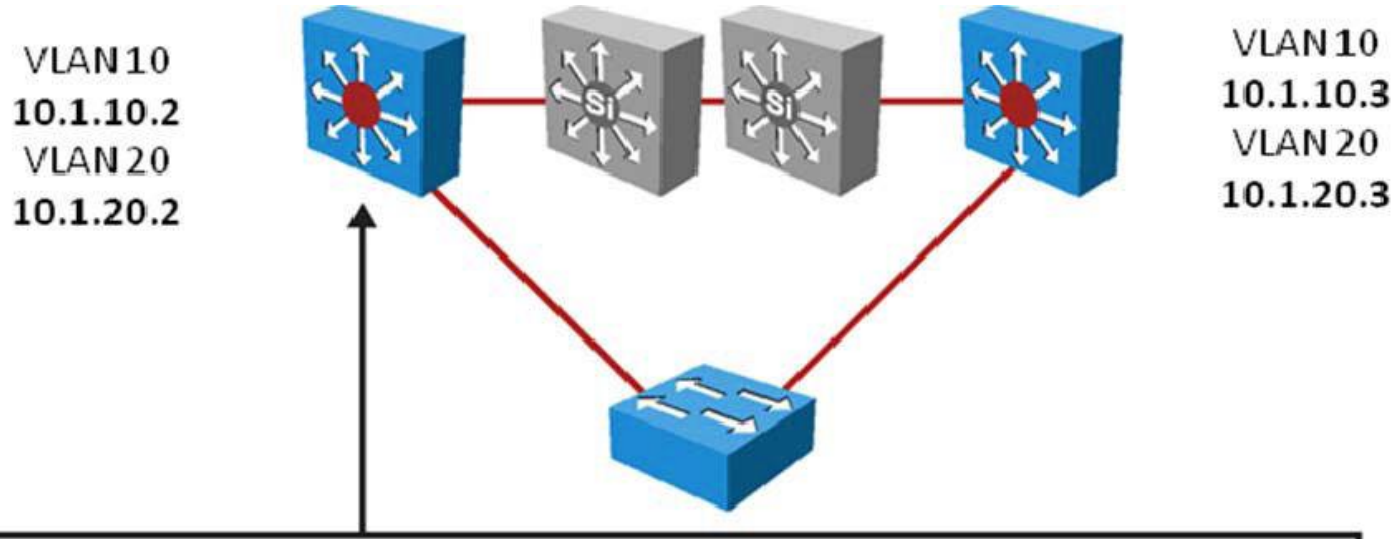


```
switch(config)# spanning-tree vlan 10 root primary
switch(config)# spanning-tree vlan 20 root secondary
switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.10.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.10.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
switch(config)# interface vlan 20
switch(config-if)# ip address 10.1.20.2 255.255.255.0
switch(config-if)# standby 20 ip 10.1.20.1
switch(config-if)# standby 20 priority 90
switch(config-if)# standby 20 preempt
```

Troubleshoot HSRP

```
show standby [brief]
show standby [IFACE [group-number]] [brief]
debug standby
```

The show standby Command

```
Sumperk# show standby
FastEthernet0/0.1 - Group 1
  State is Active
    11 state changes, last state change 00:05:16
  Virtual IP address is 172.16.10.110
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.784 secs
  Preemption enabled
  Active router is local
  Standby router is 172.16.10.2, priority 100 (expires in 9.788 sec)
  Priority 150 (configured 150)
  IP redundancy name is "hsrp-Fa0/0.1-1" (default)
FastEthernet0/0.2 - Group 2
  State is Standby
    7 state changes, last state change 01:41:07
  Virtual IP address is 172.16.20.120
  Active virtual MAC address is 0000.0c07.ac02
    Local virtual MAC address is 0000.0c07.ac02 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.988 secs
  Preemption enabled
  Active router is 172.16.20.2, priority 150 (expires in 7.796 sec)
  Standby router is local
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Fa0/0.2-2" (default)
```

Debug HSRP

- Displays all state changes to HSRP, including all hello packets (arguments minimize output):

```
DLS1# debug standby ?
  errors   HSRP errors
  events   HSRP events
  packets  HSRP packets
  terse    Display limited range of HSRP information
  <cr>
```

- Displays all HSRP errors, events, and packets, except hello and advertisement packets:

```
DLS1# debug standby terse
HSRP:
  HSRP Errors debugging is on
  HSRP Events debugging is on
    (protocol, neighbor, redundancy, track, ha, arp)
  HSRP Packets debugging is on
    (Coup, Resign)
```

The debug standby events Command

```
*Mar  3 05:38:28.502: HSRP: Vl10 Interface UP
*Mar  3 05:38:28.502: HSRP: Vl10 Starting minimum interface delay (1 secs)
*Mar  3 05:38:29.458: HSRP: Vl10 Grp 1 Active router is 172.16.10.102
*Mar  3 05:38:29.458: HSRP: Vl10 Nbr 172.16.10.102 is no longer passive
*Mar  3 05:38:29.458: HSRP: Vl10 Nbr 172.16.10.102 active for group 1
*Mar  3 05:38:29.500: HSRP: Vl10 Interface min delay expired
*Mar  3 05:38:29.500: HSRP: Vl10 Grp 1 Init: a/HSRP enabled
*Mar  3 05:38:29.500: HSRP: Vl10 Grp 1 Init -> Listen
*Mar  3 05:38:29.500: HSRP: Vl10 Grp 1 Redundancy "hsrp-Vl10-1" state Init -> Backup
*Mar  3 05:38:29.500: HSRP: Vl10 IP Redundancy "hsrp-Vl10-1" update, Init -> Backup
*Mar  3 05:38:30.507: %LINK-3-UPDOWN: Interface Vlan10, changed state to up
*Mar  3 05:38:30.515: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan10, changed
state to up
*Mar  3 05:38:32.260: HSRP: Vl10 Grp 1 Listen: h/Hello rcvd from lower pri Active router
(100/172.16.10.102)
*Mar  3 05:38:32.260: HSRP: Vl10 Grp 1 Active router is local, was 172.16.10.102
*Mar  3 05:38:32.260: HSRP: Vl10 Nbr 172.16.10.102 no longer active for group 1 (Listen)
*Mar  3 05:38:32.260: HSRP: Vl10 Nbr 172.16.10.102 Was active or standby - start passive
holddown
*Mar  3 05:38:32.260: HSRP: Vl10 Grp 1 Listen -> Active
*Mar  3 05:38:32.260: %HSRP-5-STATECHANGE: Vlan10 Grp 1 state Listen -> Active
*Mar  3 05:38:32.260: HSRP: Vl10 Grp 1 Redundancy "hsrp-Vl10-1" state Backup -> Active
*Mar  3 05:38:32.260: HSRP: Vl10 Added 172.16.10.1 to ARP (0000.0c07.ac01)
*Mar  3 05:38:32.268: HSRP: Vl10 Grp 1 Activating MAC 0000.0c07.ac01
*Mar  3 05:38:32.268: HSRP: Vl10 Grp 1 Adding 0000.0c07.ac01 to MAC address filter
*Mar  3 05:38:32.268: HSRP: Vl10 IP Redundancy "hsrp-Vl10-1" update, Backup -> Active
*Mar  3 05:38:35.254: HSRP: Vl10 IP Redundancy "hsrp-Vl10-1" update, Active -> Active
*Mar  3 05:38:42.913: HSRP: Vl10 Grp 1 Standby router is 172.16.10.102
*Mar  3 05:38:42.913: HSRP: Vl10 Nbr 172.16.10.102 is no longer passive
*Mar  3 05:38:42.913: HSRP: Vl10 Nbr 172.16.10.102 standby for group 1
```

Example: Authentication Error

```
Switch# debug standby errors
*Mar  3 05:40:49.606: HSRP: Vl1 Grp 1 Auth failed for Hello pkt
from 10.1.1.102, Text auth failed
*Mar  3 05:40:52.131: HSRP: Vl1 Grp 1 Auth failed for Hello pkt
from 10.1.1.102, Text auth failed
*Mar  3 05:40:54.715: HSRP: Vl1 Grp 1 Auth failed for Hello pkt
from 10.1.1.102, Text auth failed
```

