




"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks.com

Warning: Keep connected with our support team for latest updates

IPv6 EIGRPAS100



```

R1# show ipv6 eigrp topology
EIGRP IPv6 Topology Table for AS( 100)TD( 101121)
Codes P - Passive, A - Active, U - Update, Q - Query, R - Reply,
      r - reply Status, s - sia Status
P 2001:DB8:0:4::/64, 1 successors, FD is 28416
   via FE80:C828:DFF:FEF4:1C (28416/2816), FastEthernet3/0
P 2001:DB8:0:1::/64, 1 successors, FD is 2816
   via Connected, GigabitEthernet0/0
P 2001:DB8:0:14::/64, 1 successors, FD is 2816
   via FE80:C821:17FF:FE04:8 (2816/256), GigabitEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
   via Connected, FastEthernet3/0
P 2001:DB8:0:12::/64, 1 successors, FD is 2816
  
```

```

Branch# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS( 100)40(4 4 4 4)
Codes P - Passive, A - Active, U - Update, Q - Query, R - Reply,
      r - reply Status, s - sia Status
P 2001:068:0:4::/64, 1 successors, FD is 2816
   via Connected, GigabitEthernet0/0
P 2001:DB8:0:1764::/64, 1 successors, FD is 28416
   via FE80:C820:17FF:FE04:54 (28416/2816), FastEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
   via Connected, FastEthernet1/0
P 2001:DB8:0:12::/64, 1 successors, FD is 28416
   via FE80:C820:17FF:FE04:54 (28416/2816), FastEthernetVO
  
```

Users in the branch network of 2001:db8:0:4::/64 report that they cannot access the Internet. Which command is issued in IPv6 router EIGRP 100 configuration mode to solve this issue?

- A. Issue the eigrp stub command on R1
- B. Issue the no neighbor stub command on R2.
- C. Issue the eigrp command on R2.
- D. Issue the no eigrp stub command on R1.

Answer:


D

Explanation:

Question:

2

Refer to the exhibit.



Which configuration configures a policy on R1 to forward any traffic that is sourced from the 192.168.130.0/24 network to R2?

A. access-list 1 permit 192.168.130.0 0.0.0.255

interface GiO/2

ip policy route-map test

route-map test permit 10

match ip address 1 set ip next-hop 172.20.20.2

B. access-list 1 permit 192.168.130.0 0.0.0.255

interface Gi0/1

ip policy route-map test

route-map test permit 10

match ip address 1

set ip next-hop 172.20.40.2

c access-list 1 permit 192.168.130.0 0.0.0.255

interface GiO/2

ip policy route-map test

route-map test permit 10

match ip address 1

set ip next-hop 172.20.20.1

D access-list 1 permit 192.168.130.0 0.0.0.255

Interface Gi0/1

ip policy route-map test

route-map test permit 10

match ip address 1

set ip next-hop 172.20.40.1

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 3

R2 has a locally originated prefix 192.168.130.0/24 and has these configurations:

ip prefix-list test seq 5 permit 192.168.130.0/24 T
route-map OUT permit 10 match ip address prefix-list
test set as-path prepend 65000

What is the result when the route-map OUT command is applied toward an eBGP neighbor R1 (1.1.1.1) by using the neighbor 1.1.1.1 route-map OUT out command?

- A. R1 sees 192.168.130.0/24 as two AS hops away instead of one AS hop away.
- B. R1 does not accept any routes other than 192.168.130.0/24
- C. R1 does not forward traffic that is destined for 192.168.30.0/24
- D. Network 192.168.130.0/24 is not allowed in the R1 table

Answer: A

Explanation:

Question: 4

Which method changes the forwarding decision that a router makes without first changing the routing table or influencing the IP data plane?

- A. nonbroadcast multiaccess

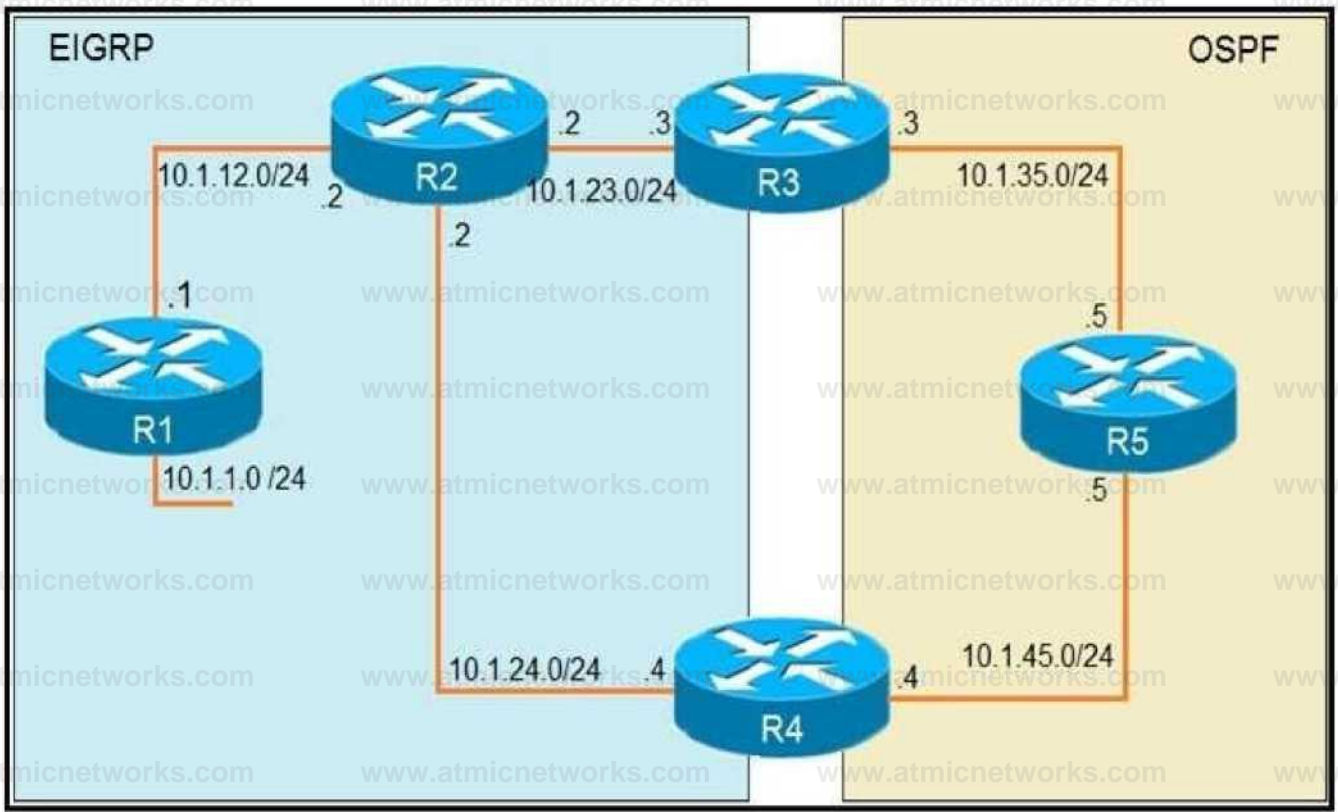
- B. packet switching
- C. policy-based routing
- D. forwarding information base

Answer: C

Explanation:

Question: 5

Refer to the exhibit.



```
R1
router eigrp 1
 redistribute connected
 network 10.1.12.1 0.0.0.0
```

```
R3
router ospf 1
 redistribute eigrp 1 subnets
 network 10.1.35.3 0.0.0.0 area 0
```

```
R4
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500
```

```
router ospf 1
 network 10.1.45.4 0.0.0.0 area 0
```

```
R5#traceroute 10.1.1.1
```

```
Type escape sequence to abort
Tracing the route to 10.1.1.1
```

```
 1  10 1 35 3 80 msec 44 msec 20 msec
 2  10 1 23 2 44 msec 104 msec 64 msec
 3  10 1 24 4 44 msec 64 msec 40 msec
 4  10.1.45 5 24 msec 40 msec 20 msec
 5  10 1 35 3 92 msec 144 msec 148 msec
 6  10 1 23 2 108 msec 76 msec 80 msec
<output truncated>
```

The output of the trace route from R5 shows a loop in the network. Which configuration

prevents this loop?

A)

R3

```
router ospf 1
 redistribute eigrp 1 subnets route-map SET-TAG
```

```
route-map SET-TAG permit 10
 set tag 1
```

R4

```
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG I
 route-map FILTER-TAG deny 10
 match tag 1

route-map FILTER-TAG permit 20
```

B)

R3

```
router eigrp 1
 redistribute OSPF 1 route-map SET-TAG
```

```
route-map SET-TAG permit 10
 set tag 1
```

R4

```
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG network
 10.1.24.4 0.0.0.0
```

```
route-map FILTER-TAG deny 10
 match tag 1
```

route-map FILTER-TAG permit 20

c)

R3

```
router ospf 1
 redistribute eigrp 1 subnets route-map SET-TAG
```

```
route-map SET-TAG permit 10
 set tag 1
```

R4

```
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
```

```
route-map FILTER-TAG permit 10
 match tag 1
```

D)

R3

```
router ospf 1
 redistribute eigrp 1 subnets route-map SET-TAG
```

```
route-map SET-TAG deny 10
 set tag 1
```

R4

```
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG f
```

```
route-map FILTER-TAG deny 10
 match tag 1
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

The reason for the loop is that R2 is forwarding the packets destined to 10.1.1.1 to R4, instead of R1.

This is because in the redistribute OSPF statement, BW metric has a higher value and delay has a

value of 1. So, R2 chooses R4 over R1 for 10.1.1.0/24 subnet causing a loop. Now, R5 learns 10.1.1.0/24 from R3 and advertises the same route to R4, that R4 redistributes back in EIGRP. If R3 sets a tag of 1 while redistributing EIGRP in OSPF, and R4 denies all the OSPF routes with tag 1 while redistributing, it will not advertise 10.1.1.0/24 back into EIGRP. Hence, the loop will be broken.

Question: 6

Refer to the exhibit.

```
Router#show running-config | include ip route
ip route 192.168.2.2 255.255.255.255 209.165.200.225 130
Router#show ip route
<output omitted>

Gateway of last resort is not set

    192.168.1.0/32 is subnetted, 1 subnets
C       192.168.1.1 is directly connected, Loopback0
    192.168.2.0/32 is subnetted, 1 subnets
O       192.168.2.2[110/11] via 192.168.12.2, 00:52:09, Ethernet0/0
    192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.12.0/24 is directly connected, Ethernet0/0
L       192.168.12.1/32 is directly connected, Ethernet0/0
    209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.200.0/24 is directly connected, Ethernet0/1
        209.165.200.226/32 is directly connected, Ethernet0/1
```

An engineer configures a static route on a router, but when the engineer checks the route to the destination, a different next hop is chosen. What is the reason for this?

- A. Dynamic routing protocols always have priority over static routes.
- B. The metric of the OSPF route is lower than the metric of the static route.

- C. The configured AD for the static route is higher than the AD of OSPF.
- D. The syntax of the static route is not valid, so the route is not considered.

Answer: C

Explanation:

The AD of static route is manually configured to 130 which is higher than the AD of OSPF router which is 110.

Question: 7

Refer to the exhibit.

```
Router#show ip route
```

```
<output omitted>
```

```
Gateway of last resort is not set
```

```
192.168.1.0/32 is subnetted, 1 subnets
O       192.168.1.1 [110/11] via 192.168.12.1,16:56:40, Ethernet0/0
192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.2.0/24 is directly connected, Loopback0
L       192.168.2.2/32 is directly connected, Loopback0
192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.3.0/24 is directly connected, Ethernet0/1
L       192.168.3.1/32 is directly connected, Ethernet0/1
192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.12.0/24 is directly connected, Ethernet0/0
L       192.168.12.2/32 is directly connected, Ethernet0/0
```

```
Router#show running-config | section ospf
```

```
router ospf 1
  summary-address 10.0.0.0 255.0.0.0
  redistribute static subnets
  network 192.168.3.0 0.0.0.255 area 0
  network 192.168.12.0 0.0.0.255 area 0
Router#
```

An engineer is trying to generate a summary route in OSPF for network 10.0.0.0/8, but the

summary route does not show up in the routing table. Why is the summary route missing?

- A. The summary-address command is used only for summarizing prefixes between areas.
- B. The summary route is visible only in the OSPF database, not in the routing table.
- C. There is no route for a subnet inside 10.0.0.0/8, so the summary route is not generated.
- D. The summary route is not visible on this router, but it is visible on other OSPF routers in the same area.

Answer: C

Explanation:

The `summary-address` is only used to create aggregate addresses for OSPF at an autonomous system boundary. It means this command should only be used on the ASBR when you are trying to summarize externally redistributed routes from another protocol domain or you have a NSSA area. But a requirement to create a summarized route is:

—The ASBR compares the summary route's range of addresses with all routes redistributed into OSPF on that ASBR to find any subordinate subnets (subnets that sit inside the summary route range). If at least one subordinate subnet exists, the ASBR advertises the summary route.

Question: 8

Refer to the exhibit.

```
Router#show access-lists
Standard IP access list 1
    10 permit 192.168.2.2 (1 match)
Router#
Router#show route-map
route-map RM-OSPF-DL, permit, sequence 10
Match clauses:
    ip address (access-lists): 1
Set clauses:
Policy routing matches: 0 packets, 0 bytes
Router#
Router#show running-config | section ospf
router ospf 1
 network 192.168.1.1 0.0.0.0 area 0
 network 192.168.12.0 0.0.0.255 area 0
 distribute-list route-map RM-OSPF-DL in
Router#|
```

An engineer is trying to block the route to 192.168.2.2 from the routing table by using the configuration that is shown. The route is still present in the routing table as an OSPF route. Which action blocks the route?

- A. Use an extended access list instead of a standard access list.
- B. Change sequence 10 in the route-map command from permit to deny.
- C. Use a prefix list instead of an access list in the route map.
- D. Add this statement to the route map: route-map RM-OSPF-DL deny 20.

Answer: B

Explanation:

Question: 9

What is a prerequisite for configuring BFD?

- A. Jumbo frame support must be configured on the router that is using BFD.
- B. All routers in the path between two BFD endpoints must have BFD enabled.
- C. Cisco Express Forwarding must be enabled on all participating BFD endpoints.
- D. To use BFD with BGP, the timers 3 9 command must first be configured in the BGP routing process.

Answer: C

Explanation:

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1043332

Question: 10

DRAG DROP

Drag and drop the OSPF adjacency states from the left onto the correct descriptions on the right.



Each router compares the DBD packets that were received from the other router.

Routers exchange information with other routers in the multiaccess network.

The neighboring router requests the other routers to send missing entries.

The network has already elected a DR and a backup BDR.

The OSPF router ID of the receiving router was not contained in the hello message.

No hellos have been received from a neighbor router.

Answer:

Explanation:

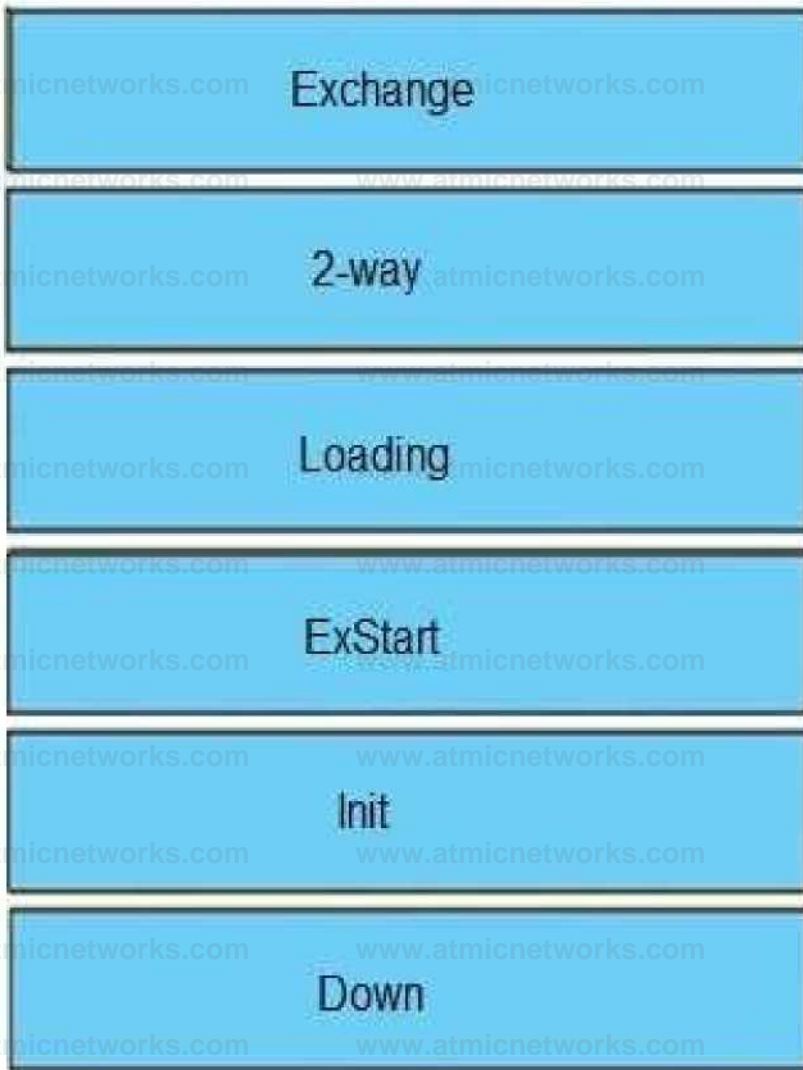


Table Description

automatically generated

(Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml)

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html>

Down

This is the first OSPF neighbor state. It means that no information (hellos) has been received from this neighbor, but hello packets can still be sent to the neighbor in this state.

During the fully adjacent neighbor state, if a router doesn't receive hello packet from a neighbor within the Router Dead Interval time (RouterDeadInterval = 4*HelloInterval by default) or if the manually configured neighbor is being removed from the configuration, then the neighbor state changes from Full

to Down.

Attempt

This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

Init

This state specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet. When a router receives a hello packet from a neighbor, it should list the sender's router ID in its hello packet as an acknowledgment that it received a valid hello packet.

2-Way

This state designates that bi-directional communication has been established between two routers. Bi-directional means that each router has seen the other's hello packet. This state is attained when the router receiving the hello packet sees its own Router ID within the received hello packet's neighbor field. At this state, a router decides whether to become adjacent with this neighbor. On broadcast media and non-broadcast multiaccess networks, a router becomes full only with the designated router (DR) and the backup designated router (BDR); it stays in the 2-way state with all other neighbors. On Point-to-point and Point-to-multipoint networks, a router becomes full with all connected routers.

At the end of this stage, the DR and BDR for broadcast and non-broadcast multiaccess networks are elected. For more information on the DR election process, refer to DR Election.

Note: Receiving a Database Descriptor (DBD) packet from a neighbor in the init state will also cause a transition to 2-way state.

Exstart

Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR. (ie. Shared or NBMA networks).

In this state, the routers and their DR and BDR establish a master-slave relationship and choose the initial sequence number for adjacency formation. The router with the higher router ID becomes the master and starts the exchange, and as such, is the only router that can increment the sequence number. Note that one would logically conclude that the DR/BDR with the highest router ID will become the master during this process of master-slave relation. Remember that the DR/BDR election might be purely by virtue of a higher priority configured on the router instead of highest router ID. Thus, it is possible that a DR plays the role of slave. And also note that master/slave election is on a per-neighbor basis.

Exchange

In the exchange state, OSPF routers exchange database descriptor (DBD) packets. Database descriptors contain link-state advertisement (LSA) headers only and describe the contents of the entire link-state database. Each DBD packet has a sequence number which can be incremented only by master which is explicitly acknowledged by slave. Routers also send link-state request packets and link-state update packets (which contain the entire LSA) in this state. The contents of the DBD received are compared to the information contained in the routers link-state database to check if new or more current link-state information is available with the neighbor.

Loading

In this state, the actual exchange of link state information occurs. Based on the information provided by the DBDs, routers send link-state request packets. The neighbor then provides the requested linkstate information in link-state update packets. During the adjacency, if a router receives an outdated or missing LSA, it requests that LSA by sending a link-state request packet. All link-state update packets are acknowledged.

Full

In this state, routers are fully adjacent with each other. All the router and network LSAs are exchanged and the routers' databases are fully synchronized.

Full is the normal state for an OSPF router. If a router is stuck in another state, it is an indication that there are problems in forming adjacencies. The only exception to this is the 2-way state, which is normal in a broadcast network.

Routers achieve the FULL state with their DR and BDR in NBMA/broadcast media and FULL state with every neighbor in the remaining media such as point-to-point and point-to-multipoint.

Note: The DR and BDR that achieve FULL state with every router on the segment will display FULL/DROTHER when you enter the show ip ospf neighbor command on either a DR or BDR. This simply means that the neighbor is not a DR or BDR, but since the router on which the command was entered is either a DR or BDR, this shows the neighbor as FULL/DROTHER.

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html>

Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.sht ml

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html>

Question: 11

Refer to the exhibit.

R1 #show ip bgp summary

BGP router identifier 192.168.1.1, local AS number 65000

<output omitted>

Neighbor	V AS	MsgRcvd	MsgSent	Tblver	InQ	OutQ	Up/Down State/PfxRcd	
192.1682.2	4 65000	28	28	22	0	0	00:21:31	0

R1#show ip bgp

BGP table version is 22, local router ID is 192.168.1.1

Status codes: s suppressed, d damped, h history, * valid > best, i - internal, r RIB-failure, s stale, m multipath, b backup-path, f RT-Filter, x best-external, a additional-path, C RIB-compressed,

Origm codes: i - IGP, e - EGP, ? - incomplete

RPKI validation codes V valid, I invalid, N Not found

Network	Next Hop	Metnc	LocPrf	Weight	Path
*> 172.16.25.0/24	209.165 200.225	0		32768	?

R1#

R2 #show ip bgp summary

BGP router identifier 192 168 2 2, local AS number 65000

<output omitted>

Neighbor	V AS	MsgRcvd	MsgSent	Tblver	InQ	OutQ	Up/Down State/PfxRcd	
192.168 1.1	4 65000	29	28	3	0	0	00:22:07	1
192.168.3.3	4 65000	7	8	3	0	0	00:02:55	0

R2#show ip bgp

BGP table version is 3, local router ID is 192.168 2.2

Status codes: s suppressed d damped, h history, * valid. > best, i - internal, r RIB-failure, s stale, m multipath, b backup-path, f RT-Filter, x best-external, a additional-path, C RIB-compressed,

Origin codes: i - IGP, e - EGP, ? - incomplete

RPKI validation codes: V valid, I invalid, N Not found

Network	Next Hop	Metnc	LocPrf	Weight	Path
* i 172.16.25.0/24	209.165 200.225	0	100	0	?

R2#

R3 #show ip bgp summary

BGP router identifier 192.168.3.3, local AS number 65000

BGP table version is 4, main routing table version 4

Neighbor V AS MsgRcvd MsgSent Tblver InQ OutQ

192.168.2.2	4 65000	8	7	4	0	0	Up/Down State/PfxRcd	
R3#							00:03:08	0

R2 is a route reflector, and R1 and R3 are route reflector clients. The route reflector learns the route to 172.16.25.0/24 from R1, but it does not advertise to R3. What is the reason the route is not advertised?

A. R2 does not have a route to the next hop, so R2 does not advertise the prefix to other clients.

- B. Route reflector setup requires full IBGP mesh between the routers.
- C. In route reflector setup, only classful prefixes are advertised to other clients.
- D. In route reflector setups, prefixes are not advertised from one client to another.

Answer: A

Explanation:

Question: 12

Refer to the exhibit.

