



## Configuring SLB Redundancy

---

This chapter describes how to configure server load-balancing redundancy on the Catalyst 4840G SLB switch. For further information about the commands used in this chapter, refer to Appendix A, “Command Reference.”

The Catalyst 4840G SLB switch could represent a point of failure and the servers could lose their connections to the backbone if power fails or if a link from a switch to the distribution-layer switch is disconnected. This section describes redundancy features you can use to reduce that risk.

This chapter includes the following sections:

- Hot Standby Router Protocol (HSRP), page 7-1
- SLB Stateful Backup, page 7-7



**Note**

---

You are at Step 4 in the suggested process for configuring your Catalyst 4840G SLB switch. See the “Switch Configuration Steps” section on page 2-1.

---

## Hot Standby Router Protocol (HSRP)

The Hot Standby Router Protocol (HSRP) provides high network availability by routing IP traffic from hosts on Ethernet networks without relying on the availability of any single Catalyst 4840G SLB switch. This feature is useful for hosts that do not support a router discovery protocol (such as the Intermediate System-to-Intermediate System [IS-IS] Interdomain Routing Protocol [IDRP]) and do not have the functionality to switch to a new load-balancing switch when their selected load-balancing switch reloads or loses power.

## Requirements

Configuring HSRP on the Catalyst 4840G SLB switch requires the following:

- Catalyst 4840G SLB switches are configured to run HSRP between interfaces on the server side.
- You must configure multiple, unique virtual IP addresses to allow multiple, active load-balancing switches.
- Each virtual IP address must be specified on a single SLB switch within a group.

## How SLB HSRP Works

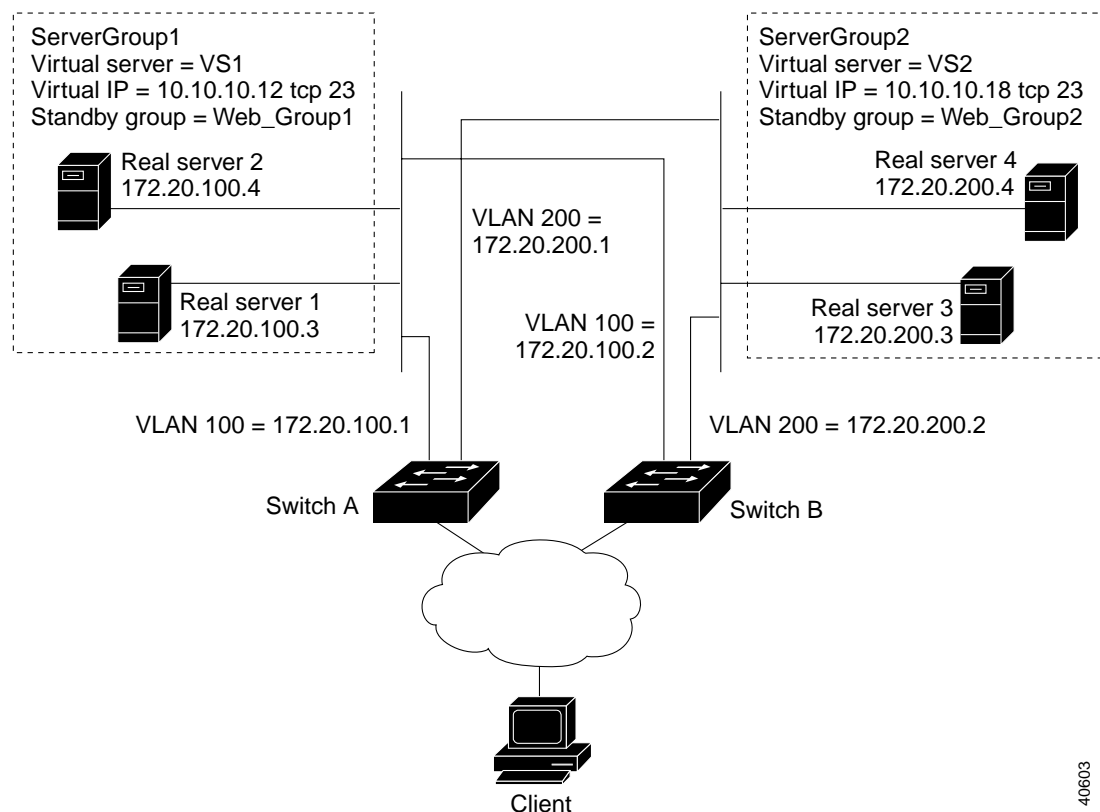
A SLB switch running the HSRP detects a failure by sending and receiving multicast User Datagram Protocol (UDP) hello packets. When the SLB switch running HSRP detects that the designated active SLB switch has failed, the selected backup SLB switch assumes control of the HSRP group MAC and IP addresses. (You can also select a new standby SLB switch at that time.)

