Cisco Certified Network Professional Enterprise (CCNP 2025)

Exam 300-410 Enterprise Advanced Routing and Services (ENARSI)

Demo Questions

CISCO

CISCO CERTIFIED CCNP

300-410: Cisco Enterprise Advanced Routing and Services

QUESTION 1

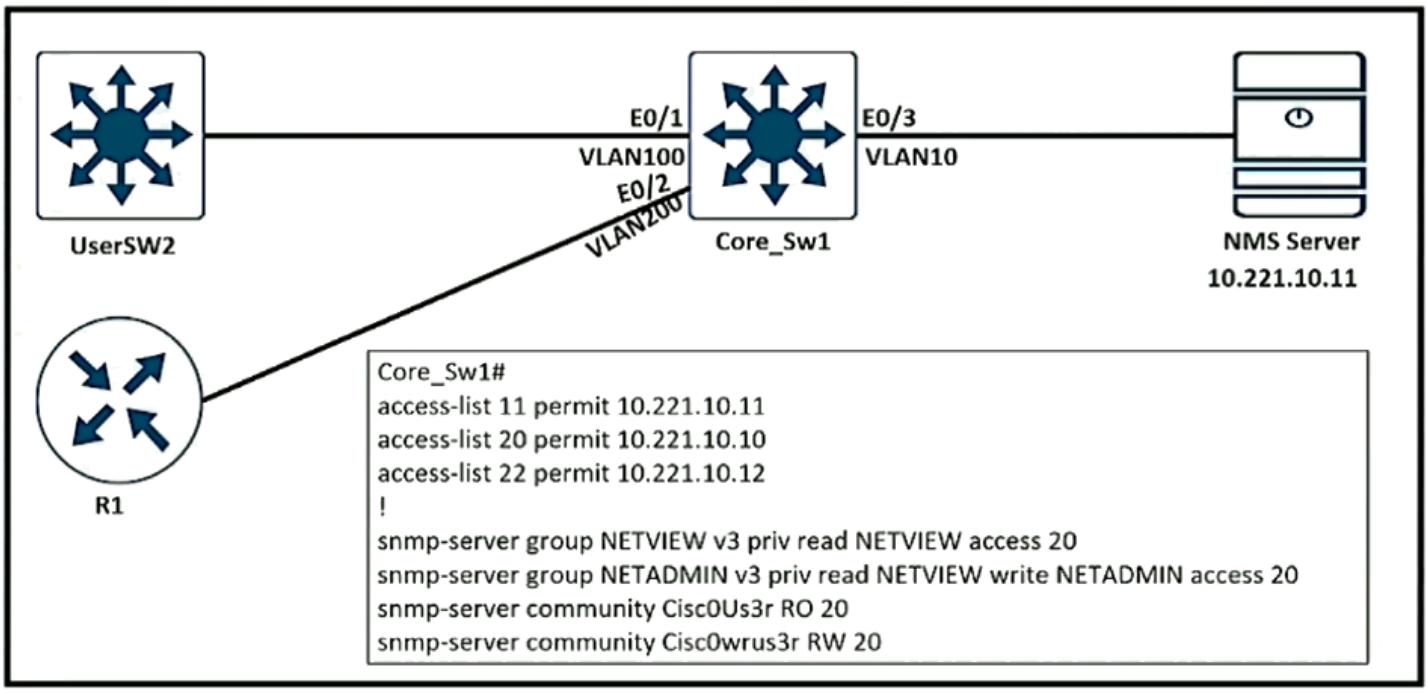

Refer to the exhibit. An engineer configured SNMP communities on the Core_Sw1, but the SNMP server cannot obtain information from Core_Sw1. Which configuration resolves this issue?

A. `snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 22`

B. `snmp-server group NETVIEW v2c priv read NETVIEW access 22`

C. `access-list 20 permit 10.221.10.11`

D. `access-list 20 permit 10.221.10.12`

Answer: C

- **The Problem:** The NMS Server, which has the IP address 10.221.10.11, is unable to communicate with Core_Sw1.
- **The Cause:** Looking at the configuration on Core_Sw1, all active SNMP commands (both the groups and the communities) are restricted by access-list 20.
- **The Fix:** Currently, access-list 20 only permits the IP 10.221.10.10. To allow the NMS Server to poll the switch, its IP address (10.221.10.11) must be added to that ACL.

Breakdown of the ACLs:

- **Access-list 11:** Permits 10.221.10.11 (but is not applied to the SNMP configuration).
- **Access-list 20:** Permits 10.221.10.10 (**Applied to SNMP**; prevents the NMS Server from connecting).
- **Access-list 22:** Permits 10.221.10.12 (not used by the SNMP community strings).

By applying `access-list 20 permit 10.221.10.11`, you grant the NMS Server the necessary permissions to access the SNMP services on the switch.

QUESTION 2

```
R2# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address         Interface
192.168.99.2       1   EXCHANGE/  -    00:00:36    192.168.99.1    Serial0/1
router-6#

R3# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address         Interface
192.168.99.1       1   EXSTART/   -    00:00:33    192.168.99.2    Serial0/1
```

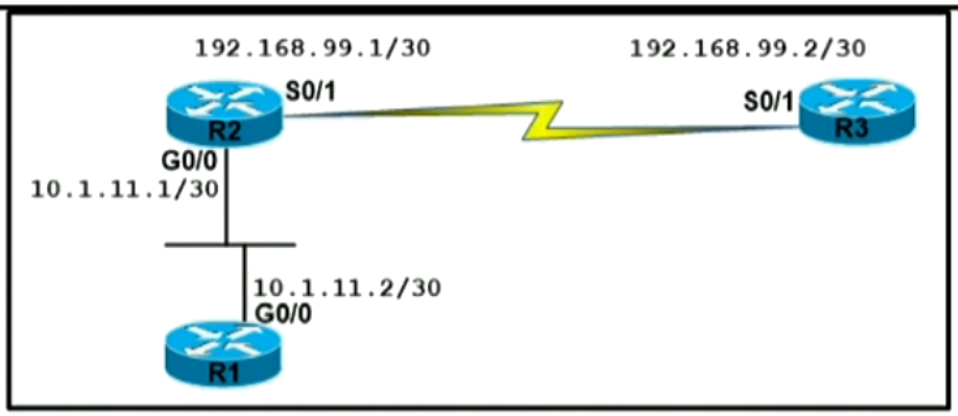

Refer to the exhibit. An OSPF neighbor relationship between R2 and R3 is showing stuck in EXCHANGE/EXSTART state. The neighbor is established between R1 and R2. The network engineer can ping from R2 to R3 and vice versa, but the neighbor is still down. Which action resolves the issue?

A. Enable OSPF on the interface, which is required.
B. Restore the Layer2/Layer3 connectivity issue in the ISP network.
C. Match MTU on both router interfaces or ignore MTU.
D. Administrative "shutdown then no shut" both router interfaces.

Answer: C

- **The Symptom:** The OSPF neighbor state is stuck in **EXSTART/EXCHANGE**. This is the classic indicator of an **MTU mismatch**.
- **The Cause:** During the **EXSTART/EXCHANGE** phase, routers exchange Database Descriptor (DBD) packets. These packets are often large. If the Maximum Transmission Unit (MTU) on one router's interface is larger than the other, the router with the smaller MTU will drop the incoming DBD packet. Because the routers cannot successfully exchange their database summaries, they remain stuck in this state.
- **The Fix:** You must ensure the MTU matches on both serial interfaces. Alternatively, you can use the interface-level command **`ip ospf mtu-ignore`** to tell OSPF to bypass the MTU check.

Why other options are incorrect:

Option A: OSPF is clearly enabled, as evidenced by the **`show ip ospf neighbor`** output showing the neighbors in the list.

Option B: The engineer can already ping between R2 and R3, meaning Layer 2 and Layer 3 connectivity is functional.

Option D: While "shut/no shut" restarts the process, it will not fix the underlying parameter mismatch; the routers will simply get stuck in EXSTART again.

