Lab 4-1 Configuring Basic Integrated IS-IS

Learning Objectives

- Configure and verify the operation of Integrated IS-IS on a router
- Configure a NET identifying a domain, area, and intermediate system
- Configure and verify Level 1 and Level 2 IS-IS adjacencies
- Verify and understand the IS-IS topology table
- Manipulate IS-IS adjacency timers
- Implement IS-IS domain and link authentication

Topology Diagram

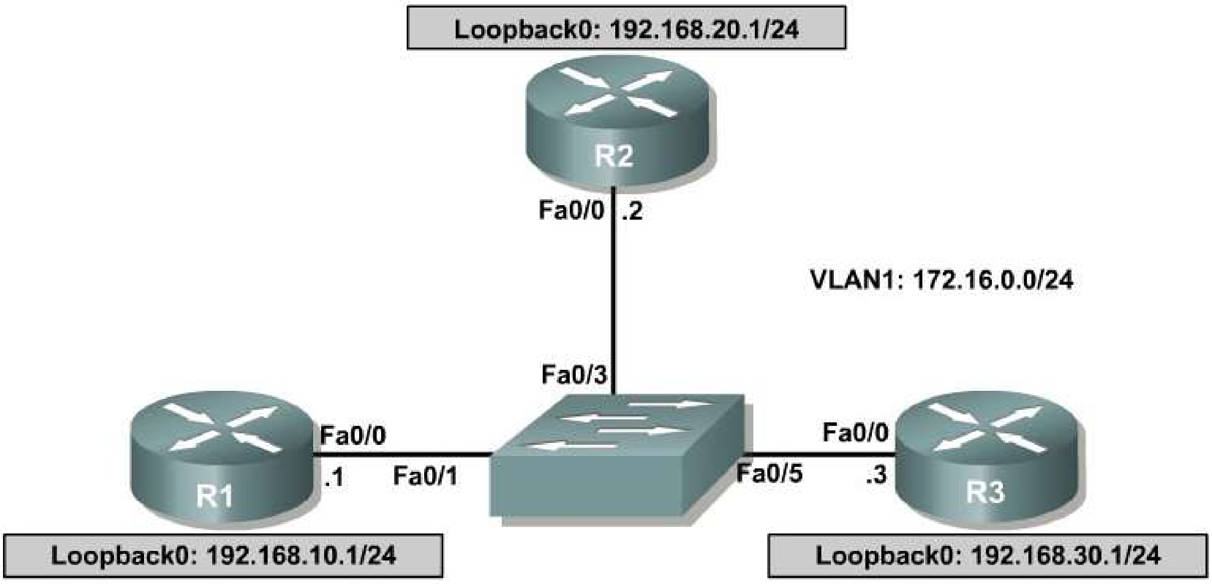

Scenario

The IS-IS routing protocol has become increasingly popular with widespread usage among service providers. The International Travel Agency (ITA) is considering implementing IS-IS because it is a link state protocol that enables very fast convergence with large scalability and flexibility. But before making a final decision, management wants a non-production network set up to test the IS-IS routing protocol.

The backbone of the production ITA WAN consists of three routers connected by an Ethernet core. Because the routers are also connected to the Internet, authentication is needed to prevent unauthorized routers from participating in the IS-IS process.

Step 1: Addressing and Basic Connectivity

Build and configure the network according to the diagram, but do not configure IS-IS yet. Configure loopback interfaces and addresses as well.

Use **ping** to test connectivity between the directly connected Fast Ethernet interfaces. You could alternatively use the following TCL script to ping across the Fast Ethernet link:

```
foreach address {
172.16.0.1
172.16.0.2
172.16.0.3 } { ping $address }
```

Step 2: Configuring Basic IS-IS

IS-IS (ISO/IEC 10589) is implemented with network service access point (NSAP) addresses consisting of three fields: area address, system ID, and NSEL (also known as N-selector, the service identifier or the process ID). The area address field can be from one to thirteen octets, the system ID field is usually six octets (must be six for Cisco IOS), and the NSEL identifies a process on the device. It is a loose equivalent to a port or socket in IP. The NSEL is not used in routing decisions.

