



First Hop Redundancy Protocols



SWITCH Module 5

Agenda

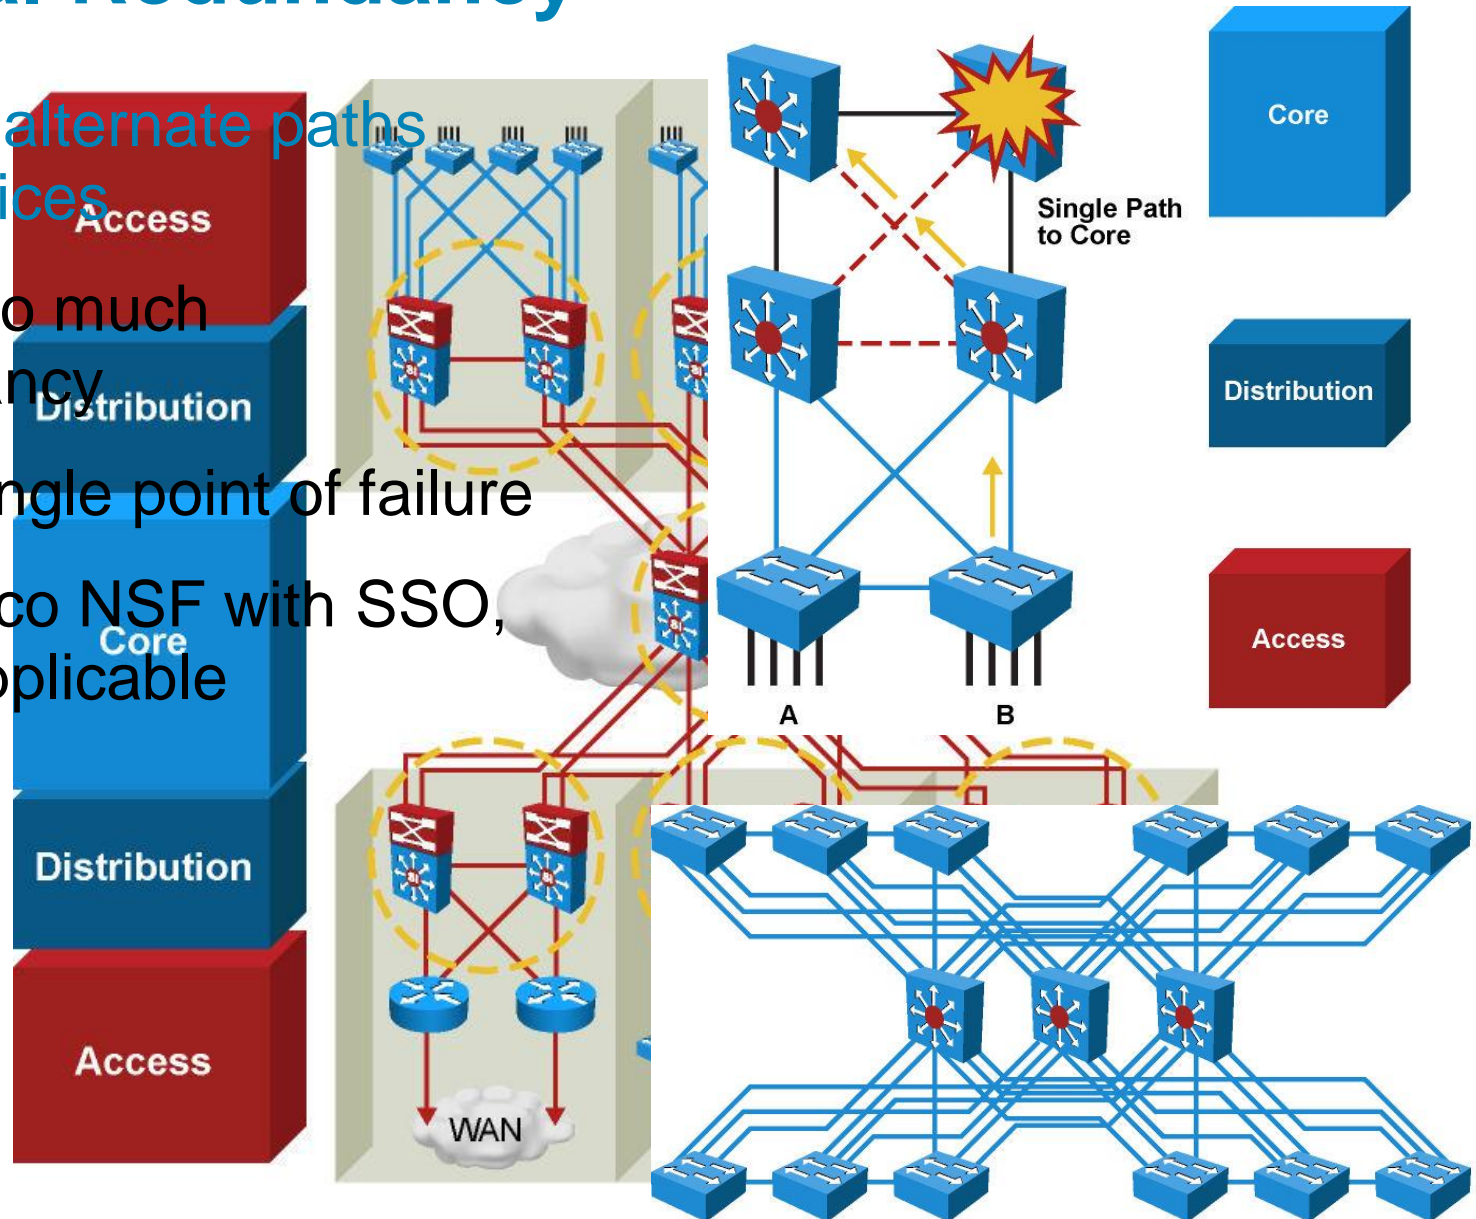
- **Motivation**
- **Hot Standby Router Protocol**
- **Virtual Router Redundancy Protocol**
- **Gateway Load Balancing Protocol**
- **Server Load Balancing**
- **Network Management**
 - IP SLAs
 - Syslog
 - SNMP

Resiliency for High Availability


- High availability is implemented with the following components
 - **Network-level resiliency**
 - Redundant links
 - Redundant devices
 - Power redundancy
 - Fast convergence
 - **System-level resiliency**
 - Integrated hardware resiliency
 - Redundant power supply
 - Stackable switches
 - **Management and monitoring**
 - Detection of failure

Optimal Redundancy

- Provide alternate paths and devices
- Avoid too much redundancy
- Avoid single point of failure
- Use Cisco NSF with SSO, when applicable

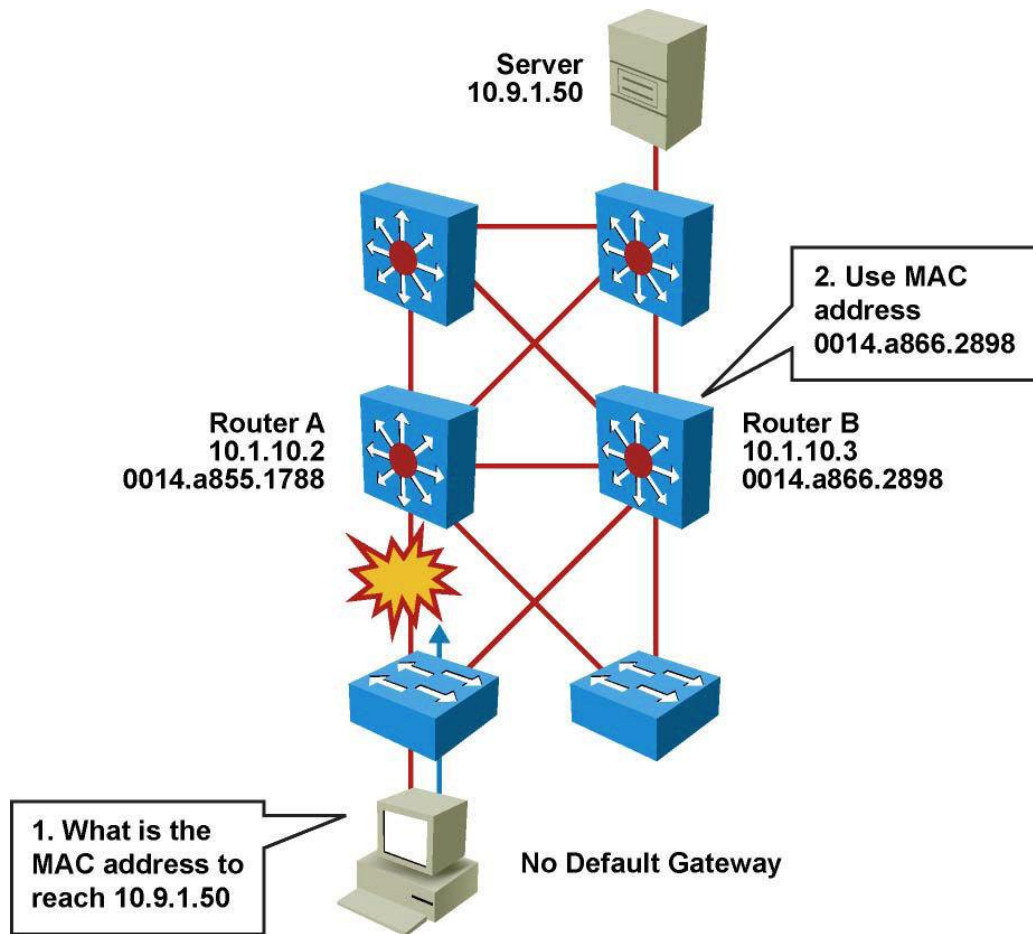


L3 Redundancy



- 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

Redundancy Features in Cisco Boxes



Redundancy Features

- Redundancy of Supervise Engines
 - Route Processor Redundancy
 - Route Processor Redundancy+
 - Stateful SwitchOver
 - Non-Stop Forwarding with SSO
- Available ONLY on Catalyst 4500/6500