The chosen MAC address and IP addresses must be unique and not conflict with any others on the same network segment. The MAC address is selected from a pool of Cisco MAC addresses. Configure the last byte of the MAC address by using the HSRP group number. When the HSRP is running, it selects an active SLB switch and instructs its device layer to listen on an additional (dummy) MAC address.

SLB switching software supports HSRP over 10/100 Ethernet, Gigabit Ethernet, FEC, GEC, and Bridge Group Virtual Interface (BVI) connections.

For example, in the HSRP network shown in Figure 7-1, Device A is the active HSRP SLB switch and is responsible for handling packets to the real servers 1 through 4. If the connection between Device A and the client accessing virtual server IP address 10.10.10.12 tcp 23 or 10.10.10.18 tcp 23 fails, fast converging routing protocols, such as the Enhanced Interior Gateway Routing Protocol (Enhanced IGRP) and Open Shortest Path First (OSPF), can respond within seconds so that Device B is prepared to transfer packets that would have gone through Device A.

**Figure 7-1 HSRP Example Network Topology**



40603

HSRP uses a priority scheme to determine which HSRP-configured SLB switch is to be the default active SLB switch. To configure a SLB switch as active, you assign it a priority that is higher than the priority of all the other HSRP-configured SLB switches. The default priority is 100, so if you configure just one SLB switch to have a higher priority, that switch becomes the default active switch.

HSRP works by the exchange of multicast messages that advertise priority among HSRP-configured SLB switches. When the active switch fails to send a hello message within a configurable period of time, the standby switch with the highest priority becomes the active switch. The transition of packet-forwarding functions between SLB switches is completely transparent to all hosts accessing the network.

HSRP-configured SLB switches exchange the following three types of multicast messages:

- Hello—The hello message conveys to other HSRP switches the switch’s HSRP priority and state information. By default, an HSRP switch sends hello messages every three seconds.
- Coup—When a standby SLB switch assumes the function of the active switch, it sends a coup message.
- Resign—The active SLB switch sends this message when it is about to shut down or when a switch that has a higher priority sends a hello message.

At any time, HSRP-configured SLB switches are in one of the following states:

- Active—The switch is performing packet-transfer functions.
- Standby—The switch is prepared to assume packet-transfer functions if the active router fails.
- Speaking and listening—The switch is sending and receiving hello messages.
- Listening—The switch is receiving hello messages.

## Configuring HSRP

To configure HSRP over VLANs between SLB switches, you must first create its environment. Perform these tasks in the order in which they appear:

- Configure the server farms, real servers, and virtual servers—See the “Required SLB Configuration Tasks” section on page 5-13.



**Note** When you configure the virtual server as `inservice`, you must use the optional **standby** command and configure an HSRP group name. See the “`inservice (server farm virtual server)`” section on page A-31.

- Configure the IP routing protocol—See the “Configuring IP Routing” section on page 8-2.
- Configure the VLAN between the switches—See the “Using VLANs in SLB” section on page 4-6.
- Enable HSRP.