QUESTION 3

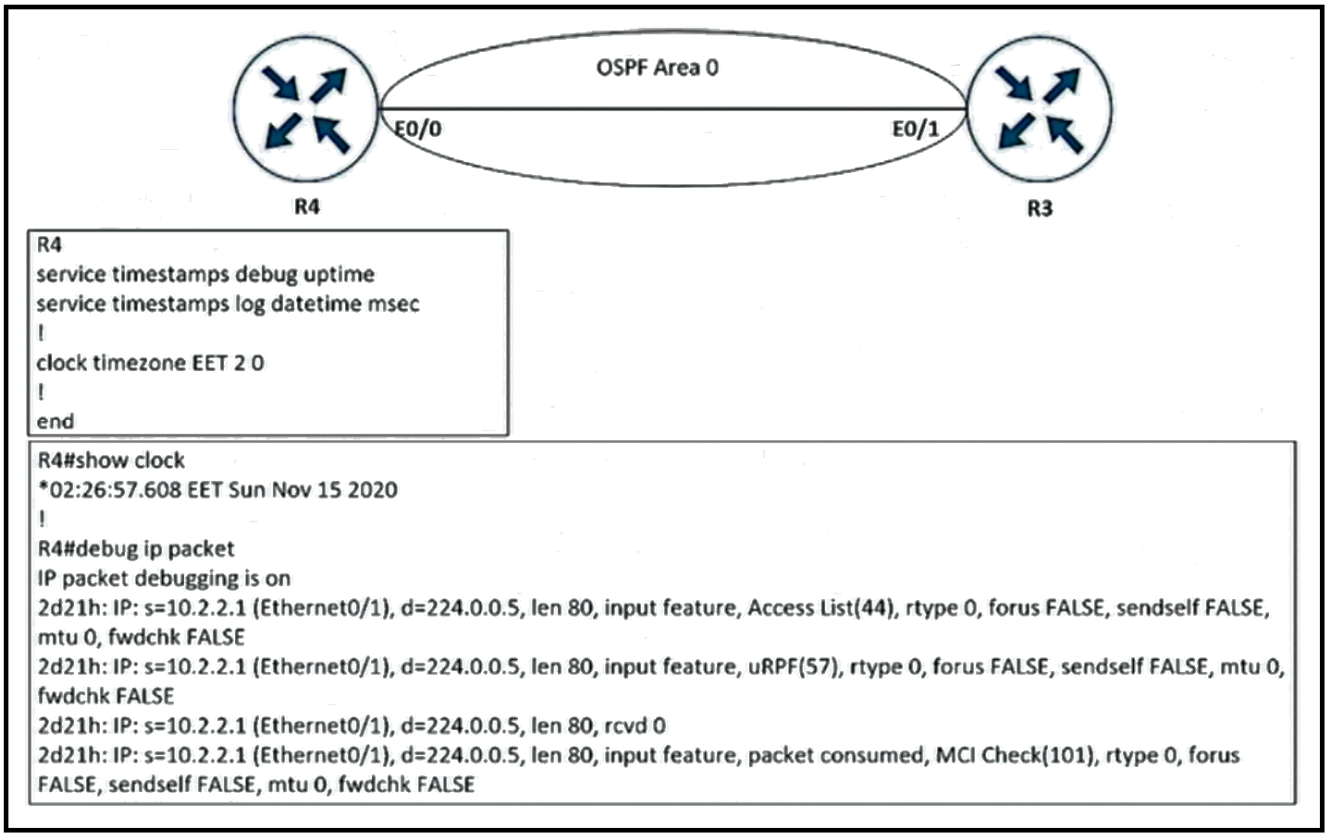

Refer to the exhibit. An engineer is troubleshooting an issue using the debug ip packet command and notices that no time stamps are shown on R4 to establish the event time. Which configuration resolves this issue by showing time stamps regardless of the time zone in R4 logs?

A. `service timestamps debug datetime localtime msec`

B. `service timestamps log datetime msec show-timezone`

C. `service timestamps log datetime localtime msec`

D. `service timestamps debug datetime msec show-timezone`

Answer: D

- **The Problem:** The user is running a debug command (debug ip packet), but the output is showing "uptime" (e.g., 2d21h) instead of a full date and time.
- **The Requirement:** The timestamps must show the actual date/time for debug logs, specifically including milliseconds (msec) and the time zone (show-timezone).
- **The Solution:**
  - `service timestamps debug:` This is required because the issue is with debug output, not standard system logs.
  - `datetime:` Switches the format from uptime to a calendar clock.
  - `msec:` Provides the requested millisecond precision.
  - `show-timezone:` Ensures the time zone is displayed as requested.

Comparison:

- **Options B and C:** These only affect log messages (syslog), not the output of debug commands.
- **Option A:** This uses localtime, which would rely on the locally configured time zone but doesn't explicitly meet the "regardless of the time zone" requirement as effectively as showing the zone name itself. Option D specifically includes the show-timezone parameter to provide that clarity in the logs.

QUESTION 4

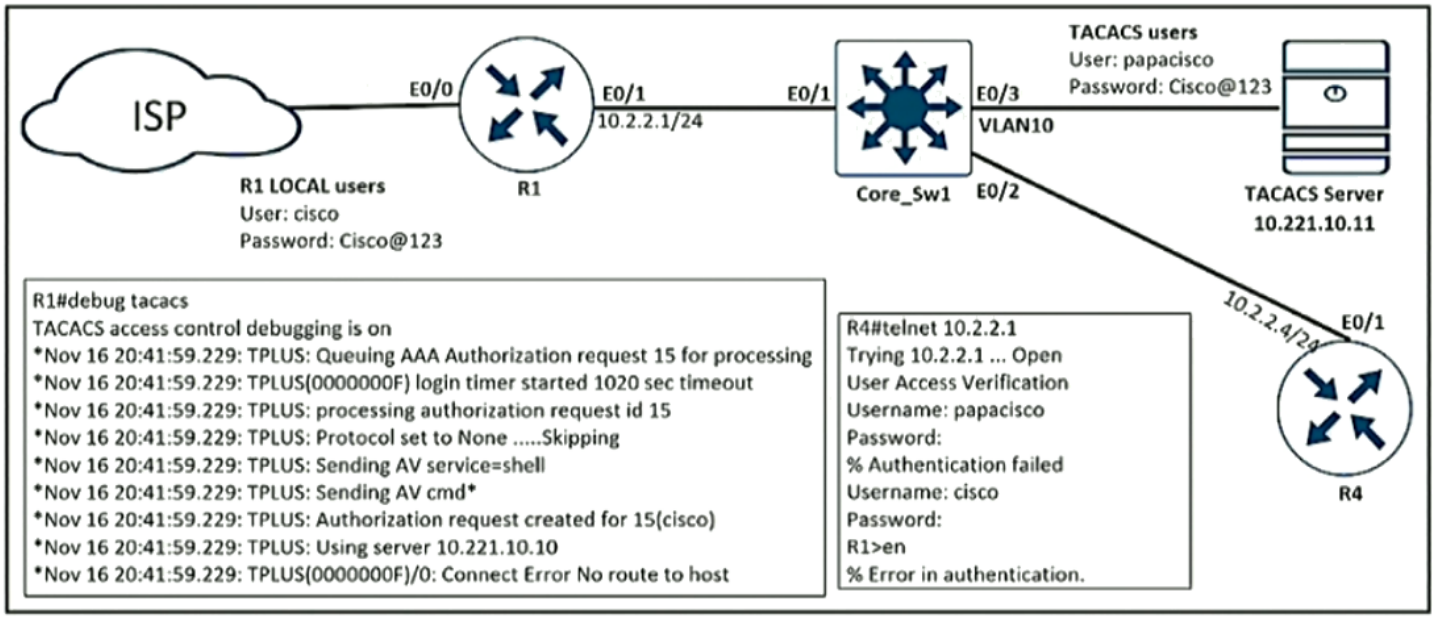

Refer to the exhibit. A network administrator is trying to switch to the privileged EXEC level on R1 but failed. Which configuration resolves the issue?