When the NSEL is set to 00, the NSAP is referred to as the network entity title (NET). NETs and NSAPs are represented in hexadecimal, and must start and end on a byte boundary, such as 49.0001.1111.1111.1111.00

Level 1, or L1, IS-IS routing is based on system ID. Therefore, each router must have a unique system ID within the area. L1 IS-IS routing equates to intra-area routing. It is customary to use either a MAC address from the router or, for Integrated IS-IS, to code the IP address of a loopback address, for example, into the system ID.

Area addresses starting with 48, 49, 50, or 51 are private addresses. This group of addresses should not be advertised to other connectionless network service (CLNS) networks. The area address must be the same for all routers in an area.

On a LAN, one of the routers is elected the designated intermediate system (DIS) based on interface priority. The default is 64. If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) address is selected. The (Ethernet) MAC address serves as the SNPA address for Ethernet LANs. The DIS serves the same purpose for IS-IS as the designated router does for OSPF. The ITA network engineer decides that R1 is the DIS, so its priority must be set higher than R2 and R3.

Now, configure Integrated IS-IS on each router and set a priority of 100 on the FastEthernet 0/0 interface of R1 as follows:

```
R1(config)# router isis
R1(config-router)# net 49.0001.1111.1111.1111.00
R1(config-router)# interface fastethernet 0/0
R1(config-if)# ip router isis
R1(config-if)# isis priority 100
R1(config-if)# interface loopback 0
R1(config-if)# ip router isis

R2(config)# router isis
R2(config-router)# net 49.0001.2222.2222.2222.00
R2(config-router)# interface fastethernet 0/0
R2(config-if)# ip router isis
R2(config-if)# interface loopback 0
R2(config-if)# ip router isis

R3(config)# router isis
R3(config-router)# net 49.0001.3333.3333.3333.00
R3(config-router)# interface fastethernet 0/0
R3(config-if)# ip router isis
R3(config-if)# interface loopback 0
R3(config-if)# ip router isis
```

1. Identify parts of the NSAP/NET addresses.

   a. Area Address:

   ____

   b. R1 System ID:

   ____

   c. R2 System ID:

   ____

   d. R3 System ID:

   ____

   e. NSEL:

   ____

Step 3: Verifying IS-IS Adjacencies and Operation

Verify IS-IS operation using **show** commands on any of the three routers. The following is output for R1:

```
R1# show ip protocols
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
```



```
  Routing for Networks:
    FastEthernet0/0
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.30.1         115      00:00:36
    192.168.20.1         115      00:00:36
  Distance: (default is 115)
```

Because you are also working with the OSI connectionless protocol suite, use the **show clns protocols** command to see the IS-IS protocol output:

```
R1# show clns protocols

IS-IS Router: <Null Tag>
  System Id: 1111.1111.1111.00  IS-Type: level-1-2
  Manual area address(es):
        49.0001
  Routing for area address(es):
        49.0001
  Interfaces supported by IS-IS:
        FastEthernet0/0 - IP
        Loopback0 - IP
  Redistribute:
    static (on by default)
  Distance for L2 CLNS routes: 110
  RRR level: none
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
R1#
```

Notice that the update timers are set to zero (0). Updates are not sent at regular intervals because they are event driven. The Last Update field indicates how long it has been since the last update in hours:minutes:seconds.

Issue the **show clns neighbors** command to view adjacencies:

```
R1# show clns neighbors

System Id     Interface    SNPA                State  Holdtime   Type Protocol
R2            Fa0/0        0004.9ad2.d0c0      Up     9          L1L2 IS-IS
R3            Fa0/0        0002.16f4.1ba0      Up     29         L1L2 IS-IS
```

Neighbor ISs (Intermediate Systems) and neighbor ESs (End Systems) are shown, if applicable. You can use the keyword **detail** to display comprehensive neighbor information:

```
R1# show clns neighbors detail

System Id     Interface    SNPA                State  Holdtime   Type Protocol
R2            Fa0/0        0004.9ad2.d0c0      Up     24         L1L2 IS-IS
  Area Address(es): 49.0001
  IP Address(es):  172.16.0.2*
  Uptime: 00:07:30
  NSF capable
R3            Fa0/0        0002.16f4.1ba0      Up     27         L1L2 IS-IS
  Area Address(es): 49.0001
```



```
  IP Address(es):  172.16.0.3*
  Uptime: 00:07:00
  NSF capable
```

The system IDs of the IS neighbors are the hostnames of the respective neighbor routers. Starting with Cisco IOS Release 12.0(5), Cisco routers support dynamic hostname mapping. The feature is enabled by default. As seen in the sample output, the configured system ID of 2222.2222.2222 has been replaced by the hostname R2. Similarly, R3 replaces 3333.3333.3333.