To enable HSRP on an interface, enable the protocol, then customize it for the interface. Perform this task in interface configuration mode:

Command	Purpose
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>ip</b> [ip-address [secondary]]	Enable HSRP.

You can customize an HSRP group attributes using these commands in interface configuration mode:

Command	Purpose
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>priority</b> priority	Set the HSRP priority used to choose the active router.
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>preempt</b>	Specify whether the local SLB switch has priority over the current active SLB switch, or whether the local SLB switch should attempt to take its place as the active SLB switch.
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>timers</b> hellotime holdtime	Configure the time between hello packets and the hold time before other SLB switches declare the active router to be down.
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>track</b> type-number [interface-priority]	Configure the HSRP interface to track other interfaces, so that if one of the other interfaces goes down, the HSRP priority for the device is lowered.
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>name</b> name	Configure the HSRP interface name.
SLB-Switch(config-if)# <b>standby</b> [group-number] <b>authentication</b> string	Select an authentication string to be carried in all HSRP messages.

This example shows how to enable the HSRP standby group 100 IP address, preempt, priority, timers, and how to configure a name and authentication for Device A in Figure 7-1:

```
SLB-Switch(config-if)# standby 100 ip 172.20.100.10
SLB-Switch(config-if)# standby 100 priority 110
SLB-Switch(config-if)# standby 100 preempt
SLB-Switch(config-if)# standby 100 timers 5 15
SLB-Switch(config-if)# standby 100 name Web_group1
SLB-Switch(config-if)# standby 100 authentication Secret
SLB-Switch(config-if)# exit
SLB-Switch#
```

## Example HSRP Configuration

Figure 7-1 shows the topology of an IP network with two Catalyst 4840G SLB switches configured for HSRP.

All hosts accessing the network use the IP address of the virtual servers (in this case, 10.10.10.12 or 10.10.10.18).



### Note

The configurations shown use the RIP routing protocol; HSRP can be used with any routing protocol supported by the Cisco IOS software. Some configurations that use HSRP still require a routing protocol to converge when a topology change occurs. The standby SLB switch becomes active, but connectivity does not occur until the protocol converges.

The following is the configuration for Switch A (active):

```
hostname Switch A
!
ip slb serverfarm ServerGroup1
  real 172.20.100.3
  inservice
  real 172.20.100.4
  inservice
!
ip slb serverfarm ServerGroup2
  real 172.20.200.3
  inservice
  real 172.20.200.4
  inservice
!
ip slb vserver VS1
  virtual 10.10.10.12 tcp 23
  serverfarm ServerGroup1
  in-service standby Web_Group1
!
ip slb vserver VS2
  virtual 10.10.10.18 tcp 23
  serverfarm ServerGroup2
  in-service standby Web_Group2
!
ip routing
router rip
network 172.20.0.0
!
interface vlan100
ip address 172.20.100.1 255.255.255.0
standby 100 ip 172.20.100.10
standby 100 priority 110
standby 100 preempt
standby 100 timers 5 15
standby 100 name Web_Group1
standby 100 authentication Secret
!
interface vlan200
ip address 172.20.200.1 255.255.255.0
standby 200 ip 172.20.200.10
standby 200 priority 110
standby 200 preempt
standby 200 timers 5 15
standby 200 name Web_Group2
standby 200 authentication Covert
!
```

The following is the configuration for Switch B (standby):

```
hostname Switch B
!
ip slb serverfarm ServerGroup1
    real 172.20.100.3
    inservice
    real 172.20.100.4
    inservice
!
ip slb serverfarm ServerGroup2
    real 172.20.200.3
    inservice
    real 172.20.200.4
    inservice
!
ip slb vserver VS1
    virtual 10.10.10.12 tcp 23
    serverfarm ServerGroup1
    in-service standby Web_Group1
!
ip slb vserver VS2
    virtual 10.10.10.18 tcp 23
    serverfarm ServerGroup2
    in-service standby Web_Group2
!
ip routing
router rip
network 172.20.0.0
!
interface vlan100
ip address 172.20.100.2 255.255.255.0
standby 100 ip 172.20.100.10
standby 100 preempt
standby 100 timers 5 15
standby 100 name Web_Group1
standby 100 authentication Secret
!
interface vlan200
ip address 172.20.200.2 255.255.255.0
standby 200 ip 172.20.200.10
standby 200 preempt
standby 200 timers 5 15
standby 200 name Web_Group2
standby 200 authentication Covert
```