A. `enable password Cisco@123`

B. `enable-password Cisco@123`

C. `tacacs server enable-password Cisco@123`

D. `tacacs-server enable-password Cisco@123`

Answer: A

- **The Problem:** The user `cisco` (a local user) successfully logged into R1 but received a "% Error in authentication" message when trying to enter privileged EXEC mode (the en command).
- **The Cause:** Looking at the `debug tacacs` output, the router is attempting to use a TACACS+ server for authorization/authentication, but it fails with "Connect Error No route to host" (it is trying to reach **10.221.10.10**, while the server is actually at **10.221.10.11**).
- **The Fallback:** When the external server is unreachable, the router falls back to its local "enable" configuration. If no enable password or secret is configured on the router, the `en` command will fail because there is no local credential to verify against.
- **The Fix:** Configuring a local enable password using `enable password Cisco@123` ensures that even if the TACACS server is offline or unreachable, the administrator can still access privileged mode using a locally defined password.

Why other options are incorrect:

Option B: `enable-password` (with a hyphen) is not a valid Cisco IOS command.

Options C & D: These are not valid commands for configuring the local enable password or for correctly pointing to a TACACS server.

QUESTION 5

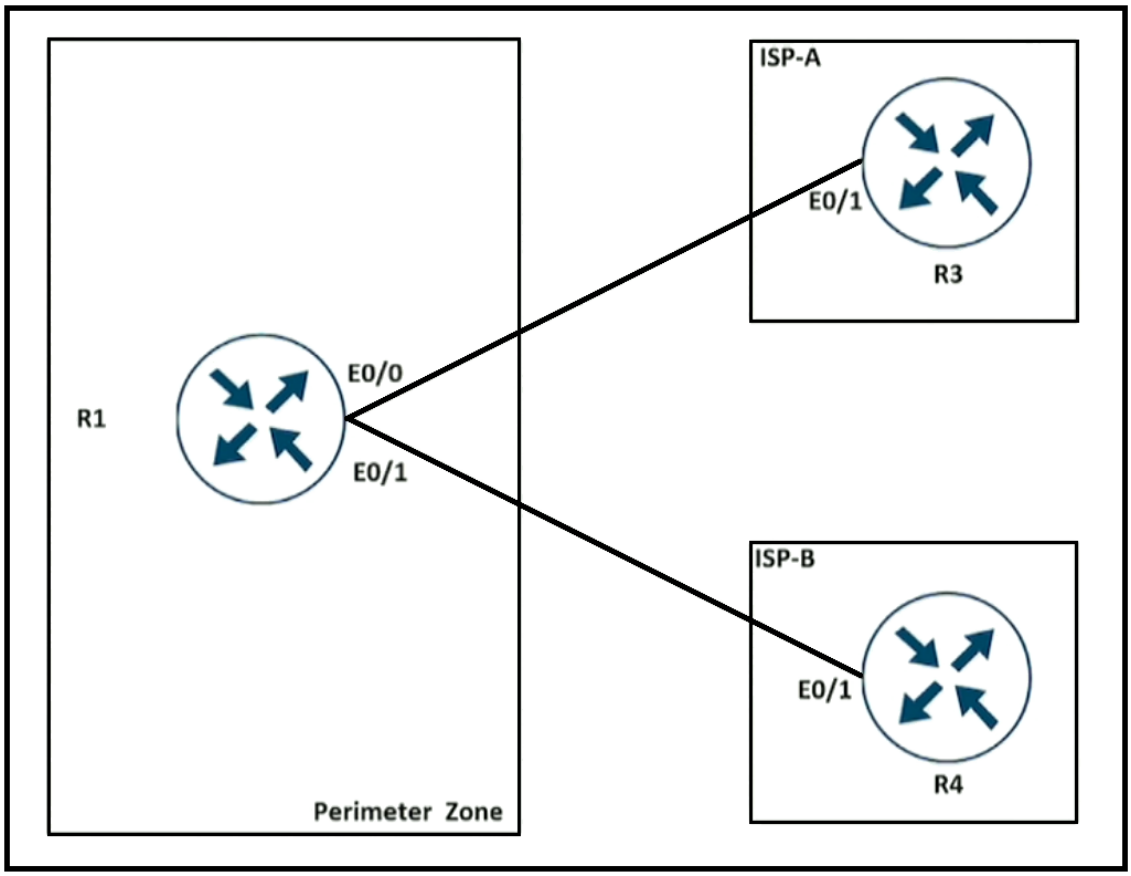

Refer to the exhibit. A network is under a cyberattack. A network engineer connected to R1 by SSH and enabled the terminal monitor via SSH session to find the source and destination of the attack. The session was flooded with messages, which made it impossible for the engineer to troubleshoot the issue. Which command resolves this issue on R1?

A. `terminal no monitor`
B. `no terminal monitor`
C. `(config)# terminal no monitor`
D. `(config)# no terminal monitor`

Answer: A

- **The Problem:** The engineer used terminal monitor to see debug messages and logs during an SSH session. Under a cyberattack, the volume of these messages is too high, "flooding" the screen and making it impossible to type or read.
- **The Command:** `terminal monitor` is a Privileged EXEC command. It is session-specific, meaning it only affects the current SSH or Telnet window you are using.
- **The Fix:** To stop the messages, you must use the negation of that command at the same EXEC prompt. The correct Cisco IOS syntax to disable this is `terminal no monitor.`

Why the others are incorrect:

Option B: `no terminal monitor` is incorrect syntax for this specific feature.

Options C & D: These are shown at the (config)# (Global Configuration) prompt. Because `terminal monitor` is a temporary session command and not a permanent router configuration, it cannot be executed or disabled from configuration mode.

QUESTION 6

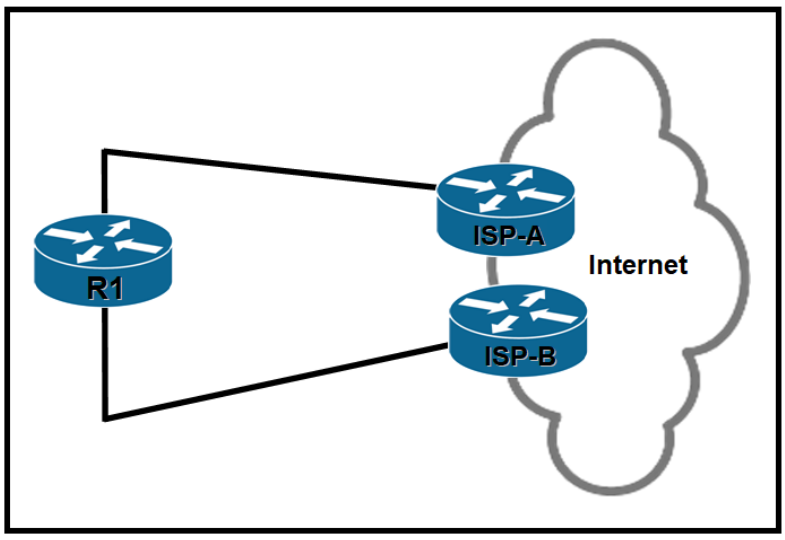

Refer to the exhibit. Router R1 peers with two ISPs using static routes to get to the internet. The requirement is that R1 must prefer ISP-A under normal circumstances and failover to ISP-B if the connectivity to ISP-A is lost. The engineer observes that R1 is load balancing traffic across the two ISPs. Which action resolves the issue but sending traffic to ISP-A only with failover to ISP-B?

A. Configure two static routes on R1, one pointing to ISP-B with more specific routes and another pointing to ISP-A with summary routes.
B. Configure two static routes on R1, one pointing to ISP-A and another point to ISP-B with 222 admin distance.
C. Change the bandwidth of the interface on R1 so that interface to ISP-A has a higher value than the interface to ISP-B.
D. Configure OSPF between R1, ISP-A, and ISP-B for dynamic failover if any ISP link to R1 fails.

Answer: B

Why Option B is the Solution:

The current issue is that R1 is load balancing because both static routes have the same default Administrative Distance (AD) of 1. To create a primary/backup scenario (Floating Static Route), you must make one route less "desirable" than the other.

- **Primary Path:** By leaving the static route to ISP-A with the default AD of 1, it remains the preferred path in the routing table.
- **Backup Path:** By configuring the static route to ISP-B with a higher AD (like 222), the router will "hide" this route as long as the ISP-A path is active.
- **Failover:** If the interface to ISP-A goes down, the primary route is removed, and the router automatically installs the backup route (AD 222) from the configuration into the routing table.

Why other options are incorrect:

Option A: More specific routes always beat summary routes. This would make ISP-B the primary, which is the opposite of the requirement.