Example: Active Election

```
DLS1# debug standby
HSRP debugging is on
DLS1#

*Mar 8 20:34:10.221: SB11: Vl11 Init: a/HSRP enabled
*Mar 8 20:34:10.221: SB11: Vl11 Init -> Listen
*Mar 8 20:34:20.221: SB11: Vl11 Listen: c/Active timer expired (unknown)
*Mar 8 20:34:20.221: SB11: Vl11 Listen -> Speak
*Mar 8 20:34:20.221: SB11: Vl11 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:23.101: SB11: Vl11 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:25.961: SB11: Vl11 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:28.905: SB11: Vl11 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:30.221: SB11: Vl11 Speak: d/Standby timer expired (unknown)
*Mar 8 20:34:30.221: SB11: Vl11 Standby router is local
*Mar 8 20:34:30.221: SB11: Vl11 Speak -> Standby
*Mar 8 20:34:30.221: SB11: Vl11 Hello out 172.16.11.111 Standby pri 100 ip 172.16.11.115
*Mar 8 20:34:30.221: SB11: Vl11 Standby: c/Active timer expired (unknown)
*Mar 8 20:34:30.221: SB11: Vl11 Active router is local
*Mar 8 20:34:30.221: SB11: Vl11 Standby router is unknown, was local
*Mar 8 20:34:30.221: SB11: Vl11 Standby -> Active
*Mar 8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Standby -> Active
*Mar 8 20:34:30.221: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
```

Example: Active Preemption

```
DLS1# debug standby
*Mar 1 00:16:41.295: %SYS-5-CONFIG_I: Configured from console by console
*Mar 1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar 1 00:16:43.099: SB: Vl11 Interface up
*Mar 1 00:16:43.099: SB11: Vl11 Init: a/HSRP enabled
*Mar 1 00:16:43.099: SB11: Vl11 Init -> Listen
*Mar 1 00:16:43.295: SB11: Vl11 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:16:43.295: SB11: Vl11 Active router is 172.16.11.112
*Mar 1 00:16:43.295: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:16:43.295: SB11: Vl11 Active router is local, was 172.16.11.112
*Mar 1 00:16:43.295: SB11: Vl11 Coup out 172.16.11.111 Listen pri 100 ip 172.16.11.115
 Mar 1 00:16:43.295
*Mar 1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Listen -> Active
*Mar 1 00:16:43.299: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:43.303: SB11: Vl11 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:44.095: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed state to up
*Mar 1 00:16:46.187: SB11: Vl11 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:46.207: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:49.095: SB11: Vl11 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:49.195: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:52.079: SB11: Vl11 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:52.147: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: Vl11 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: Vl11 Standby router is 172.16.11.112
*Mar 1 00:16:55.083: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:56.231: SB11: Vl11 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:16:58.023: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:59.223: SB11: Vl11 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:17:00.983: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:17:02.211: SB11: Vl11 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:17:03.847: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.11
```

Virtual Router Redundancy Protocol

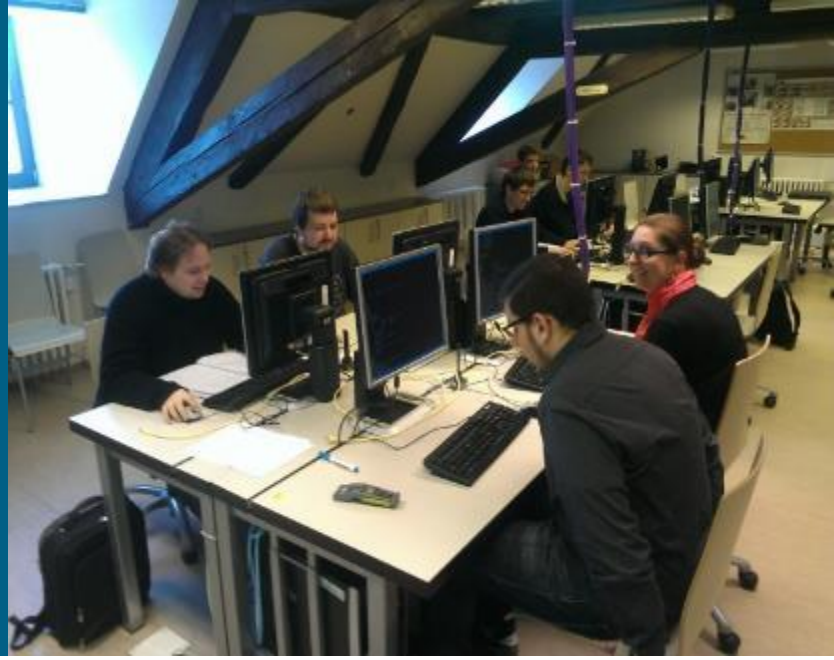

Virtual Router Redundancy Protocol

- IETF open standard
  - Own protocol number 112
  - IPv4 address 224.0.0.18
  - IPv4 address FF02::12
- VRRPv2 (RFC 3768)
- VRRPv3 (RFC 5798) – adds dual IPv4+IPv6 support
- *Almost the same as HSRP*

Similarities and Differences ①

- **VRRP group** (instead of HSRP group)
- **Master** (instead of Active)
  - Other routers in a VRRP group are called **Backup**
  - VRRP does not have Standby router
- VRRP allows virtual IP address same as real IP address of a member
  - The member is known as **IP Address Owner** and always win Master election (priority 255)
  - Election is based on priority (1-254) or higher IP address
- Virtual MAC address:
  - VRRPv2: `0000.5e00.01XX`
  - VRRPv3: `0000.5e00.02XX`
- VRRP is always preemptive by design

Similarities and Differences ②

- Instead of Hellos it sends **Advertisement** messages
  - Only Master is sending Advertisements
  - Hello timer is **Advertisement Timer** (AT)
  - By default 1 second
- Instead of Holdtime is **Master Down Interval** (MDI)
  - It IS NOT carried in Advertisement messages
  - **Skew time**: VRRPv2= $\frac{(256-priority)}{256}$, VRRPv3= $\frac{(256-priority)\times AT}{256}$
  - It is computed dynamically using following formula:

$$MDI = 3 \times AT + skew\_time$$

- Authentication of VRRP is considered deprecated
  - However, on Cisco boxes it is still supported as HSRP authentication

Operation

1) *A* sends advertisements
2) *A* fails and stops sending Advertisements
3) *B* and *C* stops receiving Advertisements and Master Down Interval on both of them starts to expire
4) Because of a scew time *B* expires it sooner (≈ 3.2 seconds) than *C* (≈ 3.6 seconds)
5) *B* transitions to master state and starts sending Advertisements
6) *C* receives Advertisement from B. Hence, it resets own Master Down Interval and continue to be backup router.