```
Router#sh ip route ospf
```

```
<output omitted>
```

```
Gateway is last resort is not set
```

```
10.0.0.0/24 is subnetted, 1 subnets
```

```
o E2 10.0.0.0 [110/20] via 192.168.12.2, 00:00:10, Ethernet0/0
```

```
o 192.168.3.0/24 [110/20] via 192.168.12.2,00:00:50, Ethernet0/0 Router#
```

```
Routertfshow ip bgp
```

```
<output omitted>
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
>*	192.168.1.1/32	0.0.0.0	0		32768	?
>*	192.168.3.0	192.168.12.2	20		32768	?
>*	192.168.12.0	0.0.0.0	0		32768	?

```
Router#show running-config | section router bgp router bgp 65000 bgp log-neighbor-changes redistribute ospf 1
```

```
Router#
```

An engineer is trying to redistribute OSPF to BGP, but not all of the routes are redistributed. What is the reason for this issue?

- A. By default, only internal routes and external type 1 routes are redistributed into BGP
- B. Only classful networks are redistributed from OSPF to BGP
- C. BGP convergence is slow, so the route will eventually be present in the BGP table
- D. By default, only internal OSPF routes are redistributed into BGP

Answer: D

Explanation:

If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default.

You can redistribute both internal and external (type-1 & type-2) OSPF routes via this command: —Router(config-router)#redistribute ospf 1 match internal external 1 external 2||

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html>

Question: 13

Refer to the exhibit.

R200#show ip bgp summary

```
BGP router identifier 10.1.11. local AS number 65000
BGP table version is 26, main routing table version 26
1 network entries using 132 bytes of memory
1 path entries using 52 bytes of memory
2/1 BGP path/bestpath attribute entries using 296 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 2) using 28 bytes of memory
BGP using 508 total bytes of memory
BGP activity 24/23 prefixes, 24/23 paths, scan interval 60 secs
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
192.0.2.2 4 65100 20335 20329 0 0 0 00:02:04 Idle (PfxCt)
R200#
```

In which circumstance does the BGP neighbor remain in the idle condition?

- A. if prefixes are not received from the BGP peer
- B. if prefixes reach the maximum limit
- C. if a prefix list is applied on the inbound direction
- D. if prefixes exceed the maximum limit

Answer: D

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/25160-bgp-maximum-prefix.html#b>

Question: 14

Which attribute eliminates LFAs that belong to protected paths in situations where links in a network are connected through a common fiber?

- A. shared-risk-link-group-disjoint
- B. linecard-disjoint
- C. lowest-repair-path-metric
- D. interface-disjoint

Answer: A

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xr-3s/asr1000/ire-xe-3s-asr1000/ire-ipfrr.html

Question: 15

Refer to the exhibit.

```
' Jun 28 14:41:57: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Down User reset
```

```
' Jun 28 14:41:57: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.2.2 IPv4 Unicast
```

topology base removed from session User reset

' Jun 28 14 41:57: %BGP-5-ADJCHANGE neighbor 192 168 2 2 Up

R1#show clock

'15:42:00.506 CET Fri Jun 28 2019

An engineer is troubleshooting BGP on a device but discovers that the clock on the device does not correspond to the time stamp of the log entries. Which action ensures consistency between the two times?

- A. Configure the service timestamps log uptime command in global configuration mode.
- B. Configure the logging clock synchronize command in global configuration mode.
- C. Configure the service timestamps log datetime localtime command in global configuration mode.
- D. Make sure that the clock on the device is synchronized with an NTP server.

Answer: C

Explanation:

https://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r3-9/system_management/command/reference/yr39xr12k_chapter4.html#wp1784026936

By default, syslog and debug messages are stamped by UTC, regardless of the time zone that device is configured. You should append localtime key word to "service timestamp {log | debug} datetime msec" global command to change that behavior.

<https://community.cisco.com/t5/networking-documents/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258>

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/service_timestamps.htm

Question: 16

Refer to the exhibit.

```
R1#show policy-map control-plane
Control Plane
  Service-policy input: CoPP-BGP
    Class-map: BGP (match all)
      2716 packets, 172071 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name BGP
      drop
    Class-map: class-default (match-any)
      5212 packets, 655966 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

What is the result of applying this configuration?

- A. The router can form BGP neighborships with any other device.
- B. The router cannot form BGP neighborships with any other device.
- C. The router cannot form BGP neighborships with any device that is matched by the access list named "BGP".
- D. The router can form BGP neighborships with any device that is matched by the access list named "BGP".

Answer: C

Explanation:

after bgp session are UP.I configured the CoPP to drop 10.3.3.3 bgp traffic (R3).

R3 bgp traffic that matched the ACL 100 is dropped and the state is in IDLE

```
access-list 100 permit tcp host 10.3.3.3 any eq bgp
```

```
access-list 100 permit tcp host 10.3.3.3 eq bgp any !
```

```
class-map match-all class-bgp
```

```
match access-group 100
```

```
!
```

```
policy-map policy-bgp
```

```
class class-bgp
```

```
drop
```

```
!
```

```
control-plane
```

```
service-policy input policy-bgp
```

The 10.3.3.3 neighbor goes to IDLE

Question: 17

Which command displays the IP routing table information that is associated with VRF-Lite?

- A. show ip vrf
- B. show ip route vrf
- C. show run vrf
- D. show ip protocols vrf

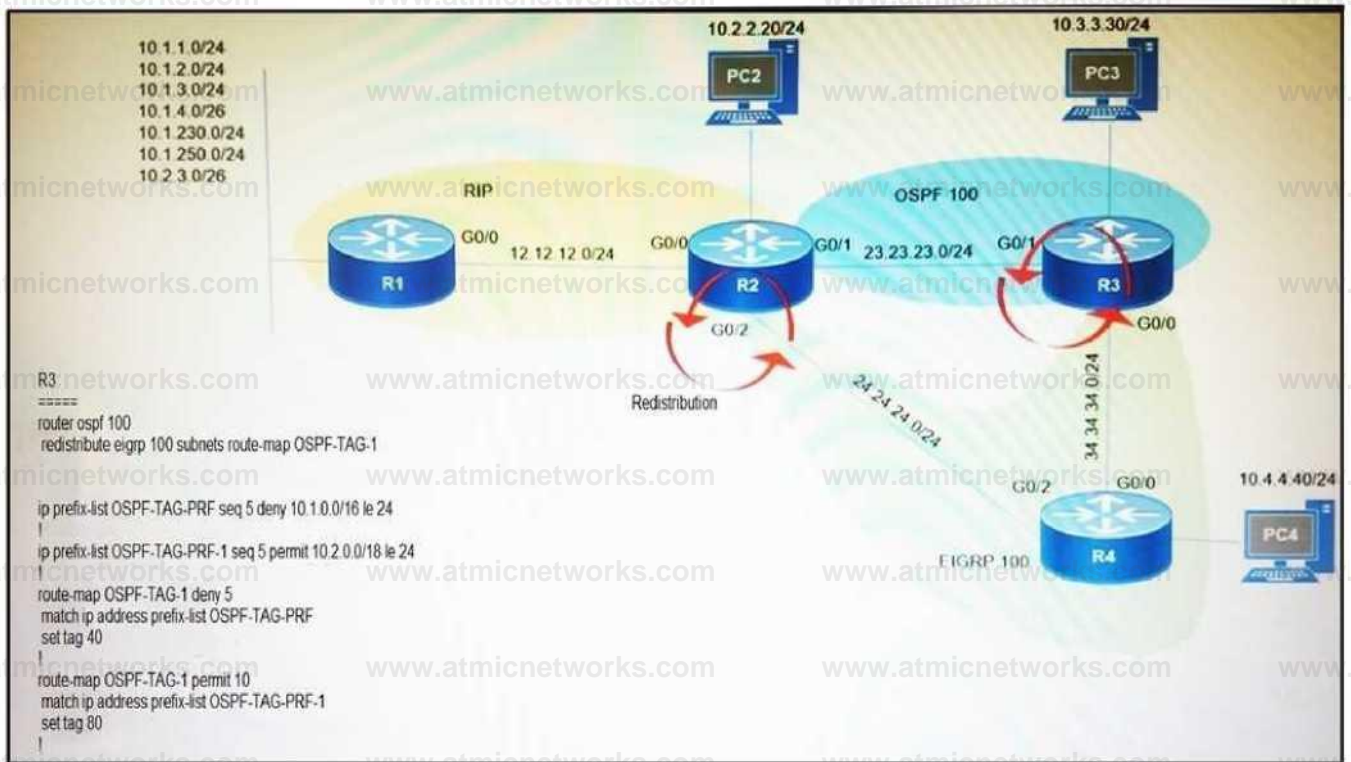
Answer: B

Explanation:

Reference: <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/50sg/configuration/guide/Wrapper-46SG/vrf.html#wp1045708>

Question: 18

Refer to the exhibit.



Which subnet is redistributed from EIGRP to OSPF routing protocols?

- A. 10.2.2.0/24
- B. 10.1.4.0/26
- C. 10.1.2.0/24
- D. 10.2.3.0/26

Answer: A

Explanation:

Question: 19

Which configuration adds an IPv4 interface to an OSPFv3 process in OSPFv3 address family

configuration?

- A. Router ospfv3 1 address-family ipv4
- B. Router(config-router)#ospfv3 1 ipv4 area 0
- C. Router(config-if)#ospfv3 1 ipv4 area 0
- D. Router ospfv3 1 address-family ipv4 unicast

Answer: C

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xen-3s/iro-xe-3s-book/ip6-route-ospfv3-add-fam-xe.html

Question: 20

Refer to the exhibit.

```
R1(config)#route-map ADD permit 20
R1(config-route-map)#set tag 1

R1(config)#router ospf1
R1(config-router)#redistribute rip subnets route-map ADD
```

Which statement about R1 is true?

- A. OSPF redistributes RIP routes only if they have a tag of one.
- B. RIP learned routes are distributed to OSPF with a tag value of one.

C. R1 adds one to the metric for RIP learned routes before redistributing to OSPF.

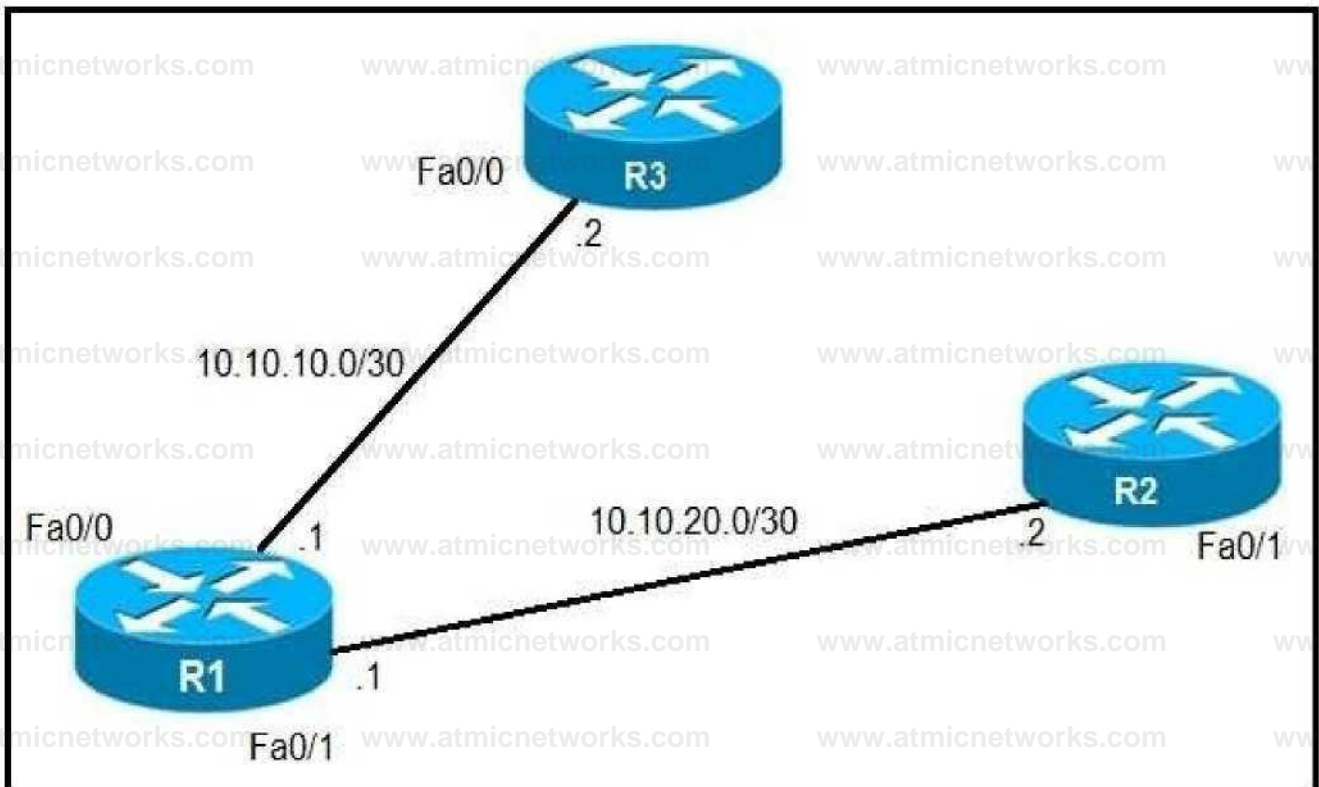
D. RIP routes are redistributed to OSPF without any changes.

Answer: B

Explanation:

Question: 21

Refer to the exhibit.



An IP SLA was configured on router R1 that allows the default route to be modified in the event that Fa0/0 loses reachability with the router R3 Fa0/0 interface. The route has changed to flow through

router R2. Which debug command is used to troubleshoot this issue?

A. debug ip flow

B. debug ip sla error

C. debug ip routing

D. debug ip packet

Answer: C

Explanation:

debug ip routing This command enables debugging messages related to the routing table.

Question: 22

Which configuration enabled the VRF that is labeled "Inet" on FastEthernet0/0?

A. R1(config)# ip vrf InetR1(config-vrf)#interface FastEthernet0/0R1(config-if)#ip vrf forwarding Inet

B. R1(config)#router ospf 1 vrf InetR1(config-router)#ip vrf forwarding FastEthernet0/0

C. R1(config)#ip vrf Inet FastEthernet0/0

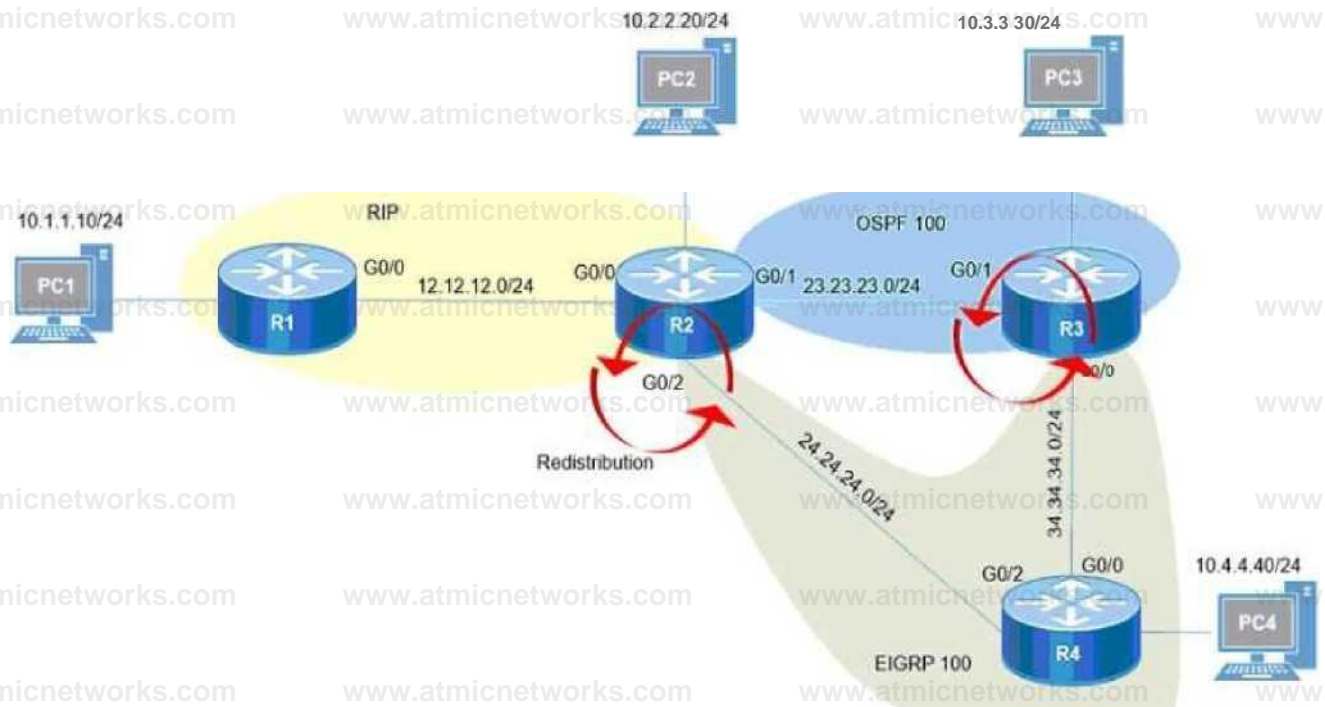
D. R1(config)# ip vrf InetR1(config-vrf)#ip vrf FastEthernet0/0

Answer: A

Explanation:

Question: 23

Refer to the exhibit.



After redistribution is enabled between the routing protocols; PC2, PC3, and PC4 cannot reach PC1. Which action can the engineer take to solve the issue so that all the PCs are reachable?

- A. Set the administrative distance 100 under the RIP process on R2.
- B. Filter the prefix 10.1.1.0/24 when redistributed from OSPF to EIGRP.
- C. Filter the prefix 10.1.1.0/24 when redistributed from RIP to EIGRP.
- D. Redistribute the directly connected interfaces on R2.

Answer: A

Explanation:

Question: 24

Which command allows traffic to load-balance in an MPLS Layer 3 VPN configuration?

- A. multi-paths eibgp 2
- B. maximum-paths 2
- C. Maximum-paths ibgp 2
- D. multi-paths 2


Answer: C

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_vpn_multipath.html

Question: 25

Refer to the exhibit.



```

on R1
R2(config)# crypto isakmp policy 10
R2(config-isakmp) # hash md5
R2(config-isakmp) # authentication pre-share
R2(config-isakmp) # group 2
R2(config-isakmp)# encryption 3des
R2(config)# crypto isakmp key cisco address 10.1.1.1
R2(config)# crypto ipsec transform-set T3ET esp-des esp-md-hmac
R2(cfg-crypto-trans)# mode transport
R2(config)# crypto ipsec profile TST R2 (ipsec-profile) # set transform-set TSET
R2(config)# interface tunnel 123
R2(config-if)# tunnel protection ipsec profile TST

```

