Chapter 5 Lab 5-2, Configure IP SLA Tracking and Path Control

Topology

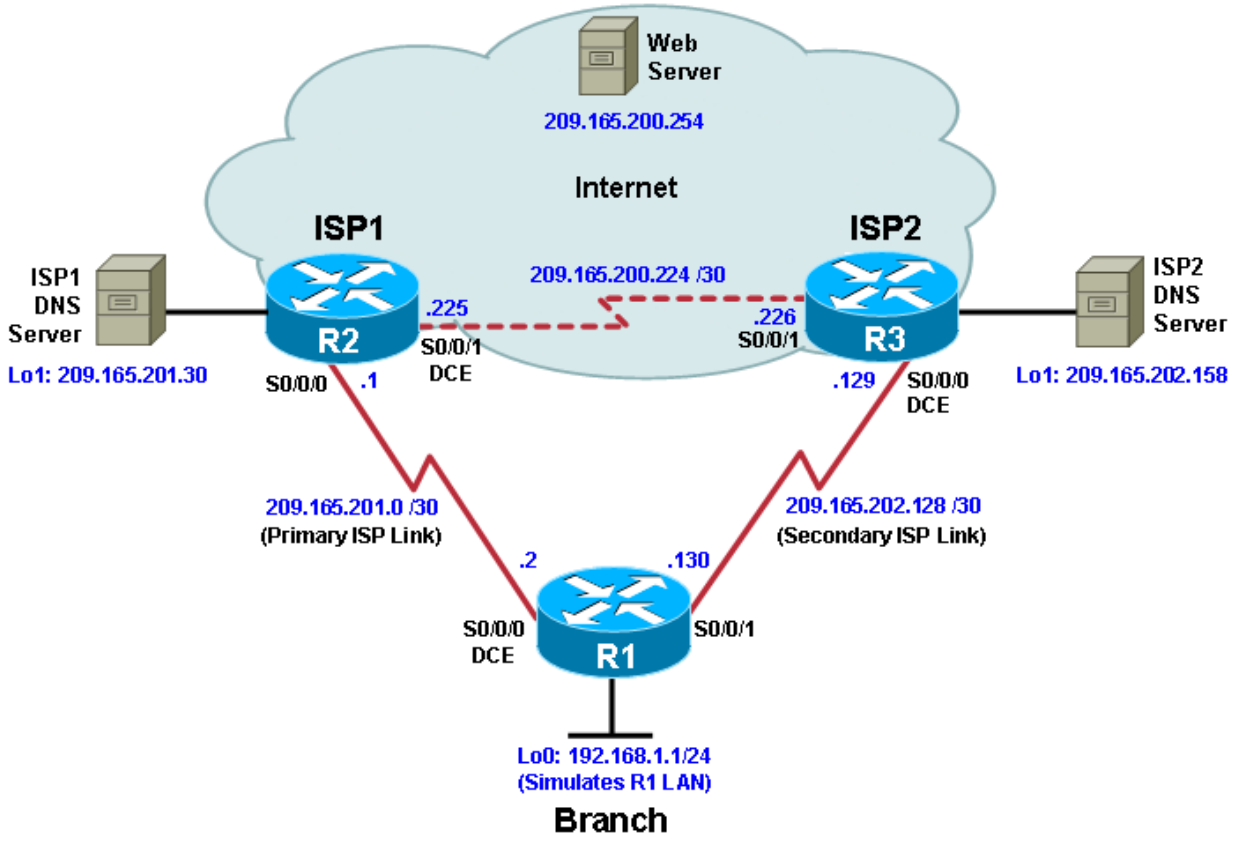

Objectives

- Configure and verify the IP SLA feature.
- Test the IP SLA tracking feature.
- Verify the configuration and operation using **show** and **debug** commands.

Background

You want to experiment with the Cisco IP Service Level Agreement (SLA) feature to study how it could be of value to your organization.

At times, a link to an ISP could be operational, yet users cannot connect to any other outside Internet resources. The problem might be with the ISP or downstream from them. Although policy-based routing (PBR) can be implemented to alter path control, you will implement the Cisco IOS SLA feature to monitor this behavior and intervene by injecting another default route to a backup ISP.