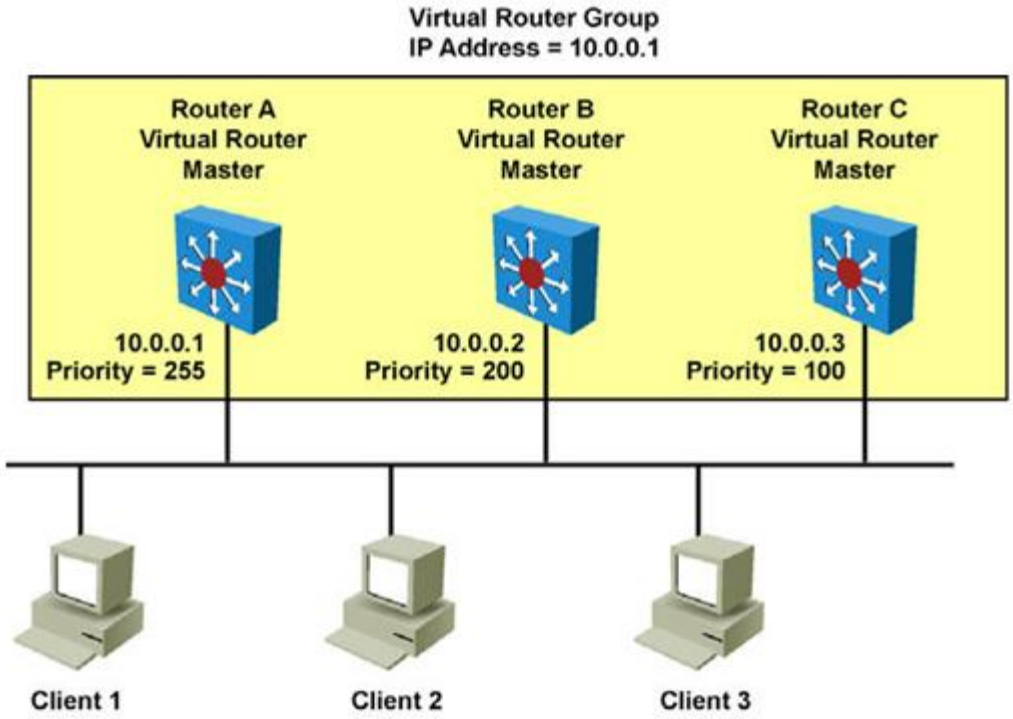

VRRP Message

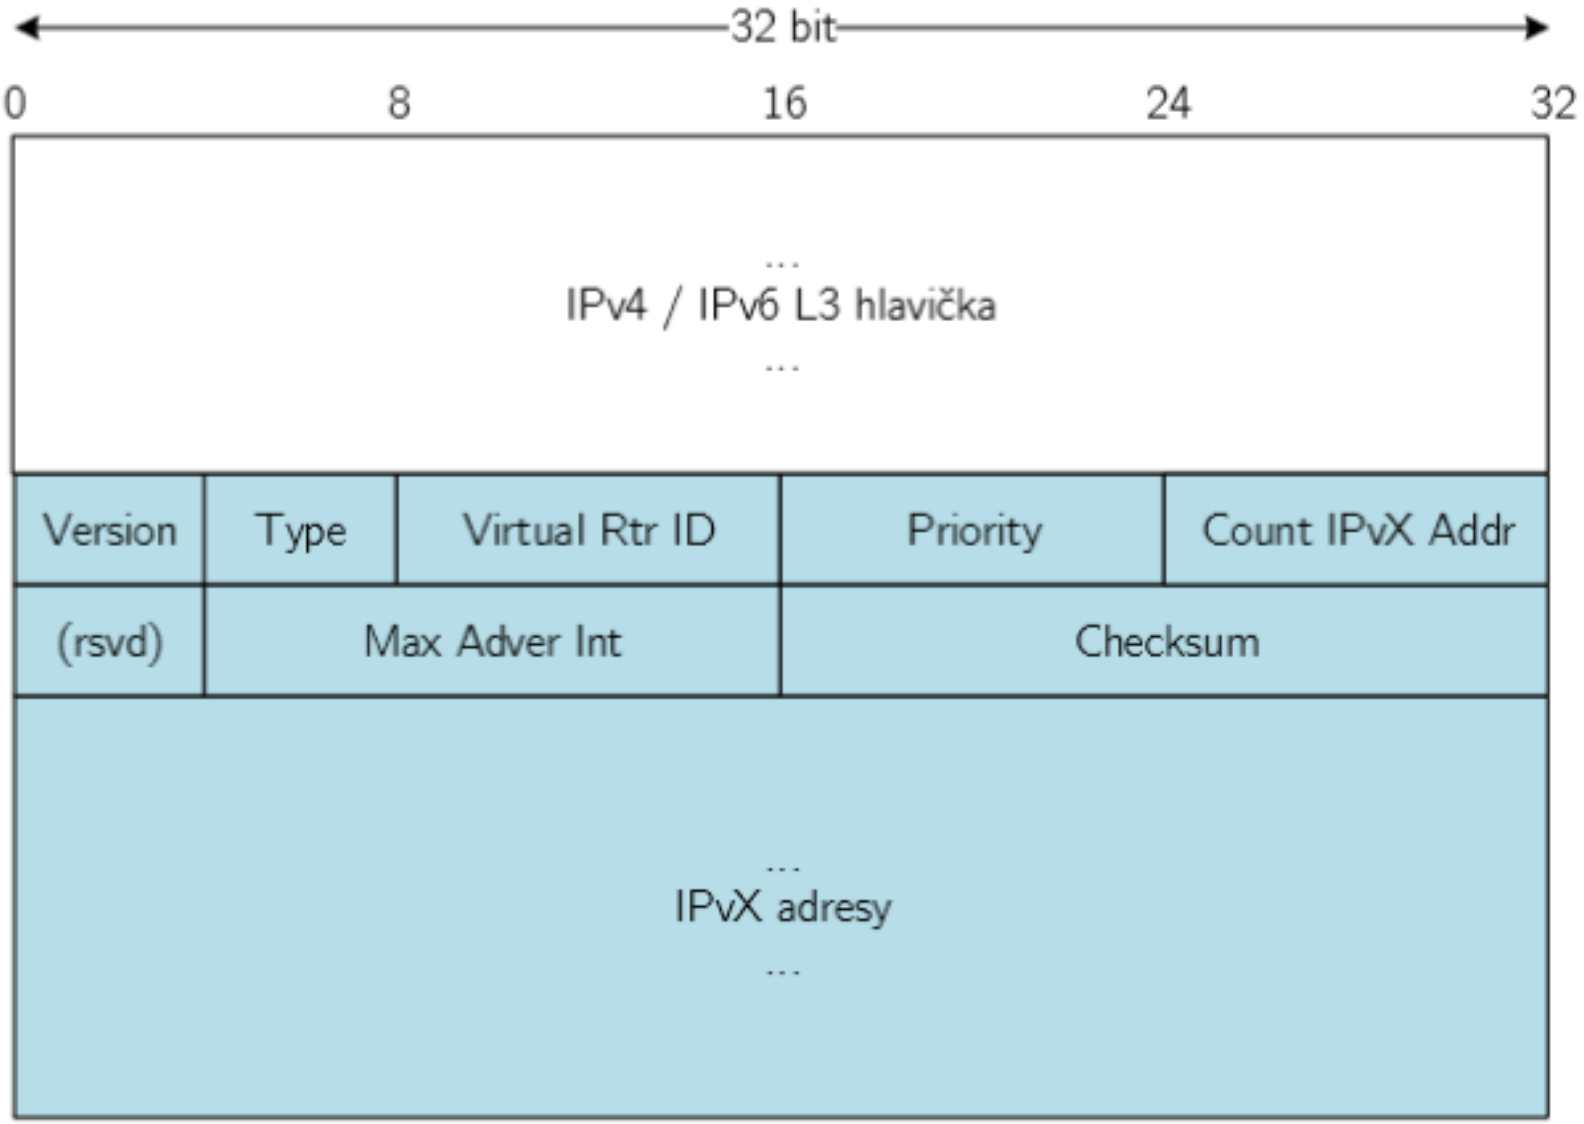

VRRP State Diagram

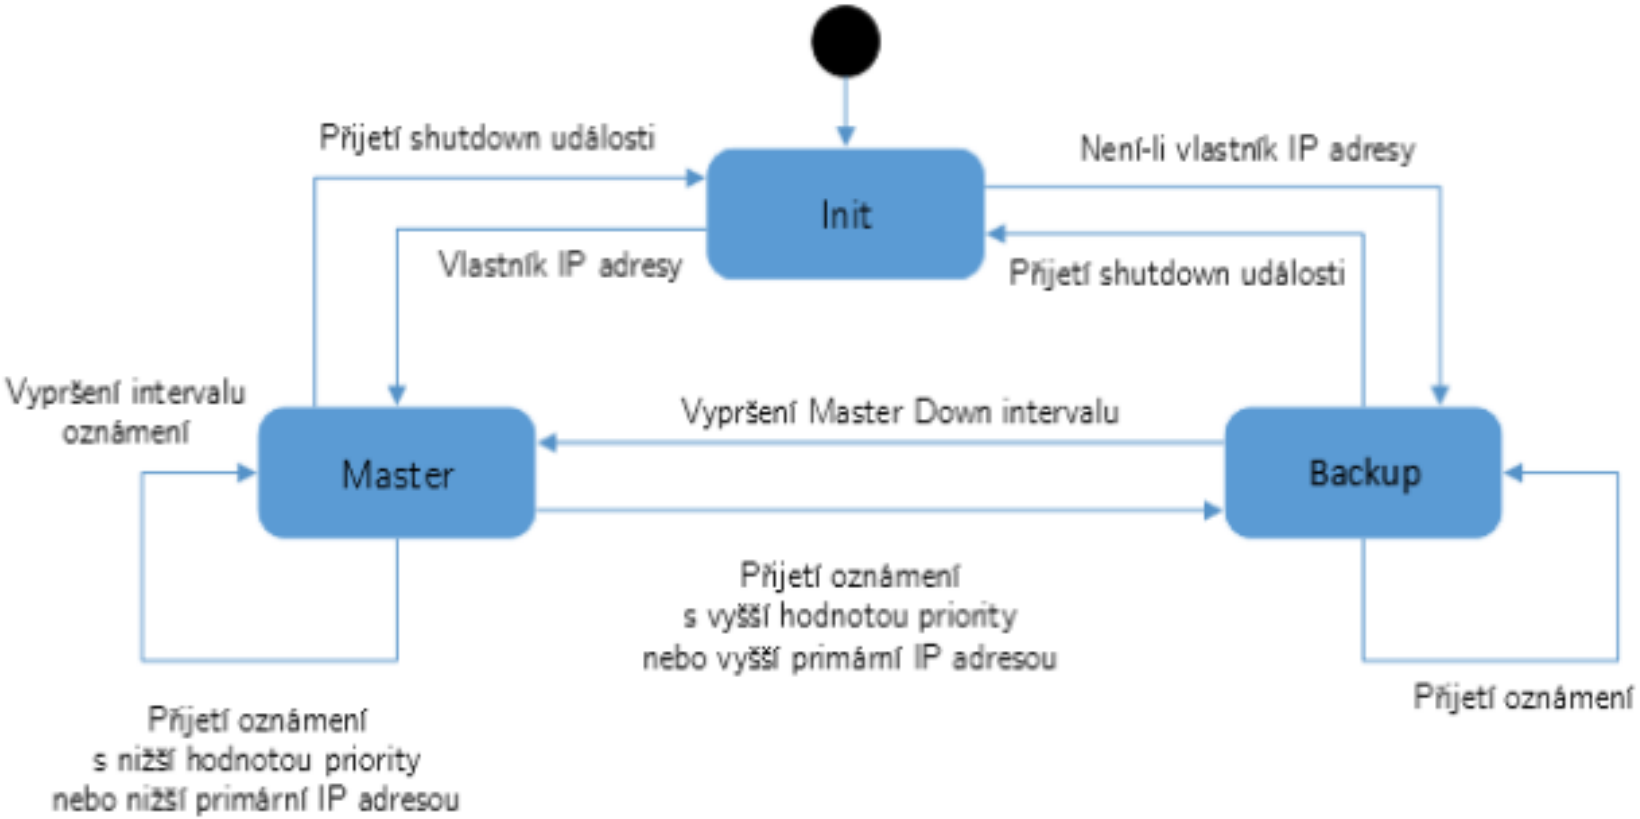

Basic Configuration

- This makes the interface a member of the virtual group identified with the IP virtual address:

```
Switch(config-if)# vrrp group-number ip virtual-ip
```

- To set a VRRP priority:
  - Default is 100
  - Virtual IP owner always has 255, Master abdicates by value 0

```
Switch(config-if)# vrrp group-number priority priority-value
```

- To change timer and indicate if it should advertise for master or just learn for backup routers

```
Switch(config-if)#
   vrrp group-number timers advertise [msec] timer-value
! Or learn them dynamically
   vrrp group-number timers learn
```

Other Configurations

- Append description to VRRP interface:

```
Switch(config-if)# vrrp group-number description string
```

- Delay preemption ability so that device has enough time to rebuild control plane:

```
Switch(config-if)#
  vrrp group-number preempt delay minimum SECONDS
  vrrp group-number preempt delay reload SECONDS
```

- Object tracking:

```
Switch(config)#
   track object-id interface IFACE line-protocol
Switch(config-if)#
   vrrp group-number track object-id decrement penalty
```

Simple Example

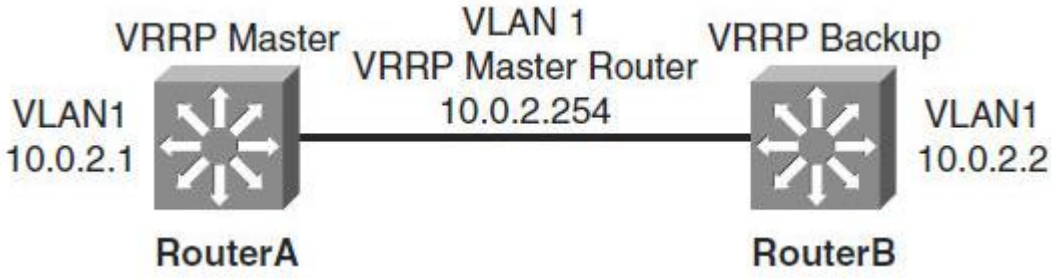


```
RouterA# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
RouterA(config)# interface vlan 1
RouterA(config-if)# ip address 10.0.2.1 255.255.255.0
RouterA(config-if)# vrrp 1 ip 10.0.2.254
RouterA(config-if)# vrrp 1 timers advertise msec 500
RouterA(config-if)# end

RouterB# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
RouterB(config)# interface vlan 1
RouterB(config-if)# ip address 10.0.2.2 255.255.255.0
RouterB(config-if)# vrrp 1 ip 10.0.2.254
RouterB(config-if)# vrrp 1 priority 90
RouterB(config-if)# vrrp 1 timers learn
RouterB(config-if)# end
```

Troubleshooting

```
show vrrp [brief]
show vrrp all
show vrrp GROUP_NUM

debug vrrp all
debug vrrp error
debug vrrp events
debug vrrp packets
debug vrrp state
```

The `show vrrp interface` Command

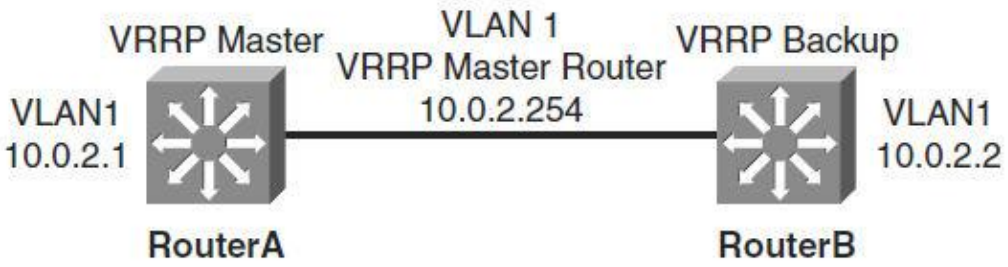


```
RouterA# show vrrp interface vlan 1
Vlan1 - Group 1
  State is Master
  Virtual IP address is 10.0.2.254
  Virtual MAC address is 0000.5e00.0101
  Advertisement interval is 0.500 sec
  Preemption is enabled
    min delay is 0.000 sec
Priority is 100
Master Router is 10.0.2.1 (local), priority is 100
Master Advertisement interval is 0.500 sec
Master Down interval is 2.109 sec

RouterB# show vrrp interface vlan 1
Vlan1 - Group 1
  State is Backup