Route Processor Redundancy (RPR)

- With **RPR**, any of the following events triggers a switchover from the active to the standby Supervisor Engine
 - Route Processor (RP) or Switch Processor (SP) crash on the active Supervisor Engine
 - A manual switchover from the CLI
 - Removal of the active Supervisor Engine
 - Clock synchronization failure between Supervisor Engines
- **RPR+** enhances Supervisor redundancy compared to RPR
 - Reduced switchover time (in the range of 30 seconds to 60 seconds)
 - No reloading of installed modules (Because both the startup configuration and the running configuration stay continually synchronized)
- *RPR is not preferred any longer!*

Configuring and Verifying RPR

- To use RPR and change its mode RPR+ issue following:

```
Router(config)# redundancy
Router(config-red)# mode rpr-plus
```

- Type following command to verify RPR status:

```
Switch# show redundancy states
      my state = 13 -ACTIVE
      peer state = 1 -DISABLED
      Mode = Simplex
      Unit = Primary
      Unit ID = 1
Redundancy Mode (Operational) = Route Processor Redundancy Plus
Redundancy Mode (Configured) = Route Processor Redundancy Plus
      Split Mode = Disabled
      Manual Swact = Disabled Reason: Simplex mode
      Communications = Down Reason: Simplex mode
```

Stateful Switchover (SSO)

- Provides minimal Layer 2 traffic disruption during Supervisor switchover
- Redundant Supervisor starts up in fully initialized state and synchronizes with startup configuration and running configuration of active Supervisor
- Standby Supervisor in SSO mode keeps in sync with active Supervisor for all changes in hardware and software states for features supported via SSO
- *Preferred solution replacing RPR!*

Features Supported by SSO

- On Cat6500 switchover is between 1 to 3 seconds, on Cat4500 it is subsecond
- Protocols that are maintained synchronized by SSO
 - 802.3x (Flow Control)
 - 802.3ad (LACP) and PAgP
 - 802.1X (Authentication) and Port security
 - 802.3af (Inline power)
 - VTP
 - Dynamic ARP Inspection/DHCP snooping/IP source guard
 - IGMP snooping (versions 1 and 2)
 - DTP (802.1Q and ISL)
 - MST/PVST+/Rapid-PVST
 - PortFast/UplinkFast/BackboneFast /BPDU Guard and filtering
 - Voice VLAN
 - Unicast MAC filtering
 - ACL (VLAN ACLs, Port ACLs, Router ACLs)
 - QOS (DBL)
 - Multicast storm control/broadcast storm control
- *Observe that mostly L2 remains synchronized, what about L3?*

Configuring and Verifying SSO

- To use SSO issue following:

```
Router(config)# redundancy
Router(config-red)# mode sso
```

- IF mode is changed THEN standby is reset
- Same command as for RPR could be used to verify SSO:

```
Switch# show redundancy states
      my state = 13 -ACTIVE
      peer state = 8 -STANDBY HOT
      Mode = Duplex
      Unit = Primary
      Unit ID = 2
Redundancy Mode (Operational) = Stateful Switchover
Redundancy Mode (Configured) = Stateful Switchover
      Split Mode = Disabled
      Manual Swact = Enabled
Communications = Up
```

Non-Stop Forwarding (NSF) with SSO

- Minimizes time that L3 network is by continuing to forward IP packets using CEF entries built from the old active SE
 - Zero or near zero packet loss
 - Supports BGP, EIGRP, OSPF, and IS-IS
 - Prevents route flapping
- *How is it done?*
 - Adjacencies must not be reset when switchover is complete; otherwise, protocol state is not maintained
 - FIB must remain unchanged during switchover
 - Current routes are marked as stale during restart and routes are refreshed after Cisco NSF convergence is complete
 - Switchover must be completed before dead or hold timer expires; otherwise, peers will reset the adjacency and reroute the traffic
 - Cisco NSF-capable routers are also aware about Cisco NSF-capable neighbours
- *The most preferred solution replacing SSO!*

Configuring NSF

- NSF is an additional configuration option when SSO is enabled
- To configure NSF for OSPF, EIGRP, and IS-IS, use the:

```
Router(config-router) # nsf router-level
```

- To configure BGP for NSF support, use the:

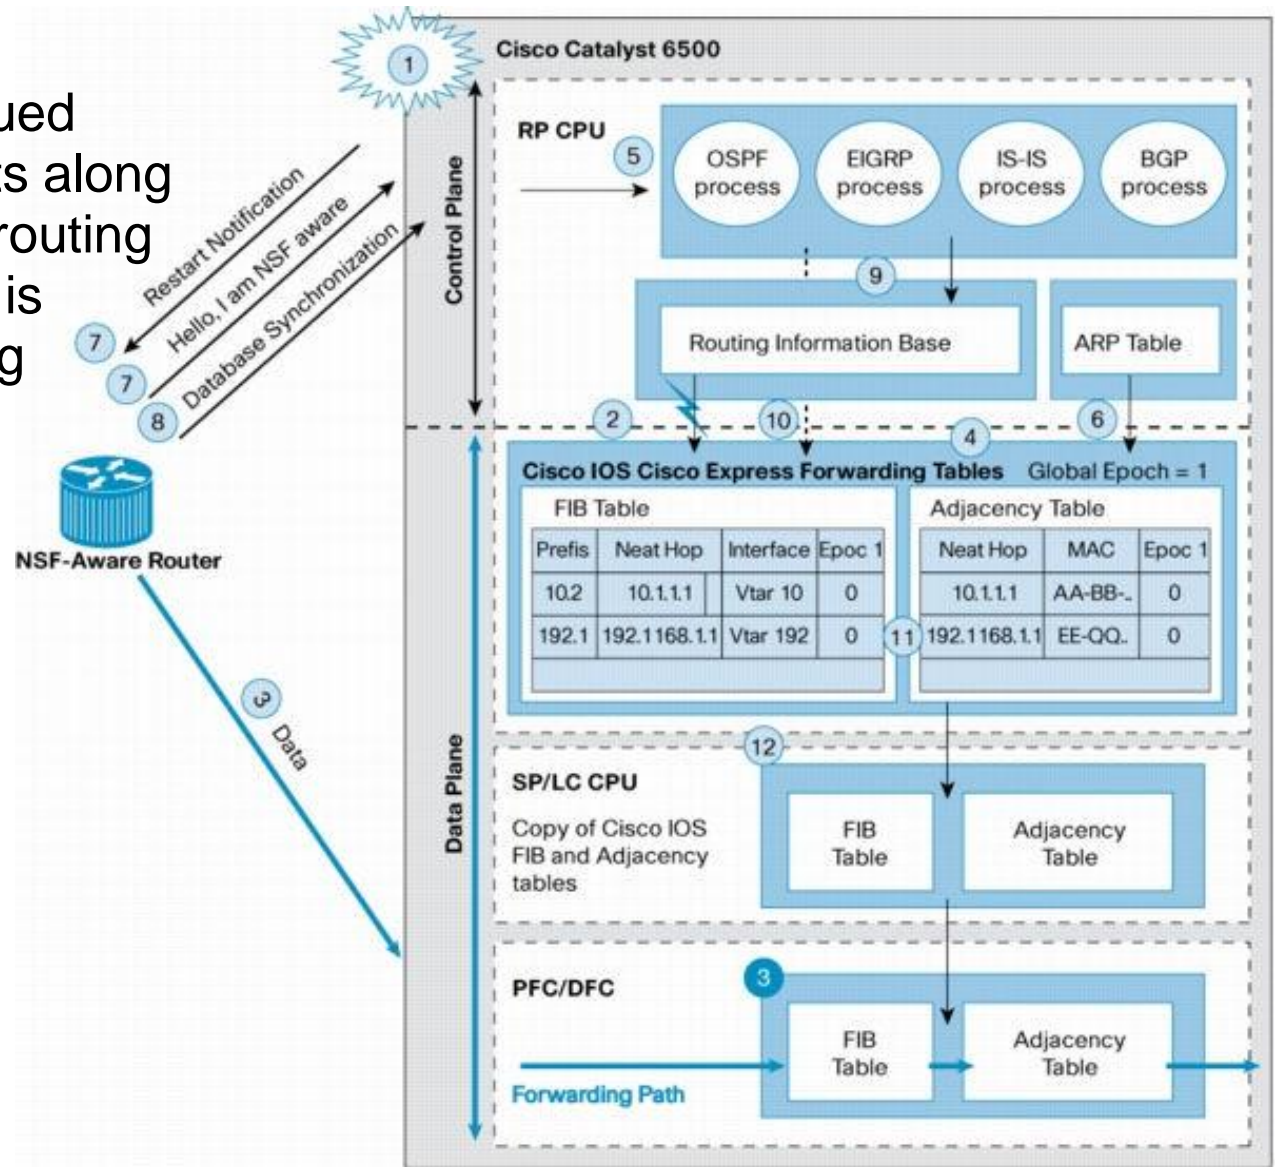
```
Router(config-router) #  
bgp graceful-restart router-level
```

Verifying NSF

```
Switch# show ip bgp neighbors 192.168.200.1
BGP neighbor is 192.168.200.1, remote AS 200, external link
BGP version 4, remote router ID 192.168.200.1
BGP state = Established, up for 00:01:23
Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh:advertised and received(new)
  Address family IPv4 Unicast:advertised and received
  Address family IPv4 Multicast:advertised and received
  Graceful Restart Capability:advertised and received
  Remote Restart timer is 120 seconds
...
Switch# show ip ospf
Routing Process "ospf 200" with ID 192.168.20.1 and Domain ID 0.0.0.1
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x0
Number of opaque AS LSA 0. Checksum Sum 0x0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled, last NSF restart 00:02:36 ago (took 34 secs)
Area BACKBONE(0)
Number of interfaces in this area is 1 (0 loopback)
```

Routing Protocols and NSF

- NSF enables continued forwarding of packets along known routes while routing protocol information is being restored during switchover



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.