```

on R3:
R3(config)# crypto isakmp policy 10
R3(conf g-isakmp) # hash md5
R3(config-isakmp) # authentication pre-share
R3(conf g-isakmp) # group 2
R3(conf gdsakmp)# encryption 3des
R3(con1ig)# crypto isakmp key cisco address 10.1.1.1
R3(config)# crypto ipsec transform-set TSET esp-des esp-md5hmac
R3(cfg-crypto-trans)# mode tunnel
R3(config)# crypto ipsec profile TST R3 (ipsec-profile) + set transform-set TSET
R3(config)# interface tunnel 123
R3(config-if)# tunnel protection ipsec profile TST

```

After applying IPsec, the engineer observed that the DMVPN tunnel went down, and both spoke-to-spoke and hub were not establishing. Which two actions resolve the issue? (Choose two.)

- A. Configure the crypto isakmp key cisco address 192.1.1.1 on R2 and R3
- B. Configure the crypto isakmp key cisco address 0.0.0.0 on R2 and R3.
- C. Change the mode from mode transport to mode tunnel on R3
- D. Change the mode from mode transport to mode tunnel on R2.
- E. Remove the crypto isakmp key cisco address 10.1.1.1 on R2 and R3

Answer: A,D

Explanation:

*When using DMVPN with IPsec, it is unnecessary to use tunnel mode. Because DMVPN uses GRE which means that a new IP header is already added by GRE. The GRE encapsulation happens on the tunnel interface before the encryption process takes place.

Question: 26

Which statement about route distinguishers in an MPLS network is true?

- A. Route distinguishers allow multiple instances of a routing table to coexist within the edge router.
- B. Route distinguishers are used for label bindings.
- C. Route distinguishers make a unique VPNv4 address across the MPLS network.
- D. Route distinguishers define which prefixes are imported and exported on the edge router.

Answer: C

Explanation:

Question: 27

Which statement about MPLS LDP router ID is true?

- A. If not configured, the operational physical interface is chosen as the router ID even if a loopback is configured.
- B. The loopback with the highest IP address is selected as the router ID.
- C. The MPLS LDP router ID must match the IGP router ID.
- D. The force keyword changes the router ID to the specified address without causing any impact.

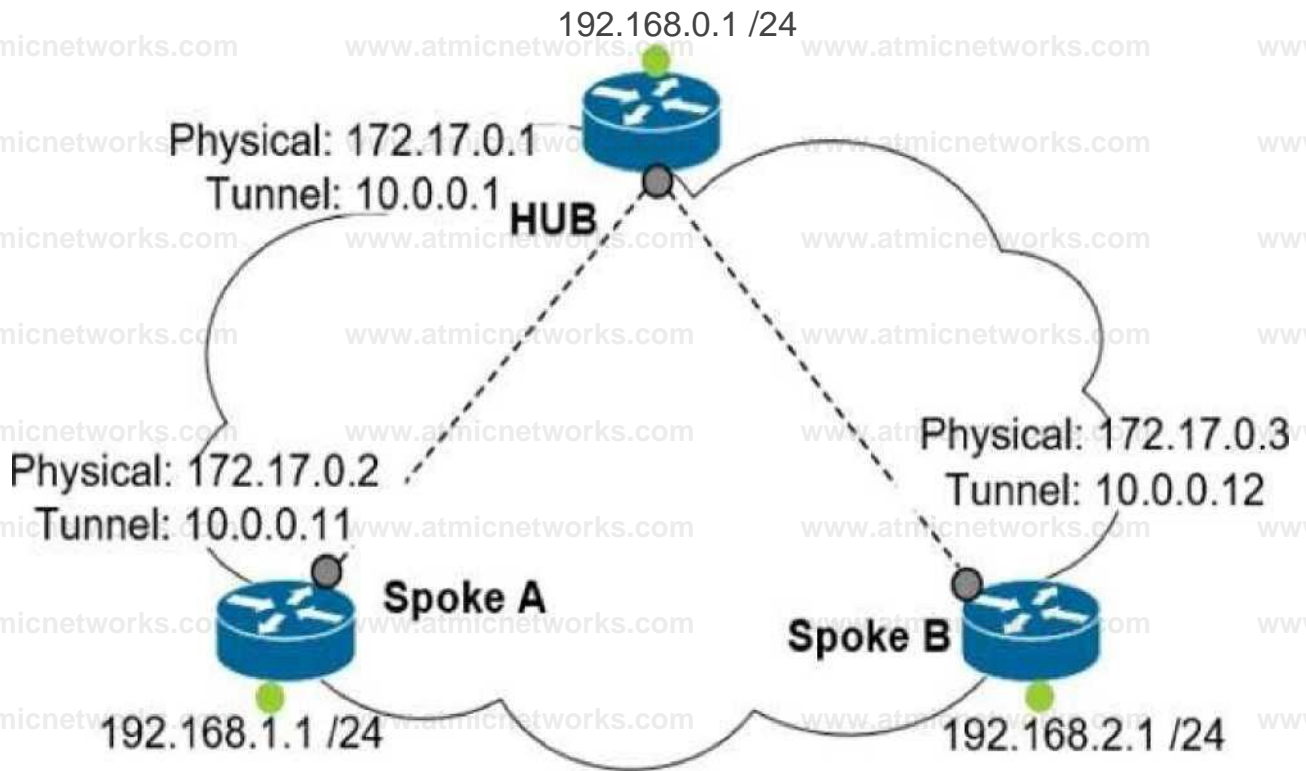
Answer: B

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/12-4m/mp_ldp-12-4mbook.pdf

Question: 28

Refer to the exhibit.



Which interface configuration must be configured on the spoke A router to enable a dynamic DMVPN tunnel with the spoke B router?

A interface Tunnel0 description mGRE - DMVPN Tunnel ip address 10.0.0.11 255.255.255.0 ip nhrp map multicast dynamic ip nhrp network-id 1 tunnel source 10.0.0.1 tunnel destination FastEthernet 0/0 tunnel mode gre multipoint

0 interface Tunnel0 ip address 10.0.0.11 255.255.255.0 ip nhrp network-id 1 tunnel source FastEthernet 0/0

tunnel mode gre multipoint ip nhrp nhs 10.0.0.1 ip nhrp
map 10.0.0.1172.17.0.1


```
interface Tunnel0
ip address 10.1.0.11 255.255.255.0
ip nhrp network-id 1 tunnel source 1.1.1.10
ip nhrp map 10.0.0.11 172.17.0.2 tunnel mode gre
```

```
u interface Tunnel0
ip address 10.0.0.11 255.255.255.0 ip nhrp map multicast static
ip nhrp network-id 1 tunnel source 10.0.0.1 tunnel mode gre
multipoint
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 29

Which list defines the contents of an MPLS label?

- A. 20-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL
- B. 32-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL

C. 20-bit label; 3-bit flow label; 1-bit bottom stack; 8-bit hop limit

D. 32-bit label; 3-bit flow label; 1-bit bottom stack; 8-bit hop limit

Answer: A

Explanation:

The first 20 bits constitute a label, which can have 2^{20} values. Next comes 3 bit value called Traffic Class. It was formerly called as experimental (EXP) field. Now it has been renamed to Traffic Class (TC). This field is used for QoS related functions. Ingress router can classify the packet according to some criterion and assign a 3 bit value to this field. If an incoming packet is marked with some IP Precedence or DSCP value and the ingress router may use such a field to assign an FEC to the packet. Next bit is Stack bit which is called bottom-of-stack bit. This field is used when more than one label is assigned to a packet, as in the case of MPLS VPNs or MPLS TE. Next byte is MPLS TTL field which serves the same purpose as that of IP TTL byte in the IP header

Reference: <https://tools.ietf.org/html/rfc5462>

Question: 30

Refer to the exhibit.

```
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

What does the imp-null tag represent in the MPLS VPN cloud?

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

Answer: A

Explanation:

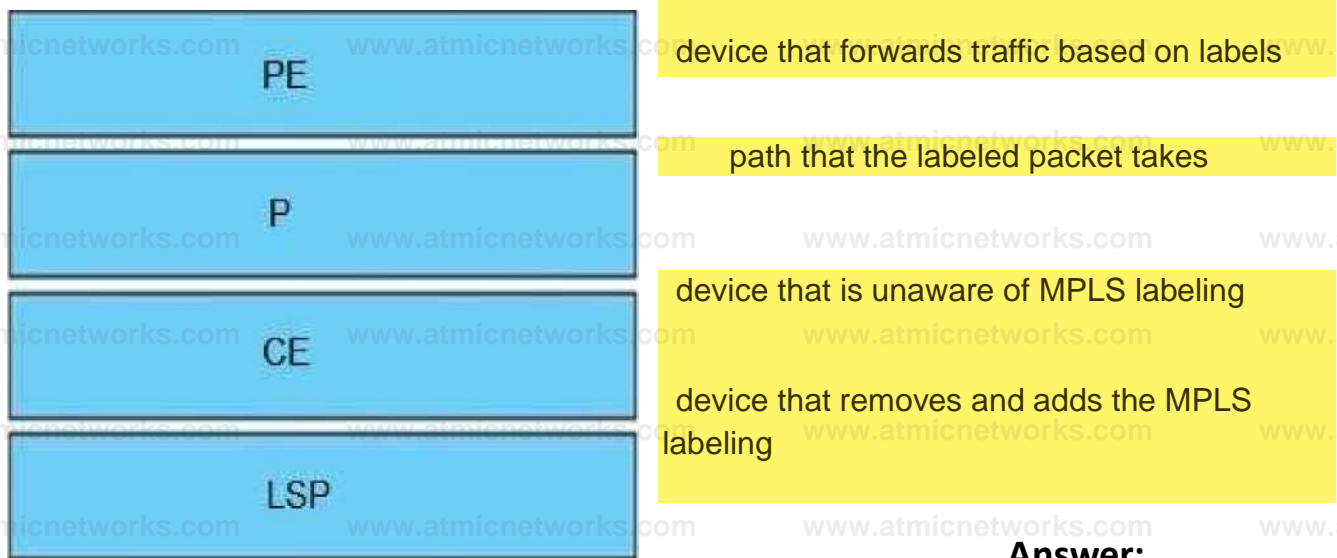
The `—imp-null` (implicit null) tag instructs the upstream router to pop the tag entry off the tag stack before forwarding the packet.

Note: pop means `—remove the top MPLS label`

Question: 31

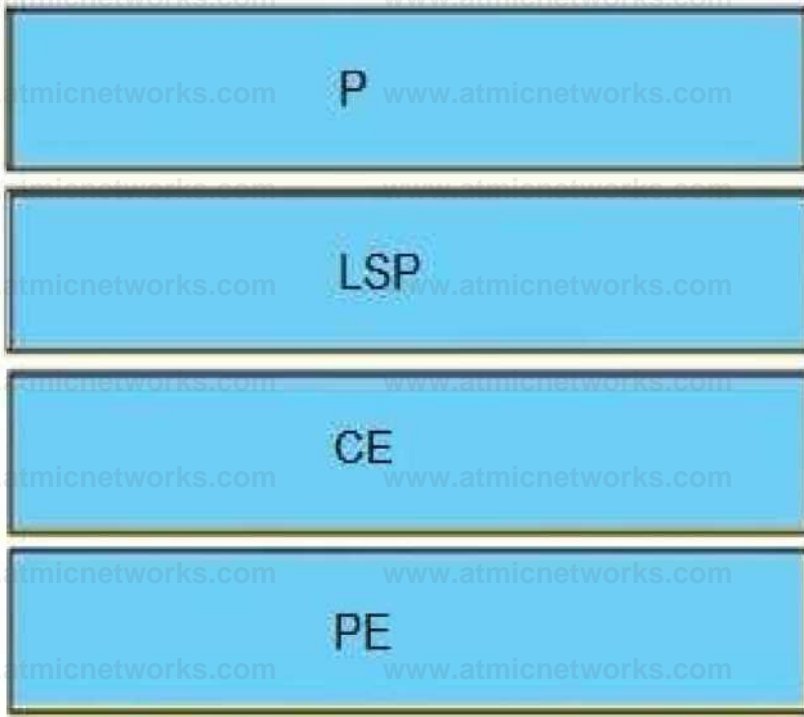
DRAG DROP

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



Answer:

Explanation:



Question: 32

Which transport layer protocol is used to form LDP sessions?

- A. UDP
- B. SCTP
- C. TCP
- D. RDP

Answer: C

Explanation:

LDP multicasts hello messages to a well-known UDP port (646) in order to discover neighbors. Once the discovery is accomplished, a TCP connection (port 646) is established and the LDP session begins. LDP keepalives ensure the health of the session. Thanks to the LDP session, LDP messages create the label mappings required for a FEC. Withdraw messages are used when FECs need to be torn down.

Question: 33

DRAG DROP

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

propagates VPN reachability information

route distinguisher

distributes labels for traffic engineering

route target

uniquely identifies a customer prefix

Resource Reservation Protocol

controls the import/export of customer prefixes

multiprotocol BGP

Explanation:

Answer:

multiprotocol BGP

Resource Reservation Protocol

route distinguisher


route target

Reference:

<https://www.rogerperkin.co.uk/featured/route-distinguisher-vs-route-target/>

Question: 34

Refer to the exhibits.



On R1:

```
R1(config)# interface tunnel 1
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# tunnel source 192.1.1.1
R1(config-if)# tunnel mode gre multipoint
R1(config-if)# ip nhrp network-id 111
```

On R2:

```
R2(config)# interface tunnel 1
R2(config-if)# ip address 10.1.1.2 255.255.255.0
R2(config-if)# tunnel source FastEthernet0/0
R2(config-if)# tunnel mode gre multipoint
R2(config-if)# ip nhrp network-id 222
R2(config-if)# ip nhrp nhs 10.1.1.1
R2(config-if)# ip nhrp map 10.1.1.1 192.1.1.1
```

On R3:

```
R3(config)# interface tunnel 1
R3(config-if)# ip address 10.1.1.3 255.255.255.0
R3(config-if)# tunnel source FastEthernet0/0
R3(config-if)# tunnel mode gre multipoint
R3(config-if)# ip nhrp network-id 333 R3(config-if)# ip nhrp nhs 10.1.1.1
R3(config-if)# ip nhrp map 10.1.1.1 192.1.1.1
```

On R4: R4(config)# interface tunnel 1

```
R4(config-if)# ip address 10.1.1.4 255.255.255.0
R4(config-if)# tunnel source FastEthernet0/0
R4(config-if)# tunnel mode gre multipoint
R4(config-if)# ip nhrp network-id 444
R4(config-if)# ip nhrp nhs 10.1.1.1
R4(config-if)# ip nhrp map 10.1.1.1 192.1.1.1
```


Phase-3 tunnels cannot be established between spoke-to-spoke in DMVPN. Which two commands are missing? (Choose two.)

- A. The ip nhrp redirect command is missing on the spoke routers.
- B. The ip nhrp shortcut command is missing on the spoke routers.
- C. The ip nhrp redirect commands is missing on the hub router.
- D. The ip nhrp shortcut commands is missing on the hub router.
- E. The ip nhrp map command is missing on the hub router.

Answer: B,C

Explanation:

Question: 35

Which protocol is used to determine the NBMA address on the other end of a tunnel when mGRE is used?

- A. NHRP
- B. IPsec
- C. MP-BGP
- D. OSPF

Answer: A

Explanation:

Question: 36

Refer to the exhibit.



Which configuration denies Telnet traffic to router 2 from 198A:0:200C::1/64?

A) `ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host 201A:0:205C::1/64 eq telnet`

`int Gi0/0`

`ipv6 traffic-filter Deny_Telnet in`

B) `ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host 201A:0:205C::1/64 eq telnet`

`int Gi0/0`

`ipv6 access-map Deny_Telnet in`

|

C)

`ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host 201A:0:205C::1/64`

|

`int Gi0/0`

`ipv6 access-map Deny_Telnet in`

D)

```
ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host  
201A:0:205C::1/64
```

```
int Gi0/0
```

```
ipv6 traffic-filter Deny_Telnet in
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 37

Refer to the exhibit.

```
access-list 100 deny tcp any any eq 465 access-list 100 deny tcp  
any eq 465 any access-list 100 permit tcp any any eq 80 access-  
list 100 permit tcp any eq 80 any access-list 100 permit udp any  
any eq 443 access-list 100 permit udp any eq 443 any
```

During troubleshooting it was discovered that the device is not reachable using a secure web browser. What is

needed to fix the problem?

- A. permit tcp port 443
- B. permit udp port 465
- C. permit tcp port 465
- D. permit tcp port 22

Answer: A

Explanation:

Question: 38

DRAG DROP

Drag and drop the packet types from the left onto the correct descriptions on the right.

data plane packets

user-generated packets that are always forwarded by network devices to other end-station devices

control plane packets

network device generated or received packets that are used for the creation of the network itself

management plane packets

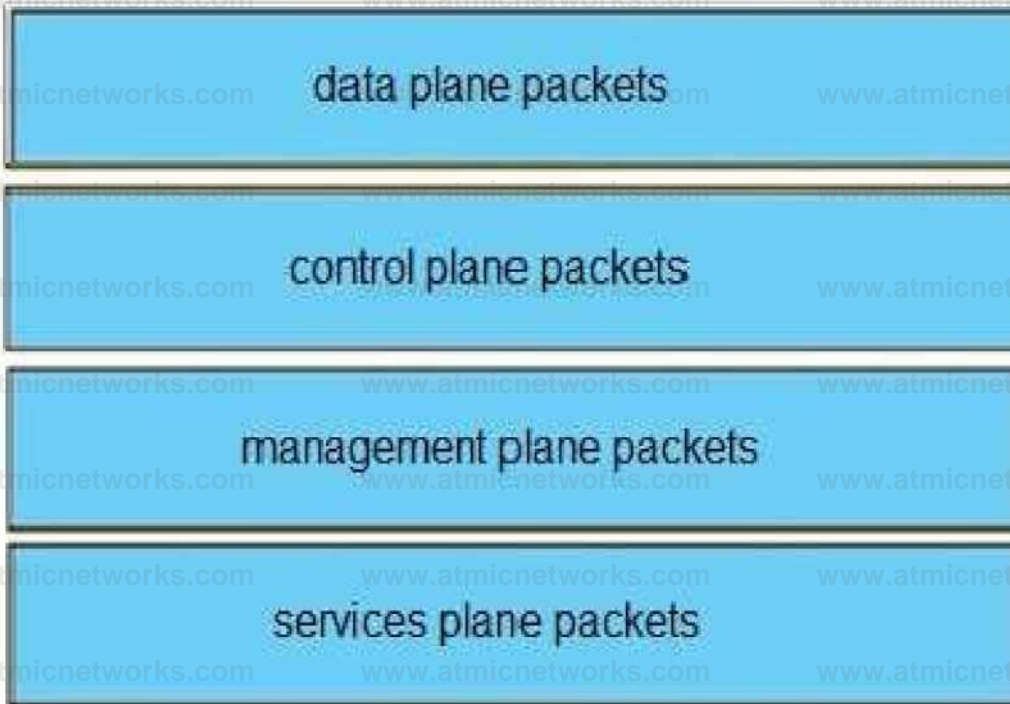
network device generated or received packets; packets that are used to operate the network

services plane packets

user-generated packets that are forwarded by network devices to other end-station devices, but that require higher priority than the normal traffic by the network devices

Answer:

Explanation:



Unlike legacy network technologies such as ISDN, Frame Relay, and ATM that defined separate data and control channels, IP carries all packets within a single pipe. Thus, IP network devices such as routers and switches must be able to distinguish between data plane, control plane, and management plane packets to treat each packet appropriately. From an IP traffic plane perspective, packets may be divided into four distinct, logical groups:

1. Data plane packets – End-station, user-generated packets that are always forwarded by network devices to other end-station devices. From the perspective of the network device, data plane packets always have a transit destination IP address and can be handled by normal, destination IP address-based forwarding processes.
2. Control plane packets – Network device generated or received packets that are used for the creation and operation of the network itself. From the perspective of the network device, control plane packets always have a receive destination IP address and are handled by the CPU in the network device route processor. Examples include protocols such as ARP, BGP, OSPF, and other protocols that glue the network together.
3. Management plane packets – Network device generated or received packets, or management station generated or received packets that are used to manage the network. From the perspective of the network device, management plane packets always have a receive destination IP

address and are handled by the CPU in the network device route processor. Examples include protocols such as Telnet, Secure Shell (SSH), TFTP, SNMP, FTP, NTP, and other protocols used to manage the device and/or network.4. Services plane packets – A special case of data plane packets, services plane packets are also user-generated packets that are also forwarded by network devices to other end-station devices, but that require high-touch handling by the network device (above and beyond normal, destination IP address-based forwarding) to forward the packet. Examples of high- touch handling include such functions as GRE encapsulation, QoS, MPLS VPNs, and SSL/IPsec encryption/decryption, etc. From the perspective of the network device, services plane packets may have a transit destination IP address, or may have a receive destination IP address (for example, in the case of a VPN tunnel endpoint).

Reference: https://tools.cisco.com/security/center/resources/copp_best_practices

Question: 39

DRAG DROP

Drag and drop the addresses from the left onto the correct IPv6 filter purposes on the right.

<pre>permit ip 2001:d8b:800:200c::/117 2001:0DBB:800:2010::/64 eq 443</pre>	<pre>Permit NTP from this source 2001:0D8B:0800:200c::1f</pre>
<pre>permit ip 2001:D88:800:200C::e/126 2001:0DBB:800:2010:764 eq 514</pre>	<pre>Permit syslog from this source 2001:0D88:0800:200c::1c</pre>
<pre>permit ip 2001 d8b 800 200c::800 /117 2001:0DBB:800:2010:764 eq 80</pre>	<pre>Permit HTTP from this source 2001:0D8B:0800:200c::0fff</pre>
<pre>permit ip 2001:D8B 800:200C::c/126 2001 ODBB 800:2010:764 eq 123</pre>	<pre>Permit HTTPS from this source 2001:0D8B:0800:200c::07ff</pre>

Answer:

Explanation:

```
permit ip 2001:D8B:800:200C::c/126  
2001:0DBB:800:2010::/64 eq 123
```

```
permit ip 2001:D88:800:200C::e/126  
2001:0DBB:800:2010::/64 eq 514
```

```
permit ip 2001:d8b:800:200c::800 /117  
2001:0DBB:800:2010::/64 eq 80
```

```
permit ip 2001:d8b:800:200c:: /117  
2001:0DBB:800:2010::/64 eq 443
```

HTTP and HTTPS run on TCP port 80 and 443, respectively and we have to remember them.

Syslog runs on UDP port 514 while NTP runs on UDP port 123 so if we remember them we can find out the matching answers easily. But maybe there is some typo in this question as 2001:d88:800:200c::c/126 only ranges from 2001:d88:800:200c:0:0:0:c to 2001:d88:800:200c:0:0:0:f (4 hosts in total). It does not cover host 2001:0D88:0800:200c::1f. Same for 2001:D88:800:200c::e/126, which also ranges from 2001:d88:800:200c:0:0:0:c to 2001:d88:800:200c:0:0:0:f and does not cover host 2001:0D88:0800:200c::1c.

Question: 40

Refer to the exhibit.

```
R1#show running-config | include aaa
aaa new-model
aaa authentication login default group tacacs+ local
aaa authentication login Console local
R1#show running-config | section line
line con 0
logging synchronous
R1#
```

An engineer is trying to configure local authentication on the console line, but the device is trying to authenticate using TACACS+. Which action produces the desired configuration?

- A. Add the aaa authentication login default none command to the global configuration.
- B. Replace the capital "C" with a lowercase "c" in the aaa authentication login Console local command.
- C. Add the aaa authentication login default group tacacs+ local-case command to the global configuration.
- D. Add the login authentication Console command to the line configuration

Answer: D

Explanation:

Reference:

<https://community.cisco.com/t5/switching/how-to-define-login-local-for-console-0/td-p/2949493>

Question: 41

Refer to the exhibit.


```
R1#show ip ssh
```

```
SSH Disabled-version 1.99
```

```
%Please create RSA keys to enable SSH (and of atleast 768 bits for SSH v2)
```

```
Authentication timeout: 120 secs: Authentication retries: 3
```

```
Minimum expected Diffie Hellman key size: 1024 bits
```

```
IOS Keys in SEOSH format (ssh-rsa, base64 encoded): NONE R1#
```

An engineer is trying to connect to a device with SSH but cannot connect. The engineer connects by using the console and finds the displayed output when troubleshooting. Which command must be used in configuration mode to enable SSH on the device?

- A. no ip ssh disable
- B. ip ssh enable
- C. ip ssh version 2
- D. crypto key generate rsa

Answer: D

Explanation:

Question: 42

Which statement about IPv6 ND inspection is true?

- A. It learns and secures bindings for stateless autoconfiguration addresses in Layer 3 neighbor tables.
- B. It learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables.
- C. It learns and secures bindings for stateful autoconfiguration addresses in Layer 3 neighbor tables.
- D. It learns and secures bindings for stateful autoconfiguration addresses in Layer 2 neighbor tables.

Answer: B

Explanation:

IPv6 ND inspection learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables. IPv6 ND inspection analyzes neighbor discovery messages in order to build a trusted binding table database, and IPv6 neighbor discovery messages that do not have valid bindings are dropped. A neighbor discovery message is considered trustworthy if its IPv6-to-MAC mapping is verifiable.

This feature mitigates some of the inherent vulnerabilities for the neighbor discovery mechanism, such as attacks on duplicate address detection (DAD), address resolution, device discovery, and the neighbor cache.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ip6f-15-s-book/ip6-snooping.pdf

Question: 43

While troubleshooting connectivity issues to a router, these details are noticed:

Standard pings to all router interfaces, including loopbacks, are successful.