  Virtual IP address is 10.0.2.254
  Virtual MAC address is 0000.5e00.0101
  Advertisement interval is 0.500 sec
  Preemption is enabled
    min delay is 0.000 sec
Priority is 90
Master Router is 10.0.2.1, priority is 100
Master Advertisement interval is 0.500 sec
Master Down interval is 2.109 sec (expires in 1.745 sec)
```

Debug Authentication ①

```
Router1# show vrrp

Ethernet0/1 - Group 1
State is Master
Virtual IP address is 10.21.0.10
Virtual MAC address is 0000.5e00.0101
Advertisement interval is 1.000 sec
Preemption is enabled
 min delay is 0.000 sec
Priority is 100
  Authentication MD5, key-string
Master Router is 10.21.0.1 (local), priority is 100
Master Advertisement interval is 1.000 sec
Master Down interval is 3.609 sec
```

Debug Authentication ②

```
Router1#: debug vrrp authentication

VRRP: Sent: 21016401FE050000AC1801FE0000000000000000
VRRP: HshC: B861CBF1B9026130DD34AED849BEC8A1

VRRP: Rcvd: 21016401FE050000AC1801FE0000000000000000
VRRP: HshC: B861CBF1B9026130DD34AED849BEC8A1
VRRP: HshR: C5E193C6D84533FDC750F85FCFB051E1
VRRP: Grp 1 Adv from 172.24.1.2 has failed MD5 auth


Router2#: debug vrrp authentication

VRRP: Sent: 21016401FE050000AC1801FE0000000000000000
VRRP: HshC: C5E193C6D84533FDC750F85FCFB051E1

VRRP: Rcvd: 21016401FE050000AC1801FE0000000000000000
VRRP: HshC: C5E193C6D84533FDC750F85FCFB051E1
VRRP: HshR: B861CBF1B9026130DD34AED849BEC8A1
VRRP: Grp 1 Adv from 172.24.1.1 has failed MD5 auth
```

HSRP vs. VRRP

HSRP	VRRP
HSRP is a Cisco proprietary protocol, created in 1994, and formalized with the RFC 2281 in March 1998	VRRP is an IEEE standard (RFC 2338 in 1998; then RFC 3768 in 2005) for router redundancy
256/4096 groups max	256 groups max
1 active, 1 standby, several candidates	1 active, several backups
Virtual IP is different from Active and Standby real IP addresses	Virtual IP can be the same as one of the group members real IP address
Uses 224.0.0.2/102 or FF02::66 for packets	Uses 224.0.0.18 or FF02:0::12 for advertisements
Default timers: hello 3 s, holdtime 10 s	The default timers are shorter in VRRP than HSRP. This often gave VRRP the reputation of being faster than HSRP
Can track interfaces or objects	Can track only objects
Uses authentication within each group by default. When authentication is not configured, a default authentication, using “cisco” as the password	Supports plaintext and HMAC/MD5 authentication methods (RFC 2338). The new VRRP RFC (RFC 3768) removes support for these methods. The consequence is that VRRP does not support authentication anymore. Nevertheless, current Cisco IOS still supports the RFC 2338 authentications mechanisms