The adjacency Type for both neighbors is L1L2. By default, Cisco IOS enables both L1 and L2 adjacency negotiation on IS-IS routers. You can use the router configuration mode command **is-type** or the interface configuration command **isis circuit-type** to specify how the router operates for L1 and L2 routing.

You can use the **show isis database** and **show clns interface fa0/0** commands to obtain DIS and related information. First, issue the **clear isis *** command on all routers to force IS-IS to refresh its link-state databases and recalculate all routes. A minute or two may be needed for all routers to update their respective IS-IS databases.

```
All_Router# clear isis *
```

Issue the **show isis database** command to view the content of the IS-IS database:

```
R1# show isis database

IS-IS Level-1 Link State Database:
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00           * 0x00000008   0x088F        1191              0/0/0
R1.01-00           * 0x00000002   0x9B60        1192              0/0/0
R2.00-00             0x00000001   0x8736        1190              0/0/0
R3.00-00             0x00000002   0x39A1        1195              0/0/0
IS-IS Level-2 Link State Database:
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00           * 0x00000017   0x4E1B        1195              0/0/0
R1.01-00           * 0x00000002   0x4D37        1192              0/0/0
R2.00-00             0x00000010   0xF4B9        1191              0/0/0
R3.00-00             0x00000002   0xD703        1195              0/0/0
```

IS-IS retains a separate database for L1 and L2 routing. Because IS-IS is a link-state protocol, the link-state database should be the same for the three routers.

As discussed earlier, if the priority for R1’s FastEthernet 0/0 interface had not been increased, the DIS would have been elected on the basis of the highest SNPA. DIS election is preemptive, unlike OSPF behavior. The **isis priority 100** command ensured that R1 would be elected the DIS, regardless of router boot order. But how can it be determined from the **show isis database** output that R1 is indeed the DIS?

Look at the entries under the link-state protocol data unit ID (LSPID) column. The first six octets form the system ID. As mentioned earlier, because of the dynamic host mapping feature, the respective router names are listed instead of the numerical system ID. Following the system ID are two octets.

The first octet is the pseudonode ID, representing a LAN. The pseudonode ID is used to distinguish LAN IDs on the same DIS. When this value is non-zero, the associated LSP is a pseudonode LSP originating from the DIS. The DIS is the only system that originates pseudonode LSPs. The DIS creates one pseudonode LSP for L1 and one for L2, as shown in the previous output.

The pseudonode ID varies upon reboot of the router as a function of the creation or deletion of virtual interfaces, such as loopback interfaces. The system ID and pseudonode ID together are referred to as the circuit ID. An example is R1.01.

A non-pseudonode LSP represents a router and is distinguished by the fact that the two-byte value in the circuit ID is 00.

The second octet forms the LSP fragmentation number. The value 00 indicates that all data fits into a single LSP. If there had been more information that did not fit into the first LSP, IS-IS would have created additional LSPs with increasing LSP numbers, such as 01, 02, and so on. The asterisk (*) indicates that the LSP was originated by the local system.

Issue the **show clns interface fastethernet 0/0** command:

```
R1# show clns interface fastethernet 0/0
FastEthernet0/0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 8 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 100, Circuit ID: R1.01
    DR ID: R1.01
    Level-1 IPv6 Metric: 10
    Number of active level-1 adjacencies: 2
    Level-2 Metric: 10, Priority: 100, Circuit ID: R1.01
    DR ID: R1.01
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 2
    Next IS-IS LAN Level-1 Hello in 803 milliseconds
    Next IS-IS LAN Level-2 Hello in 2 seconds
```

Notice that the circuit ID, R1.01, which is made up of the system and pseudonode IDs, identifies the DIS. Circuit Types, Levels, Metric, and Priority information is also displayed.