To test this, you have set up a three-router topology in a lab environment. Router R1 represents a branch office connected to two different ISPs. ISP1 is the preferred connection to the Internet, while ISP2 provides a backup link. ISP1 and ISP2 can also interconnect, and both can reach the web server. To monitor ISP1 for

failure, you will configure IP SLA probes to track the reachability to the ISP1 DNS server. If connectivity to the ISP1 server fails, the SLA probes detect the failure and alter the default static route to point to the ISP2 server.

Note: This lab uses Cisco 1841 routers with Cisco IOS Release 12.4(24)T1 and the Advanced IP Services image c1841-advipservicesk9-mz.124-24.T1.bin. You can use other routers (such as a 2801 or 2811) and Cisco IOS Software versions if they have comparable capabilities and features. Depending on the router and Cisco IOS Software version, the commands available and output produced might vary from what is shown in this lab.

Required Resources

- 3 routers (Cisco 1841 with Cisco IOS Release 12.4(24)T1 Advanced IP Services or comparable)
- Serial and console cables

Step 1: Prepare the routers and configure the router hostname and interface addresses.

a. Cable the network as shown in the topology diagram. Erase the startup configuration and reload each router to clear the previous configurations. Using the addressing scheme in the diagram, create the loopback interfaces and apply IP addresses to them as well as the serial interfaces on R1, ISP1, and ISP2.

You can copy and paste the following configurations into your routers to begin.

Note: Depending on the router model, interfaces might be numbered differently than those listed. You might need to alter them accordingly.

Router R1

```
hostname R1

interface Loopback 0
 description R1 LAN
 ip address 192.168.1.1 255.255.255.0

interface Serial0/0/0
 description R1 --> ISP1
 ip address 209.165.201.2 255.255.255.252
 clock rate 128000
 bandwidth 128
 no shutdown

interface Serial0/0/1
 description R1 --> ISP2
 ip address 209.165.202.130 255.255.255.252
 bandwidth 128
 no shutdown
```

Router ISP1 (R2)

```
hostname ISP1

interface Loopback0
 description Simulated Internet Web Server
 ip address 209.165.200.254 255.255.255.255

interface Loopback1
 description ISP1 DNS Server
 ip address 209.165.201.30 255.255.255.255
```



```
interface Serial0/0/0
 description ISP1 --> R1
 ip address 209.165.201.1 255.255.255.252
 bandwidth 128
 no shutdown

interface Serial0/0/1
 description ISP1 --> ISP2
 ip address 209.165.200.225 255.255.255.252
 clock rate 128000
 bandwidth 128
 no shutdown
```

Router ISP2 (R3)

```
hostname ISP2

interface Loopback0
 description Simulated Internet Web Server
 ip address 209.165.200.254 255.255.255.255

interface Loopback1
 description ISP2 DNS Server
 ip address 209.165.202.158 255.255.255.255

interface Serial0/0/0
 description ISP2 --> R1
 ip address 209.165.202.129 255.255.255.252
 clock rate 128000
 bandwidth 128
 no shutdown

interface Serial0/0/1
 description ISP2 --> ISP1
 ip address 209.165.200.226 255.255.255.252
 bandwidth 128
 no shutdown
```

b. Verify the configuration by using the **show interfaces description** command. The output from router R1 is shown here as an example.

```
R1# show interfaces description
Interface                      Status         Protocol Description
Fa0/0                          admin down     down
Fa0/1                          admin down     down
Se0/0/0                        up             up       R1 --> ISP1
Se0/0/1                        up             up       R1 --> ISP2
Lo0                            up             up       R1 LAN
```

All three interfaces should be active. Troubleshoot if necessary.

c. The current routing policy in the topology is as follows:

- Router R1 establishes connectivity to the Internet through ISP1 using a default static route.
- ISP1 and ISP2 have dynamic routing enabled between them, advertising their respective public address pools.
- ISP1 and ISP2 both have static routes back to the ISP LAN.