Gateway Load Balancing Protocol

Gateway Load Balancing Protocol

- HSRP/VRRP standby/backup resources are not fully utilized
  - Load balancing can be accomplished through the creation of multiple groups and through the assignment of multiple default gateways
  - Load balancing configuration creates an administrative burden
- Gateway Load Balance Protocol (U.S. Patent 7881208)
  - Document ID 81565: GLBP on Catalyst 6500 Switches Configuration Example
- Simultaneous use of multiple available routers in addition to automatic failover
  - Efficient resource utilization
  - Load sharing
- Basic characteristics
  - Max. 1024 groups
  - One virtual IP per group
  - Up to 4 virtual MAC addresses per group

Device Roles ①

- **Active virtual gateway (AVG)**
  - Elected router with the highest priority (plus the highest IP address)
  - 1 active AVG and 1 standby AVG in a GLBP group
  - Assigns a virtual MAC address (`0007.b40X.XX0[1-4]`) to each member of the GLBP group
  - Responds with virtual MAC addresses in ARP Reply
- **Active virtual forwarder (AVF)**
  - Up to 4 AVFs per GLBP group
  - AVF assumes responsibility for forwarding packets that are sent to the virtual MAC address assigned by the AVG
  - AVG may be also AVF
- **Backup AVG/AVF**
  - Other routers in a group provide backup for the AVG/AVF if they becomes unavailable

Device Roles ②

- **Virtual Forwarder (VF)**
  - Router that obtained virtual MAC address
- **Primary Virtual Forwarder (PVF)**
  - Router that with virtual MAC address assigned from AVG
- **Secondary Virtufal Forwarder (SVF)**
  - Router that learned MAC address from Hello messages of PVF

Operation

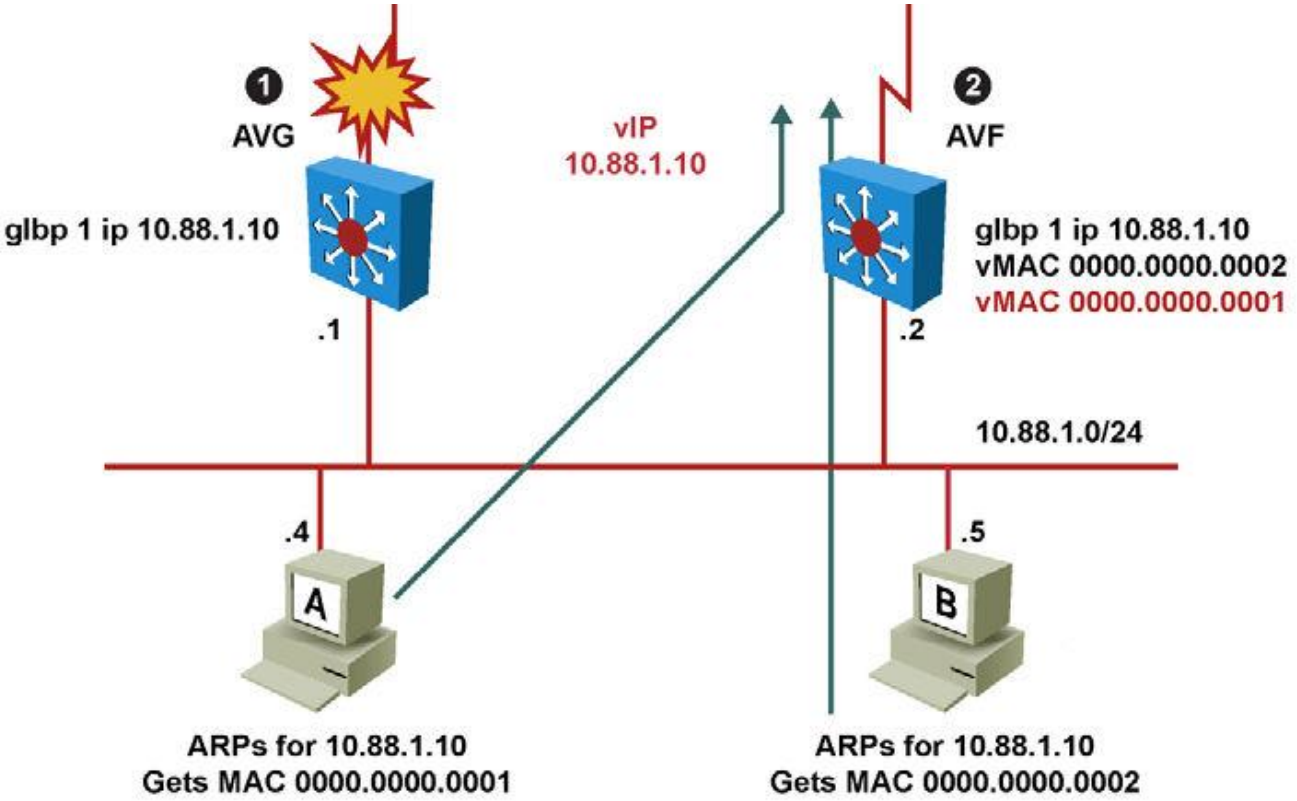
1
AVG
glbp 1 ip 10.88.1.10
.1
vIP
10.88.1.10
2
AVF
glbp 1 ip 10.88.1.10
vMAC 0000.0000.0002
vMAC 0000.0000.0001
.2
10.88.1.0/24
.4
A
.5
B
ARPs for 10.88.1.10
Gets MAC 0000.0000.0001
ARPs for 10.88.1.10
Gets MAC 0000.0000.0002

GLBP General Message Header

- Three message-like TLVs
  - **Hello**
  - **Request-Response**
  - **Auth**
- Common header