You can obtain additional information about a specific LSP ID by appending the LSP ID and **detail** keyword to the **show isis database** command, as shown in the output. The hostname is case sensitive. You can also use this command to view the IS-IS database of a neighbor router by including its hostname in the command.

```
R1# show isis database R1.00-00 detail

IS-IS Level-1 LSP R1.00-00
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00           * 0x0000000B   0x0292        831               0/0/0
  Area Address: 49.0001
  NLPID:         0xCC
  Hostname: R1
  IP Address:   192.168.10.1
  Metric: 10         IP 172.16.0.0 255.255.255.0
  Metric: 10         IP 192.168.10.0 255.255.255.0
  Metric: 10         IS R1.02
  Metric: 10         IS R1.01

IS-IS Level-2 LSP R1.00-00
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00           * 0x0000000D   0x4703        709               0/0/0
  Area Address: 49.0001
  NLPID:         0xCC
  Hostname: R1
  IP Address:   192.168.10.1
  Metric: 10         IS R1.02
  Metric: 10         IS R1.01
  Metric: 20         IP 192.168.30.0 255.255.255.0
  Metric: 10         IP 192.168.10.0 255.255.255.0
  Metric: 10         IP 172.16.0.0 255.255.255.0
  Metric: 20         IP 192.168.20.0 255.255.255.0
```

The default IS-IS metric for every link is 10, but notice that the metrics for the 192.168.20.0 and 192.168.30.0 networks are both 20. This is because the networks are not directly connected, but are directly connected to neighbor routers.

Issue the **show isis topology** command to display the paths to the other intermediate systems:

```
R1# show isis topology

IS-IS paths to level-1 routers
System Id     Metric   Next-Hop         Interface    SNPA
R1            --
R2            10       R2               Fa0/0        0004.9ad2.d0c0
R3            10       R3               Fa0/0        0002.16f4.1ba0

IS-IS paths to level-2 routers
System Id     Metric   Next-Hop         Interface    SNPA
R1            --
R2            10       R2               Fa0/0        0004.9ad2.d0c0
R3            10       R3               Fa0/0        0002.16f4.1ba0
```

The highlighted entries in the SNPA column are the MAC addresses of the R2 and R3 FastEthernet 0/0 interfaces.

Issue the **show isis route** command to view the IS-IS L1 routing table:


```
R1# show isis route

IS-IS not running in OSI mode (*) (only calculating IP routes)

(*) Use "show isis topology" command to display paths to all routers
```

This command has no useful output because it is specific to OSI routing. Remember, IP IS-IS was enabled on each router. If CLNP were configured in the network, more interesting output would appear.

Issue the **show clns route** command to view the IS-IS L2 routing table:

```
R1# show clns route
Codes: C - connected, S - static, d - DecnetIV
       I - ISO-IGRP,  i - IS-IS,  e - ES-IS
       B - BGP,       b - eBGP-neighbor

C  49.0001.1111.1111.1111.00 [1/0], Local IS-IS NET
C  49.0001 [2/0], Local IS-IS Area
```

Again, there is no useful output because this command applies to OSI routing and not IP routing.

Issue the **show ip route** command to view the IP routing table:

```
R1# show ip route
<output omitted>

Gateway of last resort is not set

i L1 192.168.30.0/24 [115/20] via 172.16.0.3, FastEthernet0/0
C    192.168.10.0/24 is directly connected, Loopback0
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.0.0 is directly connected, FastEthernet0/0
i L1 192.168.20.0/24 [115/20] via 172.16.0.2, FastEthernet0/0
```

Notice how the routes to the 192.168.30.0 and 192.168.20.0 networks were learned.

The **show clns neighbors, show isis database, show clns interface, show isis topology, show isis route,** and **show clns route** commands illustrate the somewhat confusing nature of IS-IS verification and troubleshooting. There is no clear pattern as to whether incorporation of the keyword **isis** or **clns** in a **show** command applies to IP routing or to OSI routing.