The **standby ip** interface configuration command enables HSRP and establishes 10.10.10.12 and 10.10.10.18 as the IP addresses of the virtual servers. The configurations of both switches include this command so that both switches share the same virtual IP address. The 100 establishes Hot Standby group 100. (If you do not specify a group number, the default is group 0.) The configuration for at least one switch in the Hot Standby group must specify the IP address of the virtual server; specifying the IP address of the virtual router is optional for other routers in the same group.

The **standby preempt** interface configuration command allows the switch to become the active switch when its priority is higher than all other HSRP-configured switches in this group. The configurations of both switches include this command so that each can be the standby switch for the other switch. The 100 indicates that this command applies to Hot Standby group 100. If you do not use the **standby preempt** command in the configuration for a load-balancing switch, that switch cannot become the active switch.

The **standby priority** interface configuration command sets the switch's HSRP priority to 110, which is higher than the default priority of 100. Only the configuration of Device A includes this command, which makes Device A the default active switch. The 100 indicates that this command applies to Hot Standby group 100.

The **standby authentication** interface configuration command establishes an authentication string whose value is an unencrypted eight-character string that is incorporated in each HSRP multicast message. This command is optional. If you choose to use it, each HSRP-configured switch in the group should use the same string so that each switch can authenticate the source of the HSRP messages that it receives. The 100 indicates that this command applies to Hot Standby group 100.

The **standby timers** interface configuration command sets the interval in seconds between hello messages (called the *hello time*) to five seconds and sets the duration in seconds that a switch waits before it declares the active switch to be down (called the *hold time*) to eight seconds. (The defaults are three and 10 seconds, respectively.) If you decide to modify the default values, you must configure each switch to use the same hello time and hold time. The 100 indicates that this command applies to Hot Standby group 100.

## Verifying HSRP Configuration

To verify that the HSRP feature has been configured and is operating correctly, use the following **show ip slb vsrver** commands to display information about the SLB virtual server status:

```
SLB-Switch# show ip slb vsrver
```

slb vsrver	prot	virtual	state	conns
VS1	TCP	10.10.10.12:23	INSERVICE	2
VS2	TCP	10.10.10.18:23	INSERVICE	2

```
SLB-Switch# show ip slb vsrver detail
```

```
VS1, state = INSERVICE, v_index = 10
  virtual = 10.10.10.12:23, TCP, service = NONE, advertise = TRUE
  server farm = SERVERGROUP1, delay = 10, idle = 3600
  sticky timer = 0, sticky subnet = 255.255.255.255
  sticky group id = 0
  synguard counter = 0, synguard period = 0
  conns = 0, total conns = 0, syns = 0, syn drops = 0
  standby group = None
VS2, state = INOFSERVICE, v_index = 11
  virtual = 10.10.10.18:23, TCP, service = NONE, advertise = TRUE
  server farm = SERVERGROUP2, delay = 10, idle = 3600
  sticky timer = 0, sticky subnet = 255.255.255.255
  sticky group id = 0
  synguard counter = 0, synguard period = 0
  conns = 0, total conns = 0, syns = 0, syn drops = 0
  standby group = None
```

## SLB Stateful Backup

The stateful backup feature enables SLB to back up its load-balancing decisions incrementally between primary and backup Catalyst 4840G SLB switches. The backup switch has its virtual servers in a dormant state until failover is detected by HSRP; then the backup (now primary) switch begins advertising virtual addresses and filtering traffic.