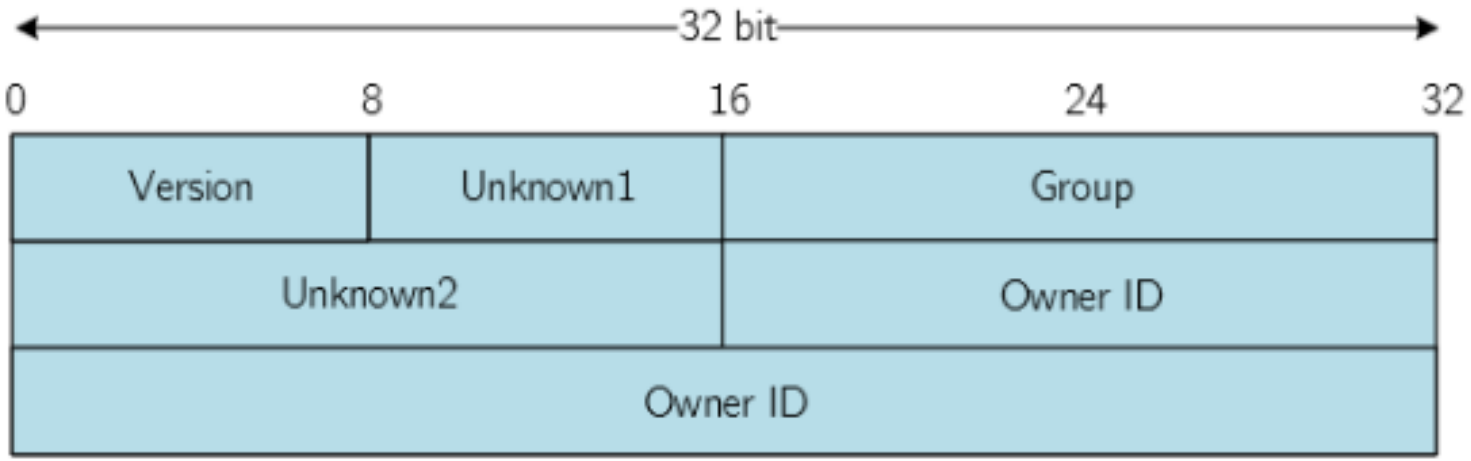

Hello TLV

- Contains fields related to AVG state
- Just one occurence

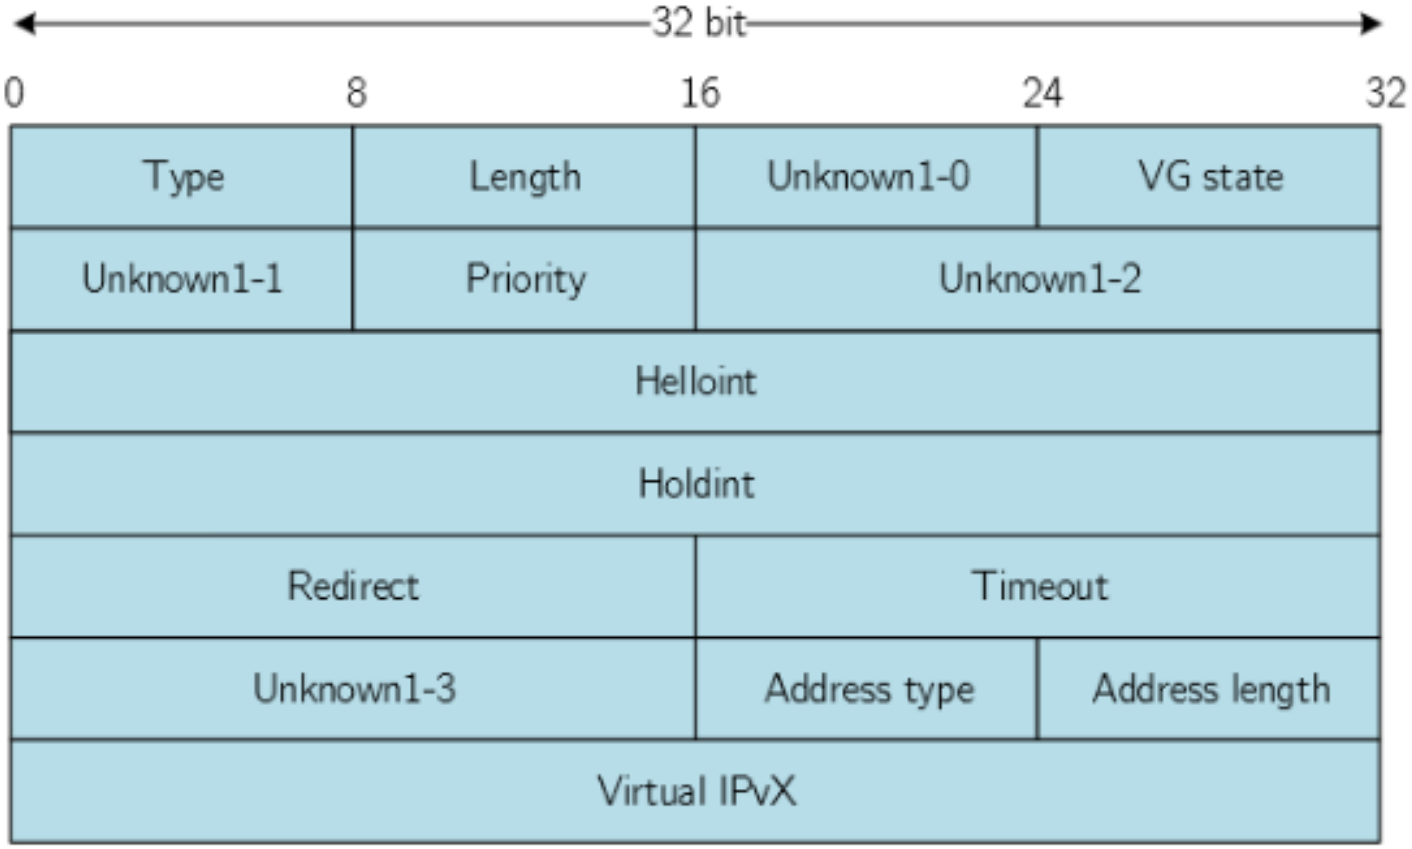

Request-Response TLV

- Contains fields related to AVF state
- Up to four occurance per each AVF instance

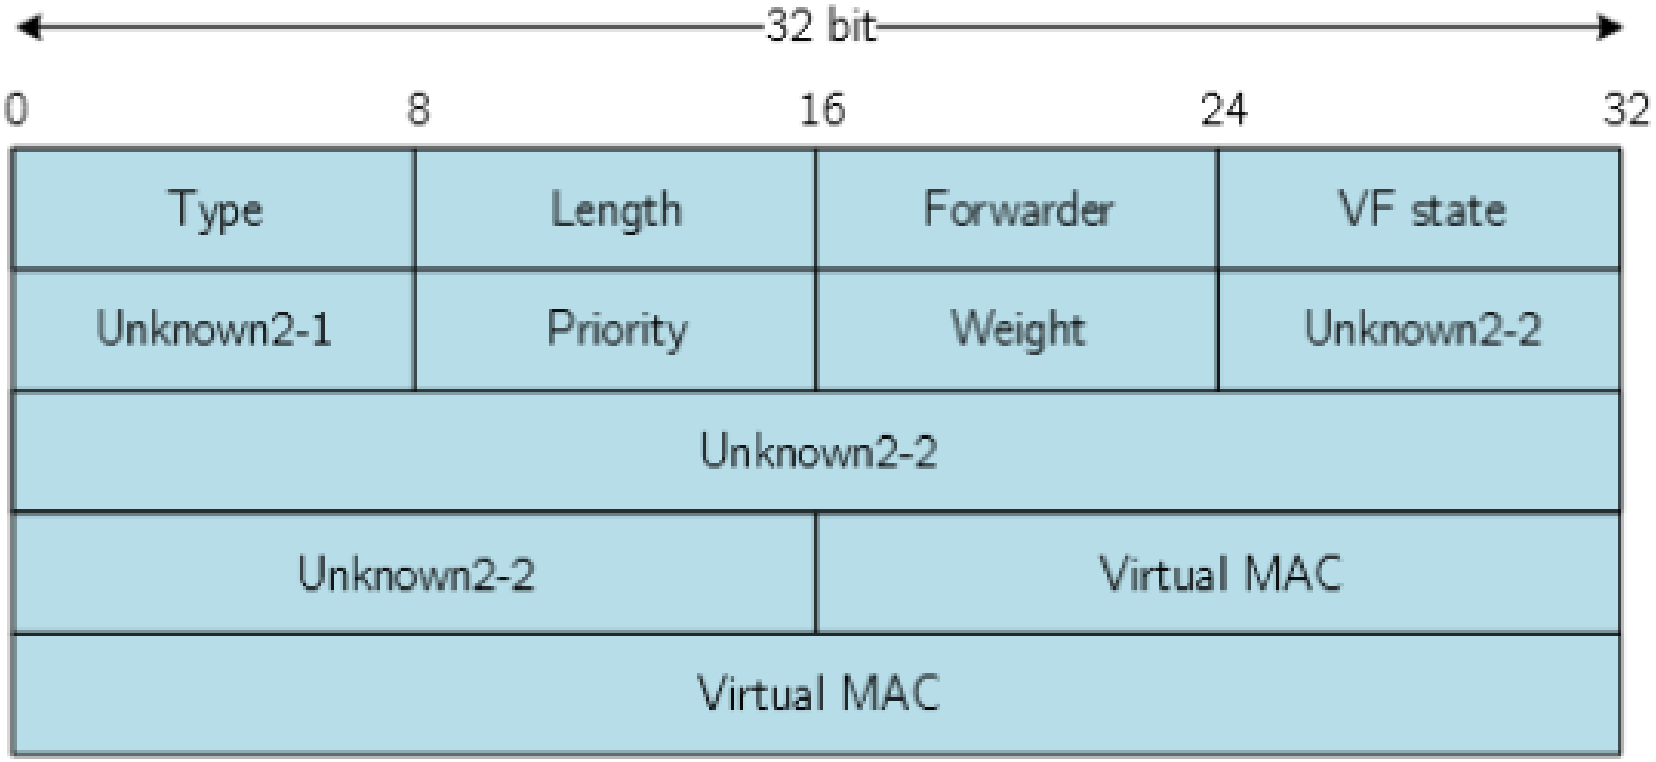

Basic Configuration

- Enable GLBP on an interface:

```
Switch(config-if)# glbp group-number ip virtual-ip
```

- Set a priority for this router within GLBP group:

```
Switch(config-if)# glbp group-number priority priority-value
```

  - The highest value wins election as active router
  - The default is 100
  - IF routers have the same GLBP priority THEN the gateway with the highest real IP address becomes the AVG

- Change timer values for hello interval and holdtime (use the argument msec to enter subsecond values):