Note: For the purpose of this lab, the ISPs have a static route to an RFC 1918 private network address on the branch router R1. In an actual branch implementation, Network Address Translation (NAT) would be configured for all traffic exiting the branch LAN. Therefore, the static routes on the ISP routers would be pointing to the provided public pool of the branch office. This is covered in Lab 7-1, "Configure Routing Facilities to the Branch Office."

Implement the routing policies on the respective routers. You can copy and paste the following configurations.

Router R1

```
ip route 0.0.0.0 0.0.0.0 209.165.201.1
```

Router ISP1 (R2)

```
router eigrp 1
 network 209.165.200.224 0.0.0.3
 network 209.165.201.0 0.0.0.31
 no auto-summary

ip route 192.168.1.0 255.255.255.0 209.165.201.2
```

Router ISP2 (R3)

```
router eigrp 1
 network 209.165.200.224 0.0.0.3
 network 209.165.202.128 0.0.0.31
 no auto-summary

ip route 192.168.1.0 255.255.255.0 209.165.202.130
```

EIGRP neighbor relationship messages on ISP1 and ISP2 should be generated. Troubleshoot if necessary.

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 209.165.200.225 (Serial0/0/1) is
up: new adjacency
```

Step 2: Verify server reachability.

The Cisco IOS IP SLA feature enables an administrator to monitor network performance between Cisco devices (switches or routers) or from a Cisco device to a remote IP device. IP SLA probes continuously check the reachability of a specific destination, such as a provider edge router interface, the DNS server of the ISP, or any other specific destination, and can conditionally announce a default route only if the connectivity is verified.

a. Before implementing the Cisco IOS SLA feature, you must verify reachability to the Internet servers. From router R1, ping the web server, ISP1 DNS server, and ISP2 DNS server to verify connectivity. You can copy the following Tcl script and paste it intoR1.

```
foreach address {
209.165.200.254
209.165.201.30
209.165.202.158
} {
ping $address source 192.168.1.1
}
```

```
R1(tcl)# foreach address {
+>(tcl)# 209.165.200.254
```



```
+>(tcl)# 209.165.201.30
+>(tcl)# 209.165.202.158
+>(tcl)# } {
+>(tcl)# ping $address source 192.168.1.1
+>(tcl)#}

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.200.254, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/15/16 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.30, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/14/16 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.202.158, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/21/24 ms
```

b. Trace the path taken to the web server, ISP1 DNS server, and ISP2 DNS server. You can copy the following Tcl script and paste it into R1.

```
foreach address {
209.165.200.254
209.165.201.30
209.165.202.158
} {
trace $address source 192.168.1.1
}
```

```
R1(tcl)# foreach address {
+>(tcl)# 209.165.200.254
+>(tcl)# 209.165.201.30
+>(tcl)# 209.165.202.158
+>(tcl)# } {
+>(tcl)# trace $address source 192.168.1.1
+>(tcl)# }

Type escape sequence to abort.
Tracing the route to 209.165.200.254

  1 209.165.201.1 20 msec 8 msec *
Type escape sequence to abort.
Tracing the route to 209.165.201.30

  1 209.165.201.1 8 msec 8 msec *
Type escape sequence to abort.
Tracing the route to 209.165.202.158

  1 209.165.201.1 8 msec 8 msec 4 msec
  2 209.165.200.226 8 msec 8 msec *
```

Through which ISP is traffic flowing?

Step 3: Configure IP SLA probes.

When the reachability tests are successful, you can configure the Cisco IOS IP SLAs probes. Different types of probes can be created, including FTP, HTTP, and jitter probes. In this scenario, you will configure ICMP echo probes.

a. Create an ICMP echo probe on R1 to the primary DNS server on ISP1 using the **ip sla** command.

Note: With Cisco IOS Release 12.4(4)T, 12.2(33)SB, and 12.2(33)SXI, the **ip sla** command has replaced the previous **ip sla monitor** command. In addition, the **icmp-echo** command has replaced the **type echo protocol ipIcmpEcho** command.