HSRP Message Format

Version	Op Code	State	Hellotime
Holdtime	Priority	Group	Reserved
Authentication Data			
Authentication Data			
Virtual IP address			

■ Op Code

- Hello
- Coup
- Resign

■ State

- Initial
- Learn
- Listen
- Speak
- Standby
- Active

■ Hellotime

- Holdtime
- Priority
- Group

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 10 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
```

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
<i>group-number</i>	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
<i>IFACE</i>	Indicates the interface type and number that will be tracked.
<i>Penalty</i>	(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 “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 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

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: V110 Interface UP
*Mar 3 05:38:28.502: HSRP: V110 Starting minimum interface delay (1 secs)
*Mar 3 05:38:29.458: HSRP: V110 Grp 1 Active router is 172.16.10.102
*Mar 3 05:38:29.458: HSRP: V110 Nbr 172.16.10.102 is no longer passive
*Mar 3 05:38:29.458: HSRP: V110 Nbr 172.16.10.102 active for group 1
*Mar 3 05:38:29.500: HSRP: V110 Interface min delay expired
*Mar 3 05:38:29.500: HSRP: V110 Grp 1 Init: a/HSRP enabled
*Mar 3 05:38:29.500: HSRP: V110 Grp 1 Init -> Listen
*Mar 3 05:38:29.500: HSRP: V110 Grp 1 Redundancy "hsrp-V110-1" state Init -> Backup
*Mar 3 05:38:29.500: HSRP: V110 IP Redundancy "hsrp-V110-1" 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: V110 Grp 1 Listen: h/Hello rcvd from lower pri Active router
(100/172.16.10.102)
*Mar 3 05:38:32.260: HSRP: V110 Grp 1 Active router is local, was 172.16.10.102
*Mar 3 05:38:32.260: HSRP: V110 Nbr 172.16.10.102 no longer active for group 1 (Listen)
*Mar 3 05:38:32.260: HSRP: V110 Nbr 172.16.10.102 Was active or standby - start passive
holddown
*Mar 3 05:38:32.260: HSRP: V110 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: V110 Grp 1 Redundancy "hsrp-V110-1" state Backup -> Active
*Mar 3 05:38:32.260: HSRP: V110 Added 172.16.10.1 to ARP (0000.0c07.ac01)
*Mar 3 05:38:32.268: HSRP: V110 Grp 1 Activating MAC 0000.0c07.ac01
*Mar 3 05:38:32.268: HSRP: V110 Grp 1 Adding 0000.0c07.ac01 to MAC address filter
*Mar 3 05:38:32.268: HSRP: V110 IP Redundancy "hsrp-V110-1" update, Backup -> Active
*Mar 3 05:38:35.254: HSRP: V110 IP Redundancy "hsrp-V110-1" update, Active -> Active
*Mar 3 05:38:42.913: HSRP: V110 Grp 1 Standby router is 172.16.10.102
*Mar 3 05:38:42.913: HSRP: V110 Nbr 172.16.10.102 is no longer passive
*Mar 3 05:38:42.913: HSRP: V110 Nbr 172.16.10.102 standby for group 1
```

Example: Authentication Error

```
Switch# debug standby errors
```

```
*Mar  3 05:40:49.606: HSRP: V11 Grp 1 Auth failed for Hello pkt  
from 10.1.1.102, Text auth failed
```

```
*Mar  3 05:40:52.131: HSRP: V11 Grp 1 Auth failed for Hello pkt  
from 10.1.1.102, Text auth failed
```

```
*Mar  3 05:40:54.715: HSRP: V11 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: V111 Init: a/HSRP enabled
*Mar 8 20:34:10.221: SB11: V111 Init -> Listen
*Mar 8 20:34:20.221: SB11: V111 Listen: c/Active timer expired (unknown)
*Mar 8 20:34:20.221: SB11: V111 Listen -> Speak
*Mar 8 20:34:20.221: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:23.101: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:25.961: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:28.905: SB11: V111 Hello out 172.16.11.111 Speak pri 100 ip 172.16.11.115
*Mar 8 20:34:30.221: SB11: V111 Speak: d/Standby timer expired (unknown)
*Mar 8 20:34:30.221: SB11: V111 Standby router is local
*Mar 8 20:34:30.221: SB11: V111 Speak -> Standby
*Mar 8 20:34:30.221: SB11: V111 Hello out 172.16.11.111 Standby pri 100 ip 172.16.11.115
*Mar 8 20:34:30.221: SB11: V111 Standby: c/Active timer expired (unknown)
*Mar 8 20:34:30.221: SB11: V111 Active router is local
*Mar 8 20:34:30.221: SB11: V111 Standby router is unknown, was local
*Mar 8 20:34:30.221: SB11: V111 Standby -> Active
*Mar 8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Standby -> Active
*Mar 8 20:34:30.221: SB11: V111 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: V111 Interface up
*Mar 1 00:16:43.099: SB11: V111 Init: a/HSRP enabled
*Mar 1 00:16:43.099: SB11: V111 Init -> Listen
*Mar 1 00:16:43.295: SB11: V111 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar 1 00:16:43.295: SB11: V111 Active router is 172.16.11.112
*Mar 1 00:16:43.295: SB11: V111 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar 1 00:16:43.295: SB11: V111 Active router is local, was 172.16.11.112
*Mar 1 00:16:43.295: SB11: V111 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: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:43.303: SB11: V111 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: V111 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:46.207: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:49.095: SB11: V111 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:49.195: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:52.079: SB11: V111 Hello in 172.16.11.112 Speak pri 50 ip 172.16.11.115
*Mar 1 00:16:52.147: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: V111 Standby router is 172.16.11.112
*Mar 1 00:16:55.083: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:56.231: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:16:58.023: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:16:59.223: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:17:00.983: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:17:02.211: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:17:03.847: SB11: V111 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
 - Protocol number 112
 - Address [224.0.0.18](#)
- 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: 0000.5e00.01XX
- 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**
 - By default 1 second
- Instead of Holdtime is **Master Down Interval**
 - It IS NOT carried in Advertisement messages
 - It is computed dynamically using following formula:

$$3 \times advertisement_time + \frac{\overbrace{256 - priority}^{\text{skew time}}}{256}$$

- 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 skew 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.



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):

```
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
16 groups max	255 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 for hello packets	Uses 224.0.0.18 for hello packets
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
- One virtual IP per GLBP group (max. 1024 groups)
 - At most 4 virtual MAC addresses per GLBP group

Device Roles

▪ Active virtual gateway (AVG)

- Router with highest priority (highest IP address)
- One AVG in a GLBP group
- Assigns a virtual MAC address (0007.b400.XX0[1-4]) to each member of the GLBP group
- Group controller: responds with virtual MAC addresses in ARP Reply

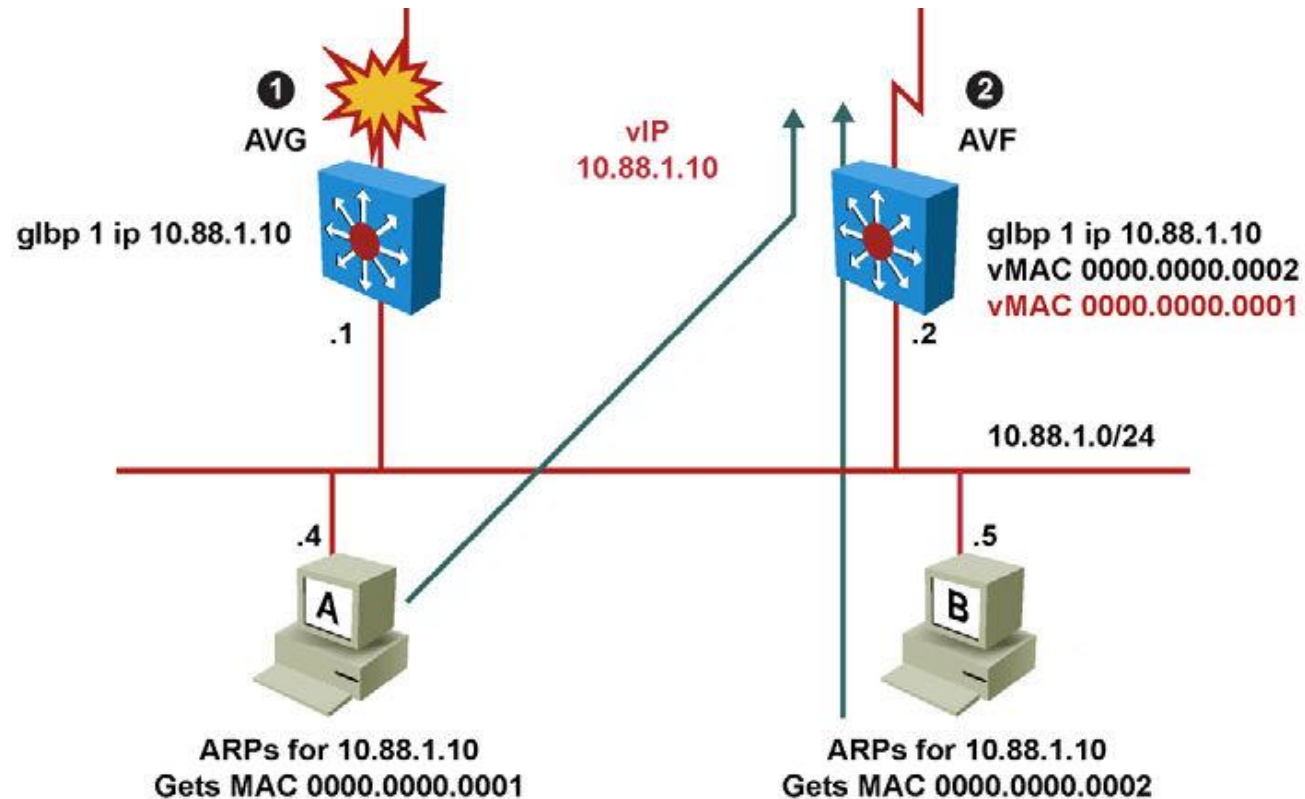
▪ Active virtual forwarder (AVF)