Hello and Holddown Timers

- **Hello timer**
  - Time between two consecutive GLBP Hello messages
  - Default value is 3 seconds
- **Holddown timer**
  - AVG is considered down after expiration of this timer
  - Default value is 10 seconds
  - Range $\langle hello, 3 \times hello + 1{,}160 \rangle$ seconds

```
Switch(config-if)#
   glbp group-number timers [msec] redirect [msec] timeout
```

Redirect and Timeout Timers

- **Redirect timer**
  - Time during AVG forwards ARP Requests to AVF
  - After expiration a given AVF is removed from AVF group
  - Default value is 600 seconds (10 minutes)
- **Timeout timer**
  - Time after which secondary AVF becomes unavailable
  - Default value is 14 400 seconds (4 hours)
  - Range $\langle redirect + 600, 64\,800 \rangle$ seconds

```
Switch(config-if)#
   glbp group-number timers redirect redirect timeout
```

Load Balancing Mechanism

- GLBP supports these operational modes for load balancing
  - **Round-robin load-balancing algorithm**
    - Default one
    - Round-robin fashion for the ARP-replay
  - **Weighted load-balancing algorithm**
    - The amount of load directed to a router is dependent upon the weighting value advertised by that router
  - **Host-dependent load-balancing algorithm**
    - A host is guaranteed use of the same virtual MAC address
- Load balancing algorithm could be configured per group:

```
Switch(config-if)#
glbp group load-balancing {host-dependent|round-robin|weighted}
```

Delay Configurations

- Both AVG and AVF could be delayed before preemption takes place:

```
Switch(config-if)#
!Configure delay for AVG functionality after recovery
  glbp group-number preempt delay minimum SECONDS
!Configure delay for AVG functionality after reboot
  glbp group-number preempt delay reload SECONDS
!Configure delay for AVF functionality after recovery
  glbp group-number preempt forwarder delay minimum SECONDS
!Configure delay for AVF functionality after reboot
  glbp group-number preempt forwarder delay reload SECONDS
```

Weight Configurations

- **Weight**
  - Weight value affects weighted load-balancing
  - Default weight value is 100
  - Default lower weight value is 1
  - Default upper weight value is 100
- Whether router is AVF or not is determined by weight that could be configured to track object:

```
Switch(config)# track object-id interface IFACE {line-protocol | ip-routing}
Switch(config)# interface type number
Switch(config-if)# glbp group weighting maximum [lower lower] [upper upper]
Switch(config-if)# glbp group weighting track object-id [decrement value]
```

Simple Example Explaining Weight

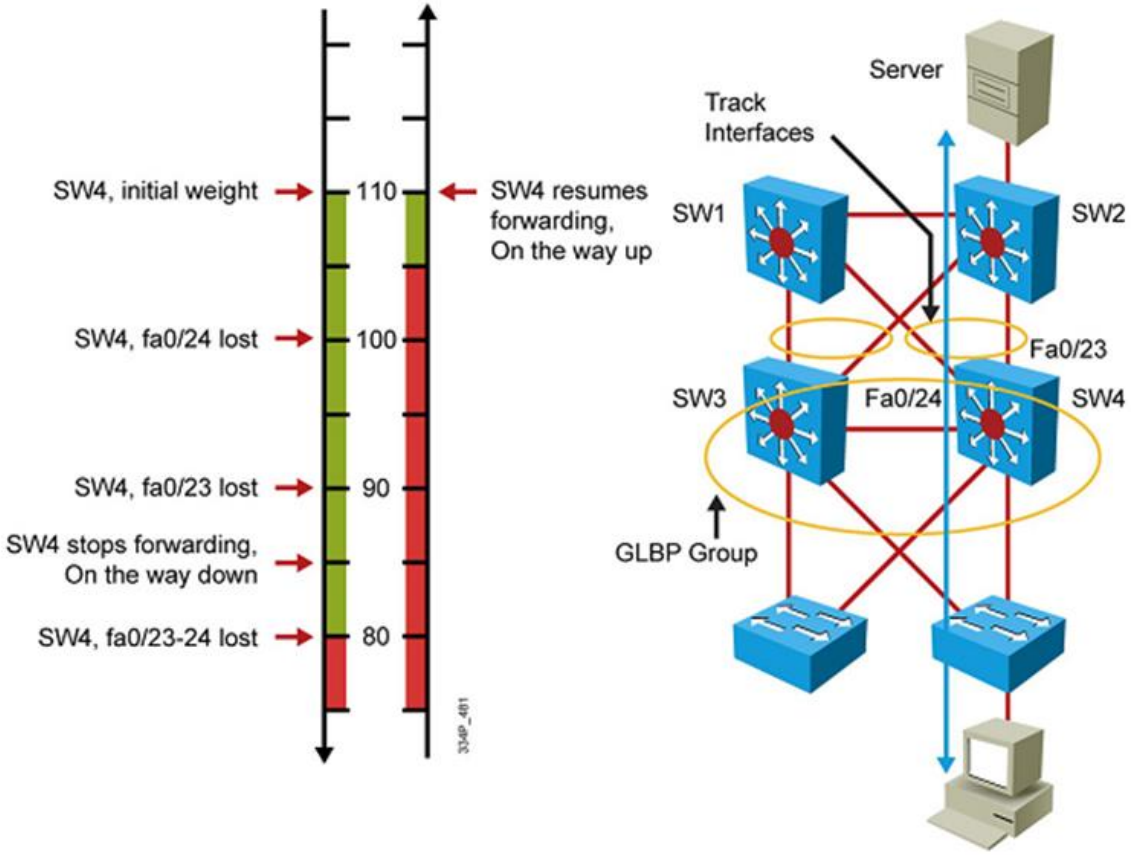


```
SW4(config)# track 90 interface fa0/24 line-protocol
SW4(config)# track 91 interface fa0/23 line-protocol
SW4(config)# interface vlan10
SW4(config-if)# ip address 10.1.10.2 255.255.255.0
SW4(config-if)# glbp 1 10.1.10.1
SW4(config-if)# glbp 1 weighting 110 lower 85 upper 105
SW4(config-if)# glbp 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
```

Troubleshooting

```
show glbp
show glbp brief
show glbp all
show glbp GROUP_NUM

debug glbp all
debug glbp error
debug glbp events
debug glbp packets
debug glbp state
```