This feature provides the switch with a one-to-one stateful or idle backup scheme. This means that only one instance of SLB is handling client or server traffic at a given time, and that there is at most one backup platform for each active load-balancing switch.


Figure 7-2 is an example of a stateful backup configuration, using HSRP on both the client and server sides to handle failover. The real servers route outbound traffic to 10.10.3.100, which is the HSRP address on the server side interfaces. The client (or access router) routes to the virtual IP address (10.10.10.12) through 10.10.2.100, the HSRP address on the client side.

Notice the loopback interfaces configured on both boxes for the exchange of these messages. Each SLB should also be given duplicate routes to the other switch loopback address to allow replication messages to flow despite an interface failure.

**Note**

To allow HSRP to function properly, you must configure **set spantree portfast** on any Layer 2 device between the load-balancing switches.

**Figure 7-2 SLB Stateful Environment**



## Configuring Stateful Backup

To configure stateful backup on the Catalyst 4840G SLB switches, perform this task beginning in global configuration mode:

	Command	Purpose
Step 1	SLB-Switch(config)# <b>ip slb vserver</b> <i>virtserver-name</i>	Configure a virtual server and enter virtual server configuration mode.
Step 2	SLB-Switch(config-slb-vserver)# <b>replicate</b> <i>casa listening-ip remote-ip port-number [interval] [password password timeout]</i>	Configure a stateful backup of SLB decision tables to a backup switch.

This example shows how to configure stateful backup for virtual server RESTRICTED\_HTTP using listening IP 10.10.3.132 and remote IP 10.10.99.3 over port 1032, and configure the password as “PASS” for Device A in Figure 7-2:

```
SLB-Switch(config)# ip slb vserver RESTRICTED_HTTP
SLB-Switch(config-slb-vserver)# virtual 10.10.10.12 tcp telnet
SLB-Switch(config-slb-vserver)# replicate casa 10.10.3.132 10.10.99.3 1024 password PASS
SLB-Switch(config-slb-vserver)# inservice standby virt
SLB-Switch(config-slb-vserver)#
.
(Information Deleted)
.
```

## Example Stateful Backup Configuration

Following is the stateful backup configuration for switch SLB1 shown in Figure 7-2:

```
!
ip slb serverfarm SF1
  nat server
  real 10.10.3.1
  inservice
  real 10.10.3.2
  inservice
  real 10.10.3.3
  inservice
!
ip slb vserver VS1
  virtual 10.10.10.12 tcp telnet
  serverfarm SF1
  replicate casa 10.10.99.132 10.10.99.99 1024 password PASS
  inservice standby virt
!
interface Loopback1
  ip address 10.10.99.132 255.255.255.255
!
!
interface FastEthernet1
  ip address 10.10.3.132 255.255.255.0
  no ip redirects
  no ip mroute-cache
  standby priority 5 preempt
  standby name out
  standby ip 10.10.3.100
  standby track FastEthernet3
!
interface FastEthernet3
  ip address 10.10.2.132 255.255.255.0
  no ip redirects
  standby priority 5 preempt
  standby name virt
```

```
standby ip 10.10.2.100
standby track FastEthernet1
!
```