```
R1(config)# ip sla 11
R1(config-ip-sla)# icmp-echo 209.165.201.30
R1(config-ip-sla-echo)# frequency 10
R1(config-ip-sla-echo)# exit
R1(config)# ip sla schedule 11 life forever start-time now
```

The operation number of 11 is only locally significant to the router. The **frequency 10** command schedules the connectivity test to repeat every 10 seconds. The probe is scheduled to start now and to run forever.

b. Verify the IP SLAs configuration of operation 11 using the **show ip sla configuration 11** command.

Note: With Cisco IOS Release 12.4(4)T, 12.2(33)SB, and 12.2(33)SXI, the **show ip sla configuration** command has replaced the **show ip sla monitor configuration** command.

```
R1# show ip sla configuration 11
IP SLAs, Infrastructure Engine-II.
Entry number: 11
Owner:
Tag:
Type of operation to perform: icmp-echo
Target address/Source address: 209.165.201.30/0.0.0.0
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Verify data: No
Vrf Name:
Schedule:
   Operation frequency (seconds): 10  (not considered if randomly scheduled)
   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 (not considered if react RTT is configured)
Distribution Statistics:
   Number of statistic hours kept: 2
   Number of statistic distribution buckets kept: 1
   Statistic distribution interval (milliseconds): 20
History Statistics:
   Number of history Lives kept: 0
   Number of history Buckets kept: 15
   History Filter Type: None
Enhanced History:
```

The output lists the details of the configuration of operation 11. The operation is an ICMP echo to 209.165.201.30, with a frequency of 10 seconds, and it has already started (the start time has already passed).

c. Issue the **show ip sla statistics** command to display the number of successes, failures, and results of the latest operations.

Note: With Cisco IOS Release 12.4(4)T, 12.2(33)SB, and 12.2(33)SXI, the **show ip sla statistics** command has replaced the **show ip sla monitor statistics** command.

```
R1# show ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 11

Latest operation start time: *21:22:29.707 UTC Fri Apr 2 2010
Latest operation return code: OK
Number of successes: 5
Number of failures: 0
Operation time to live: Forever
```

You can see that operation 11 has already succeeded five times, has had no failures, and the last operation returned an OK result.

d. Although not actually required because IP SLA session 11 alone could provide the desired fault tolerance, create a second probe, 22, to test connectivity to the second DNS server located on router ISP2. You can copy and paste the following commands on R1.

```
ip sla 22
icmp-echo 209.165.202.158
frequency 10
exit
ip sla schedule 22 life forever start-time now
```

e. Verify the new probe using the **show ip sla configuration** and **show ip sla statistics** commands.

```
R1# show ip sla configuration 22
IP SLAs, Infrastructure Engine-II.
Entry number: 22
Owner:
Tag:
Type of operation to perform: icmp-echo
Target address/Source address: 209.165.201.158/0.0.0.0
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Verify data: No
Vrf Name:
Schedule:
   Operation frequency (seconds): 10  (not considered if randomly scheduled)
   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 (not considered if react RTT is configured)
Distribution Statistics:
   Number of statistic hours kept: 2
```

```
   Number of statistic distribution buckets kept: 1
   Statistic distribution interval (milliseconds): 20
History Statistics:
   Number of history Lives kept: 0
   Number of history Buckets kept: 15
   History Filter Type: None
Enhanced History:

R1# show ip sla statistics 22
IPSLAs Latest Operation Statistics

IPSLA operation id: 22

Latest operation start time: *21:24:14.215 UTC Fri Apr 2 2010
Latest operation return code: OK
Number of successes: 4
Number of failures: 0
Operation time to live: Forever
```

The output lists the details of the configuration of operation 22. The operation is an ICMP echo to 209.165.202.158, with a frequency of 10 seconds, and it has already started (the start time has already passed). The statistics also prove that operation 22 is active.

Step 4: Configure tracking options.

Although PBR could be used, you will configure a floating static route that appears or disappears depending on the success or failure of the IP SLA.

a. Remove the current default route on R1, and replace it with a floating static route having an administrative distance of 5.