Data traffic is unaffected.

SNMP connectivity is intermittent.

SSH is either slow or disconnects frequently.

Which command must be configured first to troubleshoot this issue?

- A. show policy-map control-plane
- B. show policy-map
- C. show interface | inc drop
- D. show ip route

Answer: A

Explanation:

Question: 44

Refer to the exhibit.

**TAC+: TCP/IP open to 171.68.118.101/49 failed -
Destination unreachable; gateway or host down
AAA/AUTHEN (2546660185): status = ERROR
AAA/AUTHEN/START (2546660185): Method=LOCAL
AAA/AUTHEN (2546660185): status " FAIL
As1 CHAP: Unable to validate Response. Username chapuser: Authentication failure**

Why is user authentication being rejected?

- A. The TACACS+ server expects "user", but the NT client sends "domain/user".
- B. The TACACS+ server refuses the user because the user is set up for CHAP.
- C. The TACACS+ server is down, and the user is in the local database.
- D. The TACACS+ server is down, and the user is not in the local database.

Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-accesscontrol-system-tacacs-/13864-tacacs-pppdebug.html>


```
Cat3850-Stack-2# show policy-map
```

```
Policy Map LIMIT_EGP
```

```
Class BGP  
drop
```

```
Policy Map SHAPE_BGP
```

```
Class BGP  
Average Rate Traffic Shaping cir  
10000000 (bps)
```

```
Policy Map POLICE_BGP
```

```
Class BGP  
police cir 1000k be 1500  
conform-action transmit exceed-  
action transmit
```

```
Policy Map COPP
```

```
Class BGP  
police cir 1000k be 1500  
conform-action transmit  
exceed-action drop
```

Which control plane policy limits BGP traffic that is destined to the CPU to 1 Mbps and ignores BGP traffic that is sent at higher rate?

- A. policy-map SHAPE_BGP
- B. policy-map LIMIT_BGP
- C. policy-map POLICE_BGP
- D. policy-map COPP

Answer: D

Explanation:

Question: 46

Which statement about IPv6 RA Guard is true?

- A. It does not offer protection in environments where IPv6 traffic is tunneled.
- B. It cannot be configured on a switch port interface in the ingress direction.
- C. Packets that are dropped by IPv6 RA Guard cannot be spanned.
- D. It is not supported in hardware when TCAM is programmed.

Answer: A

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xs-3s/ip6f-xe-3s-book/ip6-ra-guard.html#GUID-589AF00C-7499-439F-AD23-51005D61CAB7

The IPv6 RA Guard feature does not offer protection in environments where IPv6 traffic is tunneled.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xs-16/ip6f-xe-16-book/ip6-ra-guard.pdf

Question: 47

An engineer is trying to copy an IOS file from one router to another router by using TFTP. Which two actions are needed to allow the file to copy? (Choose two.)

- A. Copy the file to the destination router with the copy tftp: flash: command

- B. Enable the TFTP server on the source router with the tftp-server flash: <filename> command
- C. TFTP is not supported in recent IOS versions, so an alternative method must be used
- D. Configure a user on the source router with the username tftp password tftp command
- E. Configure the TFTP authentication on the source router with the tftp-server authentication local command

Answer: A,B

Explanation:

Question: 48

Refer to the exhibit.

```
R1#show running-config | section dhcp
ip dhcp excluded-address 192.168.1.1 192.168.1.49
ip dhcp pool DHCP
  network 192.168.1.0 255.255.255.0
  default-router 192.168.1.1
  dns-server 8.8.8.8
  lease 0 12
```

Users report that IP addresses cannot be acquired from the DHCP server. The DHCP

server is configured as shown. About 300 total nonconcurrent users are using this DHCP server, but none of them are active for more than two hours per day. Which action fixes the issue within the current resources?

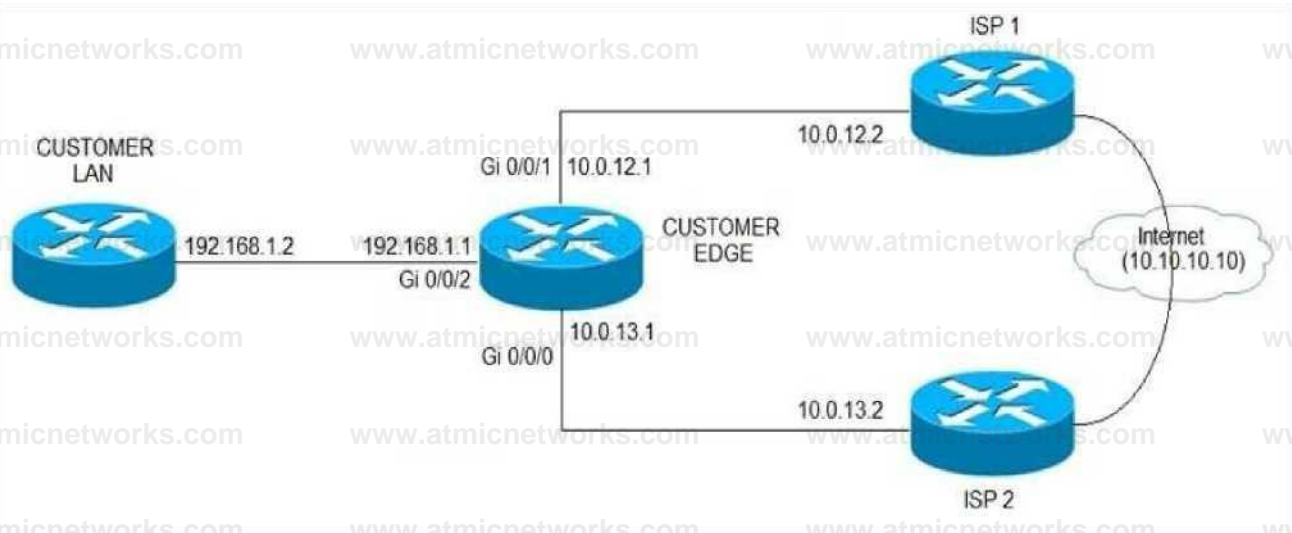
- A. Modify the subnet mask to the network 192.168.1.0 255.255.254.0 command in the DHCP pool
- B. Configure the DHCP lease time to a smaller value
- C. Configure the DHCP lease time to a bigger value
- D. Add the network 192.168.2.0 255.255.255.0 command to the DHCP pool

Answer: B

Explanation:

Question: 49

Refer to the exhibit.



ISP 1 and ISP 2 directly connect to the Internet. A customer is tracking both ISP links to

achieve redundancy and cannot see the Cisco IOS IP SLA tracking output on the router console. Which command is missing from the IP SLA configuration?

- A. Start-time 00:00
- B. Start-time 0
- C. Start-time immediately
- D. Start-time now

Answer: D

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_icmp_echo.html

Question: 50

Refer to the exhibit.

```
service timestamps debug datetime msec
service timestamps log datetime
clock timezone MST -7 0
clock summer-time MST recurring
ntp authentication-key 1 md5 00101A0B0152181206224747071E 7
ntp server 10 10 10 10
```

R1#show clock

```
*06 13 44 045 MST Sun Dec 30 2018
```

R1#conf t

Enter configuration commands, one per line End with CNTL/Z

R1 (config) #logging host 10.10.10.20

R1(config) #end

R1#

```
"Dec 30 13:15:28: %SYS-5-CONFIG_I: Configured from console by console
```

R1#

```
'Dec 30 13:15:28: %SYS-6-LOGGINGHOST_STARTSTOP: Logging to host 10.10.10.20 port 514 started - CLI initiated
```

An administrator noticed that after a change was made on R1, the timestamps on the system logs did not match the clock.

What is the reason for this error?

- A. An authentication error with the NTP server results in an incorrect timestamp.
- B. The keyword localtime is not defined on the timestamp service command.
- C. The NTP server is in a different time zone.
- D. The system clock is set incorrectly to summer-time hours.

Answer: B

Explanation:

Question: 51

DRAG DROP

Drag and drop the DHCP messages from the left onto the correct uses on the right.

DHCPACK	server-to-client communication, refusing the request for configuration parameters
DHCPINFORM	client-to-server communication, indicating that the network address is already in use
DHCPNAK	server-to-client communication with configuration parameters, including committed network address
DHCPDECLINE	client-to-server communication, asking for only local configuration parameters that the client has already externally configured as an address

Answer:

Explanation:

DHCPNAK

DHCPDECLINE

DHCPACK

DHCPINFORM

DHCPACK

The server-to-client communication with configuration parameters, including committed network address.

DHCPINFORM

The client-to-server communication, asking for only local configuration parameters that the client already has externally configured as an address.

DHCPNAK

The server-to-client communication, refusing the request for configuration parameter.

DHCPDECLINE

The client-to-server communication, indicating that the network address is already in use

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html>

DHCPINFORM: If a client has obtained a network address through some other means or has a manually configured IP address, a client workstation may use a DHCPINFORM request message to obtain other local configuration parameters, such as the domain name and Domain Name Servers (DNSs). DHCP servers receiving a DHCPINFORM message construct a DHCPACK message with any local configuration parameters appropriate for the client without allocating a new IP address. This

DHCPACK will be sent unicast to the client.

DHCPNAK: If the selected server is unable to satisfy the DHCPREQUEST message, the DHCP server will respond with a DHCPNAK message. When the client receives a DHCPNAK message, or does not receive a response to a DHCPREQUEST message, the client restarts the configuration process by going into the Requesting state. The client will retransmit the DHCPREQUEST at least four times within 60 seconds before restarting the Initializing state.

DHCPACK: After the DHCP server receives the DHCPREQUEST, it acknowledges the request with a DHCPACK message, thus completing the initialization process.

DHCPDECLINE: The client receives the DHCPACK and will optionally perform a final check on the parameters. The client performs this procedure by sending Address Resolution Protocol (ARP) requests for the IP address provided in the DHCPACK. If the client detects that the address is already in use by receiving a reply to the ARP request, the client will send a DHCPDECLINE message to the server and restart the configuration process by going into the Requesting state.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html>

Question: 52

A network engineer is investigating a flapping (up/down) interface issue on a core switch that is synchronized to an NTP server. Log output currently does not show the time of the flap. Which command allows the logging on the switch to show the time of the flap according to the clock on the device?

- A. service timestamps log uptime
- B. clock summer-time mst recurring 2 Sunday mar 2:00 1 Sunday nov 2:00
- C. service timestamps log datetime localtime show-timezone
- D. clock calendar-valid

Answer: C

Explanation:

By default, Catalyst switches add a simple uptime timestamp to logging messages. This is a cumulative counter that shows the hours, minutes, and seconds since the switch has been booted up

Question: 53

When provisioning a device in Cisco DNA Center, the engineer sees the error message “Cannot select the device. Not compatible with template”.

- What is the reason for the error?
- A. The template has an incorrect configuration.
 - B. The software version of the template is different from the software version of the device.
 - C. The changes to the template were not committed.
 - D. The tag that was used to filter the templates does not match the device tag.

Answer: D

Explanation:

If you use tags to filter the templates, you must apply the same tags to the device to which you want to apply the templates. Otherwise, you get the following error during provisioning: —Cannot select the device. Not compatible with template.||

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-2-10/user_guide/b_cisco_dna_center_ug_1_2_10/b_dnac_ug_1_2_10_chapter_0111.html

Question: 54

While working with software images, an engineer observes that Cisco DNA Center cannot upload its software image directly from the device. Why is the image not uploading?

- A. The device must be resynced to Cisco DNA Center.
- B. The software image for the device is in install mode.
- C. The device has lost connectivity to Cisco DNA Center.
- D. The software image for the device is in bundle mode

Answer: B

Explanation:

Upload Software Images for Devices in Install Mode

The Image Repository page might show a software image as being in Install Mode. When a device is in Install Mode, Cisco DNA Center is unable to upload its software image directly from the device. When a device is in install mode, you must first manually upload the software image to the Cisco DNA Center repository before marking the image as golden, as shown in the following steps.

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-andmanagement/dna-center/1-2-10/user_guide/b_cisco_dna_center_ug_1_2_10/b_dnac_ug_1_2_10_chapter_0100.html

Question: 55

An engineer configured the wrong default gateway for the Cisco DNA Center enterprise interface during the install. Which command must the engineer run to correct the configuration?

- A. `sudo maglev-config update`
- B. `sudo maglev install config update`
- C. `sudo maglev reinstall`
- D. `sudo update config install`

Answer: A

Explanation:

Question: 56

DRAG DROP

Drag and drop the SNMP attributes in Cisco IOS devices from the left onto the correct SNMPv2c or SNMPV3 categories on the right.

community string	SNMPv2c
username and password	
authentication	
no encryption	
privileged	SNMPv3
read-only	

Answer:

Explanation:

SNMP

v2c

community
string

no
encryptio

n read-
only

SNM

Pv3

username and password

authentic
ation

privileged

Graphical user interface, application Description automatically generated

Question:

57

Refer to the exhibit.

```
R1 (config) # do show running-config | section line|username
username cisco secret 5 $1Syb/o$L3G5cXODxpYMSJ70PzEyoO
line con 0
  logging synchronous
line vty 0 4
  login local
  transport input telnet
R1 (config) # logging console 7
R1 (config) # do debug aaa authentication
R1 (config)#
```

An administrator that is connected to the console does not see debug messages when remote users log in. Which action ensures that debug messages are displayed for remote logins?

- A. Enter the transport input ssh configuration command.
- B. Enter the terminal monitor exec command.
- C. Enter the logging console debugging configuration command.
- D. Enter the aaa new-model configuration command.

Answer: C

Explanation:

The `—logging console` is a default and hidden command.

Question: 58

Refer to the exhibit.

```
snmp-server community ciscotest1
snmp-server host 192.168.1.128 ciscotest
snmp-sever enable traps bgp
```

Network operations cannot read or write any configuration on the device with this configuration from the operations subnet. Which two configurations fix the issue? (Choose two.)

- A. Configure SNMP rw permission in addition to community ciscotest.
- B. Modify access list 1 and allow operations subnet in the access list.
- C. Modify access list 1 and allow SNMP in the access list.
- D. Configure SNMP rw permission in addition to version 1.
- E. Configure SNMP rw permission in addition to community ciscotest 1.

Answer: B,E

Explanation:

Question: 59

Refer to the exhibit.

```
config t
flow record v4_r1
match ipv4 tos
match ipv4 protocol
match ipv4 source address
match ipv4 destination address
match transport source-port
match transport destination-port
collect counter bytes long
collect counter packets long
!
flow exporter EXPORTER-1
destination 172.16.10.2
transport udp 90
exit
!
flow monitor FLOW-MONITOR-1
record v4_r1
exit
!
ip cef
!
interface Ethernet0/0.1
ip address 172.16.6.2 255.255.255.0
ip flow monitor FLOW-MONITOR-1 input
!
```

Why is the remote NetFlow server failing to receive the NetFlow data?

- A. The flow exporter is configured but is not used.
- B. The flow monitor is applied in the wrong direction.
- C. The flow monitor is applied to the wrong interface.
- D. The destination of the flow exporter is not reachable.

Answer: A

Explanation:

Question: 60

Refer to the exhibit.

```
!
neighbor 10.222.1.1 route-map SET-WEIGHT in
neighbor 10.222.1.1 remote-as 1
!
ip as-path access-list 200 permit ^690$
ip as-path access-list 200 permit ^1800
!
route-map SET-WEIGHT permit 10
match as-path 200
set local-preference 250
set weight 200
```

A router receiving BGP routing updates from multiple neighbors for routers in AS 690. What is the reason that the router still sends traffic that is destined to AS 690 to a neighbor other than 10.222.1.1?

- A. The local preference value in another neighbor statement is higher than 250.
- B. The local preference value should be set to the same value as the weight in the route map.

- C. The route map is applied in the wrong direction.
- D. The weight value in another neighbor statement is higher than 200.

Answer: C

Explanation:

Question: 61

Refer to the exhibit.

```
router# show ip route
****
D 192.168.32.0/19 [90/25789217] via 10.1.1.1
R 192.168.32.0/24 [120/4] via 10.1.1.2
O 192.168.32.0/26 [110/229840] via 10.1.1.3
```

Refer to the exhibit. an engineer is trying to get 192.168.32.100 forwarded through 10.1.1.1, but it was forwarded through 10.1.1.2. What action forwards the packets through 10.1.1.1?

- A. Configure EIGRP to receive 192.168.32.0 route with lower admin distance.
- B. A. Configure EIGRP to receive 192.168.32.0 route with longer prefix than /19.
- C. A. Configure EIGRP to receive 192.168.32.0 route with lower metric.
- D. A. Configure EIGRP to receive 192.168.32.0 route with equal or longer prefix than /24.

Answer: D

Explanation:

Question: 62

What is a limitation of IPv6 RA Guard?

- A. It is not supported in hardware when TCAM is programmed
- B. It does not offer protection in environments where IPv6 traffic is tunneled.
- C. It cannot be configured on a switch port interface in the ingress direction
- D. Packets that are dropped by IPv6 RA Guard cannot be spanned

Answer: B

Explanation:

Restrictions for IPv6 RA Guard

The IPv6 RA Guard feature does not offer protection in environments where IPv6 traffic is tunneled.

This feature is supported only in hardware when the ternary content addressable memory (TCAM) is programmed.

This feature can be configured on a switch port interface in the ingress direction.

This feature supports host mode and router mode.

This feature is supported only in the ingress direction; it is not supported in the egress direction.

This feature is not supported on EtherChannel and EtherChannel port members.

This feature is not supported on trunk ports with merge mode.

This feature is supported on auxiliary VLANs and private VLANs (PVLANS). In the case of PVLANS, primary VLAN features are inherited and merged with port features.

Packets dropped by the IPv6 RA Guard feature can be spanned.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16-10/ip6f-xe-16-10-book/ip6-ra-guard.html#GUID-589AF00C-7499-439F-AD23-51005D61CAB7

Question: 63

Refer to the exhibit.

```
R1(config)# ip route 0.0.0.0 0.0.0.0 1.1.1.1
R1(config)# ip route 0.0.0.0 0.0.0.0 2.2.2.2 10
R1(config)# ip sla 1
R1(config)# icmp-echo 1.1.1.1 source-interface FastEthernet0/0
R1(config)# ip sla schedule 1 life forever start-time now

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

An IP SLA is configured to use the backup default route when the primary is down, but it is not working as desired. Which command fixes the issue?

- A. R1(config)# ip route 0.0.0.0 0.0.0.0 2.2.2.2 10 track 1
- B. R1(config)# ip route 0.0.0.0 0.0.0.0 2.2.2.2
- C. R1(config)# ip sla track 1
- D. R1(config)# ip route 0.0.0.0 0.0.0.0 1.1.1.1 track 1

Answer: D

Explanation:

Reference:

Note: By default Static Router AD value-1 hence ip route 0.0.0.0. 0.0.0.0. 1.1.1.1 track 1 means AD-1 which must be less than of back up route AD.

Define the backup route to use when the tracked object is unavailable. !--- The administrative distance of the backup route must be greater than !--- the administrative distance of the tracked route.!--- If the primary gateway is unreachable, that route is removed!--- and the backup route is installed in the routing table!--- instead of the tracked route.

<https://www.cisco.com/c/en/us/support/docs/ip/ip-routing/200785-ISP-Failover-with-default-routes-using-l.html>

<https://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/118962-configure-asa-00.html>

Question: 64

Which label operations are performed by a label edge router?

- A. SWAP and POP
- B. SWAP and PUSH
- C. PUSH and PHP
- D. PUSH and POP

Answer: D

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.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/mpls/configuration/guide/mpls_cg/mp_mpls_overview.pdf

Question: 65

Refer to the exhibit.

```
BRANCH-RTR#
router eigrp 100
 network 10.4.31.0 0.0.0.7
 network 10.100.100.1 0.0.0.0
 distribute-list route-map FILTER-IN in FastEthernet0/0
 eigrp router-id 10.100.100.1
!
ip prefix-list 102 seq 10 permit 10.1.1.100/32
!
route-map FILTER-IN deny 10
 match ip address prefix-list 102
!
```

A junior engineer updated a branch router configuration. Immediately after the change, the engineer receives calls from the help desk that branch personnel cannot reach any network destinations. Which configuration restores service and continues to block 10.1.1.100/32?

- A. route-map FILTER-IN deny 5
- B. ip prefix-list 102 seq 15 permit 0.0.0.0/32 le 32

C. ip prefix-list 102 seq 5 permit 0.0.0.0/32 le 32

D. route-map FILTER-IN permit 20

Answer: D

Explanation:

By using “deny” keyword in a route-map, we can filter out the prefix specified in the prefix-list. But there is an implicit

“deny all” statement in the prefix-list so we must permit other prefixes with “permit” keyword in the route-

map.

Question: 66

An engineer configured a leak-map command to summarize EIGRP routes and advertise specifically loopback 0 with an IP of 10.1.1.1.255.255.255.252 along with the summary route. After finishing configuration, the customer complained not receiving summary route with specific loopback address. Which two configurations will fix it? (Choose two.)

```
router eigrp 1
```

```
route-map Leak-Route deny 10
```

```
interface Serial 0/0
```

```
ip summary-address eigrp 110.0.0.0 255.0.0.0 leak-map Leak-Route
```

A. Configure access-list 1 permit 10.1.1.0.0.0.3.

B. Configure access-list 1 permit 10.1.1.1.0.0.252.

C. Configure access-list 1 and match under route-map Leak-Route.

D. Configure route-map Leak-Route permit 10 and match access-list 1.

E. Configure route-map Leak-Route permit 20.

Answer: A,D

Explanation:

When you configure an EIGRP summary route, all networks that fall within the range of your summary are suppressed and no longer advertised on the interface. Only the summary route is advertised. But if we want to advertise a network that has been suppressed along with the summary route then we can use leak-map feature. The below

commands will fix the configuration in this

question:

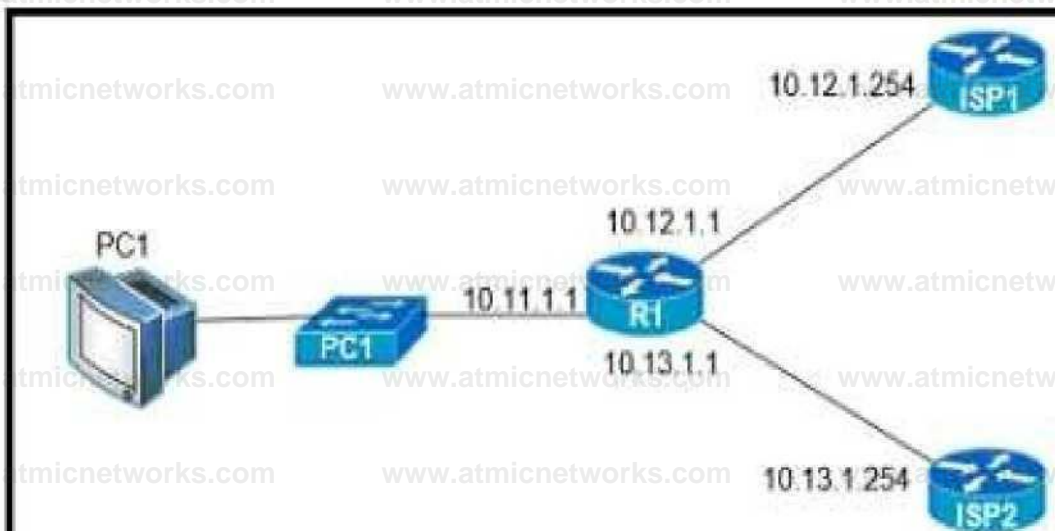
```
R1(config)#access-list 1 permit 10.1.1.0 0.0.0.3
```

```
R1(config)#route-map Leak-Route permit 10 // this command will also remove the "route_map  
Leak-Route deny 10" command.
```

```
R1(config-route-map)#match ip address 1
```

Question: 67

Refer to the exhibit.