Example

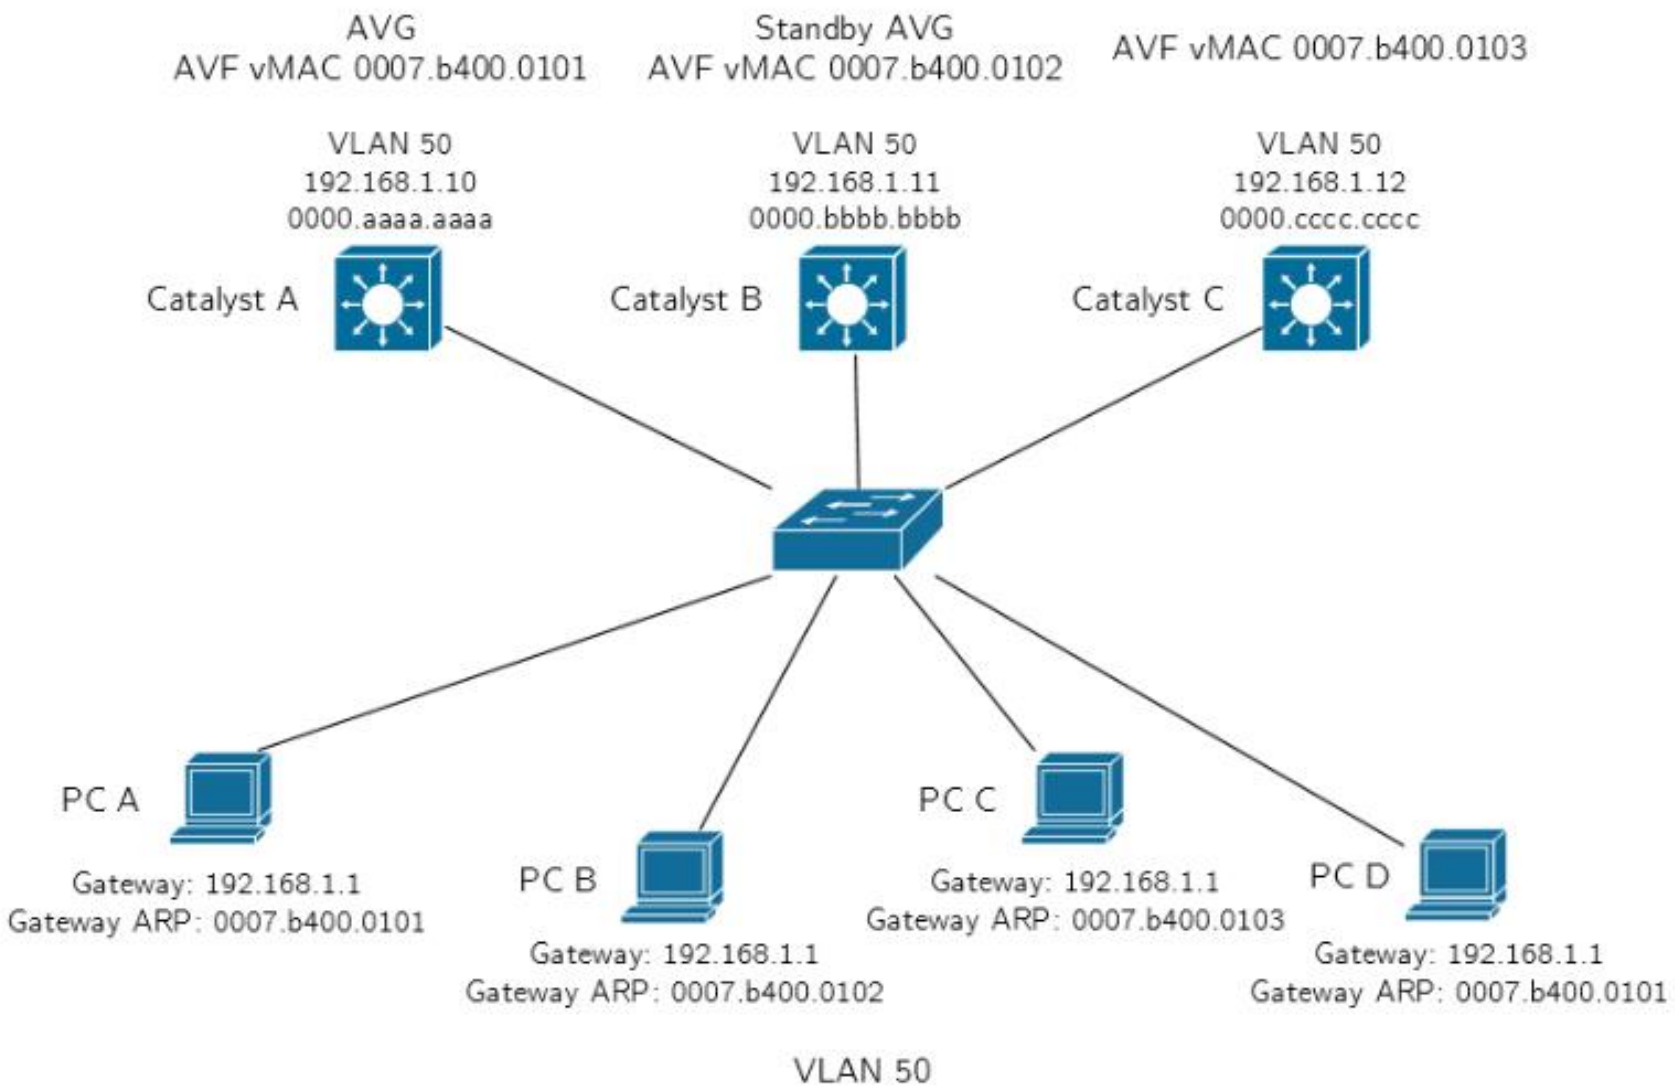
AVG
AVF vMAC 0007.b400.0101
Standby AVG
AVF vMAC 0007.b400.0102
AVF vMAC 0007.b400.0103
VLAN 50
192.168.1.10
0000.aaaa.aaaa
VLAN 50
192.168.1.11
0000.bbbb.bbbb
VLAN 50
192.168.1.12
0000.cccc.cccc
Catalyst A
Catalyst B
Catalyst C
PC A
Gateway: 192.168.1.1
Gateway ARP: 0007.b400.0101
PC B
Gateway: 192.168.1.1
Gateway ARP: 0007.b400.0102
PC C
Gateway: 192.168.1.1
Gateway ARP: 0007.b400.0103
PC D
Gateway: 192.168.1.1
Gateway ARP: 0007.b400.0101
VLAN 50

Example: Configuration

```
CatalystA(config)# interface vlan 50
CatalystA(config-if)# ip address 192.168.1.10 255.255.255.0
CatalystA(config-if)# glbp 1 priority 200
CatalystA(config-if)# glbp 1 preempt
CatalystA(config-if)# glbp 1 ip 192.168.1.1
```

```
CatalystB(config)# interface vlan 50
CatalystB(config-if)# ip address 192.168.1.11 255.255.255.0
CatalystB(config-if)# glbp 1 priority 150
CatalystB(config-if)# glbp 1 preempt
CatalystB(config-if)# glbp 1 ip 192.168.1.1
```

```
CatalystC(config)# interface vlan 50
CatalystC(config-if)# ip address 192.168.1.12 255.255.255.0
CatalystC(config-if)# glbp 1 priority 100
CatalystC(config-if)# glbp 1 ip 192.168.1.1
```

Example: Verification

```
CatalystA# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active router   Standby router
Vl50        1    -   200 Active   192.168.1.1     local           192.168.1.11
Vl50        1    1   7   Active   0007.b400.0101  local           -
Vl50        1    2   7   Listen   0007.b400.0102  192.168.1.11    -
Vl50        1    3   7   Listen   0007.b400.0103  192.168.1.13    -
CatalystB# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active router   Standby router
Vl50        1    -   150 Standby  192.168.1.1     192.168.1.10    local
Vl50        1    1   7   Listen   0007.b400.0101  192.168.1.10    -
Vl50        1    2   7   Active   0007.b400.0102  local           -
Vl50        1    3   7   Listen   0007.b400.0103  192.168.1.13    -
CatalystB#
```

```
CatalystC# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active router   Standby router
Vl50        1    -   100 Listen   192.168.1.1     192.168.1.10    192.168.1.11
Vl50        1    1   7   Listen   0007.b400.0101  192.168.1.10    -
Vl50        1    2   7   Listen   0007.b400.0102  192.168.1.11    -
Vl50        1    3   7   Active   0007.b400.0103  local           -
CatalystC#
```

Example: Verification on CatalystA

```
CatalystA# show glbp
Vlan50 - Group 1
State is Active
7 state changes, last state change 03:28:05
Virtual IP address is 192.168.1.1
Hello time 3 sec, hold time 10 sec
Next hello sent in 1.672 secs
Redirect time 600 sec, forwarder time-out 14400 sec
Preemption enabled, min delay 0 sec
Active is local
Standby is 192.168.1.11, priority 150 (expires in 9.632 sec)
Priority 200 (configured)
Weighting 100 (default 100), thresholds: lower 1, upper 100
Load balancing: round-robin
There are 3 forwarders (1 active)
Forwarder 1
State is Active
3 state changes, last state change 03:27:37
MAC address is 0007.b400.0101 (default)
…
```

HSRP vs. GLBP

HSRP	GLBP
Cisco Proprietary, 1994	Cisco Proprietary, 2005
256 groups max	1024 groups max
1 active 1 standby several candidates	1 active AVG, 1 standby AVG up to 4 AVF AVG load balances traffic among AVF
Virtual IP is different from the real IP addresses of HSRP routers	Virtual IP is different from the real IP addresses of GLBP routers
1 Virtual MAC address for each group	1 Virtual MAC address per AVF in each group
Uses 224.0.0.2/102 for hello packets	Uses 224.0.0.102 for hello packets
Default timers: hello 3 s, holdtime 10 s	Default timers: hello 3 s, holdtime 10 s
Can track interfaces or objects	Can track only objects
Authentication supported	Authentication supported

Where to go next?

- Catalyst 3560 Command Reference

  www.cisco.com/en/US/partner/docs/switches/lan/catalyst3560/software/release/12.2_55_se/command/reference/3560_cr.html

- Configuring HSRP:

  www.cisco.com/en/US/partner/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/swhsrp.html

- Configuring VRRP:

  www.cisco.com/en/US/partner/docs/ios/ipapp/configuration/guide/ipapp_vrrp.html

- Configuring GLBP:

  www.cisco.com/en/US/partner/docs/ios/ipapp/configuration/guide/ipapp_glbp.htm

- Configuring Enhanced Object Tracking:

  www.cisco.com/en/US/partner/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/sweot.html

Slides adapted by Matěj Grégr and tuned by Vladimír Veselý
partially from official course materials
but the most of the credit goes to CCIE#23527 Ing. Peter Palúch, Ph.D.

The last update: 2016-11-03