```
R1(config)# no ip route 0.0.0.0 0.0.0.0 209.165.201.1
R1(config)# ip route 0.0.0.0 0.0.0.0 209.165.201.1 5
R1(config)# exit
```

b. Verify the routing table.

```
R1# show ip route
*Apr  2 20:00:37.367: %SYS-5-CONFIG_I: Configured from console by console
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 209.165.201.1 to network 0.0.0.0

     209.165.201.0/30 is subnetted, 1 subnets
C       209.165.201.0 is directly connected, Serial0/0/0
     209.165.202.0/30 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, Serial0/0/1
C    192.168.1.0/24 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [5/0] via 209.165.201.1
```

Notice that the default static route is now using the route with the administrative distance of 5. The first tracking object is tied to IP SLA object 11.

c. Use the **track 1 ip sla 11 reachability** command to enter the config-track subconfiguration mode.

Note: With Cisco IOS Release 12.4(20)T, 12.2(33)SXI1, and 12.2(33)SRE and Cisco IOS XE Release 2.4, the **track ip sla** command has replaced the **track rtr** command.

```
R1(config)# track 1 ip sla 11 reachability
R1(config-track)#
```

d. Specify the level of sensitivity to changes of tracked objects to 10 seconds of down delay and 1 second of up delay using the **delay down 10 up 1** command. The delay helps to alleviate the effect of flapping objects—objects that are going down and up rapidly. In this situation, if the DNS server fails momentarily and comes back up within 10 seconds, there is no impact.

```
R1(config-track)# delay down 10 up 1
R1(config-track)# exit
R1(config)#
```

e. Configure the floating static route that will be implemented when tracking object 1 is active. To view routing table changes as they happen, first enable the **debug ip routing** command. Next, use the **ip route 0.0.0.0 0.0.0.0 209.165.201.1 2 track 1** command to create a floating static default route via 209.165.201.1 (ISP1). Notice that this command references the tracking object number 1, which in turn references IP SLA operation number 11.

```
R1# debug ip routing
IP routing debugging is on
R1#
*Apr  2 21:26:46.171: RT: NET-RED 0.0.0.0/0

R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ip route 0.0.0.0 0.0.0.0 209.165.201.1 2 track 1
R1(config)#
*Apr  2 21:27:02.851: RT: closer admin distance for 0.0.0.0, flushing 1
routes
*Apr  2 21:27:02.851: RT: NET-RED 0.0.0.0/0
*Apr  2 21:27:02.851: RT: add 0.0.0.0/0 via 209.165.201.1, static metric
[2/0]
*Apr  2 21:27:02.851: RT: NET-RED 0.0.0.0/0
*Apr  2 21:27:02.851: RT: default path is now 0.0.0.0 via 209.165.201.1
*Apr  2 21:27:02.855: RT: new default network 0.0.0.0
*Apr  2 21:27:02.855: RT: NET-RED 0.0.0.0/0
*Apr  2 21:27:07.851: RT: NET-RED 0.0.0.0/0
```

Notice that the default route with an administrative distance of 5 has been immediately flushed because of a route with a better admin distance. It then adds the new default route with the admin distance of 2.

f. Repeat the steps for operation 22, track number 2, and assign the static route an admin distance higher than track 1 and lower than 5. On R1, copy the following configuration, which sets an admin distance of 3.

```
track 2 ip sla 22 reachability
delay down 10 up 1
exit
ip route 0.0.0.0 0.0.0.0 209.165.202.129 3 track 2
```

g. Verify the routing table again.

```
R1# show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```



```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 209.165.201.1 to network 0.0.0.0

     209.165.201.0/30 is subnetted, 1 subnets
C       209.165.201.0 is directly connected, Serial0/0/0
     209.165.202.0/30 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, Serial0/0/1
C    192.168.1.0/24 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [2/0] via 209.165.201.1
```

Although a new default route was entered, its administrative distance is not better than 2. Therefore, it does not replace the previously entered default route.

Step 5: Verify IP SLA operation.