```

R1
ip sla 100
icmp-echo 10.12.1.254
|
track 10 ip sla 100 reachability
|
ip route 0.0.0.0 0.0.0.0 10.12.1.254 track 10
ip route 0.0.0.0 0.0.0.0 10.13.1.254 10

R1#show ip route
(Output Omitted)
Gateway of last resort is 10.13.1.254 to network 0.0.0.0

S* 0.0.0.0/0 [10/0] via 10.13.1.254
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C   10.11.1.0/24 is directly connected, GigabitEthernet0/1
L   10.11.1.1/32 is directly connected, GigabitEthernet0/1
C   10.12.1.0/24 is directly connected, GigabitEthernet0/0
L   10.12.1.1/32 is directly connected, GigabitEthernet0/0
C   10.13.1.0/24 is directly connected, GigabitEthernet0/2
L   10.13.1.1/32 is directly connected, GigabitEthernet0/2

```

An engineer is monitoring reachability of the configured default routes to ISP1 and ISP2. The default route from ISP1 is preferred if available. How is this issue resolved?

A. Use the icmp-echo command to track both default routes

- B. Use the same AD for both default routes
- C. Start IP SLA by matching numbers for track and ip sla commands
- D. Start IP SLA by defining frequency and scheduling it

Answer: D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/ip-routing/200785-ISP-Failover-with-default-routes-using-l.html>

In the above configuration we have not had activated our IP SLA operation. We can start it with this command:

```
R1(config)#ip sla schedule 100 life forever start-time now
```

Also we should specify the rate of ICMP echo:

```
R1(config-ip-sla-echo)#frequency 5 // Send ICMP echo every 5 seconds
```

Question: 68

After some changes in the routing policy, it is noticed that the router in AS 45123 is being used as a transit AS router for several service providers. Which configuration ensures that the branch router in AS 45123 advertises only the local networks to all SP neighbors?

A)

```
ip as-path access-list 1 permit ^43123
```

```
router bgp 46123  
neighbor SP-Neighbors filter-list 1 out
```

B)

```
ip as-path access-list 1 permit /
```

```
router bgp 45123
```

neighbor SP-Neighbors filter-list 1 out

c)

ip as-path access-list 1 permit *45123\$

router bgp 45123
neighbor SP-Neighbors filter-^# 1 out

D)

ip as-path access-list 1 permit *S

router bgp 45123
neighbor SP-Neighbors filter-bst 1 cut

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

By default BGP advertises all prefixes to external BGP neighbors. This means that if you are multi-homed (connected to two or more ISPs) then you might become a transit AS. For example, ISP 2 in AS 200 can send traffic to your

router in AS 100 to reach ISP 3 in AS 300 because you

advertised prefixes in ISP 3 to ISP 2.

This is what will be seen in the BGP routing table of ISP1:

```
ISP1#show ip bgp
--output omitted--
Network                Next Hop                Metric LocPrf Weight Path
* > 3.3.3.0/24          192.168.12.1           0 100 300 i
```

Question: 69

DRAG DROP

Drag and drop the operations from the left onto the locations where the operations are performed on the right.

assigns labels to unlabeled packets

handles traffic between multiple VPNs

reads the labels and forwards the packet based on the labels

performs penultimate hop popping

Label Switch Router

Label Edge Router

Answer:

Explanation:

Label Switch Router 1. Reads labels and forwards the packet based on the based on the label.

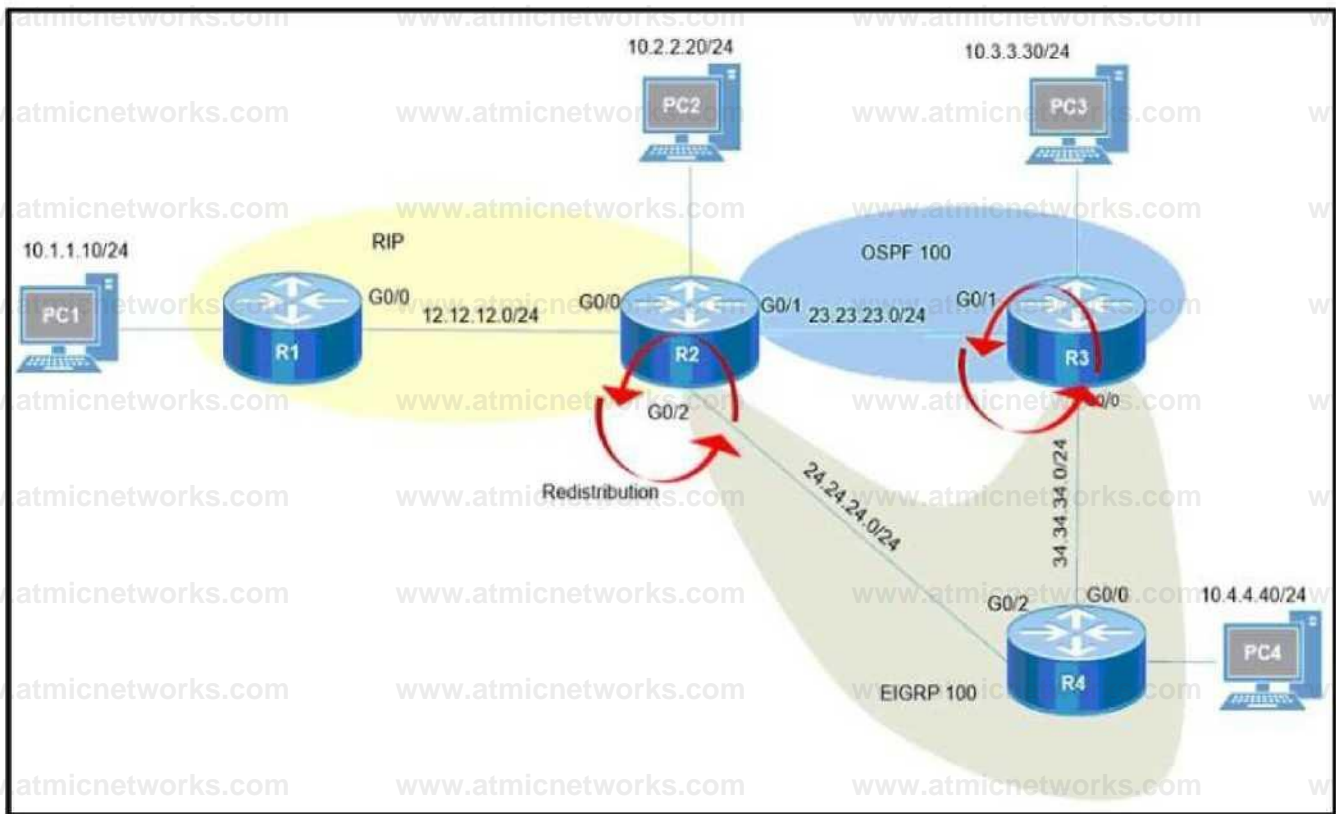
2. Performs PHP

Label Edge Router: 1 Assigns labels and unlabeled packets.

3. Handles traffic between multiple VPNs

Question: 70

Refer to the exhibit.



Redistribution is enabled between the routing protocols, and now PC2, PC3, and PC4 cannot reach PC1. What are the two solutions to fix the problem? (Choose two.)

- A. Filter RIP routes back into RIP when redistributing into RIP in R2
- B. Filter OSPF routes into RIP FROM EIGRP when redistributing into RIP in R2.
- C. Filter all routes except RIP routes when redistributing into EIGRP in R2.
- D. Filter RIP AND OSPF routes back into OSPF from EIGRP when redistributing into OSPF in R2
- E. Filter all routes except EIGRP routes when redistributing into OSPF in R3.

Answer: A,C

Explanation:

Even PC2 cannot reach PC1 so there is something wrong with RIP redistribution in R2. Because RIP has higher

Administrative Distance (AD) value than OSPF and EIGRP so it will be looped when doing mutual redistribution.

Question: 71

Refer to the exhibit.

```
R1#show policy-map control-plane
Control Plane
Class-map: NMS (match-all)
 500461 packets, 24038351 bytes
 5 minute offered rate 1390000 bps, drop rate 0 bps
police:
  cir 50000 bps, bc 5000 bytes
  conformed 50444 packets, 24031001 bytes; actions:
  transmit
  exceeded 990012 packets, 94030134 bytes; actions
  drop conformed 4000 bps, exceed 0 bps
R1#
```

A company is evaluating multiple network management system tools. Trending graphs generated by SNMP data are returned by the NMS and appear to have multiple gaps. While troubleshooting the issue, an engineer noticed the relevant output. What solves the gaps in the graphs?

- A. Remove the exceed-rate command in the class map.
- B. Remove the class map NMS from being part of control plane policing.
- C. Configure the CIR rate to a lower value that accommodates all the NMS tools
- D. Separate the NMS class map in multiple class maps based on the specific protocols with appropriate CoPP actions

Answer: D

Explanation:

Reference:

https://tools.cisco.com/security/center/resources/copp_best_practices

The class-map NMS in the exhibit did not classify traffic into specific protocols so many packets were dropped. We

should create some class-map to classify the receiving traffic. It is also a recommendation of CoPP/CP policy:

“Developing a CPP policy starts with the classification of the control plane traffic. To that end, the control plane traffic needs to be first identified and separated into different class maps.”

Question: 72

What is a role of route distinguishers in an MPLS network?

- A. Route distinguishers define which prefixes are imported and exported on the edge router
- B. Route distinguishers allow multiple instances of a routing table to coexist within the edge router.
- C. Route distinguishers are used for label bindings.
- D. Route distinguishers make a unique VPNv4 address across the MPLS network

Answer: D

Explanation:

Question: 73

Refer to the exhibit.

```
Global RADIUS shared secret:*****
retransmission count:5
timeout value:10
following RADIUS servers are configured:
  myradius.cisco.users.com:
    available for authentication on port:1814
    available for accounting on port:1813
  10.1.1.1:
    available for authentication on port:1814
    available for accounting on port:1813
    RADIUS shared secret:*****
  10.2.2.3:
    available for authentication on port:1814
    available for accounting on port:1813
    RADIUS shared secret:*****
```

AAA server 10.1.1.1 is configured with the default authentication and accounting settings, but the switch cannot communicate with the server Which action resolves this issue?

- A. Match the authentication port
- B. Match the accounting port
- C. Correct the timeout value.
- D. Correct the shared secret.

Answer: A

Explanation:

Command Default

Accounting port: 1813

Authentication port: 1812

Accounting: enabled

Authentication: enabled

Retransmission count: 1

Idle-time: 0

Server monitoring: disabled

Timeout: 5 seconds

Test username: test

Test password: test

Reference:

https://www.cisco.com/c/m/en_us/techdoc/dc/reference/cli/n5k/commands/radius-server-host.html

By default, RADIUS uses UDP port 1812 for authentication and port 1813 for accounting. In the exhibit above we see

port 1814 is being used for authentication to AAA server at 10.1.1.1 which is not the default port so we must adjust the

authentication port to the default value 1812.

Question: 74

DRAG DROP

Refer to the exhibit.

add new-model

aaa authentication login default none aaa authentication login telnet local

r 4

```
username cisco password 0 ocsic
```

```
j
```

```
line vty 0
```

```
password LecMeIn
```

```
login authentication telnet
```

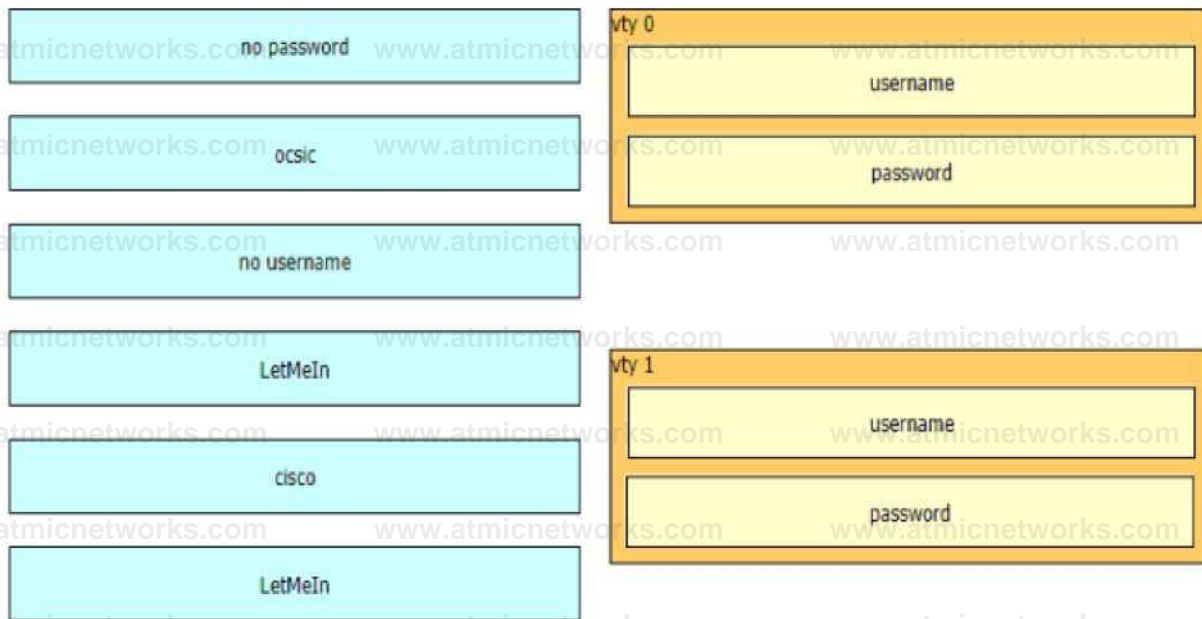
```
transp nut Input telnet
```

```
line vty 1
```

```
password LetNeIn
```

```
transport input telnet
```

Drag and drop the credentials from the left onto the remote login information on the right to resolve a failed login attempt to vtys. Not all credentials are of SLA by defining frequency and scheduling



Answer:

Explanation:

vty 0:

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

+ cisco

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

+ no

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

vtty 1:

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

+ no username

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

+ no password

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

The command "aaa authentication login default none" means no authentication is required when access to the device

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

via Console/VTY/AUX so if one interface does not specify another login authentication method (via the "login

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

authentication ..." command), it will allow to access without requiring username or password. In this case VTY 1 does

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

not specify another authentication login method so it will use the default method (which is "none" in this case).

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Question: 75

Refer to the exhibit.

```
Router Configuration:
ip vrf customer_a
  rd 1:1
  route-target export 1:1
  route-target import 1:1
!
!
interface FastEthernet0.1
  encapsulation dot1Q 2
  ip vrf forwarding customer_a
  ip address 192.168.4.1 255.255.255.0
!
router ospf 1
  log-adjacency-changes
!
router ospf 2 vrf customer_a
  log-adjacency-changes
  network 192.168.4.0 0.0.0.255 area 0
!
end
```

The network administrator configured VRF lite for customer

A. The technician at the remote site misconfigured VRF on the router. Which configuration will resolve connectivity for both sites of customer_a?


```
ip vrf customer a rd 1:1
route-target export 1:2 route-target import 1:2
```

```
ip vrf customer a rd 1:1
route-target import 1:1 route-target export 1:2
```

```
ip vrf customer a rd 1:2
route-target both 1:2
```

```
ip vrf customer a
rd 1:2
route-target both 1:1
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

From the exhibit, we learned:

+ VRF customer_a was exported with Route target (RT) of 1:1 so at the remote site it must be

imported with the same RT 1:1.

+ VRF customer_a was imported with Route target (RT) of 1:1 so at the remote site it must be

exported with the same RT 1:1.

Therefore at the remote site we must configure the command "route-target both 1:1" (which is equivalent to two commands "route-target import 1:1" & "route-target export 1:1").

Question: 76

What is a function of IPv6 ND inspection?

- A. It learns and secures bindings for stateless autoconfiguration addresses in Layer 3 neighbor tables
- B. It learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables
- C. It learns and secures bindings for stateful autoconfiguration addresses in Layer 2 neighbor tables.
- D. It learns and secures bindings for stateful autoconfiguration addresses in Layer 3 neighbor tables.

Answer: B

Explanation:

Question: 77

Exhibit:

```
policy-map COPP-7600
class COPP-CRITICAL-7600
  police cir 2000000 bc 62500
  conform-action transmit
  exceed-action transmit
!
class class-default
  police cir 200000 bc 6250
  conform-action transmit
  exceed-action drop
!
class-map match-all COPP-CRITICAL-7600
  match access-group name COPP-CRITICAL-7600
!
ip access-list extended COPP-CRITICAL-7600
  permit ip any any eq http
  permit ip any any eq https
```

BGP is flapping after the Copp policy is applied. What are the two solutions to fix the issue?

(Choose two)

- A. Configure BGP in the COPP-CRITICAL-7600 ACL
- B. Configure a higher value for CIR under the default class to allow more packets during peak traffic
- C. Configure a higher value for CIR under the class COPP-CRITICAL-7600
- D. Configure a three-color policer instead of two-color policer under class COPP-CRITICAL-7600
- E. Configure IP CEF to CoPP policy and BGP to work

Answer: A,B

Explanation:

The policy-map COPP-7600 only rate-limit HTTP & HTTPS traffic (based on the ACL conditions) so any BGP packets will be processed in the class "class-default", which drops exceeded BGP packets. Therefore we have two ways to solve this problem:

+ Add BGP to the ACL with the statement "permit tcp any any eq bgp"

+ Configure higher value for CIR in default class as 2Mbps is too low for web traffic (http & https)

Question: 78

What is an advantage of using BFD?

- A. It detects local link failure at layer 1 and updates routing table.
- B. It detects local link failure at layer 2 and updates routing protocols.
- C. It has sub-second failure detection for layer 1 and layer 3 problems.
- D. It has sub-second failure detection for layer 1 and layer 2 problems.

Answer: D

Explanation:

Question: 79

```
R3#show policy-map control-plane
Control Plane
Service-policy output: R3_CoPP

Class-map: mgmt (match-all)
 361 packets, 73858 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
 Match: access-group 120
 police:
  cir 8000 bps, bc 1500 bytes, be 1500 bytes
  conformed 8 packets, 1506 bytes; actions:
   transmit
  exceeded 353 packets, 72352 bytes; actions:
   drop
  violated 0 packets, 0 bytes; actions:
   drop
  conformed 0 bps, exceed 0 bps, violate 0 bps

Class-map: class-default (match-any)
 124 packets, 10635 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
 Match: any
R3#show access-lists 120
Extended IP access list 120
 10 permit udp any any eq snmptrap (361 matches)
```

Which action resolves intermittent connectivity observed with the SNMP trap packets?

- A. Decrease the committed burst Size of the mgmt class map
- B. Increase the CIR of the mgmt class map
- C. Add a new class map to match TCP traffic
- D. Add one new entry in the ACL 120 to permit the UDP port 161

Answer: B

Explanation:

Question: 80

Which component of MPLS VPNs is used to extend the IP address so that an engineer is able to identify to which VPN it belongs?

- A. VPNv4 address family
- B. RD
- C. RT
- D. LDP

Answer: B

Explanation:

E. Specify the correct route distinguisher used for that VPN. This is used to extend the IP address so that you can identify which VPN it belongs to

rd <VTN route dis

Question: 81

During the maintenance window an administrator accidentally deleted the Telnet-related configuration that permits a Telnet connection from the inside network (Eth0/0) to the outside of the networking between Friday – Sunday night hours only. Which configuration resolves the issue?

A)

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit udp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic Friday Saturday Sunday 22:00 to 05:00
```

B)

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic 22:00 to 05:00
```

C)

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic Friday Saturday Sunday 22:00 to 05:00
```

D)


```
interface Ethernet0/0
```

```
ip address 10.1.1.1 255.255.255.0
```

```
ip access-group 101 in
```

```
access-list 101 permit udp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255  
eq telnet time-range changewindow
```

```
!
```

```
timA-rannA channAwindnw
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 82

An engineer configured a company's multiple area OSPF head office router and Site A cisco routers with VRF lite. Each site router is connected to a PE router of an MPLS backbone.

Head & Site A

```
ip cel  
ip wf abc
```

rd 101 iQt

```
interface FwtEthemetO/D
ip vrf forwarding abc
ip address 172.16.16.X 255.255.255.252
```

```
router ospf 1 vrf abc
log adjacency-changes
network 172.16.10.0/24 area 1
```

After finishing both site router configurations, none of the LSA 3,4, 5, and 7 are installed at Site A router. Which configuration resolves this issue?

- A. configure capability vrf-lite on Site A and its connected PE router under router ospf 1 vrf abc
- B. configure capability vrf-lite on Head Office and its connected PE router under router ospf 1 vrf abc
- C. configure capability vrf-lite on both PE routers connected to Head Office and Site A routers under router ospf 1 vrf abc
- D. configure capability vrf-lite on Head Office and Site A routers under router ospf 1 vrf abc

Answer: C

Explanation:

Question: 83

Refer to the exhibit.

```
ip access-list extended FILTER
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 22
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 23
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 80
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 443
permit tcp host 192.168.10.10 host 192.168.100.10 eq ssh
permit ip any any
!
interface GigabitEthernet0/1
ip address 192.168.10.1 255.255.255.0
ip access-group FILTER in
!
```

The ACL is placed on the inbound Gigabit 0/1 interface of the router. Host

192.168.10.10 cannot SSH to host 192.168.100.10 even though the flow is permitted. Which action resolves the issue without opening full access to this router?