- At the most 4 per GLBP group
- AVF assumes responsibility for forwarding packets that are sent to the virtual MAC address assigned by the AVG
- AVG is also AVF
- AVG/AVF exchange Hello messages every 3 sec to the multicast address 224.0.0.102, UDP port 3222

▪ Backup AVG/AVF

- Other routers in a group provide backup for the AVG/AVF if they becomes unavailable

Operation



Basic Configuration

- Enable GLBP on an interface:

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

- Set a GLBP priority for this router for this GLBP group. The highest value wins election as active router. The default is 100. If routers have the same GLBP priority, the gateway with the highest real IP address becomes the AVG:

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

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

```
Switch(config-if) #  
glbp group-number timers advertise [msec] hello holdtime
```

Load Balancing Mechanism

- GLBP supports these operational modes for load balancing
 - **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
 - **Round-robin load-balancing algorithm**
 - Default one
 - Round-robin fashion for the ARP-replay
- Load balancing algorithm could be configured per group:

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

Other Configurations


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

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

- 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-track) # exit  
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]  
Switch(config-if) # end
```

Simple Example Explaining Weight



Troubleshooting

```
show glbp
```

```
show glbp all
```

```
show glbp GROUP_NUM
```

```
debug glbp all
```

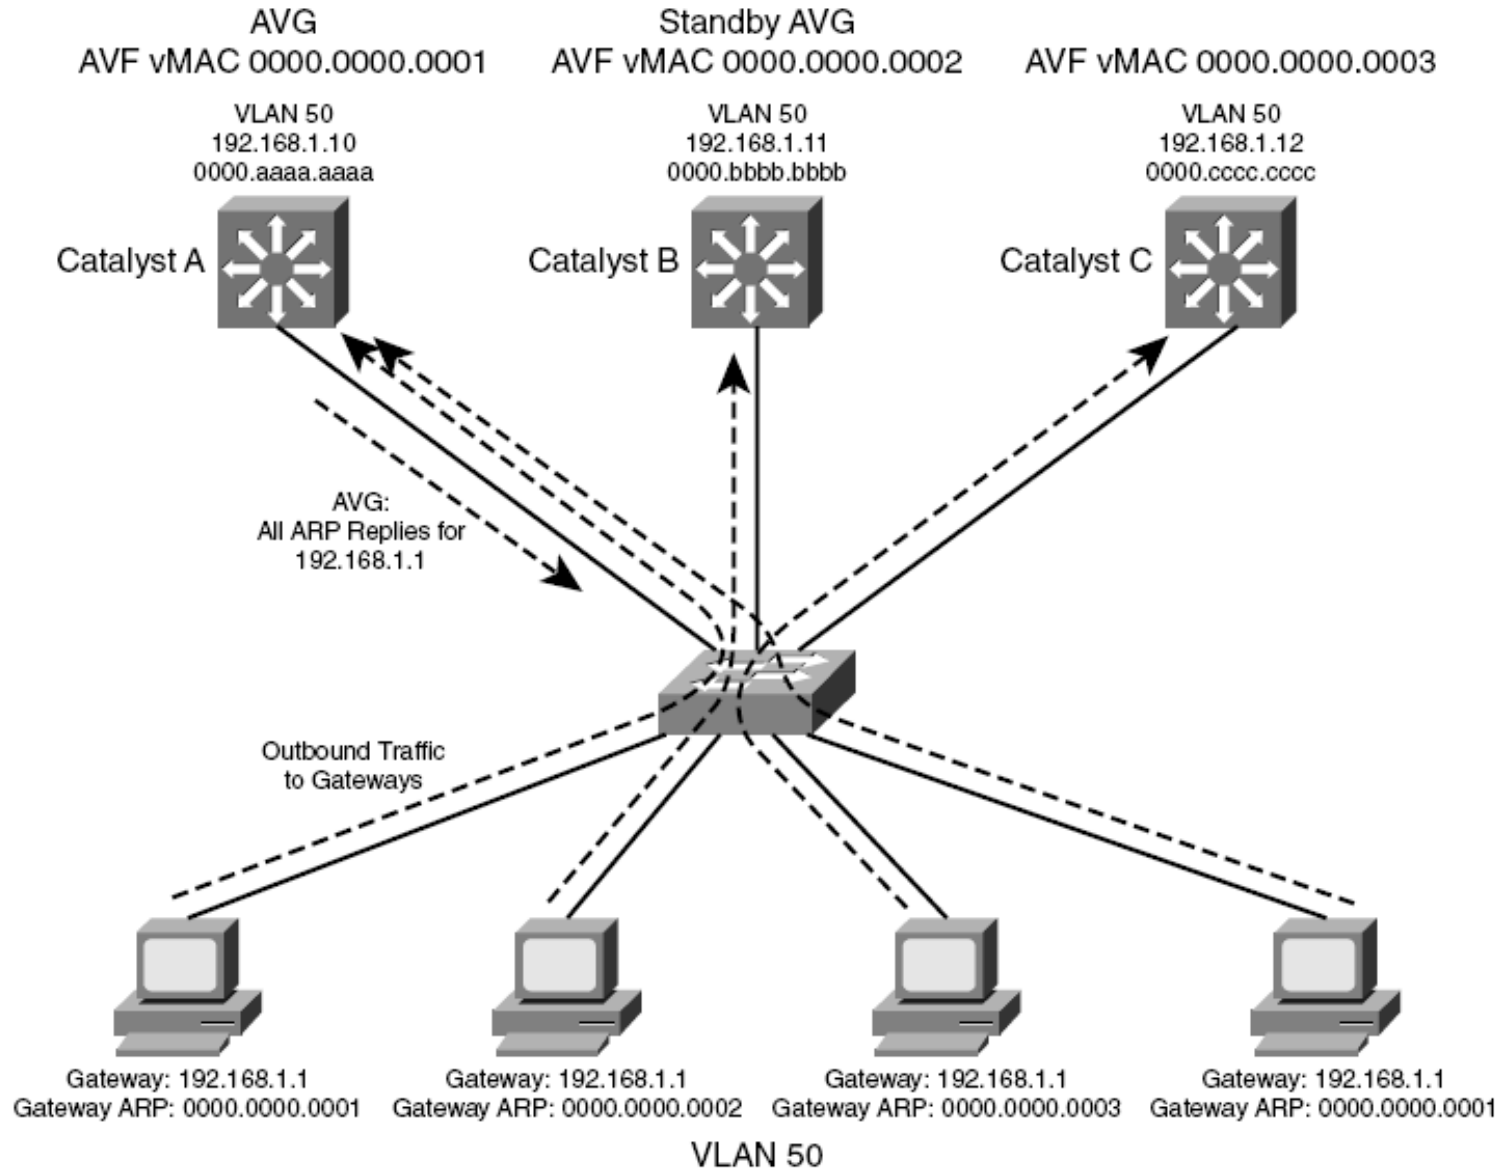
```
debug glbp error
```

```
debug glbp events
```

```
debug glbp packets
```

```
debug glbp state
```

Example



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

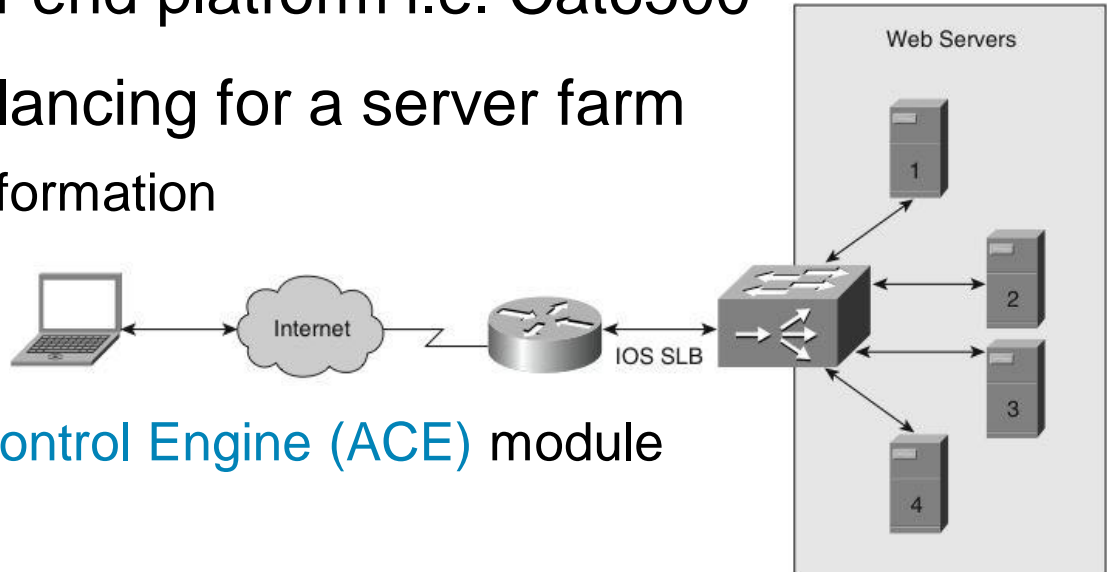
HSRP	GLBP
Cisco Proprietary, 1994	Cisco Proprietary, 2005
16 groups max	1024 groups max
1 active, 1 standby, several candidates	1 AVG, several AVF, AVG load balances traffic among AVF and AVGs
Virtual IP is different from Active and Standby real IP addresses	Virtual IP is different from AVG and AVF real IP addresses
1 Virtual MAC address for each group	1 Virtual MAC address per AVF/AVG in each group
Uses 224.0.0.2 for hello packets	Uses 224.0.0.102 for hello packets
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
Default timers: hello 3 s, holdtime 10 s	Default timers: hello 3 s, holdtime 10 s
Authentication supported	Authentication supported

IOS Server Load Balancing



Server Load Balancing


- Available only on high-end platform i.e. Cat6500
- SLB provides load balancing for a server farm
 - According to L4 - L7 information
 - SW
 - HW
 - Cisco Application Control Engine (ACE) module



- Advantages
 - Reducing server load
 - Increased security – real IP address is not visible
 - Reducing downtime (switch detects down servers)

Virtual Server and Server Farm

- Cisco IOS SLB enables users to represent a group of network servers (a server farm in a data center) as a single server instance so called **virtual server**
 - Balance the traffic and limit it to individual servers
 - Any request to virtual server is served by **real servers**



Cisco IOS SLB modes

▪ Dispatched mode

- Each of the real servers is configured with the virtual server address as a loopback address or secondary IP address
- Packets are redirected to the real servers at the MAC layer
 - Packet targeted to the virtual IP address is encapsulated into the frame with MAC address corresponding to the real server IP address
- Servers must be in same network (Layer2 adjacent)

▪ Directed mode

- Each of the real servers has own real IP address
- Server does not know virtual IP address of a server farm
- Packets are translated using NAT

Configuring the Server Farm with Real Servers

1) Define the server farm:

```
Switch(config)# ip slb serverfarm SERVERFARM-NAME
```

2) Associate the real server with the server farm:


```
Switch(config-slb-sfarm)# real A.B.C.D
```

3) Enable the real server in a server farm:

```
Switch(config-slb-real)# inservice
```


Example: Server Farm


- Two server farms in a data center, PUBLIC and RESTRICTED
- PUBLIC: three real servers: 10.1.1.1, 10.1.1.2 a 10.1.1.3
- RESTRICTED: two real servers: 10.1.1.20 a 10.1.1.21



```
Switch(config)# ip slb serverfarm PUBLIC
Switch(config-slb-sfarm)# nat server ! Directed Mode
Switch(config-slb-sfarm)# real 10.1.1.1
Switch(config-slb-real)# inservice
Switch(config-slb-real)# real 10.1.1.2
Switch(config-slb-real)# inservice
Switch(config-slb-real)# real 10.1.1.3
Switch(config-slb-real)# inservice
!
Switch(config)# ip slb serverfarm RESTRICTED
Switch(config-slb-sfarm)# nat server ! Directed Mode
Switch(config-slb-sfarm)# real 10.1.1.20
Switch(config-slb-real)# inservice
Switch(config-slb-real)# real 10.1.1.21
Switch(config-slb-real)# inservice
```

SLB Verification

- Displaying the status of the server farms
 - Associated servers
 - State of real servers
 - Load balancing mode



```
Switch# show ip slb real
real
```

real	farm name	weight	state	cons
10.1.1.1	PUBLIC	8	OPERATIONAL	0
10.1.1.2	PUBLIC	8	OPERATIONAL	0
10.1.1.3	PUBLIC	8	OPERATIONAL	0
10.1.1.20	RESTRICTED	8	OPERATIONAL	0
10.1.1.21	RESTRICTED	8	OPERATIONAL	0

```
Switch# show ip slb serverfarm
```

server farm	predictor	nat	reals	bind id
PUBLIC	ROUNDROBIN	none	3	0
RESTRICTED	ROUNDROBIN	none	2	0

Configuring Virtual Servers

- 1) Define the virtual server:

```
Switch(config)# ip slb vserver vserver-name
```

- 2) Configure the IP address of the virtual server:

```
Switch(config-slb-vserver)# virtual ip-address [network-mask]  
{tcp | udp} [port-number | wsp | wsp-wtp | wsp-wtls | wsp-wtp-wtls]  
[service service-name]
```

- 3) Associate the primary and secondary server farm to the virtual server:

```
Switch(config-slb-vserver)# serverfarm primary-servfarm-name  
[backup backup-serverfarm-name [sticky]]
```


- 4) Enable the virtual server:

```
Switch(config-slb-vserver)# inservice
```

- 5) Specify the clients allowed to access the virtual server:

```
Switch(config-slb-vserver)# client ip-address network-mask
```

Example: Virtual Servers



```
Switch(config)# ip slb vsrver PUBLIC_HTTP
Switch(config-slb-vserver)# virtual 10.1.1.100 tcp www
Switch(config-slb-vserver)# serverfarm PUBLIC
Switch(config-slb-vserver)# inservice
Switch(config)# ip slb vsrver RESTRICTED_HTTP
Switch(config-slb-vserver)# virtual 10.1.1.200 tcp www
Switch(config-slb-vserver)# client 10.4.4.0 255.255.255.0
Switch(config-slb-vserver)# serverfarm RESTRICTED
Switch(config-slb-vserver)# inservice
```

Virtual Server Verification

```
Switch# show ip slb vserver
```

slb vserver	prot	virtual	state	cons
PUBLIC_HTTP	TCP	10.1.1.100:80	OPERATIONAL	0
RESTRICTED_HTTP	TCP	10.1.1.200:80	OPERATIONAL	0

! Check the connections

```
Switch# show ip slb connections
```

vserver	prot	client	real	state	nat
RESTRICTED_HTTP	TCP	10.4.4.0:80	10.1.1.20	CLOSING	none

Troubleshooting

- Display detailed info Information for an SLB Client

```
show ip slb connections client
```

- Display the statistics

```
show ip slb stats
```

```
Switch# show ip slb connections client 10.4.4.0 detail
VSTEST_UDP, client = 10.4.4.0:80
state = CLOSING, real = 10.1.1.20, nat = none
v_ip = 10.1.1.200:80, TCP, service = NONE
client_syms = 0, sticky = FALSE, flows attached = 0
```

```
Switch# show ip slb stats
Pkts via normal switching: 0
Pkts via special switching: 6
Connections Created: 1
Connections Established: 1
Connections Destroyed: 0
Connections Reassigned: 0
Zombie Count: 0
Connections Reused: 0
```

Bidirectional Forward Detection




Bidirectional Forwarding Detection

- [RFC 5880](#)
- **Bidirectional Forwarding Detection (BFD)** provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers
- Once a BFD session has been established and timer negotiations are complete, BFD peers send BFD control packets that act in the same manner as an IGP hello protocol to detect liveness
- BFDv0 and BFDv1 do exist, both supported on Cisco boxes
- Prerequisites
 - CEF and IP routing enabled on all BFD neighbors
 - Each routing protocol MUST be configured to benefit from BFD
- [“Bidirectional Forwarding Detection”, Cisco IOS Release 12.2SR](#)

Features

- BFD detects a failure, but the IGP/BGP/FHRP must take action to bypass a failed peer
- BFD can provide **failure detection in less than one second**
 - Reducing the IGP/BGP/FHRP timers can result in minimum detection timer of one to two seconds
- BFD can be used as a generic and consistent failure detection mechanism
- BFD can be **less CPU-intensive**
 - Some parts of BFD can be distributed to the data plane
 - Reduced IGP/BGP/FHRP timers exist wholly at the control plane



Configuration

- On interface issue following command:

```
Router(config-if)# bfd interval send-timer  
min_rx receive-timer multiplier interval-multiplier
```

- **interval**: period between two consecutive BFD control messages
- **min_rx**: minimum interval between packets accepted from BFD peers
- **multiplier**: specifies the minimum number of consecutive packets that can be missed before a BFD session is declared down and neighbor dead (default is 3)

Supported Protocols

```
(conf-router) # bfd all-interfaces
```

■ IGP

■ EIGRP

```
(conf-router) # bfd interface
```

■ OSPF

```
(conf-if) # ip ospf bfd [disable]
```

■ IS-IS

```
(conf-if) # isis bfd [disable]
```

■ EGP

■ BGP

```
(conf-router) # neighbor ip-address fall-over bfd
```

■ FHRP

■ HSRP

```
(conf-if) # standby bfd
```

■ VRRP

```
(conf-if) # vrrp bfd
```

■ PIM

Verifying

show ip bfd neighbors [detail]

R1# **show bfd neighbor**

OurAddr	NeighAddr	LD/RD	RH/RS	Holddown(mult)	State	Int	
10.1.3.1	10.1.3.3	1/2	Up	0	(3)	Up	Fa0/1

Example

```
R1# show ip ospf neighbor
```

NeighborID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	FULL/DR	00:00:37	10.1.2.2	FastEthernet0/0
3.3.3.3	1	FULL/BDR	00:00:37	10.1.3.3	FastEthernet0/1

```
R1(config)# int fa 0/0
```

```
R1(config-if)# sh
```

```
19:52:13.115: %LINK-5-CHANGED: Interface FastEthernet0/0, changed state to  
administratively down
```

```
R2#
```

```
19:52:42.643: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on FastEthernet0/0 from  
FULL to DOWN, Neighbor Down: Dead timer expired
```

```
...
```

```
R1(config)#int fa 0/1
```

```
R1(config-if)#shut
```

```
20:04:10.204: %OSPF-5-ADJCHG: Process 1, Nbr on FastEthernet0/1 from FULL to  
DOWN, Neighbor Down: Interface down or detached
```

```
20:04:12.202: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to  
administratively down
```