In this step you observe and verify the dynamic operations and routing changes when tracked objects fail. The following summarizes the process:

- Disable the DNS loopback interface on ISP1 (R2).
- Observe the output of the **debug** command on R1.
- Verify the static route entries in the routing table and the IP SLA statistics of R1.
- Re-enable the loopback interface on ISP1 (R2) and again observe the operation of the IP SLA tracking feature.

```
ISP1(config)# interface loopback 1
ISP1(config-if)# shutdown
ISP1(config-if)#
*Apr  2 15:53:14.307: %LINK-5-CHANGED: Interface Loopback1, changed state to
administratively down
*Apr  2 15:53:15.307: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Loopback1, changed state to down
```

a. Shortly after the loopback interface is administratively down, observe the debug output being generated on R1.

```
R1#
*Apr  2 21:32:33.323: %TRACKING-5-STATE: 1 ip sla 11 reachability Up->Down
*Apr  2 21:32:33.323: RT: del 0.0.0.0 via 209.165.201.1, static metric [2/0]
*Apr  2 21:32:33.323: RT: delete network route to 0.0.0.0
*Apr  2 21:32:33.323: RT: NET-RED 0.0.0.0/0
*Apr  2 21:32:33.323: RT: NET-RED 0.0.0.0/0
*Apr  2 21:32:33.323: RT: add 0.0.0.0/0 via 209.165.202.129, static metric
[3/0]
*Apr  2 21:32:33.323: RT: NET-RED 0.0.0.0/0
*Apr  2 21:32:33.323: RT: default path is now 0.0.0.0 via 209.165.202.129
*Apr  2 21:32:33.323: RT: new default network 0.0.0.0
*Apr  2 21:32:33.327: RT: NET-RED 0.0.0.0/0
*Apr  2 21:32:46.171: RT: NET-RED 0.0.0.0/0
```

The tracking state of track 1 changes from up to down. This is the object that tracked reachability for IP SLA object 11, with an ICMP echo to the ISP1 DNS server at 209.165.201.30.

R1 then proceeds to delete the default route with the administrative distance of 2 and installs the next highest default route to ISP2 with the administrative distance of 3.

b. Verify the routing table.

```
R1# show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 209.165.202.129 to network 0.0.0.0

     209.165.201.0/30 is subnetted, 1 subnets
C       209.165.201.0 is directly connected, Serial0/0/0
     209.165.202.0/30 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, Serial0/0/1
C    192.168.1.0/24 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [3/0] via 209.165.202.129
```

The new static route has an administrative distance of 3 and is being forwarded to ISP2 as it should.

c. Verify the IP SLA statistics.

```
R1# show ip sla statistics
IPSLAs Latest Operation Statistics

PSLA operation id: 11
Type of operation: icmp-echo
        Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: *15:36:42.871 UTC Fri Apr 2 2010
Latest operation return code: No connection
Number of successes: 84
Number of failures: 13
Operation time to live: Forever


IPSLA operation id: 22
Type of operation: icmp-echo
        Latest RTT: 8 milliseconds
Latest operation start time: *15:36:46.335 UTC Fri Apr 2 2010
Latest operation return code: OK
Number of successes: 81
Number of failures: 1
Operation time to live: Forever
```

Notice that the latest return code is **No connection** and there have been 12 failures on IP SLA object 11.

d. Initiate a trace to the web server from the internal LAN IP address.

```
R1# trace 209.165.200.254 source 192.168.1.1

Type escape sequence to abort.
Tracing the route to 209.165.200.254

  1 209.165.202.129 8 msec 8 msec *
```

This confirms that traffic is leaving router R1 and being forwarded to the ISP2 router.

e. To examine the routing behavior when connectivity to the ISP1 DNS is restored, re-enable the DNS address on ISP1 (R2) by issuing the **no shutdown** command on the loopback 1 interface on ISP2.

```
ISP1(config-if)# no shutdown
*Apr  2 15:56:24.655: %LINK-3-UPDOWN: Interface Loopback1, changed state to
up
*Apr  2 15:56:25.655: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Loopback1, changed state to up
```