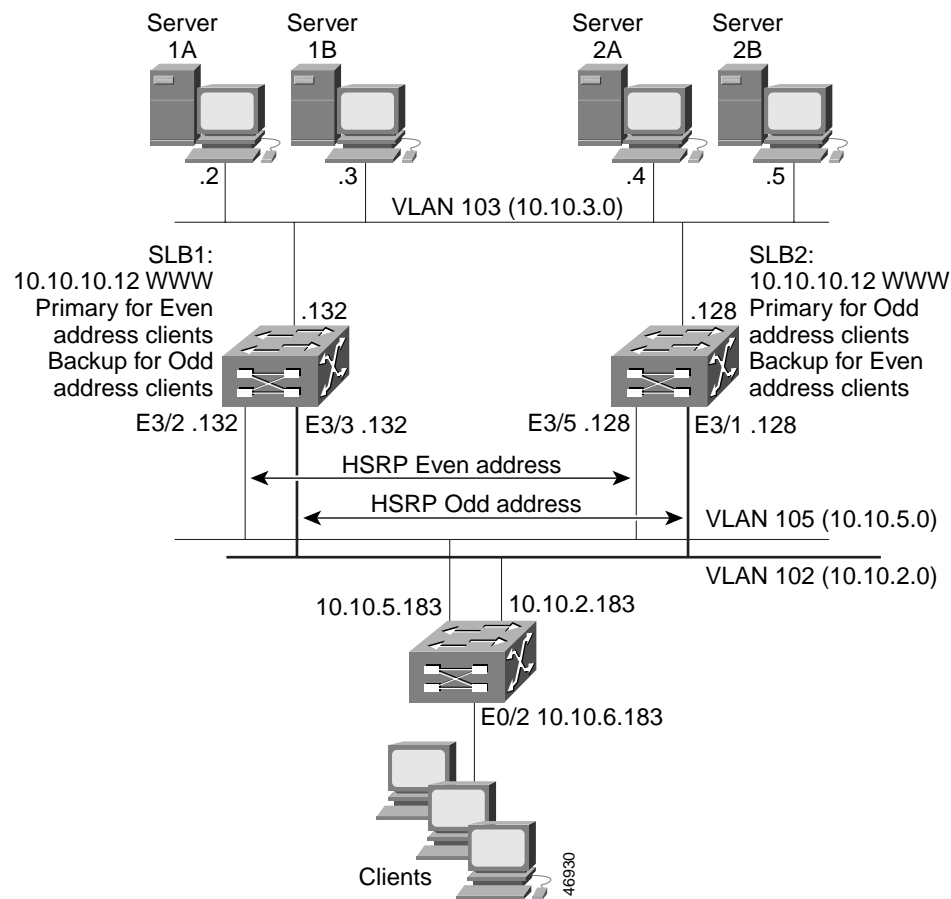
Following is the stateful backup configuration for switch SLB2 shown in Figure 7-2:

```
ip slb serverfarm SF1
  nat server
  real 10.10.3.1
  inservice
  real 10.10.3.2
  inservice
  real 10.10.3.3
  inservice
!
ip slb vserver VS1
  virtual 10.10.10.12 tcp telnet
  serverfarm SF1
  replicate casa 10.10.99.99 10.10.99.132 1024 password PASS
  inservice standby virt
!
!
interface Loopback1
  ip address 10.10.99.99 255.255.255.255
!
interface FastEthernet2
  ip address 10.10.2.99 255.255.255.0
  no ip redirects
  no ip route-cache
  no ip mroute-cache
  standby priority 10 preempt
  standby name virt
  standby ip 10.10.2.100
  standby track FastEthernet3
!
interface FastEthernet3
  ip address 10.10.3.99 255.255.255.0
  no ip redirects
  no ip route-cache
  no ip mroute-cache
  standby priority 10 preempt
  standby name out
  standby ip 10.10.3.100
  standby track FastEthernet2
!
```

## Example of IOS SLB with Active Standby

Figure 7-3 shows an IOS SLB network configured for active standby, with two IOS SLB devices load-balancing the same virtual IP address while backing up each other. If either device fails, the other takes over its load via normal HSRP failover and IOS SLB stateless redundancy.

Figure 7-3 IOS SLB Active Standby



The sample network configuration in Figure 7-3 has the following characteristics:

- SLB 1 balances servers 1A and 1B and SLB 2 balances 2A and 2B.
- A single virtual IP address (10.10.10.12 for Web) is supported across the two IOS SLB devices.
- Client traffic is divided in an access router, sending clients with even IP addresses to HSRP1 (10.10.5.100) and clients with odd IP addresses to HSRP2 (10.10.2.100). SLB 1 is configured as primary for clients with odd IP addresses, and SLB 2 is primary for clients with even IP addresses.
- The IOS SLB devices balance the traffic to disjoint sets of real servers. (If client NAT was used in this example, this would not be a requirement).
- Each set of real servers has a default gateway configured to its IOS SLB device.
- The HSRP address on VLAN 105 is 10.10.5.100. The HSRP address on VLAN 102 is 10.10.2.100.