```
R3#
```


```
20:04:10.511: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on FastEthernet0/1 from  
FULL to DOWN, Neighbor Down: BFD node down
```

Network Management



Network Management Overview

- Ability to verify the network is working well and behaving in the planned manner
- Ability to characterize the performance of the network
- Ability to understand how much traffic is flowing and where it is flowing in the network
- Ability to troubleshoot the network

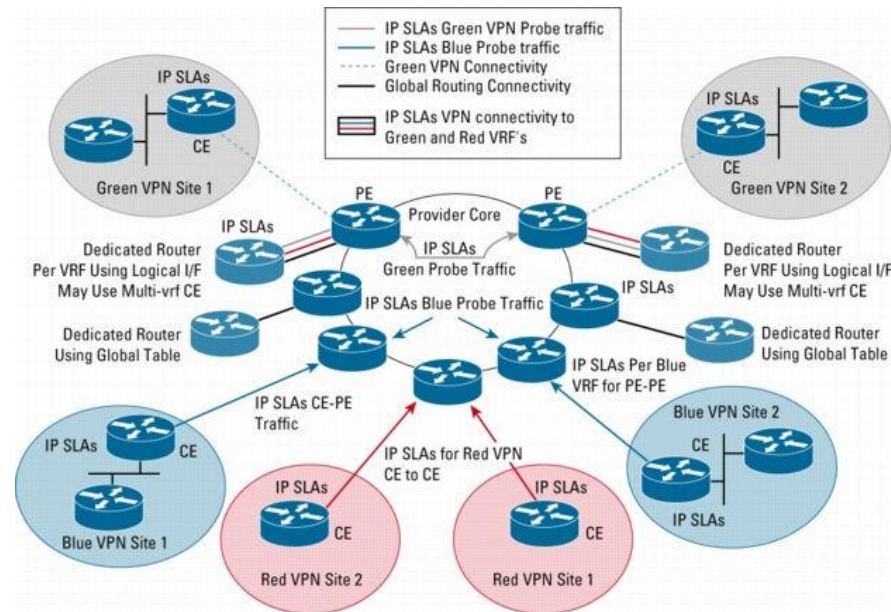


IP Service Level Agreements



IP Service Level Agreements

- Cisco IOS **IP Service Level Agreements (IP SLAs)** use active traffic monitoring
- Cisco IOS IP SLAs tests send simulated data and measure performance between network locations
- Administrator can set conditions needed to satisfy the test



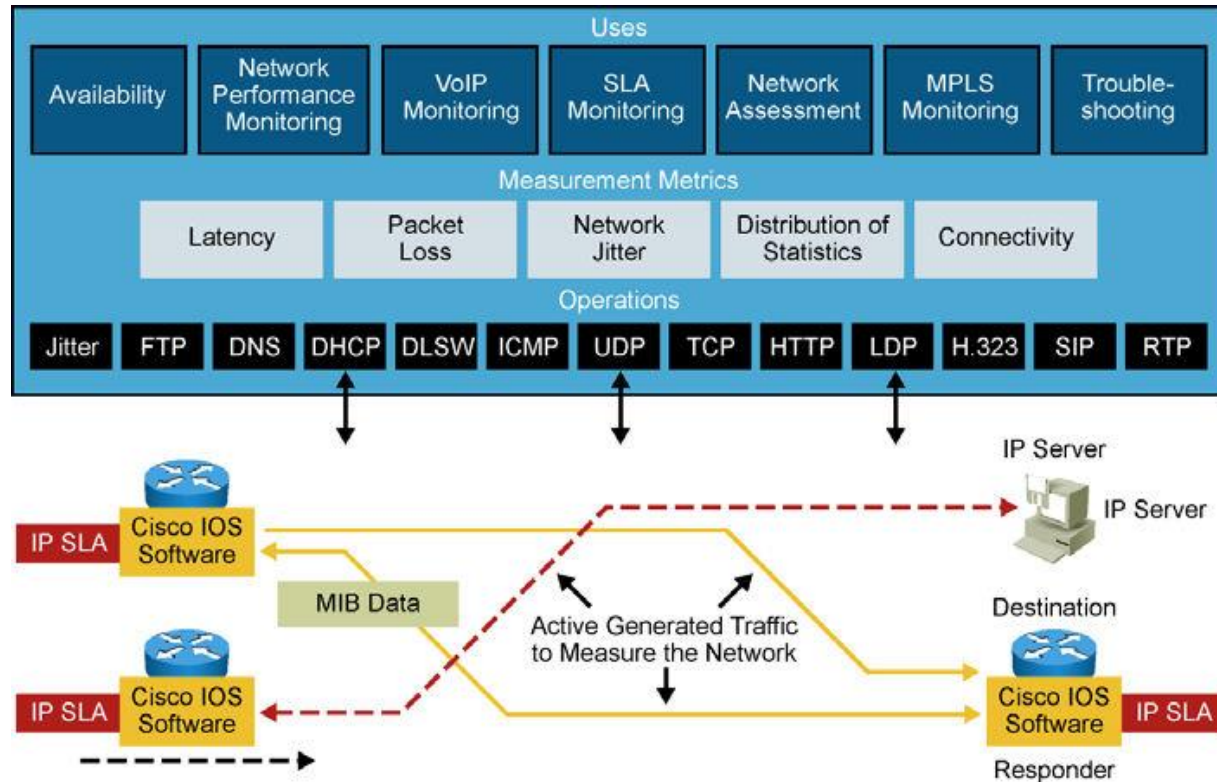
Cisco IOS IP SLAs

- Possible parameters:
 - Network resource availability
 - Response time
 - One-way latency
 - Jitter
 - Packet lost
 - Voice quality scoring
 - Application performance

IP SLA Sources, Responders and Operation

- **IP SLA source** sends testing data to destination
 - All test are configured on SLA source
 - SLA source use **control protocol** to communication with responder before the test starts (agreement on TCP/UDP ports, test type etc.)
 - Source and responder have to use time synchronization (NTP)
- **IP SLA responder** allows to anticipate and respond to IP SLAs request packets
 - To increase security on IP SLA measurements control messages, responder can utilize MD5 authentication
- **IP SLA operation** is measurement that includes protocol, frequency, traps and thresholds
 - IP SLA operations are defined by capabilities of target devices

IP SLA Operations



- IP SLA against device **without SLA responder** (web server or IP station)
 - *E.g.* ping test
- IP SLA against device **with SLA responder** running
 - More powerful test or more accurate results

Configuration Steps

- 1) Define SLA operation (probe)
- 2) Scheduling IP SLA operation
- 3) Define at least one tracking object and define action associated with the tracking object
- 4) IF responder is Cisco device THEN enable wide range of IP SLA capabilities

```
Switch(config)# ip sla {monitor} responder
```

■ Note:

- Starting from IOS 12.4(4)T, 12.2(33)SB and 12.2(33)SXI , command **ip sla monitor** is replaced by command **ip sla**

Step 1 – Define SLA Operation ①

1) Defining SLA operation

```
Router(config)# ip sla operation-number
```

- Parameter *operation-number* is ID of operation

```
R1(config)# ip sla 1  
R1(config-ip-sla)# ?  
IP SLAs entry configuration commands:  
  dhcp          DHCP Operation  
  dns           DNS Query Operation  
  exit          Exit Operation Configuration  
  icmp-echo     ICMP Echo Operation  
  icmp-jitter   ICMP Jitter Operation  
  ...  
R1(config-ip-sla)#
```

Step 1 – Define SLA Operation ②

- PING probe against non-responder node:

```
Router(config-ip-sla) #  
icmp-echo {destination-ip-address | destination-hostname}  
[source-ip {ip-address | hostname} | source-interface IFACE]
```

Parameter	Description
<i>destination-ip-address</i> <i>destination-hostname</i>	Destination IPv4/IPv6 address
source-ip { <i>ip-address</i> <i>hostname</i> }	(Optional) Sets source IPv4/IPv6 address
source-interface <i>interface-name</i>	(Optional) Sets the interface IP address as a source

- Note:

- Starting from IOS 12.4(4)T, 12.2(33)SB and 12.2(33)SXI command **type echo protocol ipIcmpEcho** is replaced by command **icmp-echo**

The icmp-echo Command

```
R1(config-ip-sla)# icmp-echo 209.165.201.30
R1(config-ip-sla-echo)# ?
```

IP SLAs echo Configuration Commands:

default	Set a command to its defaults
exit	Exit operation configuration
frequency	Frequency of an operation
history	History and Distribution Data
no	Negate a command or set its defaults
owner	Owner of Entry
request-data-size	Request data size
tag	User defined tag
threshold	Operation threshold in milliseconds
timeout	Timeout of an operation
tos	Type Of Service
verify-data	Verify data
vrf	Configure IP SLAs for a VPN Routing/Forwarding in-stance

```
R1(config-ip-sla-echo)#
```

■ Bare minimum:

```
Router(config-ip-sla-echo)#
```

```
frequency seconds
```

```
timeout milliseconds
```


Step 2 – Scheduling IP SLA Operation

2) IP SLA operation needs to be scheduled