Notice the output of the **debug ip routing** command on R1.

```
R1#
*Apr  2 21:35:34.327: %TRACKING-5-STATE: 1 ip sla 11 reachability Down->Up
*Apr  2 21:35:34.327: RT: closer admin distance for 0.0.0.0, flushing 1
routes
*Apr  2 21:35:34.327: RT: NET-RED 0.0.0.0/0
*Apr  2 21:35:34.327: RT: add 0.0.0.0/0 via 209.165.201.1, static metric
[2/0]
*Apr  2 21:35:34.327: RT: NET-RED 0.0.0.0/0
*Apr  2 21:35:34.327: RT: default path is now 0.0.0.0 via 209.165.201.1
*Apr  2 21:35:34.327: RT: new default network 0.0.0.0
*Apr  2 21:35:34.327: RT: NET-RED 0.0.0.0/0
*Apr  2 21:35:39.327: RT: NET-RED 0.0.0.0/0
*Apr  2 21:35:46.171: RT: NET-RED 0.0.0.0/0
```

Now the IP SLA 11 operation transitions back to an up state and reestablishes the default static route to ISP1 with an administrative distance of 2.

f. Again examine the IP SLA statistics.

```
R1# show ip sla statistics
IPSLAs Latest Operation Statistics

Type of operation: icmp-echo
        Latest RTT: 8 milliseconds
Latest operation start time: *15:40:42.871 UTC Fri Apr 2 2010
Latest operation return code: OK
Number of successes: 88
Number of failures: 35
Operation time to live: Forever


IPSLA operation id: 22
Type of operation: icmp-echo
        Latest RTT: 16 milliseconds
Latest operation start time: *15:40:46.335 UTC Fri Apr 2 2010
Latest operation return code: OK
Number of successes: 105
Number of failures: 1
Operation time to live: Forever
```

The IP SLA 11 operation is active again, as indicated by the OK return code, and the number of successes is incrementing.

g. Verify the routing table.

```
R1# show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```



```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 209.165.201.1 to network 0.0.0.0

     209.165.201.0/30 is subnetted, 1 subnets
C       209.165.201.0 is directly connected, Serial0/0/0
     209.165.202.0/30 is subnetted, 1 subnets
C       209.165.202.128 is directly connected, Serial0/0/1
C    192.168.1.0/24 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [2/0] via 209.165.201.1
```

The default static through ISP1 with an administrative distance of 2 is reestablished.

There are many possibilities available with object tracking and Cisco IOS IP SLAs. As shown in this lab, a probe can be based on reachability, changing routing operations, and path control based on the ability to reach an object. However, Cisco IOS IP SLAs also allow paths to be changed based on network conditions such as delay, load, and other factors.

Before deploying a Cisco IOS IP SLA solution, the impact of the additional probe traffic being generated should be considered, including how that traffic affects bandwidth utilization, and congestion levels. Tuning the configuration (for example, with the **delay** and **frequency** commands) is critical to mitigate possible issues related to excessive transitions and route changes in the presence of flapping tracked objects.

The benefits of running IP SLAs should be carefully evaluated. The IP SLA is an additional task that must be performed by the router's CPU. A large number of intensive SLAs could be a significant burden on the CPU, possibly interfering with other router functions and having detrimental impact on the overall router performance. The CPU load should be monitored after the SLAs are deployed to verify that they do not cause excessive utilization of the router CPU.

Router Interface Summary Table

Router Interface Summary				
Router Model	Ethernet Interface #1	Ethernet Interface #2	Serial Interface #1	Serial Interface #2
1700	Fast Ethernet 0 (FA0)	Fast Ethernet 1 (FA1)	Serial 0 (S0)	Serial 1 (S1)
1800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
2600	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0 (S0/0)	Serial 0/1 (S0/1)
2800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)

Note: To find out how the router is configured, look at the interfaces to identify the type of router and how many interfaces the router has. Rather than list all combinations of configurations for each router class, this table includes identifiers for the possible combinations of Ethernet and serial interfaces in the device. The table does not include any other type of interface, even though a specific router might contain one. For example, for an ISDN BRI interface, the string in parenthesis is the legal abbreviation that can be used in Cisco IOS commands to represent the interface.