- A. Move the SSH entry to the beginning of the ACL
- B. Temporarily move the permit ip any any line to the beginning of the ACL to see if the flow works
- C. Temporarily remove the ACL from the interface to see if the flow works
- D. Run the show access-list FILTER command to view if the SSH entry has any hit statistic associated with it

Answer: A

Explanation:

Question: 84

Which security feature can protect DMVPN tunnels?

- A. IPsec
- B. TACACS+
- C. RTBH
- D. RADIUS

Answer: A

Explanation:

Question: 85

Which two methods use IPsec to provide secure connectivity from the branch office to the headquarters office?

(Choose two.)

- A. DMVPN
- B. MPLS VPN

C. Virtual Tunnel Interface (VTI)

D. SSL VPN

E. PPPoE

Answer: A,C

Explanation:

Question: 86

Which protocol is used in a DMVPN network to map physical IP addresses to logical IP addresses?

A. BGP

B. LLDP

C. EIGRP

D. NHRP

Answer: D

Explanation:

Question: 87

Which Cisco VPN technology can use multipoint tunnel, resulting in a single GRE tunnel interface on the hub, to support multiple connections from multiple spoke devices?

- A. DMVPN
- B. GETVPN
- C. Cisco Easy VPN
- D. FlexVPN

Answer:

A

Explanation:

Question:

88

Which option is the best for protecting CPU utilization on a device?

- A. fragmentation
- B. COPP
- C. ICMP redirects
- D. ICMP unreachable messages

Answer:

B

Explanation:

Question:

89

What is the role of a route distinguisher via a VRF-Lite setup implementation?

- A. It extends the IP address to identify which VFP instance it belongs to.
- B. It manages the import and export of routes between two or more VRF instances
- C. It enables multicast distribution for VRF-Lite setups to enhance EGP routing protocol capabilities
- D. It enables multicast distribution for VRF-Lite setups to enhance IGP routing protocol capabilities

Answer: A

Explanation:

Question: 90

Refer to the following output:

```
Router#show ip nhrp detail
```

```
10.1.1.2/8 via 10.2.1.2, Tunnel1 created 00:00:12, expire 01:59:47
```

```
TypeE. dynamic, Flags: authoritative unique nat registered used
```

```
NBMA address: 10.12.1.2
```

What does the authoritative flag mean in regards to the NHRP information?

- A. It was obtained directly from the next-hop server.

- B. Data packets are process switches for this mapping entry.
- C. NHRP mapping is for networks that are local to this router.
- D. The mapping entry was created in response to an NHRP registration request.
- E. The NHRP mapping entry cannot be overwritten.

Answer:

A

Explanation:

Question:

91

Which two protocols can cause TCP starvation? (Choose two)

- A. TFTP
- B. SNMP
- C. SMTP
- D. HTTPS
- E. FTP

Answer:

A,B

Explanation:

Question:

92

Which two statements about VRF-Lite configurations are true? (Choose two.)

- A. They support the exchange of MPLS labels
- B. Different customers can have overlapping IP addresses on different VPNs
- C. They support a maximum of 512,000 routes
- D. Each customer has its own dedicated TCAM resources
- E. Each customer has its own private routing table.
- F. They support IS-IS

Answer: B,E

Explanation:

Question: 93

A network engineer needs to verify IP SLA operations on an interface that shows an indication of excessive traffic.

Which command should the engineer use to complete this action?

- A. show frequency
- B. show track
- C. show reachability
- D. show threshold

Explanation:

Question: 94

Answer: B

Which protocol does VRF-Lite support?

- A. IS-IS
- B. ODR
- C. EIGRP
- D. IGRP

Explanation:

Question: 95

Answer: C

Refer to Exhibit.

```
router ospf 10 router-id 192.168.1.1 log-adj aaency-
changes redistribute bgp 1 subnets route-map BGP-TO-OSPF I
route-map BGP-TO-OSPF deny 10 match ip address 50
route-map BGP-TO-OSPF permit 20
access-list 50 permit 172.16.1.0 0.0.0.255
```

Which statement about redistribution from BGP into OSPF process 10 is true?

- A. Network 172.16.1.0/24 is not redistributed into OSPF.
- B. Network 10.10.10.0/24 is not redistributed into OSPF.
- C. Network 172.16.1.0/24 is redistributed with administrative distance of 1.
- D. Network 10.10.10.0/24 is redistributed with administrative distance of 20.

Answer: A

Explanation:

Question: 96

Which two statements about redistributing EIGRP into OSPF are true? (Choose two)

- A. The redistributed EIGRP routes appear as type 3 LSAs in the OSPF database
- B. The redistributed EIGRP routes appear as type 5 LSAs in the OSPF database
- C. The administrative distance of the redistributed routes is 170
- D. The redistributed EIGRP routes appear as OSPF external type 1
- E. The redistributed EIGRP routes as placed into an OSPF area whose area ID matches the EIGRP autonomous system number
- F. The redistributed EIGRP routes appear as OSPF external type 2 routes in the routing table

Answer: B,F

Explanation:

Question: 97

Refer to the exhibit.

```
router eigrp 1
 redistribute ospf 5 match external route-map OSPF-TO-EIGRP
 metric 10000 2000 255 1 1500
 route-map OSPF-TO-EIGRP
 match ip address TO-OSPF
```

Which routes from OSPF process 5 are redistributed into EIGRP?

- A. E1 and E2 subnets matching access list TO-OSPF
- B. E1 and E2 subnets matching prefix list TO-OSPF
- C. only E2 subnets matching access list TO-OSPF
- D. only E1 subnets matching prefix list TO-OSPF

Answer: A

Explanation:

Question: 98

Users were moved from the local DHCP server to the remote corporate DHCP server. After the move, none of the users were able to use the network.

Which two issues will prevent this setup from working properly? (Choose two)

- A. Auto-QoS is blocking DHCP traffic.
- B. The DHCP server IP address configuration is missing locally
- C. 802.1X is blocking DHCP traffic
- D. The broadcast domain is too large for proper DHCP propagation
- E. The route to the new DHCP server is missing

Answer: B,E

Explanation:

Question: 99

Which command is used to check IP SLA when an interface is suspected to receive lots of traffic with options?

- A. show track
- B. show threshold
- C. show timer
- D. show delay

Answer: A

Explanation:

Question: 100

Which SNMP verification command shows the encryption and authentication protocols that are used in

SNMPV3?

- A. show snmp group
- B. show snmp user
- C. show snmp
- D. show snmp view

Answer: B

Explanation:

Question: 101

Which is statement about IPv6 inspection is true?

- A. It teams and secures bindings for stateless autoconfiguration addresses in Layer 3 neighbor tables
- B. It learns and secures bindings for stateful autoconfiguration addresses in Layer 3 neighbor tables
- C. It teams and secures bindings for stateful autoconfiguration addresses in Layer 2 neighbor tables
- D. It team and secures binding for stateless autoconfiguration addresses in Layer 2 neighbor tables.

Answer: D

Explanation:

Question: 102

What is the output of the following command:

```
show ip vrf
```

- A. Show's default RD values
- B. Displays IP routing table information associated with a VRF
- C. Show's routing protocol information associated with a VRF.
- D. Displays the ARP table (static and dynamic entries) in the specified VRF

Answer: A

Explanation:

Question: 103

Refer to the exhibit.

```
ip dhcp pool 1
```

```
network 200.30.30.0/24
```

```
default-router 200.30.30.100
```

```
lease 40
```

```
ip dhcp pool 2
```

```
network 200.30.40.0/24 default-router 200.30.40.100
```

```
lease 40
```

The server for the finance department is not reachable consistently on the 200.30.40.0/24 network and after every second month it gets a new IP address. Which two actions must be taken to resolve this Issue? (Choose two.)

- A. Configure the server to use DHCP on the network with default gateway 200.30.40.100.
- B. Configure the server with a static IP address and default gateway.
- C. Configure the router to exclude a server IP address.
- D. Configure the server to use DHCP on the network with default gateway 200.30.30.100.
- E. Configure the router to exclude a server IP address and default gateway.

Answer: B,C

Explanation:

Question: 104

Refer to the exhibit.

```
Spoke# show dmvpn
Tunnel0, Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 172.18.16.2 192.168.1.1 UP 01:05:35 S
1 172.18.46.2 192.168.1.4 UP 00:00:25 D
```

An engineer has configured DMVPN on a spoke router. What is the WAN IP address of another spoke router within the DMVPN network?

- A. 172.18.46.2
- B. 192.168.1.4

c. 172.18.16.2

d. 192.168.1.1

Answer: A


Explanation:

From the output we can see there are 2 NHRP Peers. The first one with the NBMA Address of 172.18.16.2 and the

“Attribute” (Attrb) of Static (S) so we can deduce it is the Hub device.

Therefore the second one must be the remaining Spoke device with the attribute of Dynamic (D).

HeadQuarter



-> S - Static, D - Dynamic, I - Incomplete

N - NATed, L - Local, X - No Socket

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: Tunnel1, IPv4 NHRP Details

Type:Spoke, NHRP Peers:2,

Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb

1 44.44.44.4 192.168.100.254 UP 00:03:40 S

1 12.12.12.2 192.168.100.2 UP 00:03:20 D

Question: 105

DRAG DROP

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

Customer (C) device

devc? m TH acre of me p wrier IKW-
rnatswtr-ifc: bins [HdiFF

CE oevks

dis rr rw .stflr.be'. and draches ; i- 7PM Inbels- ki tr .Tarkan ri the
HiiwIdei iisWMIll

PE (Wee

kvia 1 r+w pnt>Fr, i;w nefurc k tip; cnm^F TO .other ri Ft orrer
device

'icwdei P) deMite

t|iu|¥4- ltiecd^eif the cnLniprsu iietunk Hlat.c.uiii"wiJa.tu HIE SP
rethork

Answer:

Explanation:

Provider (P) device



Question:

106

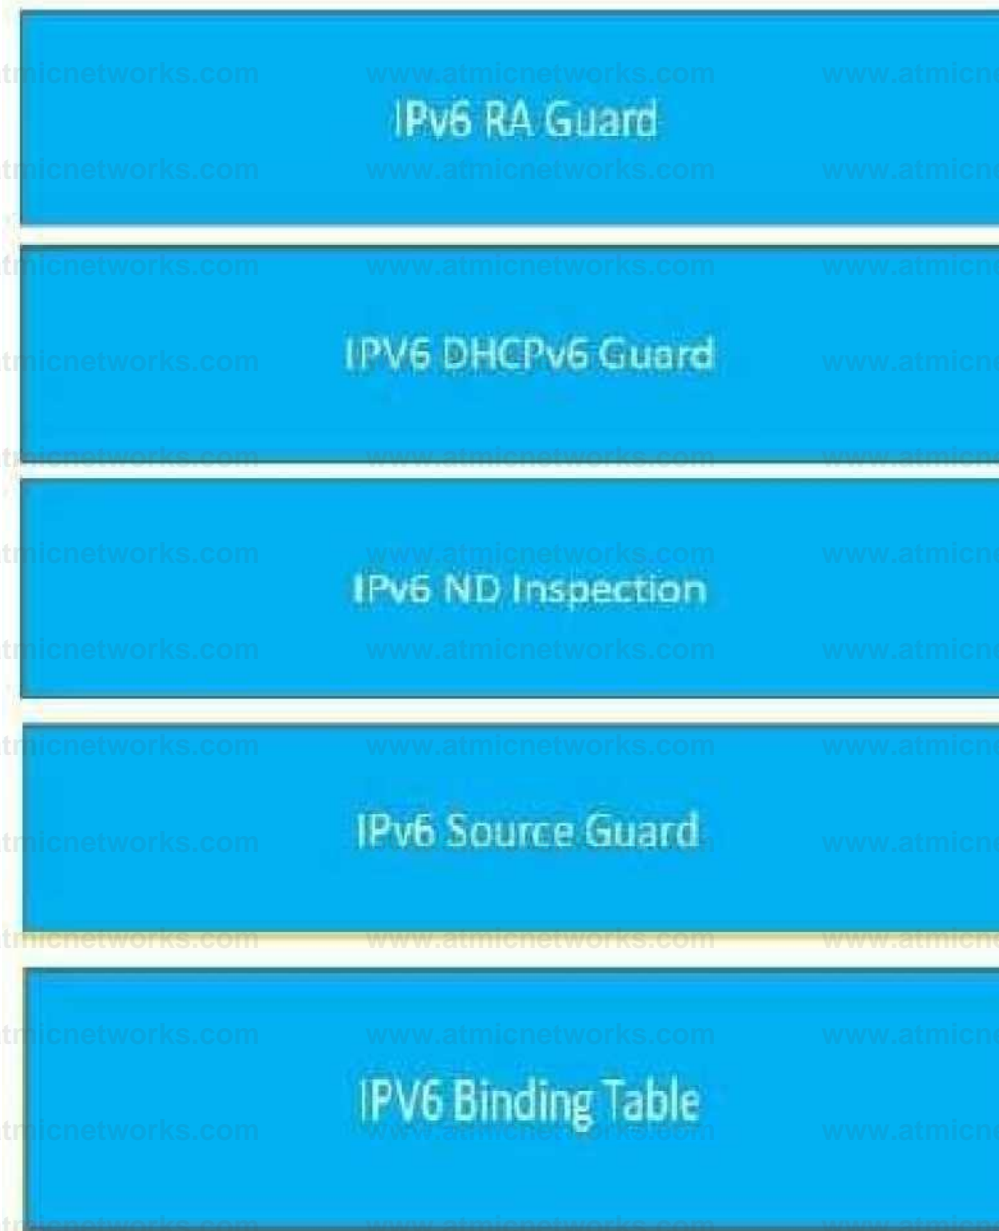
DRAG DROP

Drag and Drop the IPv6 First-Hop Security features from the left onto the definitions on the right.

IPv6 DHCPv6 Guard	Block a malicious host and permit the router from a legitimate route.
IPv6 Binding Table	Block reply and advertisement messages from unauthorized DHCP servers and relay agents.
IPv6 Source Guard	Create a binding table that is based on NS and NA messages.
IPv6 RA Guard	Filter inbound traffic on Layer 2 switch ports that are not in the IPv6 binding table.
IPv6 ND Inspection	Create IPv6 neighbors connected to the device from information sources such as NDP snooping.

Answer:

Explanation:



Graphical user

interface, chart Description automatically generated

Question: 107

An engineer is configuring a network and needs packets to be forwarded to an interface for any destination address that is not in the routing table. What should be configured to accomplish this task?

- A. set ip next-hop
- B. set ip default next-hop
- C. set ip next-hop recursive

D. set ip next-hop verify-availability

Answer: B

Explanation:

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...

- if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.
- if the destination IP address **does not exist**, the command policy routes the packet by **sending it to the specified next hop**

Question: 108

Which protocol does MPLS use to support traffic engineering?

- A. Tag Distribution Protocol (TDP)
- B. Resource Reservation Protocol (RSVP)
- C. Border Gateway Protocol (BGP)
- D. Label Distribution Protocol (LDP)

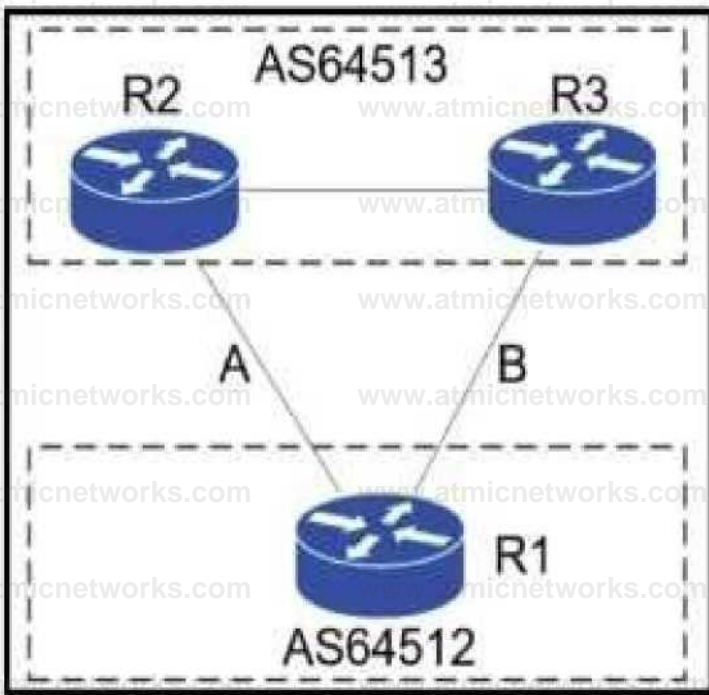
Answer: B

Explanation:

MPLS TE provides a way to integrate TE capabilities (such as those used on Layer 2 protocols like ATM) into Layer 3 protocols (IP). MPLS TE uses an extension to existing protocols (Intermediate System-to-Intermediate System (IS-IS), Resource Reservation Protocol (RSVP), OSPF) to calculate and establish unidirectional tunnels that are set according to the network constraint. Traffic flows are mapped on the different tunnels depending on their destination.

Question: 109

Refer to the exhibit.



A network engineer for AS64512 must remove the inbound and outbound traffic from link A during maintenance without closing the BGP session so that there a backup link over link A toward the ASN. Which BGP configuration on R1 accomplishes this goal?

A)

```

route-map link-a-in permit 10
set weight 200
route-map link-a-out permit 10
set as-path prepend 64512
route-map link-b-in permit 10
set weight 100
route-map link-b-out permit 10

```

B)

```
route-map link-a-in permit 10 set local-preference 200 route-map link-a-out permit 10 route-map link-b-in permit 10 route-map link-b-out permit 10 set as-path prepend 64512
```

C)

```
route-map link-a-in permit 10 route-map link-a-out permit 10 set as-path prepend 64512 route-map link-b-in permit 10
```

```
set local-preference 200 route-map link-b-out permit 10
```

D)

```
route-map link-a-in permit 10 set weight 200 route-map link-a-out permit 10 route-map link-b-in permit 10 set weight 100
```

```
route-map link-b-out permit 10 set as-path prepend 64512
```

A. Option A

B. Option B

C. Option C

D. Option D

Explanation:

Topic 2, Exam Pool B

Question:

110

Refer to the exhibit.

Answer:

C


```
R1#show policy-map control-plane
```

Control Plane

Service-policy output: CoPP

Class-map: SNMP-Out (match-all)

124 packets, 3693 bytes

5 minute offered rate 0000 bps, drop rate 0000 bps

Match: access-group name SNMP

police:

 cir 800 0 bps, be 1500 bytes

 conformed 0 packets, 0 bytes; actions: transmit

 exceeded 0 packets, 0 bytes; actions: drop

 conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)

10 packets, 1003 bytes

5 minute offered rate 0000 bps, drop rate 0000 bps

Match: any

```
R1#show ip access-list SNMP
```

Extended IP access list SNMP

```
10 permit udp any eq snmp any
```

R1 is being monitored using SNMP and monitoring devices are getting only partial information. What action should be taken to resolve this issue?

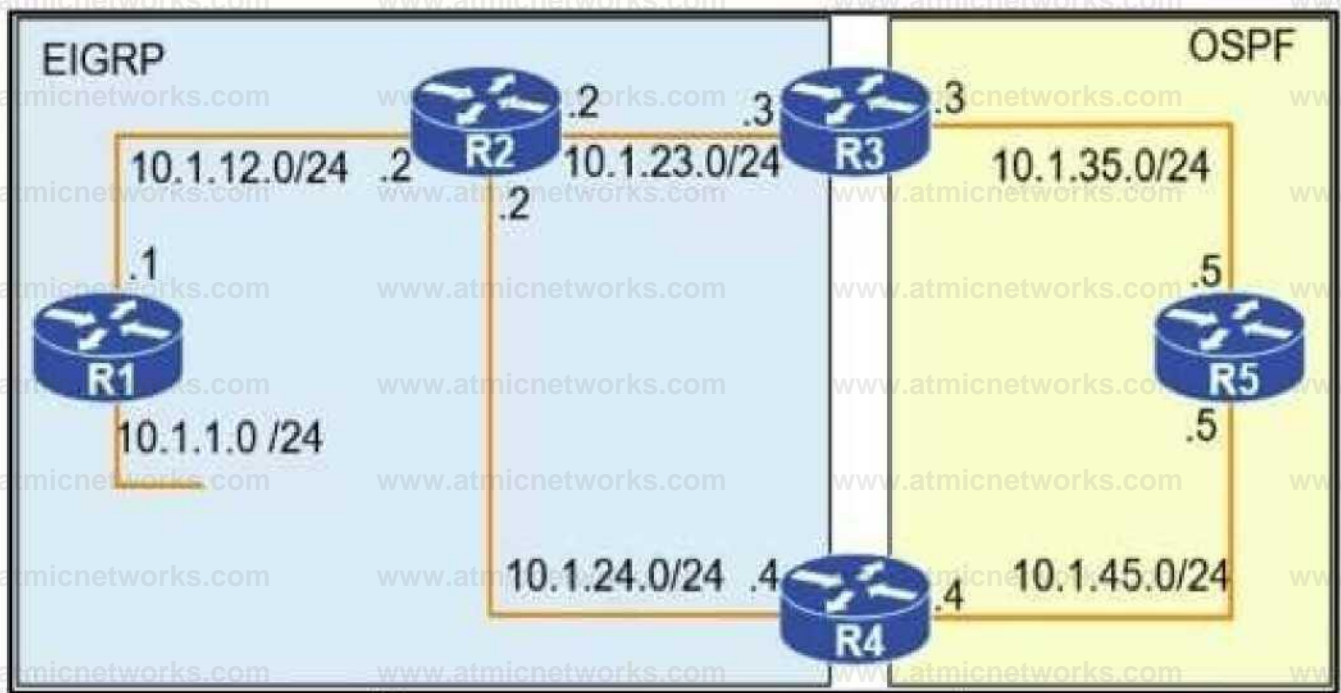
- A. Modify the CoPP policy to increase the configured exceeded limit for SNMP.
- B. Modify the access list to include snmptrap.
- C. Modify the CoPP policy to increase the configured CIR limit for SNMP.
- D. Modify the access list to add a second line to allow udp any any eq snmp

Answer: D

Explanation:

Question: 111

Refer to the exhibit.



```

R1
router eigrp 1
 redistribute connected
 network 10.1.12.1 0.0.0.0
 default-metric 1000000 10 255 1 1500

R3
router eigrp 1
 network 10.1.23.3 0.0.0.0
!
router ospf 1
 redistribute eigrp 1 subnets
 network 10.1.35.3 0.0.0.0 area 0
  
```

To provide reachability to network 10.1.1.0/24 from R5, the network administrator redistributes EIGRP into OSPF on R3 but notices that R4 is now taking a ... path through R5 to reach 10.1.1.0/24 network. Which action fixes the issue while keeping the reachability from R5 to 10.1.1.0/24 network?