```
Router(config)#  
  ip sla schedule operation-number [life {forever | seconds}]  
  [start-time {hh:mm[:ss] [month day | day month] | pending |  
  now | after hh:mm:ss}] [ageout seconds] [recurring]]
```

■ Note:

- Starting from IOS 12.4(4)T, 12.2(33)SB and 12.2(33)SXI the command **ip sla monitor schedule** is replaced by **ip sla schedule**

Step 3 – Creating Tracking Object

3) Creating tracking object:

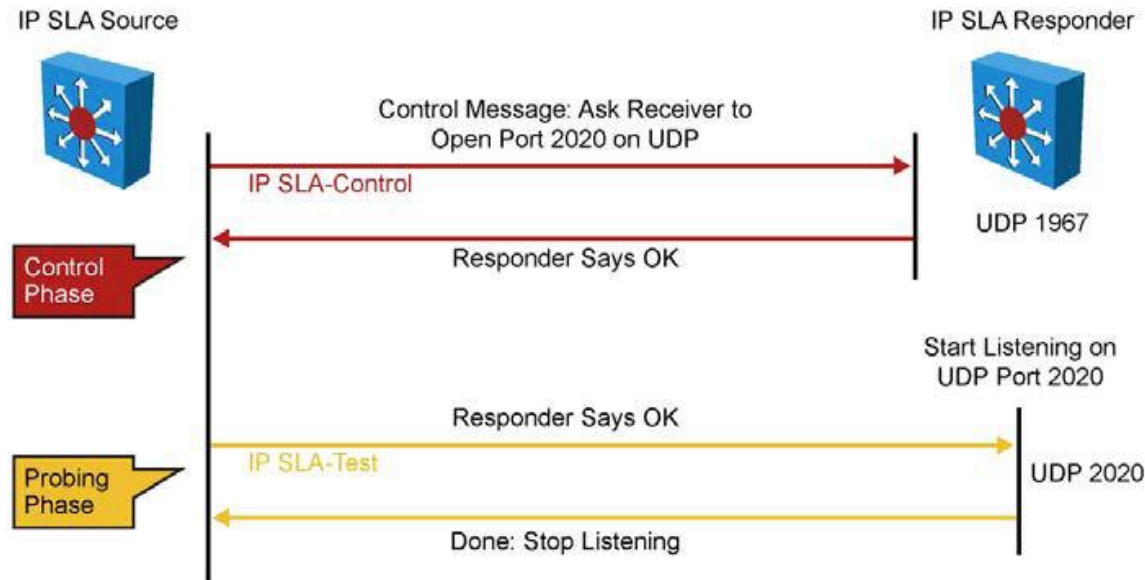
```
Router(config) #  
track object-number ip sla operation-number {state |  
reachability}
```

Parameter	Description
<i>object-number</i>	Tracking object number from 1 to 500
<i>operation-number</i>	Number used for the identification of the IP SLAs operation you are tracking
state	Tracks the operation return code
reachability	Tracks the reachability

■ Note:


- Starting from IOS 12.4(20)T, 12.2(33)SXI1 and 12.2(33)SRE the command **track rtr** is replaced by **track ip sla**

IP SLA Operation with Responder




- 1) IP SLA source sends a control message with the configured IP SLA operation information (protocol, port number, and duration) to IP SLA **control port UDP 1967** on responder
- 2) IF the responder processes the control message THEN it sends an OK message to the source router and listens on the port specified in the control message for a specified duration
 - IF the responder cannot process the control message THEN it returns an error
- 3) IF the return code of control message is OK THEN the IP SLA operation moves to the probing phase, where it sends one or more test packets to the responder for response time computations.
 - The return code is available in IP SLA statistics
- 4) The responder accepts the test packets and responds
 - Based on the type of operation, the responder might add an "in" timestamp and an "out" timestamp in the response packet payload to account for CPU time spent in measuring unidirectional packet loss, latency, and jitter to a Cisco device.

IP SLA Responder Timestamps




- IP SLA responder timestamps are used in round-trip calculations
- IP SLA source sends test packet at time T1
- IP SLA responder includes receipt time (T2) and transmitted time (T3)

Example: IS SLAs on Multi-homed Connection



```
R1(config)# ip sla 11
R1(config-ip-sla)# icmp-echo 10.1.3.3
R1(config-ip-sla-echo)# frequency 10
R1(config-ip-sla-echo)# exit ! 2x
R1(config)# ip sla 22
R1(config-ip-sla)# icmp-echo 172.16.3.3
R1(config-ip-sla-echo)# frequency 10
R1(config-ip-sla-echo)# exit ! 2x
R1(config)# track 1 ip sla 11 reachability
R1(config-track)# delay down 10 up 1
R1(config-track)# exit
R1(config)# track 2 ip sla 22 reachability
R1(config-track)# delay down 10 up 1
R1(config-track)# exit
R1(config)# ip sla schedule 11 life forever start-time now
R1(config)# ip sla schedule 22 life forever start-time now
R1(config)# ip route 0.0.0.0 0.0.0.0 10.1.1.1 2 track 1
R1(config)# ip route 0.0.0.0 0.0.0.0 172.16.1.1 3 track 2
```

Example: HSRP and IP SLA Tracking



```
sw(config)# ip sla 18
sw(config-sla)# icmp-echo 10.9.9.1
sw(config)# ip sla schedule 18 start-time now life forever
sw(config)# track 90 rtr 18 state
sw(config)# interface vlan10
sw(config-if)# ip address 10.1.1.2 255.255.255.0
sw(config-if)# standby 10 ip 10.1.1.1
sw(config-if)# standby 10 priority 110
sw(config-if)# standby 10 preempt
sw(config-if)# standby 10 track 90 decrement 20
```

Verify IP SLA Operation

- When IP SLA is configured, the test is conducted as per the scheduled configuration - the test might succeed or fail
- IF you do not monitor the test results THEN it might fail silently

```
Switch# show ip sla statistics
Round Trip Time (RTT) for Index 1
Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: 11:11:22.533 eastern Thu Jul 9 2010
Latest operation return code: Timeout
Over thresholds occurred: FALSE
Number of successes: 177
Number of failures: 6
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

Verify IP SLA Configuration


```
Switch# show ip sla configuration
IP SLAs, Infrastructure Engine-II
Entry number: 1
Owner:
Tag:
Type of operation to perform: echo
Target address/Source address: 10.1.3.10/10.1.253.1
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Verify data: No
Vrf Name:
Schedule:
    Operation frequency (seconds): 5
    Next Scheduled Start Time: Start Time already passed
    Group Scheduled : FALSE
    Randomly Scheduled : FALSE
    Life (seconds): Forever
    Entry Ageout (seconds): never
    Recurring (Starting Everyday): FALSE
    Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
<output omitted>
```


Syslog



Syslog

- Networking protocol which provides comprehensive reporting mechanism
- UDP port 514
- Logs are sent in plain text
- Supported almost on all network devices (also on Windows/MAC/Linux servers)



Syslog Severity Level


- The lower number, the more serious the situation

Syslog severity terminology	Syslog severity Cisco terminology
Emergency	Level 0, the most serious
Alert	Level 1
Critical	Level 2
Error	Level 3
Warning	Level 4
Notice	Level 5
Informational	Level 6
Debugging	Level 7, the least serious

Syslog Message Format

- **Facility:** A code consisting of two or more letters that indicates the message origin (hardware device, protocol, system software etc.)
 - Cisco IOS has more than 500 facilities (e.g. IP, OSPF, SYS, IPsec, RSP, IF)
- **Severity:** A single-digit code from 0 to 7 that reflects the severity
- **Mnemonic:** A code that uniquely identifies the error message
- **Message-text:** Text string describing the condition

%FACILITY-SUBFACILITY-SEVERITY-MNEMONIC: Message-text



The diagram illustrates the mapping of the example message to the Syslog format fields. It uses colored lines to connect the fields in the example message to the corresponding fields in the format string above it: a yellow line connects '%SYS' to '%FACILITY-SUBFACILITY', a blue line connects '5' to 'SEVERITY', a red line connects 'CONFIG_I' to 'MNEMONIC', and a purple line connects the rest of the message to 'Message-text'.

```
%SYS-5-CONFIG_I: Configured from console by  
cwr2000 on vty0 (192.168.64.25)
```

Configuring Syslog

- Setup Syslog server and severity level:

```
! Syslog configuration
Switch(config)# logging host 192.0.2.1
Switch(config)# logging trap ?
<0-7>          Logging severity level
alerts         Immediate action needed          (severity=1)
critical       Critical conditions               (severity=2)
debugging      Debugging messages               (severity=7)
emergencies    System is unusable                (severity=0)
errors         Error conditions                  (severity=3)
informational  Informational messages            (severity=6)
notifications  Normal but significant conditions (severity=5)
warnings       Warning conditions                (severity=4)

Switch(config)# logging trap informational
```

- Setup circular local Syslog buffer:

```
! Local logs configuration
! Parameters: local log size and the severity level that has to be logged
Switch(config)# logging buffered 10000 6
```

Verifying

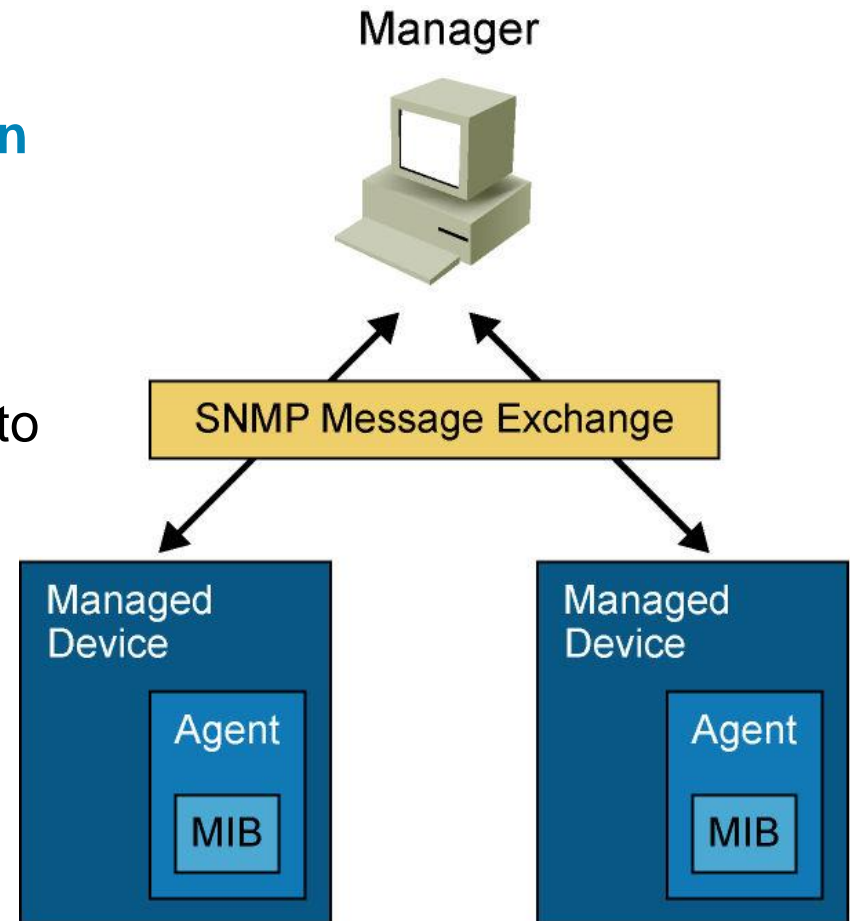
```
Switch# show logging
Syslog logging: enabled (11 messages dropped, 0 messages rate-limited,
                  0 flushes, 0 overruns, xml disabled, filtering disabled)
  Console logging: level debugging, 174 messages logged, xml disabled,
                  filtering disabled
  Monitor logging: level debugging, 0 messages logged, xml disabled,
                  filtering disabled
  Buffer logging: level informational, 3 messages logged, xml disabled,
                  filtering disabled
  Logging Exception size (4096 bytes)
  Count and timestamp logging messages: disabled
No active filter modules.
  Trap logging: level informational, 43 message lines logged
    Logging to 192.0.2.1(global) (udp port 514, audit disabled, link up),
    2 message lines logged, xml disabled, filtering disabled
Log Buffer (10000 bytes):
*Mar  1 01:40:47.395: %SYS-5-CONFIG_I: Configured from console by console
...
```