**Configuration on SLB 1**

```

ip slb serverfarm EVEN
  nat server
  real 10.10.3.2
  inservice
  real 10.10.3.3
  inservice
!
ip slb serverfarm ODD
  nat server
  real 10.10.3.2
  inservice
  real 10.10.3.3
  inservice
!
ip slb vserver EVEN                                ; Same EVEN virtual server as in SLB 2
  virtual 10.10.10.12 tcp www
  serverfarm EVEN
  client 0.0.0.0 0.0.0.1
  inservice standby STANDBY_EVEN                 ; See standby name in Ethernet 3/3 below
!
ip slb vserver ODD                                ; Same ODD virtual server as in SLB 2
  virtual 10.10.10.12 tcp www
  serverfarm ODD
  client 0.0.0.1 0.0.0.1
  inservice standby STANDBY_ODD                 ; See standby name in Ethernet 3/2 below
!
interface Ethernet3/2
  ip address 10.10.5.132 255.255.255.0
  standby priority 20 preempt
  standby name STANDBY_ODD                       ; See standby name in SLB 2, Ethernet 3/5
  standby ip 10.10.5.100
!
interface Ethernet3/3
  ip address 10.10.2.132 255.255.255.0
  standby priority 10
  standby name STANDBY_EVEN                       ; See standby name in SLB 2, Ethernet 3/1
  standby ip 10.10.2.100

```

**Configuration on SLB 2**

```

ip slb serverfarm EVEN
nat server
real 10.10.3.4
inservice
real 10.10.3.5
inservice
!
ip slb serverfarm ODD
nat server
real 10.10.3.4
inservice
real 10.10.3.5
inservice
!
ip slb vserver EVEN ; Same EVEN virtual server as in SLB 1
virtual 10.10.10.12 tcp www
serverfarm EVEN
client 0.0.0.0 0.0.0.1
inservice standby STANDBY_EVEN ; See standby name in Ethernet 3/1 below
!
ip slb vserver ODD ; Same ODD virtual server as in SLB 1
virtual 10.10.10.12 tcp www
serverfarm ODD
client 0.0.0.1 0.0.0.1
inservice standby STANDBY_ODD ; See standby name in Ethernet 3/5 below
!
interface Ethernet3/1
ip address 10.10.2.128 255.255.255.0
standby priority 20 preempt
standby name STANDBY_EVEN ; See standby name in SLB 1, Ethernet 3/3
standby ip 10.10.2.100
!
interface Ethernet3/5
ip address 10.10.5.128 255.255.255.0
standby priority 10 preempt
standby name STANDBY_ODD ; See standby name in SLB 1, Ethernet 3/2
standby ip 10.10.5.100

```

**Configuration on the Access Router**

```
interface Ethernet0/0
ip address 10.10.5.183 255.255.255.0
no ip directed-broadcast
no ip route-cache
no ip mroute-cache
!
interface Ethernet0/1
ip address 10.10.2.183 255.255.255.0
no ip directed-broadcast
no ip route-cache
no ip mroute-cache
!
interface Ethernet0/2
ip address 10.10.6.183 255.255.255.0
no ip directed-broadcast
no ip route-cache
no ip mroute-cache
ip policy route-map virts
!
access-list 100 permit ip 0.0.0.1 255.255.255.254 host 10.10.10.12
access-list 101 permit ip 0.0.0.0 255.255.255.254 host 10.10.10.12
route-map virts permit 10
match ip address 100
set ip next-hop 10.10.5.100
!
route-map virts permit 15
match ip address 101
set ip next-hop 10.10.2.100
!
```