Step 4: Converting to the IS-IS Backbone

L1 routers communicate with other L1 routers in the same area, while L2 routers route between L1 areas, forming an interdomain routing backbone. This lab scenario does not illustrate the typical multi-area composition of the set of L2 routers in an IS-IS domain, because the routers all reside in Area 49.0001. Since the main function of the San Jose routers is to route between areas in the ITA internetwork, they should be configured as L2-only routers as follows:


```
R1(config)# router isis
R1(config-router)# is-type level-2-only

R2(config)# router isis
R2(config-router)# is-type level-2-only

R3(config)# router isis
R3(config-router)# is-type level-2-only
```

To see the effect of the **is-type** command, reenter the previous commands: **show ip protocols**, **show clns neighbors**, **show isis database**, **show clns interface fastethernet 0/0**, **show isis database R1.00-00 detail**, **show isis topology**, and **show ip route**. Here are the sample outputs:

```
R1# show ip protocols
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    Loopback0
    FastEthernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.30.1         115      00:08:48
    192.168.20.1         115      00:00:09
  Distance: (default is 115)

R1# show clns neighbors

System Id      Interface    SNPA                State  Holdtime   Type Protocol
R2             Fa0/0        0004.9ad2.d0c0      Up     26         L2   IS-IS
R3             Fa0/0        0002.16f4.1ba0      Up     22         L2   IS-IS

R1# show isis database

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000001   0x623C        1086              0/0/0
R1.01-00            * 0x0000000F   0x3344        1092              0/0/0
R2.00-00              0x00000001   0x13AA        1091              0/0/0
R3.00-00              0x00000002   0xD703        1096              0/0/0
```

If the LSP ID is seen with an LSP Holdtime of 0 followed by a parenthetical value, that rogue entry can be purged with the **clear isis *** command.

```
R1# show clns interface fastethernet 0/0
FastEthernet0/0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 16 seconds
  Routing Protocol: IS-IS
```



```
    Circuit Type: level-1-2
    DR ID: R1.02
    Level-2 IPv6 Metric: 10
    Interface number 0x0, local circuit ID 0x1
    Level-2 Metric: 10, Priority: 100, Circuit ID: R1.01
    Number of active level-2 adjacencies: 2
    Next IS-IS LAN Level-2 Hello in 2 seconds
```

Even though the Circuit Type is level-1-2, the entries following the Circuit Type show that only L2 operations are taking place.

```
R1# show isis database R1.00-00 detail

IS-IS Level-2 LSP R1.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000001   0x623C        892               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: R1
  IP Address:   192.168.10.1
  Metric: 10         IS R1.02
  Metric: 10         IS R1.01
  Metric: 10         IP 192.168.10.0 255.255.255.0
  Metric: 10         IP 172.16.0.0 255.255.255.0
```

The output shows that the IDs, R1.02 and R.01, are used to number the router interfaces participating in IS-IS. This is also seen in the **show clns interface** output.

```
R1# show isis topology

IS-IS paths to level-2 routers
System Id     Metric  Next-Hop          Interface   SNPA
R1            --
R2            10      R2                Fa0/0       0004.9ad2.d0c0
R3            10      R3                Fa0/0       0002.16f4.1ba0

R1# show ip route
<output omitted>

Gateway of last resort is not set

i L2 192.168.30.0/24 [115/20] via 172.16.0.3, FastEthernet0/0
C    192.168.10.0/24 is directly connected, Loopback0
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.0.0 is directly connected, FastEthernet0/0
i L2 192.168.20.0/24 [115/20] via 172.16.0.2, FastEthernet0/0
```

What types of routes are being placed into the routing table?

Step 5: Manipulating the IS-IS Interface Timers

The default value of the hello interval is 10 seconds, and the default value of the hello multiplier is 3. The hello multiplier specifies the number of IS-IS hello

PDUs a neighbor must miss before the router declares the adjacency as down. With the default hello interval of 10 seconds, it takes 30 seconds for an adjacency to be declared down due to missed hello PDUs. The analogous OSPF settings are controlled by the **ip ospf hello-interval** and **ip ospf dead-interval** interface commands.