- A. Change the administrative distance of the external EIGRP to 90.
- B. Apply the outbound distribution list on R5 toward R4 in OSPF.
- C. Change the administrative distance of OSPF to 200 on R5.
- D. Redistribute OSPF into EIGRP on R4

Answer: A

Explanation:

Question: 112

Refer to the exhibit.

```
'Jun 24 08:54:51 530 IF-EvD(GigabdEtherneM): IP Routing reports state transition from DOWN to DOWN
'Jun 24 08:54 52 525 %LINEPROTO-5-UPDOWN Line protocol on Interface GigabitEthernetO/O, changed state to down
• Jun 24 08:54:52.528 IF-EvD(GigabitEthernetO/O): IP Routing reports state transition from DOWN to DOWN
'Jun 24 08:54 53 215 IF-EvD(GigabitEthernetO/O): IP Routing reports state transition from DOWN to DOWN
'Jun 24 08:54:54.998: %LINK-3-UPDOWN Interface GigabitEthernetO/O. changed state to up
• Jun 24 08:54 55.006 IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to UP
• Jun 24 08 54 55 998 %LINEPROTO-5-UPDOWN Line protocol on Interface GigabitEthernetO/O, changed state to up
```

R1 is connected with R2 via GigabitEthernet0/0, and R2 cannot ping R1. What action will fix the issue?

- A. Fix route dampening configured on the router.
- B. Replace the SFP module because it is not supported.
- C. Fix IP Event Dampening configured on the interface.
- D. Correct the IP SLA probe that failed.

Answer: C

Explanation:

The **IP Event Dampening** feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping.

Question: 113

Which IGPs are supported by the MPLS LDP autoconfiguration feature?

- A. RIPv2 and OSPF
- B. OSPF and EIGRP
- C. OSPF and ISIS
- D. ISIS and RIPv2

Answer: C

Explanation:

The MPLS LDP Autoconfiguration feature enables you to globally enable Label Distribution Protocol (LDP) on even' interface associated with an Interior Gateway Protocol (IGP) instance. This feature is supported on Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) IGPs. It provides

Question: 114

An engineer needs dynamic routing between two routers and is unable to establish OSPF adjacency. The output of the show ip ospf neighbor command shows that the neighbor state is EXSTART/EXCHANGE. Which action should be taken to resolve this issue?

- A. match the passwords
- B. match the hello timers
- C. match the MTUs
- D. match the network types

Answer: C

Explanation:

Neighbors Stuck in Exstart/Exchange State

The problem occurs most frequently when attempting to run OSPF between a Cisco router and another vendor's router. The problem occurs when the **maximum transmission unit (MTU)** settings for neighboring router interfaces **don't match**. If the router with the higher MTU sends a packet larger than the MTU set on the neighboring router, the neighboring router ignores the packet. When

Question: 115

Refer to the exhibit.

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
|
time-range Office-hour
periodic weekdays 08:00 to 17:00
|
access-list 101 permit tcp 10.0.0.0 0.0.0.0 172.16.1.0 0.0.0.255 eq ssh time-range Office-hour
```

An IT staff member comes into the office during normal office hours and cannot access devices through SSH. Which action should be taken to resolve this issue?

- A. Modify the access list to use the correct IP address.
- B. Configure the correct time range.
- C. Modify the access list to correct the subnet mask.

D. Configure the access list in the outbound direction.

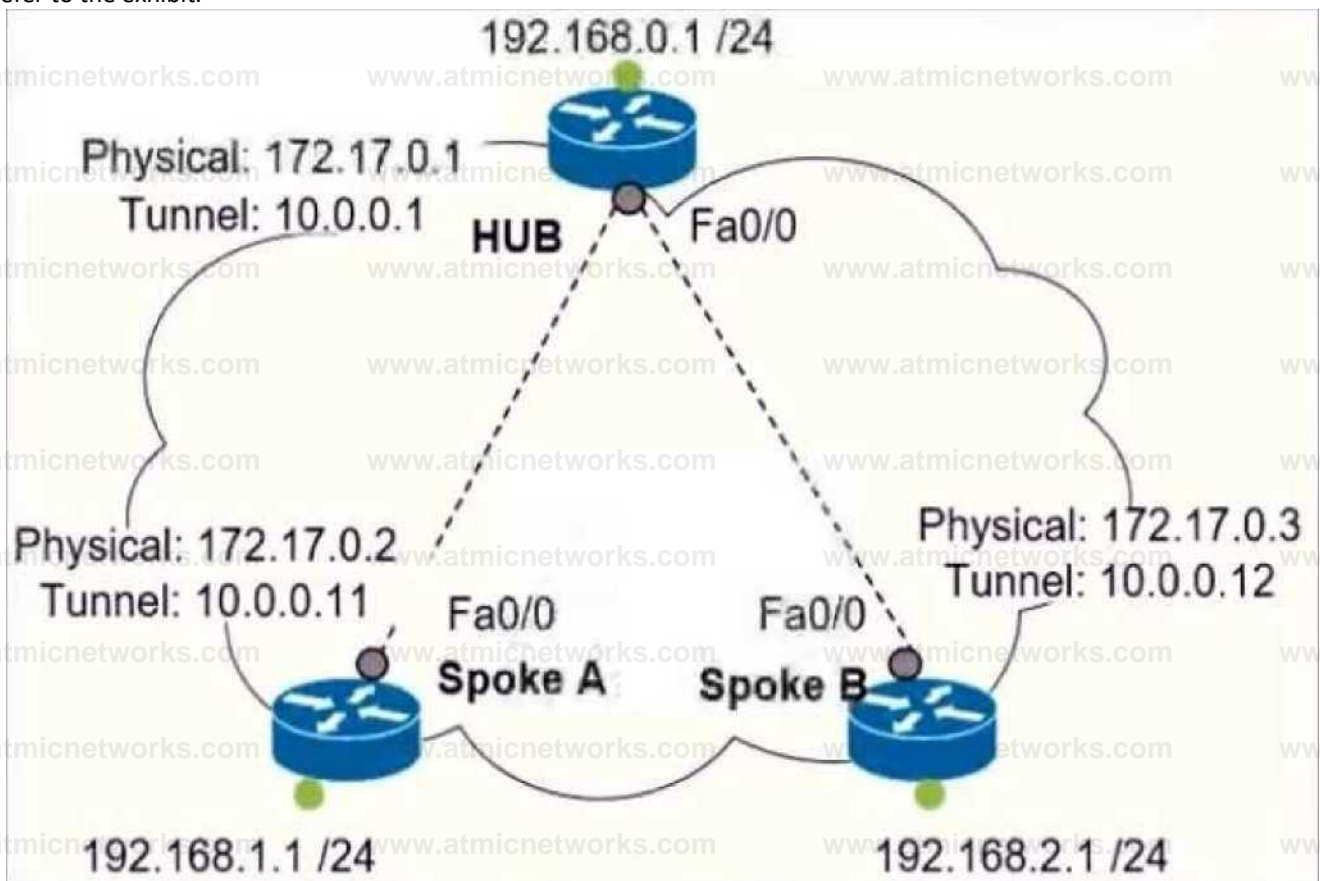
Answer: A

Explanation:

To ACL should be permit tcp 101 10.1.1.1 0.0.0.0

Question: 116

Refer to the exhibit.



Which interface configuration must be configured on the HUB router to enable MVPN with mGRE mode?

```
interface Tunnel0 description mGRE ■ DMVPN Tunnel ip address 10.1 0,1  
255.255 255.0 ip nhrp map multicast dynamic ip nhrp network-id 1 tunnel  
source 172.17 0.1 ip nhrp map 10.0 0 11172.17.0.2 ip nhrp map 10.0 0 12172.17
```

0 3 tunnel mode gre

interface TunneK)

```
description mGRE * DMVPN Tunnel ip address 10 0 0 1 255.255.255 0 ip nhrp  
map multicast dynamic ip nhrp networked 1
```

```
tunnel source 10.0.0.1
```

```
tunnel mode gre multipoint
```

interface TunnetO

```
description mGRE - DMVPN Tunnel ip address 10 0 0.1 255.255 255 0 ip nhrp  
networked 1 tunnel source 172,17,0.1 tunnel mode gre multipoint
```

interface TunnetO

```
description mGRE - DMVPN Tunnel Ip address 10.0.0.1 2 5 5.2 55 255.0 ip nhrp  
map multicast dynamic ip nhrp network -id I tunnel source 10.0.0.1 tunnel  
destination 172.17.0.2 tunnel mode gre multipoint
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html

Question: 117

What does IPv6 Source Guard utilize to determine if IPv6 source addresses should be forwarded?

- A. ACE
- B. ACLS
- C. DHCP
- D. Binding Table

Answer: D

Explanation:

IPv6 source guard is an interface feature between the populated binding table and data traffic filtering. This feature enables the device to deny traffic when it is originated from an address that is not stored in the binding table. IPv6 source guard does not inspect ND or DHCP packets; rather, it works

Question: 118

Refer to the exhibit.

```
R1#show run | begin line
```

```
line con 0
```

```
exec-timeout 0 0
```

```
privilege level 15
```

```
logging synchrous
```

```
transport preferred telnet
```

```
transport output none
```

```
stopbits 0 4
```

```
line vty 0 4
```

```
login
```

```
transport referred telnet
```

```
transport input none
```

```
transport output telnet
```

```
R1#
```

```
R1#ssh -1 cisco 192.168.12.2
```

```
% ssh connections not permitted from this terminal
```

```
R1#
```

An engineer receives this error message when trying to access another router in-band from the serial interface connected to the console of R1. Which configuration is needed on R1 to resolve this issue?

```
P1(ronfig)#line console 0
```

```
R1(config-line)#transport preferred ssh
```

```
R1 (config)#line vty 0
```

```
R1(config-line)# transport output ssh
```

```
Pl(cwtfig)#line vty 0
```

```
R1 (config-line)# transport output ssh
```

```
R1 (config-line)# transport preferred ssh
```

```
R1(config)#line console 0
```

```
R1(config-line)# transport output ssh
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

<https://community.cisco.com/t5/other-network-architecture/out-of-band-router-access/td-p/333295>

The “transport output none” command prevents any protocol connection made from R1.

Therefore our SSH connection to 192.168.12.2 was refused. In order to fix this problem we can

configure “transport output ssh” under “line console 0” of R1.

Note: The parameter “-l” specifies the username to log in as on the remote machine.

Question: 119

Refer to the exhibit.

```
Routertfshow access-lists
```

```
Standard IP access list 1
```

```
10 permit 192.168.2.2 (1 match)
```

```
Router#
```

```
Routertfshow route-map
```

```
route-map RM-OSPF-DL, deny, sequence 10
```

```
Match clauses:
```

```
ip address (access-ls); 1
```

```
Set clauses:
```

```
Policy routing matches: 0 packets, 0 bytes
```

```
Router#
```

```
Routertfshow running-config | section ospf
```

```
router ospf 1
```

```
network 192.168.1.1 0.0.0.0 area 0
```

```
network 192.168.12.0 0.0.0.255 area 0
```

```
distribute-list route-map RM-OSPF-DL in
```

```
Router#|
```

Which two actions should be taken to access the server? (Choose two.)

- A. Modify the access list to add a second line of permit ip any
- B. Modify the access list to deny the route to 192.168.2.2.
- C. Modify distribute list seq 10 to permit the route to 192.168.2.2.

D. Add a sequence 20 in the route map to permit access list 1.

E. Add a floating static route to reach to 192.168.2.2 with administrative distance higher than OSPF.

Answer: B,E

Explanation:

Question: 120

Refer to the exhibit.

```
router# show running-config
Building configuration
|
<output omitted ---->
|
hostname R1
|
ip domain-name cisco.com
|
crypto key generate rsa modulus 2048
|
username admin privilege 15 secret cisco123
|
access-list 1 permit 10.1.1.0 0.0.0.255
access-list 1 deny any log
|
line vty 0 15
access-class 1 in
login local
|
<output omitted ---->
|
end
```

A user cannot SSH to the router. What action must be taken to resolve this issue?

- A. Configure transport input ssh
- B. Configure transport output ssh
- C. Configure ip ssh version 2
- D. Configure ip ssh source-interface loopback0

Answer: A

Explanation:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/security/configuration_guide/b_sec_152ex_2960-x_cg/b_sec_152ex_2960-x_cg_chapter_01001.html

Question: 121

Refer to the exhibit.

```
MASS-RTR#show running-config
!
hostname MASS-RTR
!
aaa new-model
!
aaa authentication login default local
aaa authorization exec default local
aaa authorization commands 15 default local
!
username admin privilege 15 password 7 0236244818115F3348
username cisco privilege 15 password 7 0607072C494A5B
archive
 log config
  logging enable
  logging size 1000
!
interface GigabitEthernet0/0
 ip address dhcp
 duplex auto
 speed auto
!
line vty 0 4
!
MASS-RTR#show archive log config all
  idx  sess      user@line      Logged command
  ---  ---      -
  1    1      console@console |interface GigabitEthernet0/0
  2    1      console@console | no shutdown
  3    1      console@console | ip address dhcp
  4    2      admin@vty0     |username cisco privilege 15 password cisco
  5    2      admin@vty0     |!config: USER TABLE MODIFIED
```

A client is concerned that passwords are visible when running this show archive log config all.

Which router configuration is needed to resolve this issue?

- A. MASS-RTR(config-archive-log-cfg)#password encryption aes
- B. MASS-RTR(config)#aaa authentication arap
- C. MASS-RTR(config)#service password-encryption
- D. MASS-RTR(config-archive-log-cfg)#hidekeys

Answer: D

Explanation:

Step 7 hidekeys

Example:

```
Device(config-archive-log-config) # hidekeys
```

(Optional) Suppresses the display of password information in configuration log files.

Note

Enabling the **hidekeys** command increases security by preventing password information from being displayed in configuration log files.

Question: 122

Refer to the exhibit.

R1

```
ip prefix-list ccnp1 seq 5 permit 10.1.48.0/24 le 24
ip prefix-list ccnp2 seq 5 permit 10.1.80.0/24 le 32
ip prefix-list ccnp3 seq 5 permit 10.1.64.0/24 le 24

route-map ospf-to-eigrp permit 10
  match ip address prefix-list ccnp1
  set tag 30
route-map ospf-to-eigrp permit 20
  match ip address prefix-list ccnp2
  set tag 20
route-map ospf-to-eigrp permit 30
  match ip address prefix-list ccnp3
  set tag 10
```

An engineer wanted to set a tag of 30 to route 10.1.80.65/32 but it failed. How is the issue fixed?

- A. Modify route-map ospf-to-eigrp permit 30 and match prefix-list ccnp2.
- B. Modify route-map ospf-to-eigrp permit 10 and match prefix-list ccnp2.
- C. Modify prefix-list ccnp3 to add 10.1.64.0/20 le 24
- D. Modify prefix-list ccnp3 to add 10.1.64.0/20 ge 32

Answer:

B

Explanation:

Question: 123

Refer to the exhibit.

```
ipv6 access-list inbound  
permit tcp any any  
deny ipv6 any any log  
!  
interface gi0/0  
ipv6 traffic-filter inbound out
```

A network administrator configured an IPv6 access list to allow TCP return frame only, but it is not working as expected. Which changes resolve this issue?

```
ipv6 access-list inbound permit tcp any any established deny ipv6 any any
log
```

```
interface giOX*
```

```
  ipv6 traffic-filter inbound out
```

```
  ipv6 access-list inbound permit tcp any any syn deny |pv6 any any log
```

```
interface giOrt)
```

```
  ipv6 traffic-filter inbound out
```

```
  ipv6 access-list inbound permit tcp any any established deny |pv6S any any
```

```
log
```

```
interface giO/D
```

```
  ipv6 traffic-filter inbound in
```

```
  ipv6 access-list inbound permit tcp any any syn deny ipv6 any any log
```

```
interface giM ipvG traffic-filter inbound In
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

[https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/122_55_se/co
nfiguration/guide/scg3750/swv6acl.html](https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/122_55_se/configuration/guide/scg3750/swv6acl.html)

Question: 124

What does the PE router convert the IPv4 prefix to within an MPLS VPN?

- A. VPN-IPv4 prefix combined with the 64-bit route distinguisher
- B. 48-bit route combining the IP and PE router-id
- C. prefix that combines the ASN, PE router-id, and IP prefix
- D. eBGP path association between the PE and CE sessions

Answer: A

Explanation:

The IP prefix is a member of the IPv4 address family. After the PE device learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family. It uniquely identifies the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses. The route distinguisher used to generate the VPN-IPv4 prefix is specified by a configuration command associated with the virtual routing and forwarding (VRF) instance on the PE device.

Question: 125

How are MPLS Layer 3 VPN services deployed?

- A. The RD and RT values must match under the VRR
- B. The RD and RT values under a VRF must match on the remote PE router
- C. The import and export RT values under a VRF must always be the same.
- D. The label switch path must be available between the local and remote PE routers.

Answer: D

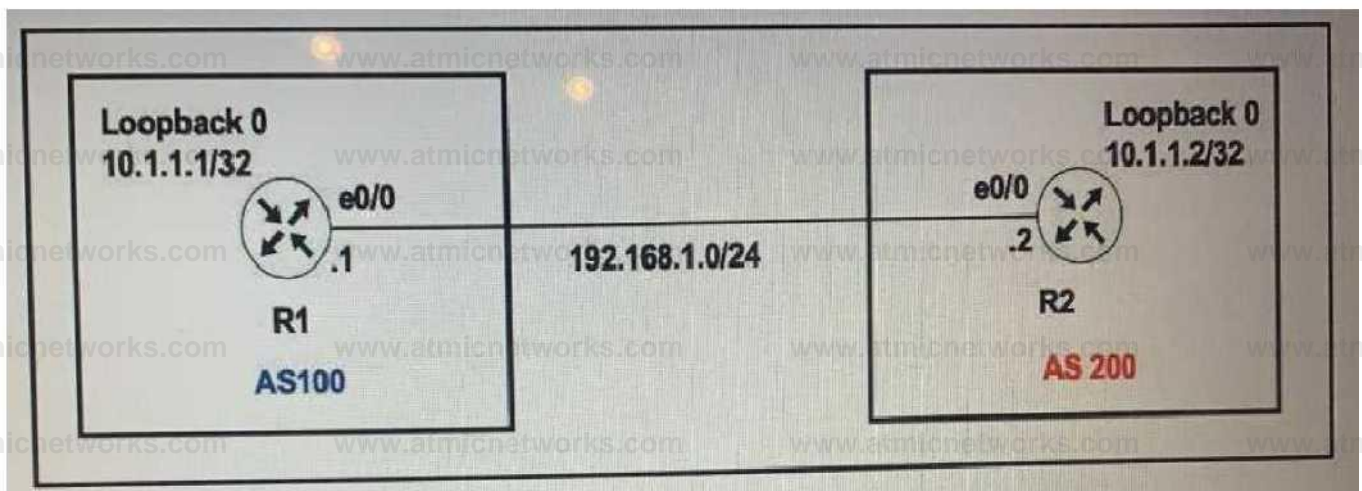
Explanation:

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/65x/b-l3vpn-cg-ncs5500-65x/b-l3vpn-cg-ncs5500-65x_chapter_010.html

The ingress PE router must be able to reach the egress PE router for a packet to be relayed to its destination.

Question: 126

Refer to the exhibit.



The R1 and R2 configurations are:

R1

```
router bgp 100
  neighbor 10.1.1.2 remote-as 200
```

R2

```
router bgp 200
  neighbor 10.1.1.1 remote-as 100
```

The neighbor is not coming up. Which two sets of configurations bring the neighbors up? (Choose two.)

- A. R2ip route 10.1.1.1 255.255.255.255 192.168.1.1 ! router bgp 200neighbor 10.1.1.1 ttl-security hops 1neighbor 10.1.1.1 update-source loopback 0
- B. R2ip route 10.1.1.1 255.255.255.255 192.168.1.1 ! router bgp 200neighbor 10.1.1.1 disable- connected- checkneighbor 10.1.1.1 update-source loopback 0
- C. R2ip route 10.1.1.2 255.255.255.255 192.168.1.2 ! router bgp 100neighbor 10.1.1.2 ttl-security hops 1neighbor 10.1.1.2 update-source loopback 0
- D. R1ip route 10.1.1.2 255.255.255.255 192.168.1.2 ! router bgp 100neighbor 10.1.1.1 ttl-security hops 1neighbor 10.1.1.2 update-source loopback 0
- E. R1ip route 10.1.1.2 255.255.255.255 192.168.1.2 ! router bgp 100neighbor 10.1.1.2 disable- connected-checkneighbor 10.1.1.2 update-source Loopback0

Answer: B,E

Explanation:

The neighbor disable-connected-check command is used to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address.

Disable-connected-check enables a directly connected eBGP neighbor to peer using a loopback address without adjusting the default TTL of 1. In disable connected check the router does not decrease the TTL of an IP packet that is destined to itself so it only counts or considers as one hop between the two loopbacks of the routers.

Question: 127

Refer to the exhibit.

```

R1#
interface GigabitEthernet0/0
 ip address 209.165.201.2 255.255.255.252
 !
interface GigabitEthernet0/1
 ip address 209.165.201.6 255.255.255.252
 !
router bgp 65401
  bgp log-neighbor-changes
  redistribute static
  neighbor 209.165.201.1 remote-as 65402
  neighbor 209.165.201.5 remote-as 65403
  !
 ip route 209.165.200.224 255.255.255.224 Null0
 ip route 209.165.202.128 255.255.255.224 Null0
 !

```

A company with autonomous system number AS65401 has obtained IP address block

209.165.200.224/27 from ARIN. The company needed more IP addresses and was assigned block 209.165.202.128/27 from ISP2. An engineer at ISP1 reports they are receiving ISP2 routes from AS65401. Which configuration on R1 resolves the issue?