Option C: Static routes do not look at interface bandwidth to make decisions; they only care about Administrative Distance.

Option D: While OSPF can handle failover, the requirement specifically mentions using static routes. Furthermore, you usually cannot run an IGP like OSPF directly with an ISP for internet edge routing.

QUESTION 7

```
R2(config)#router ospf 1
R2(config-router)#network 10.101.1.0 0.0.0.0 area 1
R2(config-router)#end

%OSPF-4-ERRRCV: Received invalid packet: mismatched area ID from backbone area
from 10.101.1.1, GigabitEthernet0/0
```

Refer to the exhibit. Which solution ensures that messages have local device time mentioned?

A. `service timestamps log uptime`

B. `service timestamps log datetime`

C. `service timestamps debug datetime`

D. `service timestamps debug uptime`

Answer: B

Explanation

- `log uptime` → shows how long the device has been running, not the actual clock time.
- `debug uptime` → same idea, but for debug messages.
- `debug datetime` → applies to debug messages only, not logs.
- `log datetime` → ensures that log messages include the local device's date and time, which is exactly what the QUESTION asks.

So the correct solution is `service timestamps log datetime` because it ensures log messages are stamped with the router's local time.

QUESTION 8

Which Layer 3 VPN attribute installs customer routes in the VRF?

A. RT

B. RD

C. MPLS label

D. extended-community

Answer: A

- **Route Target (RT):** This is a BGP extended community that functions as a "tag" to define the import and export policies. When a PE router receives a BGP VPNv4 route, it looks at the RT value. If the RT matches an "import" statement in a local VRF, the route is installed in that VRF's routing table.
- **Route Distinguisher (RD):** While often confused with the RT, the RD's only job is to make a standard IPv4 address unique within the MPLS core (creating a VPNv4 address). It does not determine which VRF a route belongs to on a receiving router.
- **MPLS Label:** This is used for data plane forwarding to reach the correct egress interface/VRF, but it is not used to "install" the route in the control plane.
- **Extended-community:** This is a broad category of BGP attributes. While the Route Target is an extended community, the RT is the specific attribute within that category that performs the installation logic.

QUESTION 9

Which method provides failure detection in BFD?

A. short duration, high overhead
B. short duration, low overhead
C. long duration, high overhead
D. long duration, low overhead

Answer: B

- **Short Duration:** BFD (Bidirectional Forwarding Detection) is designed for sub-second failure detection. It can detect a link or neighbor failure in milliseconds, which is much faster than the default timers of routing protocols like OSPF or BGP.
- **Low Overhead:** BFD uses very small, fixed-length "Hello" style control packets. Because it is a simple "keepalive" mechanism that can often be offloaded to the hardware (the data plane) rather than the CPU (the control plane), it places very little strain on the router's resources compared to more complex protocols.

QUESTION 10

```
R1#show ipv6 ospf neighbor
R1#debug ipv6 ospf hello
OSPFv3 hello events debugging is on
OSPFv3: Rcv hello from 10.1.1.1 area 10 from FastEthernet0/1 FE80::C801:FFF:FE94:1C
interface ID 4
OSPFv3: Hello from FE80::C801:FFF:FE94:1C with mismatched NSSA option bit
OSPFv3: Send hello to FF02::5 area 10 on FastEthernet0/1 from FE80::C004:22FF:FE78:1
interface ID 5
```

Look at the exhibit, R1 is not forming an adjacency on a point to point interface, what can be determined from the debug output and what corrective action is required?

Refer to the exhibit. R1 is not forming adjacency on a point-to-point interface. Which action resolves the issue?

A. The area numbers must be configured the same on each router.
B. The area types must be configured the same on each router.
C. The **`no-summary`** command must be included in the area configuration on R1.
D. The **`no-summary`** command must be included in the area configuration on R2.

Answer: B

- **The Problem:** The debug output explicitly states: OSPFv3: Hello from FE80::C801:FFF:FE94:1C with mismatched NSSA option bit.
- **The Cause:** For OSPF neighbors to form an adjacency, certain parameters in the Hello packet must match. One of these is the **Area Type** (e.g., Normal, Stub, or NSSA). The "mismatched NSSA option bit" indicates that one router is configured for a Not-So-Stubby Area (NSSA) while the other is not.
- **The Fix:** You must ensure the area types are configured identically on both routers.

Options C and D refer to the **`no-summary`** command, which is used to create a Totally Stubby or Totally NSSA area. While this affects the type of area, it doesn't address the fundamental mismatch shown in the debug—both sides must agree on whether the area is an NSSA or a standard area first.

QUESTION 11

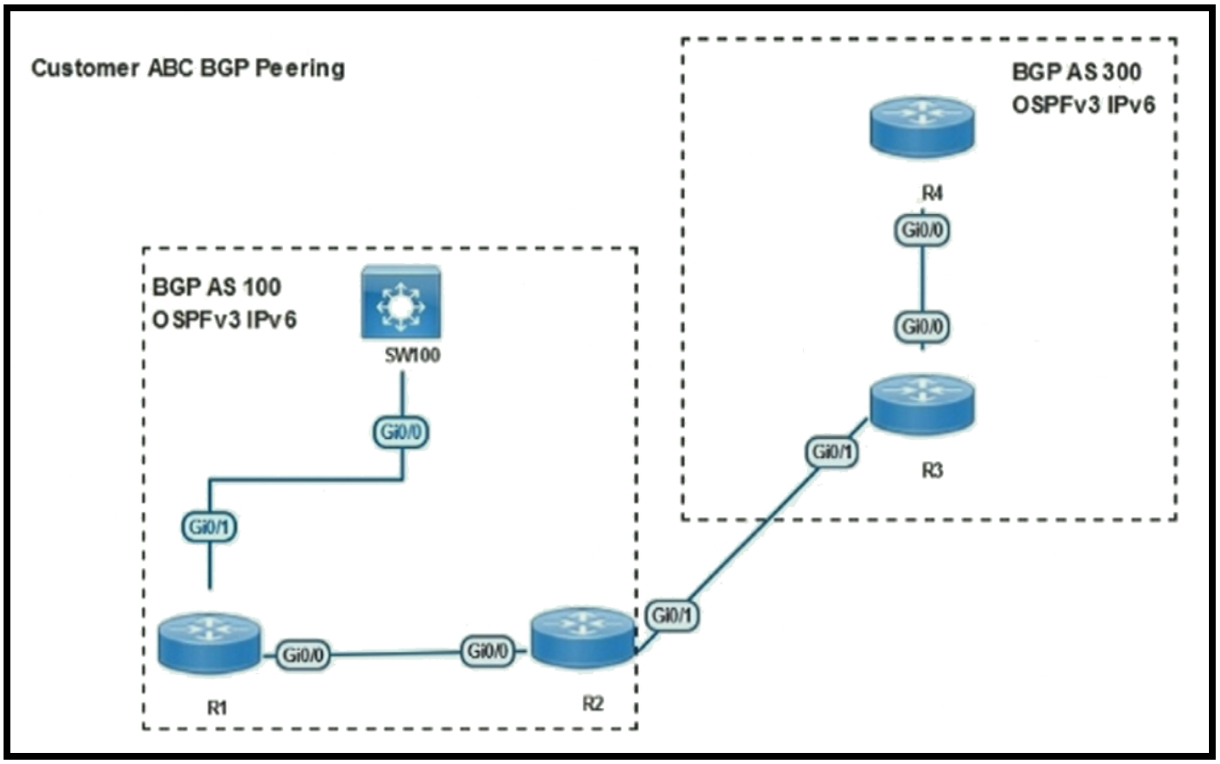


```
SW100#sh ip bgp ipv6 uni summ
BGP router identifier 100.0.0.1, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V      AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:ABC:AABB:1100:1122:1111:2222:AAA1
                4      100       6       5        1    0    0 00:00:58            0

SW100#sh ip bgp ipv6 unicast
SW100#

R1#sh ip bgp ipv6 uni
BGP table version is 4, local router ID is 1.1.1.1
      Network          Next Hop      Metric LocPrf Weight Path
 * i  2001::4/128      2001::4            0    100      0 300 i
 *>i  2002::2/128      2001::2            0    100      0 i
R1#
R1#sh ipv6 route
O   2001::2/128 [110/1]
     via FE80::5200:C3FF:FE01:E600, GigabitEthernet0/0
B   2002::2/128 [200/0]
     via 2001::2
```

Refer to the exhibit. SW100 cannot receive routes from R1. Which configuration resolves the issue?