A decision is made to adjust the IS-IS timers so that the core routers detect network failures in less time. This will increase traffic, but this is much less of a concern on the high-speed core Ethernet segment than on a busy WAN link. It is determined that the need for quick convergence on the core outweighs the negative effect of extra control traffic. Change the hello interval to 5 on all FastEthernet 0/0 interfaces, as shown below for the R1 router:

```
R1(config)# interface fastethernet 0/0
R1(config-if)# isis hello-interval 5
```

3. How long will it take for an adjacency to be declared down with the new hello interval of 5?

Step 6: Implementing IS-IS L2 Core Authentication

There should not be any unauthorized routers forming adjacencies within the IS-IS core. Adding authentication to each IS-IS enabled interface can help to ensure this.

Configure interface authentication on R1:

```
R1(config)# interface FastEthernet 0/0
R1(config-if)# isis password cisco level-2
```

This command prevents unauthorized routers from forming level-2 adjacencies with this router.

Important: Be sure to add the keyword **level-2**, which refers to the level-2 database, not an encryption level. If you do not specify a keyword, the default is level-1. Keep in mind that the passwords are exchanged in clear text and provide only limited security.

Wait 20 seconds and then issue the **show clns neighbors** command on R1.

4. Does R1 still show that it has IS-IS neighbors? Why or why not?

Issue the **debug isis adj-packets** command to verify that R1 does not recognize its neighbors, because it requires authentication that has not been configured on R2 and R3 yet.

```
R1# debug isis adj-packets
IS-IS Adjacency related packets debugging is on
03:22:28: ISIS-Adj: Sending L2 LAN IIH on FastEthernet0/0, length 1497
03:22:29: ISIS-Adj: Sending L2 LAN IIH on Loopback0, length 1514
03:22:30: ISIS-Adj: Sending L2 LAN IIH on FastEthernet0/0, length 1497
03:22:31: ISIS-Adj: Rec L2 IIH from 0004.9ad2.d0c0 (FastEthernet0/0), cir type
L2, cir id 1111.1111.1111.01, length 1497
03:22:31: ISIS-Adj: Authentication failed
```

IS-IS routers do not communicate unless the authentication parameters match. However, many other interface-specific IS-IS parameters can vary on a given segment without disrupting communication, such as those set by the commands **isis hello-interval, isis hello-multiplier, isis retransmit-interval, isis retransmit-throttle-interval,** and **isis csnp-interval**. Of course, it makes sense for these parameters to coincide on a given segment.

Correct the authentication mismatch by configuring interface authentication on R2 and R3. After the configurations are complete, verify that the routers can communicate by using the **show clns neighbors** command on R1.

```
R2(config)# interface FastEthernet 0/0
R2(config-if)# isis password cisco level-2

R3(config)# interface FastEthernet 0/0
R3(config-if)# isis password cisco level-2

R1# show clns neighbors

System Id Interface   SNPA                State  Holdtime   Type Protocol
R2        Fa0/0       0004.9ad2.d0c0      Up     23         L2   IS-IS
R3        Fa0/0       0002.16f4.1ba0      Up     26         L2   IS-IS
```

In time, the system IDs resolve to the router names. This is done through the dynamic hostname mapping feature automatically enabled on Cisco routers. In the interim, the output may appear with the actual numerical ID for that system.

Step 7: Implementing IS-IS Domain Authentication

IS-IS provides two additional layers of authentication, area passwords for L1 and domain passwords for L2, to prevent unauthorized adjacencies between routers. The interface, area, and domain password options all use plain text authentication and, therefore, are of limited use. However, beginning with Cisco IOS Release 12.2(13)T, MD5 authentication is available for IS-IS.

The command for L1 password authentication is **area-password** *password*. Using this command on all routers in an area prevents unauthorized routers from injecting false routing information into the L1 database.

The command for L2 password authentication is **domain-password** *password*. Using this command on all L2 routers in a domain prevents unauthorized routers from injecting false routing information into the L2 database. Since the core routers are operating at L2, implement domain password authentication as follows:

```
R1(config)# router isis
R1(config-router)# domain-password cisco
```

The password is case-sensitive. Time permitting, intentionally configure mismatched interface passwords. Do the same for area, and domain passwords. By seeing the way in which the router responds, it will be easier for you to spot this error when you unintentionally mismatch passwords in a production network.