Simple Network Management Protocol



Simple Network Management Protocol (SNMP)

- SNMP contains three elements:
 - **Network Management Application**
 - **SNMP Agents** – running inside a managed device
 - **MIB database** – describes the information that the agent can use to populate the data
- SNMP modes:
 - **Pull model** – manager periodically polls the SNMP agents
 - **Push model** – agents inform the manager of certain events

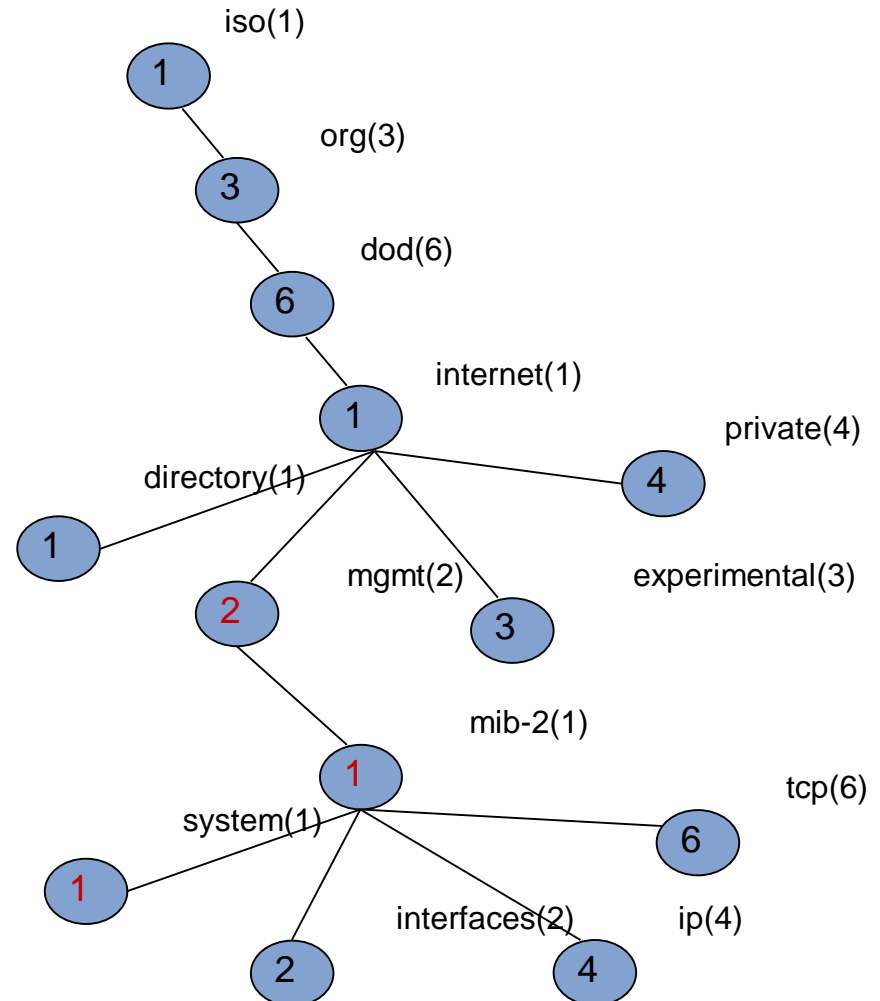


Management Information Base (MIB)

- Agent's objects have own identifiers **OID (Object Identifier)**
 - Hierarchical tree structure
 - Number + name
 - Object's address is a path from the root node


- E.g.:* OID = 1.3.6.1.2.1.1

iso(1)
org(3)
dod(6)
internet(1)
mgmt(2)
mib-2(1)
system (1)



Ports

- UDP Port 161 - SNMP Messages
- UDP Port 162 - SNMP Trap Messages



SNMP Version 1 (SNMPv1)

- [RFC 1157](#)

- Five basic SNMP messages

- **Get Request (Get)**

- Used to request the value of a specific MIB variable

- **Get Next Request (GetNext)**

- Used after the initial Get Request to retrieve the next object instance from a table or a list

- **Set Request (Set)**

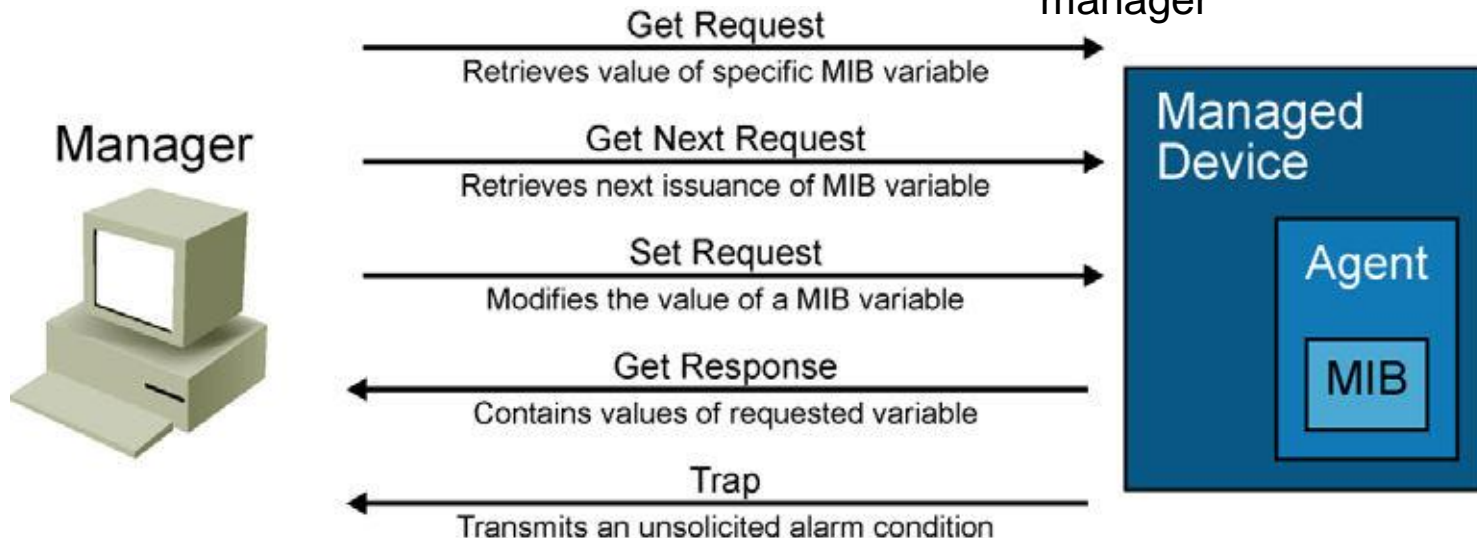
- Used to set a MIB variable on an agent

- **Get Response (Response)**

- Used by an agent to respond to a Get Request or Get Next Request from a manager

- **Trap**

- Transmit an unsolicited alarm to manager



SNMP Version 2 (SNMPv2)

- **Internet standard Management Framework** ([RFC 1441](#))
 - Lack of security – problems with standardization
 - Only experimental implementations, adds 64 bits counters
- Community-based **SNMPv2 (SNMPv2c)**
 - [RFC 1901](#)
 - The most common implementation of SNMPv2
 - SNMPv2c uses community strings for administrative access
 - **READ-ONLY**
 - **READ-WRITE**
 - **TRAP**
- SNMPv2 introduces two new message types:
 - **Get Bulk Request**
 - Reduces repetitive requests
 - Improve performance when you are retrieving large amounts of data
 - **Inform Request**
 - Alert an SNMP manager of specific conditions
 - Message is confirmed by Inform Response

SNMP Version 3

- [RFC 3410](#), [3411](#), [3412](#), [3413](#), [3414](#), [3415](#)
- It adds methods to ensure the secure transmission
- SNMPv3 introduces three levels of security
 - **noAuthNoPriv**
 - No authentication is required
 - No privacy is provided
 - **authNoPriv**
 - Authentication using MD5 or SHA
 - No privacy
 - **authPriv**
 - Authentication using MD5 or SHA
 - Privacy (encryption) using DES, 3DES or AES
- Cisco supports User-based Security Model (Authentication with name and password)

Recommendations

- SNMPv1 and SNMPv2 use community strings in clear text
 - Community strings should be changed at regular intervals
 - IF SNMP is used only to monitor devices
THEN use read-only communities
 - Use ACL to prevent SNMP messages from going beyond the required devices
- SNMPv3 is recommended because it provides authentication and encryption

Configuring SNMPv2 and SNMPv3

- Configure SNMP access lists
- Configure SNMP community strings
- Configure SNMP Trap receiver
- Configure SNMP traps

```
Switch(config)# access-list 1 permit 10.1.1.0 0.0.0.255
Switch(config)# snmp-server community cisco RO 1
Switch(config)# snmp-server community xyz123 RW 1
Switch(config)# snmp-server host 10.1.1.50
Switch(config)# snmp-server enable traps ?
```

```
Switch(config)# snmp-server group MYGROUP v3 priv
Switch(config)# snmp-server user MYUSER MYGROUP v3 auth md5
MYPASS123 priv aes 128 MYKEY123
Switch(config)# snmp-server host 10.9.99.50 priv
Switch(config)# snmp-server enable traps snmp linkdown linkup
coldstart warmstart hsrp vrrp
```

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 NSF with SSO:

www.cisco.com/en/US/partner/docs/switches/lan/catalyst6500/ios/12.2SXF/native/configuration/guide/nfsso.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/sweet.html

- Configuring Server Load Balancing:

www.cisco.com/en/US/partner/docs/ios/ipapp/configuration/guide/ipapp_slb.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: 2015-03-11