A. R1

```
router bgp 100
 address-family ipv6
  neighbor 2001::2 route-reflector-client
  neighbor 2001::ABC:AABB:1100:1122:1111:2222:AAA2 route-reflector-client
```

R2

```
router bgp 100
 address-family ipv6
  neighbor 2001::2
  neighbor 2001::1 as-override
```

B. R1

```
router bgp 100
 address-family ipv6
  neighbor 2001::2 route-reflector-client
  neighbor 2001::ABC:AABB:1100:1122:1111:2222:AAA2 route-reflector-client
```

R2

```
router bgp 100
 address-family ipv6
  neighbor 2001::2
  neighbor 2001::1 next-hop-self
```

C. R1

```
router bgp 100
 address-family ipv6
  redistribute connected
```

R2

```
router bgp 100
 address-family ipv6
  redistributed connected
```

D. R1

```
router bgp 100
 address-family ipv6
  no synchronization
```

R2

```
router bgp 100
 address-family ipv6
  no synchronization
SW100
router bgp 100
 address-family ipv6
  no synchronization
```

Answer: B

- SW100 is in AS 100, and it's receiving BGP routes from R3/R4 in AS 300.
- R1 and R2 are in AS 100, and likely form an iBGP mesh.
- In iBGP, the next-hop IP is not changed by default, which can cause issues if the next-hop is not reachable by the receiving router.

Key Fixes in Option B:

- R1 Configuration:
- Marks R2 and SW100 as route-reflector clients, allowing R1 to reflect routes between them
- This avoids the need for a full iBGP mesh.

R2 Configuration:

- Uses next-hop-self for the neighbor (SW100), ensuring that the next-hop IP in advertised routes is R2's own IP, which SW100 can reach.
- This is crucial because SW100 may not have a route to the original next-hop (e.g., R1's interface), causing route installation to fail.

QUESTION 12

```
Router# show tag-switching tdp bindings
(...)
tib entry: 10.10.10.1/32, rev 31
      local binding: tag: 18
      remote binding: tsr: 10.10.10.1:0, tag: imp-null
      remote binding: tsr: 10.10.10.2:0, tag: 18
      remote binding: tsr: 10.10.10.6:0, tag: 21
tib entry: 10.10.10.2/32, rev 22
      local binding: tag: 17
      remote binding: tsr: 10.10.10.2:0, tag: imp-null
      remote binding: tsr: 10.10.10.1:0, tag: 19
      remote binding: tsr: 10.10.10.6:0, tag: 22
```

Refer to the exhibit. What does the imp-null tag represent in the MPLS VPN cloud?

A. Exclude the EXP bit.
B. Include the EXP bit.
C. Impose the label.
D. Pop the label.

Answer: D

- **Mechanism:** In MPLS, **`imp-null`** is used for **Penultimate Hop Popping (PHP).**
- **The Process:** An egress LSR (Label Switch Router) advertises an implicit-null label to its upstream neighbor for a specific prefix it owns. This tells the upstream neighbor to remove (pop) the MPLS label before forwarding the packet to the egress router.
- **The Benefit:** By popping the label one hop early, the egress router receives a pure IP packet and only needs to perform a single lookup in its routing table, rather than performing an MPLS label lookup followed by an IP lookup.

QUESTION 13

Which OSI model is used to insert an MPLS label?

A. between Layer 1 and Layer 2
B. between Layer 2 and Layer 3
C. between Layer 3 and Layer 4
D. between Layer 5 and Layer 6

Answer: B

- **The "Layer 2.5" Protocol:** MPLS is often called a "Layer 2.5" protocol because of its placement. It sits between the Data Link layer (Layer 2, like Ethernet) and the Network layer (Layer 3, like IP).
- **Insertion Process:** When a standard IP packet enters an MPLS network, the router (LER) inserts the MPLS header after the Layer 2 header but before the Layer 3 header.
- **Data Flow:**
  - **Layer 2 Header** (Ethernet MAC)
  - **MPLS Label** (The "Shim" Header)
  - **Layer 3 Header** (IP Packet)
  - **Layer 4-7** Data

QUESTION 143

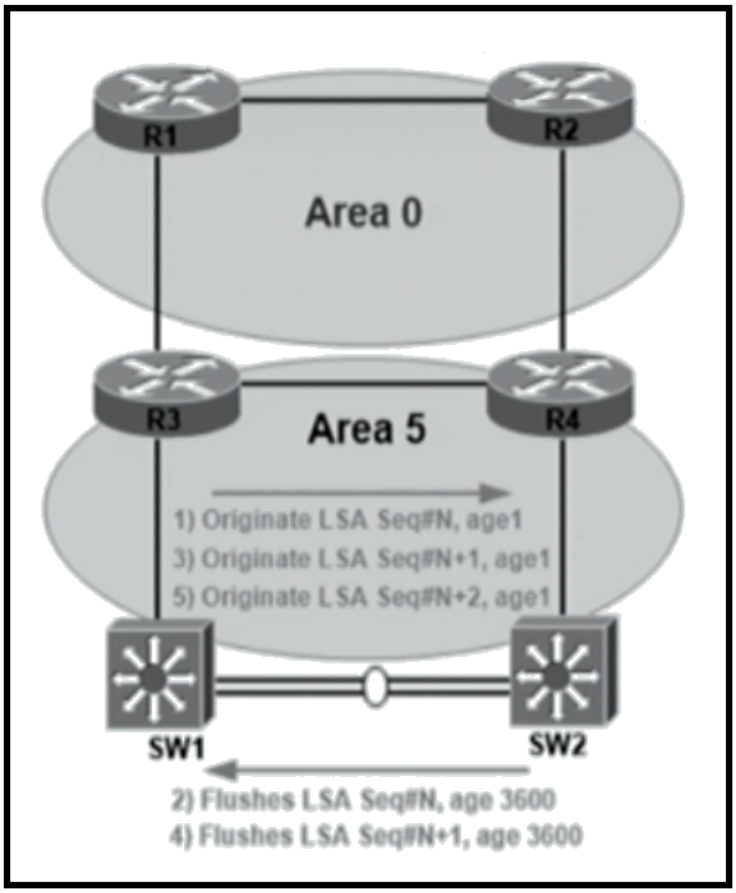

Refer to the exhibit. An error message "OSPF-4-FLOOD_WAR" is received on SW2 from SW1. SW2 is repeatedly receiving its own link-state advertisement and flushes it from the network. Which action resolves the issue?

A. Resolve different subnet mask issue on the link.
B. Resolve duplicate IP address issue in the network.
C. Change area 5 to a normal area from a nonstub area.
D. Configure Layer 3 port channel on interfaces between switches.

Answer: B

Explanation

The error message "OSPF-4-FLOOD_WAR" indicates that SW2 is repeatedly receiving its own Link-State Advertisement (LSA) back from SW1. This typically happens when there is a duplicate IP address configured in the OSPF domain. Because OSPF uses router IDs and interface IPs to identify neighbors, a duplicate IP causes LSAs to loop back, leading to flooding and instability.

- Different subnet mask issue → would cause adjacency problems, but not repeated self-LSA flooding.
- Changing area type (stub vs. normal) → unrelated to LSA flooding.
- Layer 3 port channel → improves redundancy but does not fix duplicate IP conflicts.

Therefore, the correct resolution is to eliminate duplicate IP addresses in the OSPF network.

Drag and Drop Questions

QUESTION 1

Drag and drop the MPLS VPN device types from the left onto the definitions on the right.