Refresh the IS-IS link-state database and recalculate all routes using the **clear isis *** command on all routers. It may take a minute or two for all routers to update their databases.

```
All_Router# clear isis *
```

Use the **show isis database** command to view the changes to the R1 link-state database:

```
R1# show isis database

IS-IS Level-2 Link State Database:
LSPID            LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00       * 0x00000004   0xDCB5        1155              0/0/0
R1.01-00       * 0x00000007   0xB4C1        1156              0/0/0
```

Change the other routers to reflect the new authentication policy:

```
R2(config)# router isis
R2(config-router)# domain-password cisco

R3(config)# router isis
R3(config-router)# domain-password cisco
```

View the R1 link-state database to verify that the LSPs were propagated:

```
R1# show isis database

IS-IS Level-2 Link State Database:
LSPID            LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00       * 0x00000001   0xE2B2        1189              0/0/0
R1.01-00       * 0x00000002   0xBEBC        1195              0/0/0
R2.00-00         0x00000002   0x5A59        1190              0/0/0
R3.00-00         0x00000002   0xF3DD        1185              0/0/0
```

The configuration of basic Integrated IS-IS routing protocol is now complete. In addition to enabling Integrated IS-IS, L2-specific routing was enabled, and the hello interval was changed to enable IS-IS to detect network failures faster. Two types of password authentication, interface and domain, were enabled to prevent unauthorized routers from forming adjacencies with these core routers.

Run the TCL script to verify full connectivity after implementing L2 authentication:

```
foreach address {
192.168.10.1
172.16.0.1
192.168.20.1
172.16.0.2
192.168.30.1
172.16.0.3 } { ping $address }
```

Save the R1 and R2 configurations for use with the next lab.

Appendix A: TCL Script Output

```
R1# tclsh
R1(tcl)#foreach address {
+>(tcl)#192.168.10.1
+>(tcl)#172.16.0.1
+>(tcl)#192.168.20.1
+>(tcl)#172.16.0.2
+>(tcl)#192.168.30.1
+>(tcl)#172.16.0.3 } { ping $address }

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.30.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1(tcl)# tclquit

R2# tclsh
R2(tcl)#foreach address {
+>(tcl)#192.168.10.1
+>(tcl)#172.16.0.1
+>(tcl)#192.168.20.1
+>(tcl)#172.16.0.2
+>(tcl)#192.168.30.1
+>(tcl)#172.16.0.3 } { ping $address }
```



```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.30.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
R2(tcl)# tclquit

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.30.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
R3(tcl)# tclquit
```

Final Configurations

```
R1# show run
Building configuration...

Current configuration : 1290 bytes
!
version 12.4
!
hostname R1
!
```



```
interface Loopback0
 ip address 192.168.10.1 255.255.255.0
 ip router isis
!
interface FastEthernet0/0
 ip address 172.16.0.1 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 isis password cisco level-2
 isis priority 100
 isis hello-interval 5
!
router isis
 net 49.0001.1111.1111.1111.00
 is-type level-2-only
 domain-password cisco
!
end


R2# show run
Building configuration...

Current configuration : 1044 bytes
!
version 12.4
!
hostname R2
!
interface Loopback0
 ip address 192.168.20.1 255.255.255.0
 ip router isis
!
interface FastEthernet0/0
 ip address 172.16.0.2 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 isis password cisco level-2
 isis priority 100
 isis hello-interval 5
!
router isis
 net 49.0001.2222.2222.2222.00
 is-type level-2-only
 domain-password cisco
!
end

R3# show run
Building configuration...

Current configuration : 1182 bytes
!
version 12.4
!
hostname R3
!
interface Loopback0
 ip address 192.168.30.1 255.255.255.0
 ip router isis
!
```



```
interface FastEthernet0/0
 ip address 172.16.0.3 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 isis password cisco level-2
 isis priority 100
 isis hello-interval 5
!
router isis
 net 49.0001.3333.3333.3333.00
 is-type level-2-only
 domain-password cisco
!
end
```