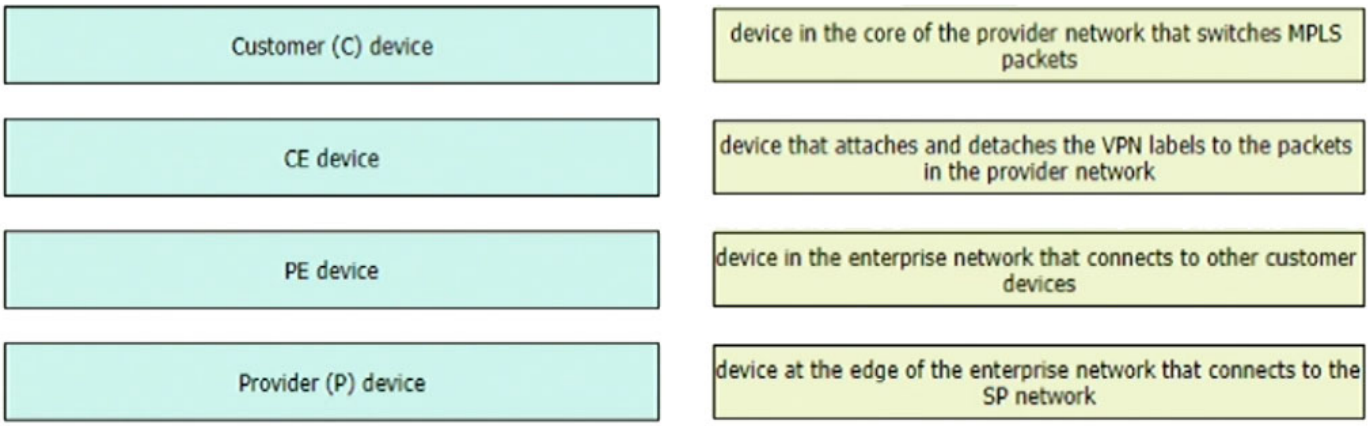

Answer:

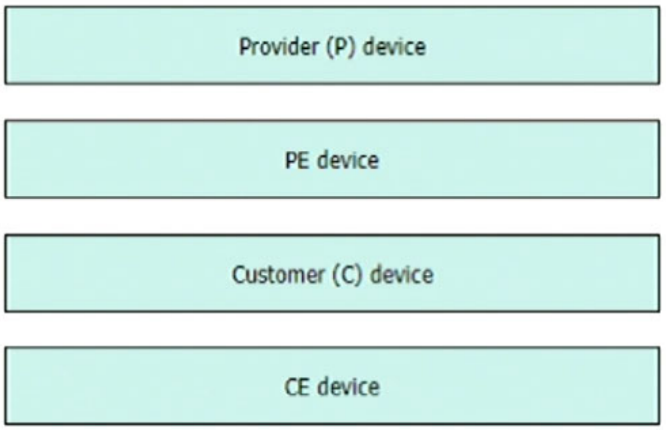

QUESTION 2

Drag and drop the MPLS concepts from the left onto the descriptions on the right.

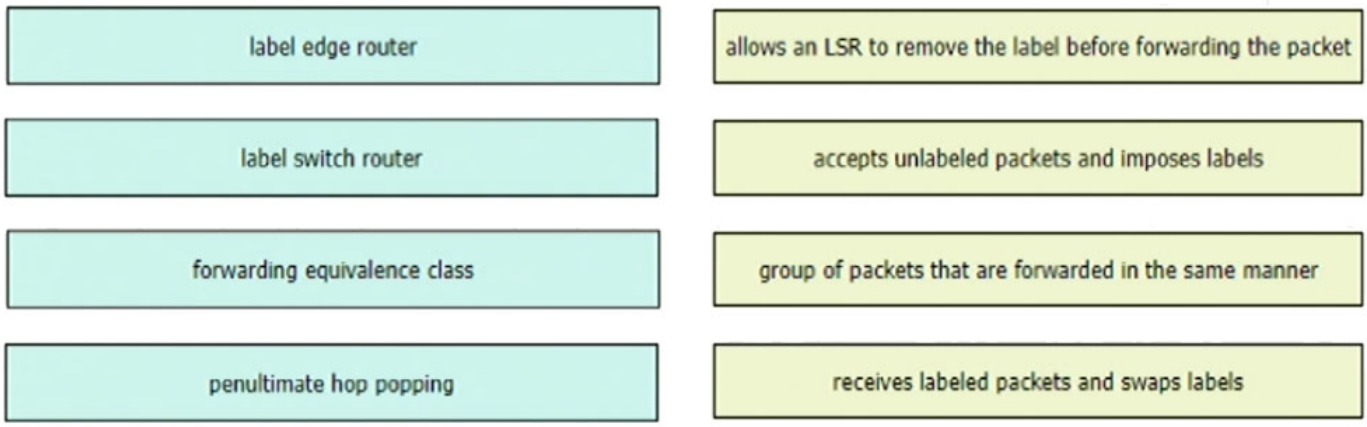

Answer:

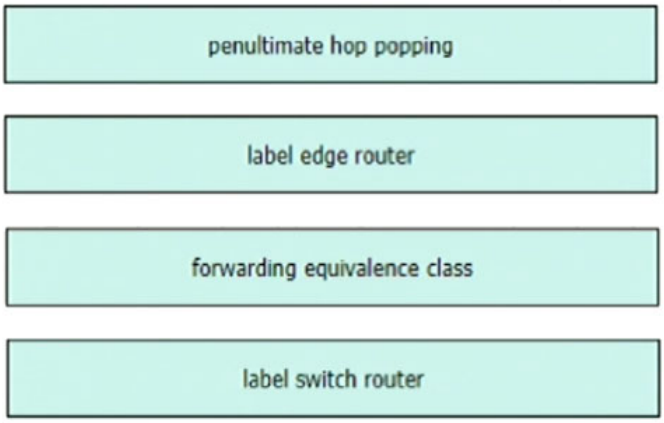

Explanation

A label edge router (LER, also known as edge LSR) is a router that operates at the edge of an MPLS network and acts as the entry and exit points for the network. LERs push an MPLS label onto an incoming packet and pop it off an outgoing packet. A forwarding equivalence class (FEC) is a term.

QUESTION 3

Drag and drop the MPLS terms from the left onto the correct definitions on the right.

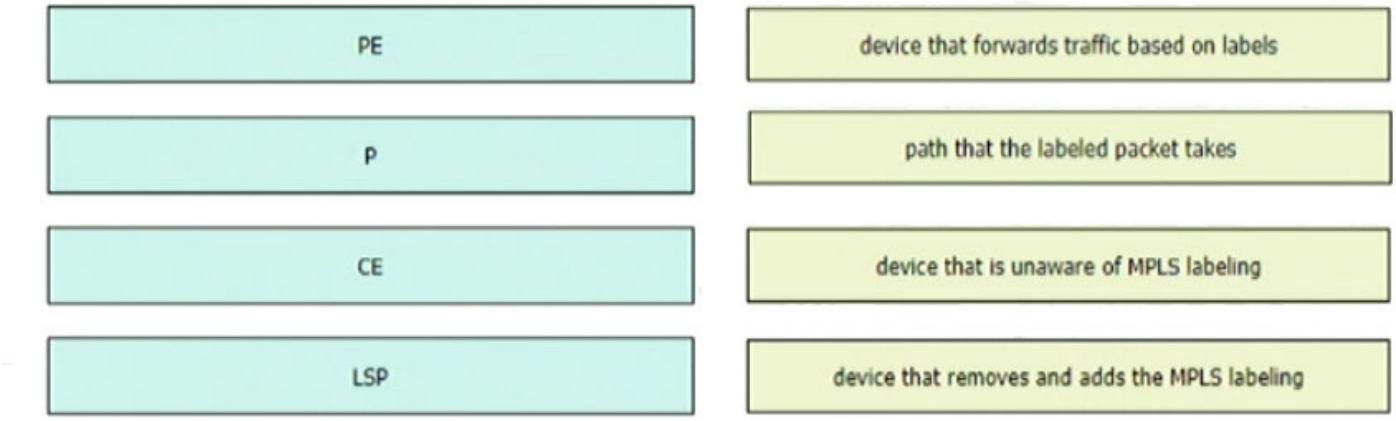

Answer:

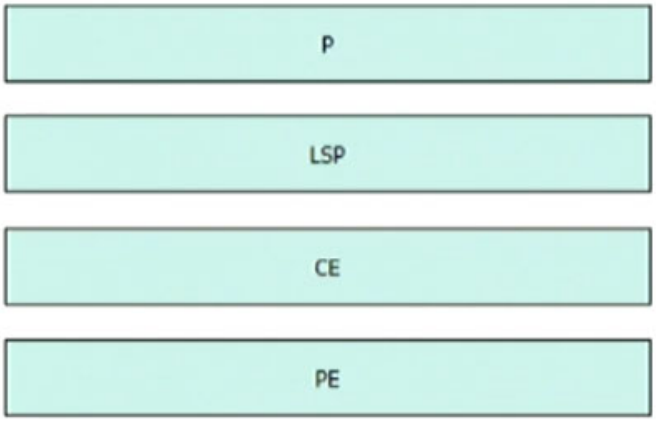

Simulation 1 ★★★★

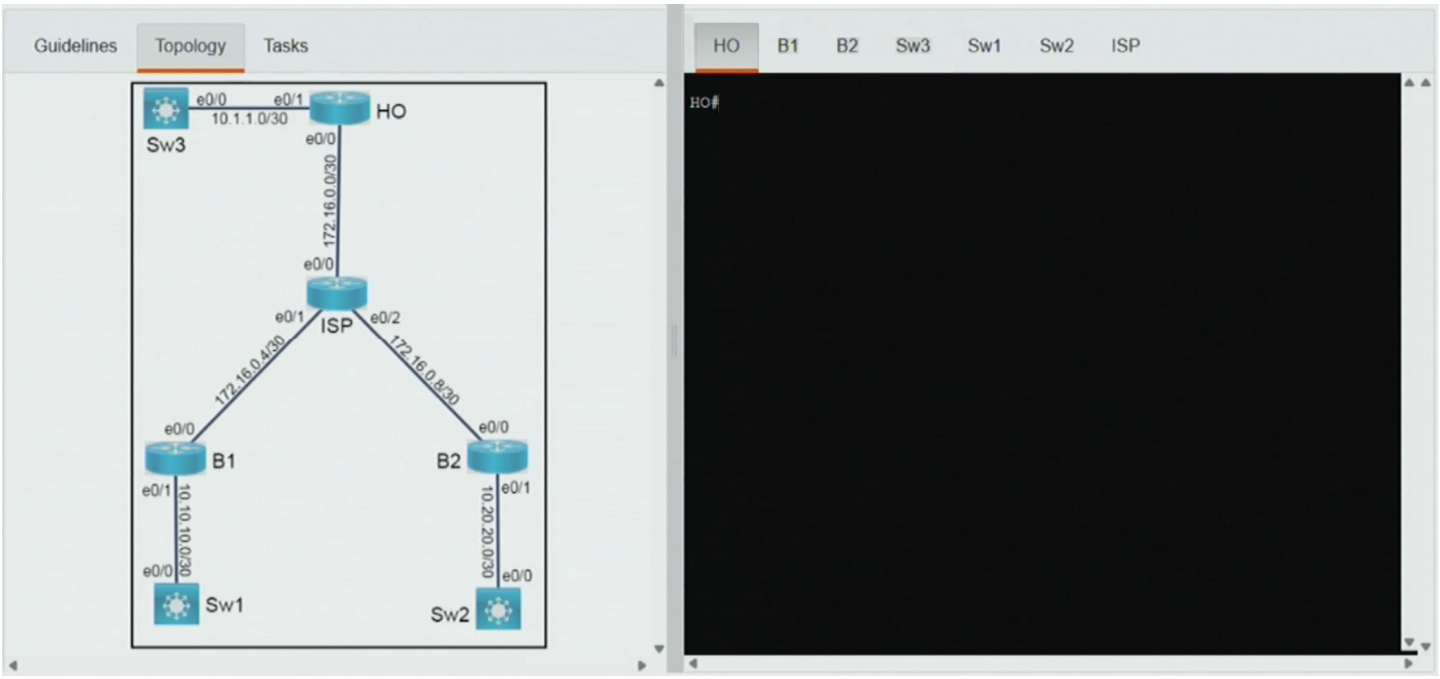

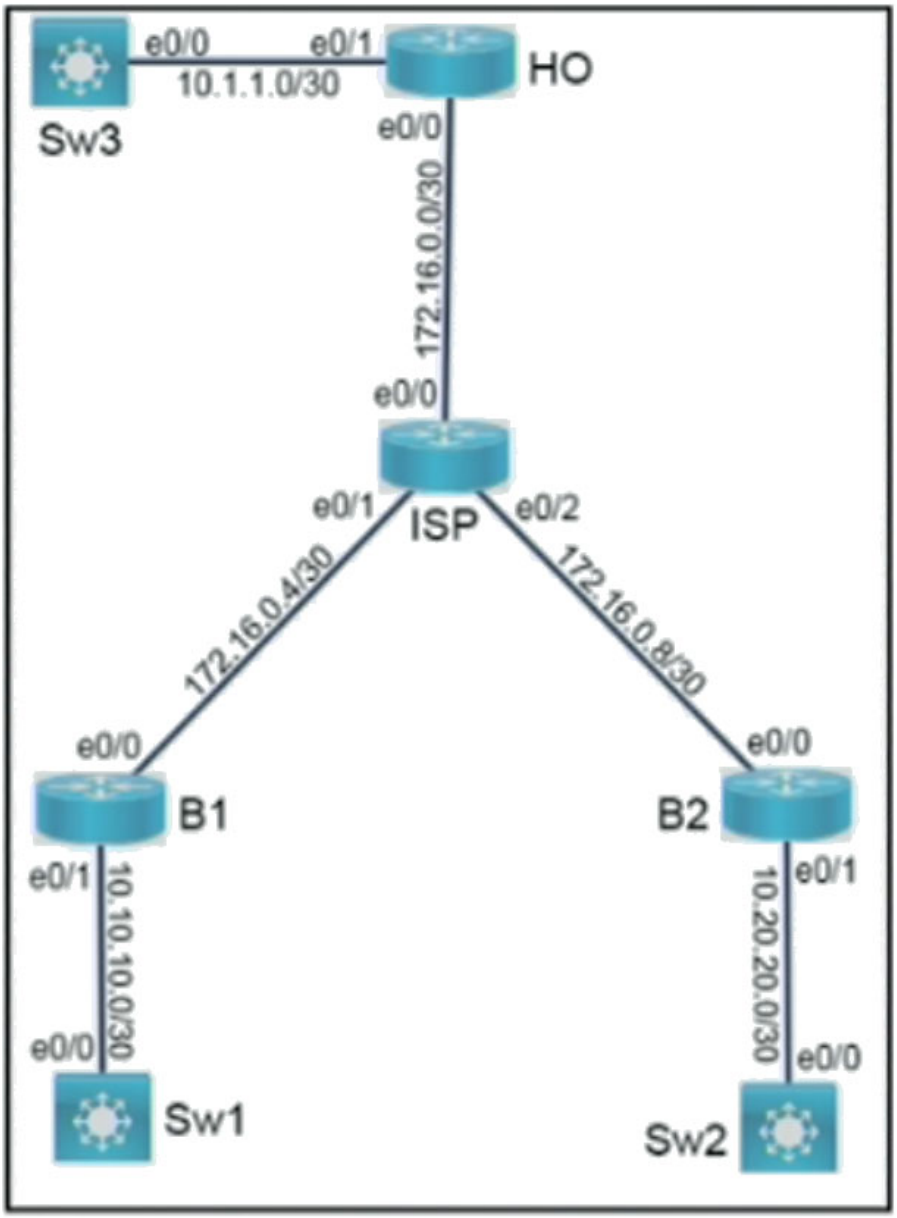

Guidelines | Topology | Tasks

Task 1:

Complete the configuration on the B1 and B2 routers to establish a DMVPN tunnel using mGRE between HO, B1 and B2. B1 and B2 routers should be able to establish a dynamic tunnel for direct traffic between them. The output from ***show DMVPN*** should match the output below, except for the UpDn Tm.

- Sw1 and Sw2 must be able to ping each other to match desired output (SW1#ping 10.20.20.1 repeat 500)

```
B1#show dmvpn
 ======================================================================
 Interface: Tunnel10, IPv4 NHRP Details Type:Spoke, NHRP Peers:2,
 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 172.16.0.2      192.168.0.1     UP    00:04:44 S
     1 172.16.0.10     192.168.0.20    UP    00:04:44 D
B2#show dmvpn
 ======================================================================
 Interface: Tunnel10, IPv4 NHRP Details Type:Spoke, NHRP Peers:2,
 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 172.16.0.2    192.168.0.1       UP    00:01:35 S
     1 172.16.0.6    192.168.0.10      UP    00:00:18 D
```

Task 2:

Complete the configuration on the HO router to encrypt end-to-end traffic between HO and B1 and B2 using these parameters:

1. Crypto IPsec Parameters
    - ESP Encryption Transform: esp-aes 256
    - ESP Authentication Transform: esp-sha384-hmac
    - Profile: AdRt_PROFILE
    - Transform-set: AdRt_TRANSFORM
    - Mode: Transport
    - Tunnel Protection: ipsec profile AdRt_PROFILE

Note: Complete Task 2 to see the results. Bounce tunnel interface on all routers after completing configurations to see both tunnels up on HO router if required.

Task 1:

Complete the configuration on the B1 and B2 routers to establish a DMVPN tunnel using mGRE between HO, B1 and B2. B1 and B2 routers should be able to establish a dynamic tunnel for direct traffic between them. The output from *show DMVPN* should match the output below, except for the UpDn Tm.

*** Sw1 and Sw2 must be able to ping each other to match desired output (SW1#ping 10.20.20.1 repeat 500)**

```
B1# show dmvpn
======================================================================
Interface: Tunnel10, IPv4 NHRP Details Type:Spoke, NHRP Peers:2,
# Ent  Peer NBMA Addr  Peer Tunnel Add  State   UpDn Tm   Attrb
-----  --------------  ---------------  -----   -------   -----
    1  172.16.0.2       192.168.0.1      UP     00:04:44  S
    1  172.16.0.10      192.168.0.20     UP     00:04:44  D

B2# show dmvpn
======================================================================
Interface: Tunnel10, IPv4 NHRP Details Type:Spoke, NHRP Peers:2,
# Ent  Peer NBMA Addr  Peer Tunnel Add  State   UpDn Tm   Attrb
-----  --------------  ---------------  -----   -------   -----
   1   172.16.0.2       192.168.0.1      UP     00:01:35  S
   1   172.16.0.6       192.168.0.10     UP     00:00:18  D
```

Task 2:

Complete the configuration on the HO router to encrypt end-to-end traffic between HO and B1 and B2 using these parameters:

1. Crypto IPsec Parameters
   - ESP Encryption Transform: esp-aes 256
   - ESP Authentication Transform: esp-sha384-hmac
   - Profile: AdRt_PROFILE
   - Transform-set: AdRt_TRANSFORM
   - Mode: Transport
   - Tunnel Protection: ipsec profile AdRt_PROFILE

Note: Complete Task 2 to see the results. Bounce tunnel interface on all routers after completing configurations to see both tunnels up on HO router if required.

Task 1: B1 and B2 DMVPN Spoke Configuration

The spokes are currently configured with tunnel destination, which creates a point-to-point tunnel. For dynamic spoke-to-spoke traffic, they must use mGRE (tunnel mode gre multipoint).

On B1:

```
conf t
interface Tunnel10
 no tunnel destination 172.16.0.2
 tunnel mode gre multipoint
 ip nhrp shortcut
end
```

On B2:

```
conf t
interface Tunnel10
 no tunnel destination 172.16.0.2
 tunnel mode gre multipoint
 ip nhrp shortcut
end
```

Task 2: HO Router IPsec Configuration

The HO router has ISAKMP (Phase 1) but is missing the IPsec Transform-set and Profile (Phase 2) to protect the tunnel.

On HO:

```
conf t
crypto ipsec transform-set AdRt_TRANSFORM esp-aes 256 esp-sha384-hmac
 mode transport
!
crypto ipsec profile AdRt_PROFILE
 set transform-set AdRt_TRANSFORM
!
interface Tunnel10
 tunnel protection ipsec profile AdRt_PROFILE
 ip nhrp redirect
 ip nhrp map multicast dynamic
end
```

Verification and Troubleshooting

1. Bounce Interfaces: To ensure the new security parameters take effect, flap the tunnels on all routers: **interface Tunnel10 -> shutdown -> no shutdown**

```
HO(config)# interface Tunnel10
HO(config-if)# shutdown
HO(config-if)# no shutdown

*Apr 25 10:45:52.562: %OSPF-5-ADJCHG: Process 10, Nbr 10.202.202.1 on Tunnel10
from LOADING to FULL, Loading Done
*Apr 25 10:45:52.563: %OSPF-5-ADJCHG: Process 10, Nbr 10.101.101.1 on Tunnel10
from LOADING to FULL, Loading Done
```

2. Verify DMVPN: Run show dmvpn. You should see the Static (S) entry for the Hub and a Dynamic (D) entry appear for the other spoke once you initiate the ping.

```
HO# show dmvpn
Legend:    Attrb --> S - Static, D - Dynamic, I - Incomplete
           N - NATed, L - Local, X - No Socket
           T1 - Route Installed, T2 - Nexthop-override
           C - CTS Capable, I2 - Temporary
           # Ent --> Number of NHRP entries with same NBMA peer
           NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
           UpDn Time --> Up or Down Time for a Tunnel
==========================================================================
Interface: Tunnel10, IPv4 NHRP Details
Type:Hub, NHRP Peers:2,

# Ent  Peer NBMA Addr   Peer Tunnel Add  State  UpDn Tm  Attrb
-----  --------------   ---------------  -----  -------  -----
    1  172.16.0.6       192.168.0.10       UP   00:00:13     D
    1  172.16.0.10      192.168.0.20       UP   00:00:13     D
```

```
B1# show dmvpn
Legend:   Attrb --> S - Static, D - Dynamic, I - Incomplete
          N - NATed, L - Local, X - No Socket
          T1 - Route Installed, T2 - Nexthop-override
          C - CTS Capable, I2 - Temporary
          # Ent --> Number of NHRP entries with same NBMA peer
          NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
          UpDn Time --> Up or Down Time for a Tunnel
======================================================================
Interface: Tunnel10, IPv4 NHRP Details
Type:Hub, NHRP Peers:2,

# Ent  Peer NBMA Addr   Peer Tunnel Add  State  UpDn Tm  Attrb
-----  --------------   ---------------  -----  -------  -----
    2  172.16.0.10      192.168.0.20       UP   00:00:05     D
                        192.168.0.20       UP   00:00:05   DT1
    1  172.16.0.2        192.168.0.1       UP   00:01:41     S



B2# show dmvpn
Legend:   Attrb --> S - Static, D - Dynamic, I - Incomplete
          N - NATed, L - Local, X - No Socket
          T1 - Route Installed, T2 - Nexthop-override
          C - CTS Capable, I2 - Temporary
          # Ent --> Number of NHRP entries with same NBMA peer
          NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
          UpDn Time --> Up or Down Time for a Tunnel
======================================================================
Interface: Tunnel10, IPv4 NHRP Details
Type:Hub, NHRP Peers:2,

# Ent  Peer NBMA Addr   Peer Tunnel Add  State  UpDn Tm  Attrb
-----  --------------   ---------------  -----  -------  -----
    2  172.16.0.6       192.168.0.10       UP   00:00:19     D
                        192.168.0.10       UP   00:00:19   DT1
    1  172.16.0.2        192.168.0.1       UP   00:00:45     S
```

3. Connectivity Test: From **Sw1**, run the ping to **Sw2's** interface:

```
Sw1# ping 10.20.20.1 repeat 500
Type escape sequence to abort.
Sending 500, 100-byte ICMP Echos to 10.20.20.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (500/500), round-trip min/avg/max = 1/1/5 ms
Sw1#
```

Connectivity Test: From **Sw2**, run the ping to **Sw1's** interface:

```
Sw2# ping 10.10.10.1 repeat 500
Type escape sequence to abort.
Sending 500, 100-byte ICMP Echos to 10.10.10.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (500/500), round-trip min/avg/max = 1/1/5 ms
Sw2#
```

4. Routing Check: Ensure OSPF or a routing protocol is active on B2 and HO (the provided B2 config was missing the OSPF section) so that the Sw1 and Sw2 subnets are reachable across the tunnel.