A)

```
access-list 10 deny 209.165.202.128 0.0.0.31
access-list 10 permit any
!
router bgp 65401
neighbor 209.165.201.1 distribute-list 10 out
```

B)

```
access-list 10 deny 209.165.202.128 0.0.0.31
access-list 10 permit any
!
router bgp 65401
neighbor 209.165.201.1 distribute-list 10 in
```

C)

```
ip route 209.165.200.224 255.255.255.224 209.165.201.1
ip route 209.165.202.128 255.255.255.224 209.165.201.5
```

D)



```
ip route 0.0.0.0 0.0.0.0 209.165.201.1
```

```
ip route 0.0.0.0 0.0.0.0 100 209.165.201.5
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/23675-27.html>

Question: 128

Refer to the exhibit.

Filter

Priority	Issue Type	Device Role	Category	Issue Count	Site Count (Area)	Device Count
High	Layer 2 loop symptoms	DISTRIBUTION	Connectivity	48	1	2

Layer 2 loop symptoms

2 Open Issues

1 Area
1 Buildings, 0 Floors

2 DISTRIBUTION

Issue	Site	Device	Device Type	Issue Count
Issue Type: Network Down (1) (NPD)	USA/DF	SF-D9300-1	Cisco Catalyst 9300 Switch	28
Issue Type: Network Down (1) (NPD)	USA/DF	SF-D9300-2	Cisco Catalyst 9300 Switch	20

Potential Loop Details

Filter

Device	Role	Port to Loop	Degree	VLAN to Loop
SF-D9300-1	DISTRIBUTION	GigabitEthernet1/0/13	Full	30-33
SF-D9300-2	DISTRIBUTION	GigabitEthernet1/0/13	Full	30-33
SF-D9300-1	DISTRIBUTION	GigabitEthernet1/0/23	Full	30-33
SF-D9300-2	ACCESS	GigabitEthernet1/0/23	Full	30-33

```

interface GigabitEthernet1/0/13
switchport trunk allowed vlan 30-33
switchport mode trunk
!
interface GigabitEthernet1/0/23
switchport trunk allowed vlan 30-33
switchport mode trunk

```

An engineer identifies a Layer 2 loop using DNAC. Which command fixes the problem in the SF-D9300-1 switch?

- A. no spanning-tree uplinkfast
- B. spanning-tree loopguard default
- C. spanning-tree backbonesfast
- D. spanning-tree portfast bpduguard

Answer: D

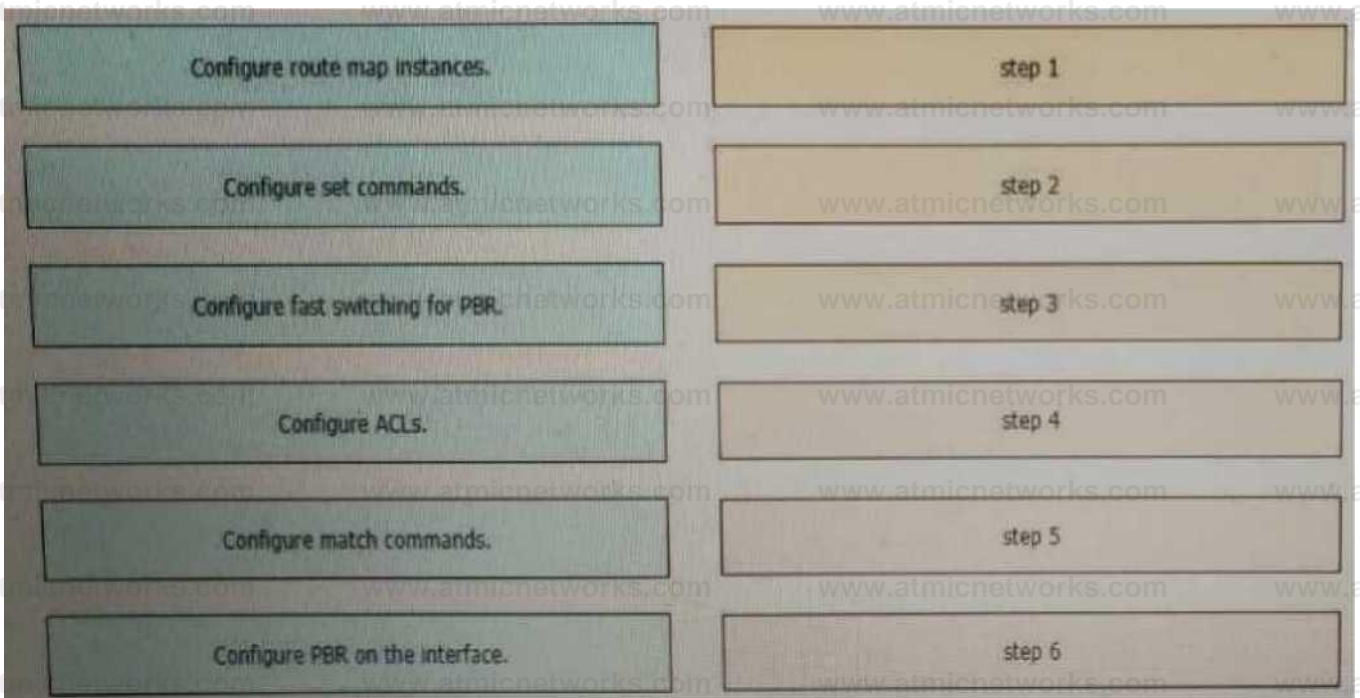
Explanation:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dnacenter/tech_notes/b_dnac_sda_lan_automation_deployment.html

Question: 129

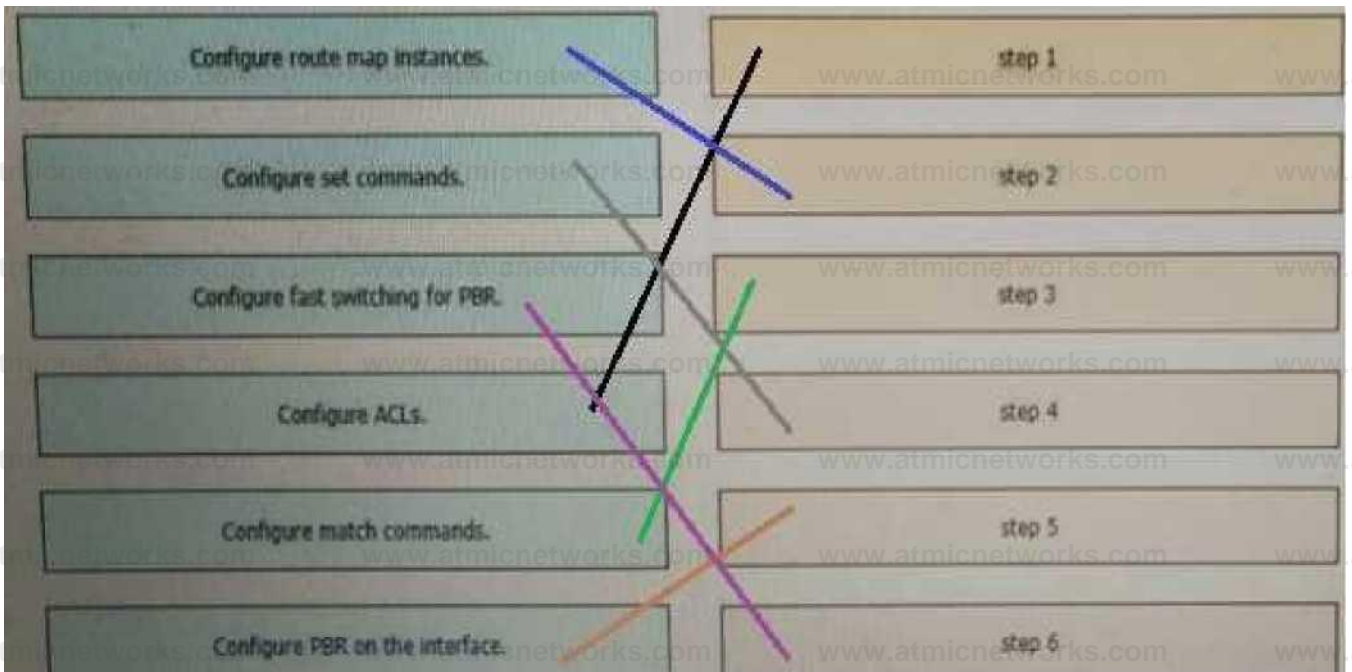
DRAG DROP

Drag and drop the actions from the left into the correct order on the right to configure a policy to **avoid** following packet forwarding based on the normal routing path.



Answer:

Explanation:



<https://community.cisco.com/t5/networking-documents/how-to-configure-pbr/ta-p/3122774>

Question:

130

What are two functions of LDP? (Choose two.)

- A. It is defined in RFC 3038 and 3039.
- B. It requires MPLS Traffic Engineering.
- C. It advertises labels per Forwarding Equivalence Class.
- D. It must use Resource Reservation Protocol.
- E. It uses Forwarding Equivalence Class

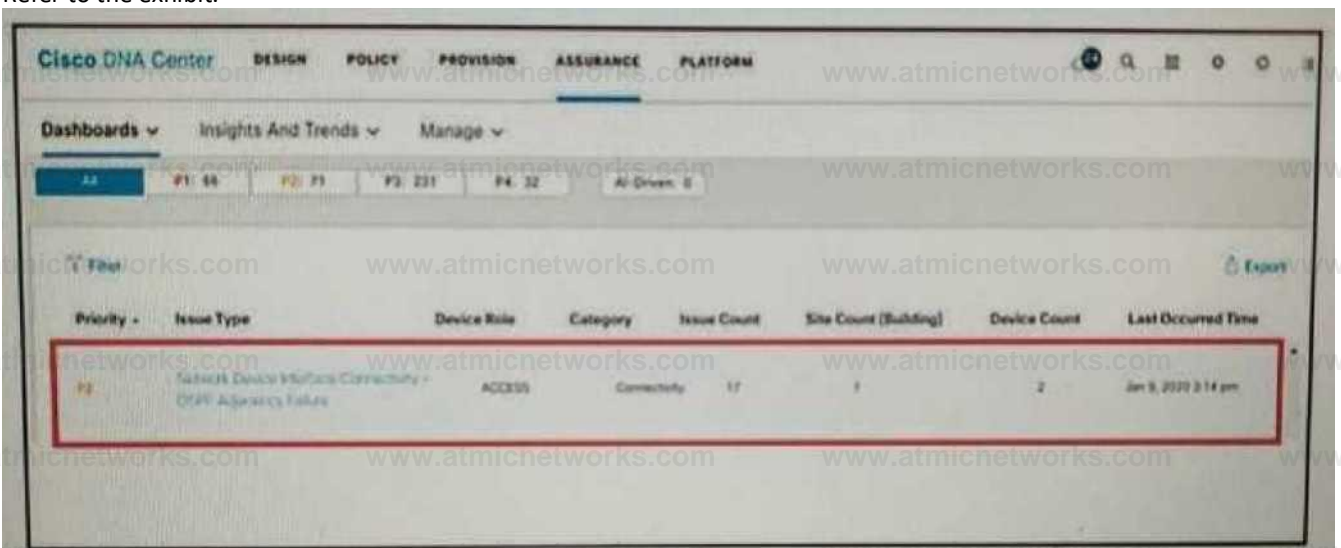
Answer: C,E

Explanation:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_mpls_overview.pdf

Question: 131

Refer to the exhibit.



A network administrator is using the DNA Assurance Dashboard panel to troubleshoot an OSPF adjacency that failed between Edge_NYC interface GigabitEthernet1/3 with Neighbor Edge_SNJ. The administrator

observes that the neighborship is stuck in exstart state. How does the administrator fix this issue?

- A. Configure to match the OSPF interface speed and duplex settings on both routers.
- B. Configure to match the OSPF interface MTU settings on both routers.
- C. Configure to match the OSPF interface unique IP address and subnet mask on both routers.
- D. Configure to match the OSPF interface network types on both routers.

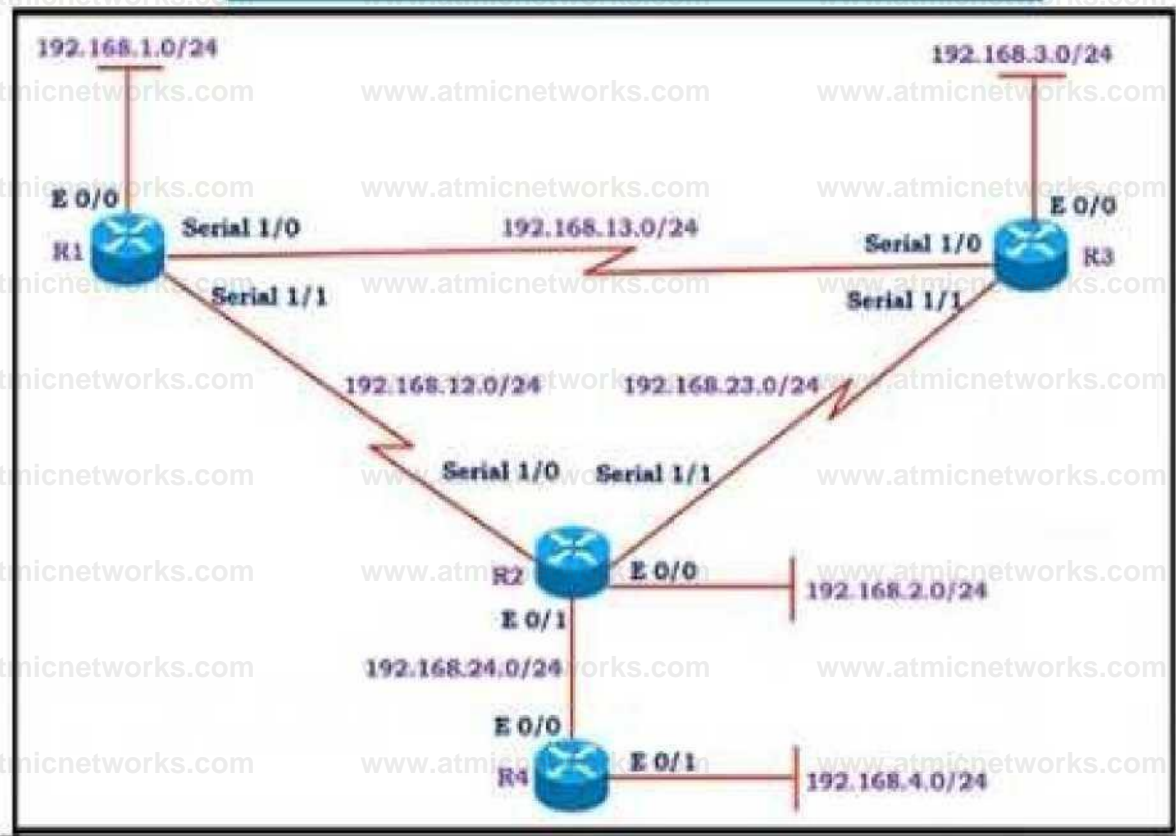
Answer: B

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13684-12.html>

Question: 132

Refer to the exhibit.



d Show IP route on R1

192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks

C 192.168.1.0/24 is directly connected. Ethernet0/0

L 192.168.1.1/32 is directly connected. Ethernet0/0

O 192.168.2.0/24 [90/2297856] via 192.168.12.2, 00:02:14. Serial1/1

S 192.168.3.0/24 [1/0] via 192.168.12.2

192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks

C 192.168.12.0/24 is directly connected. Serial1/1

L 192.168.12.1/32 is directly connected. Serial1/1

192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks

C 192.168.13.0/24 is directly connected. Serial1/0

L 192.168.13.1/32 is directly connected, Serial1/0

D 192.168.23.0/24 [90/2681856] via 192.168.13.3, 00:06:38 Serial1/0

[90/2681858] via 192.168.12.2, 00:06:30, Serial1/1

All the serial between R1, R2, and R3 have the same bandwidth. User on the 192.168.1.0/24 network report slow response times while they access resource on network 192.168.3.0/24. When a traceroute is run on the path. It shows that the packet is getting forwarded via R2 to R3 although the link between R1 and R3 is still up. What must the network administrator do to fix the slowness?

- A. Change the Administrative Distance of EIGRP to 5.
- B. Add a static route on R1 using the next hop of R3.
- C. Remove the static route on R1.
- D. Redistribute the R1 route to EIGRP

Answer: C

Explanation:

Question: 133

An engineer configured a Cisco router to send reliable and encrypted notifications for any events to the management server. It was noticed that the notification messages are reliable but not encrypted. Which action resolves the issue?

- A. Configure all devices for SNMPv3 informs with priv.
- B. Configure all devices for SNMPv3 informs with auth.
- C. Configure all devices for SNMPv3 traps with auth.
- D. Configure all devices for SNMPv3 traps with priv.

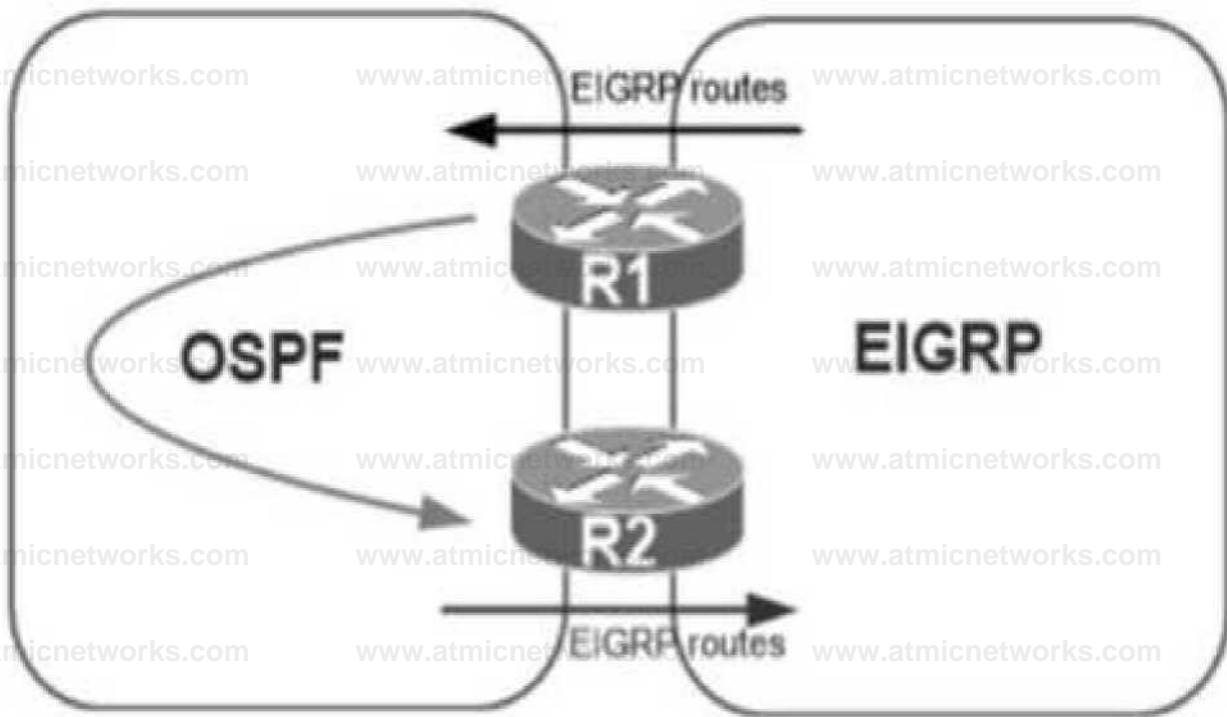
Answer: A

Explanation:

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when this device receives traps. "Send reliable and encrypted notifications for any events" so it is SNMP notifications. For encryption we need to configure "priv".

Question: 134

Refer to the exhibit.



A network administrator configured mutual redistribution on R1 and R2 routers, which caused instability in the network. Which action resolves the issue?

- A. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to allow when redistributing OSPF into EIGRP.
- B. Apply a prefix list of EIGRP network routes in OSPF domain on R1 to propagate back into the EIGRP routing domain.
- C. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to deny when redistributing OSPF into EIGRP.
- D. Advertise summary routes of EIGRP to OSPF and deny specific EIGRP routes when redistributing into OSPF.

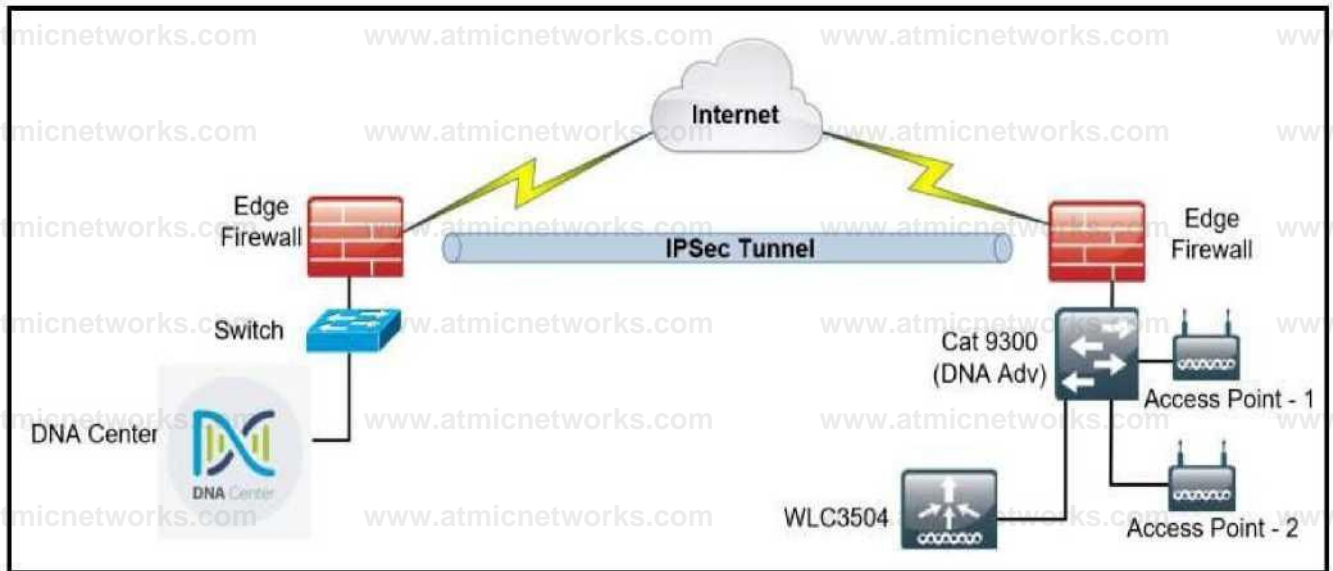
Answer: C

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html>

Question: 135

Refer to the exhibit.



A network administrator is discovering a Cisco Catalyst 9300 and a Cisco WLC 3504 in Cisco DNA Center. The Catalyst 9300 is added successfully. However, the WLC is showing [error "uncontactable"] when the administrator tries to add it in Cisco DNA Center. Which action discovers WLC in Cisco DNA Center successfully?

- A. Copy the .cert file from the Cisco DNA Center on the USB and upload it to the WLC 3504.
- B. Delete the WLC 3504 from Cisco DNA Center and add it to Cisco DNA Center again.
- C. Add the WLC 3504 under the hierarchy of the Catalyst 9300 connected devices.
- D. Copy the .pern file from the Cisco DNA Center on the USB and upload it to the WLC 3504.

Answer: D

Explanation:

<https://www.cisco.com/c/en/us/support/docs/wireless/4400-series-wireless-lan-controllers/109597-csr->

Question: 136

Which feature drops packets if the source address is not found in the snooping table?

- A. IPv6 Source Guard
- B. IPv6 Destination Guard
- C. IPv6 Prefix Guard
- D. Binding Table Recovery

Answer: A

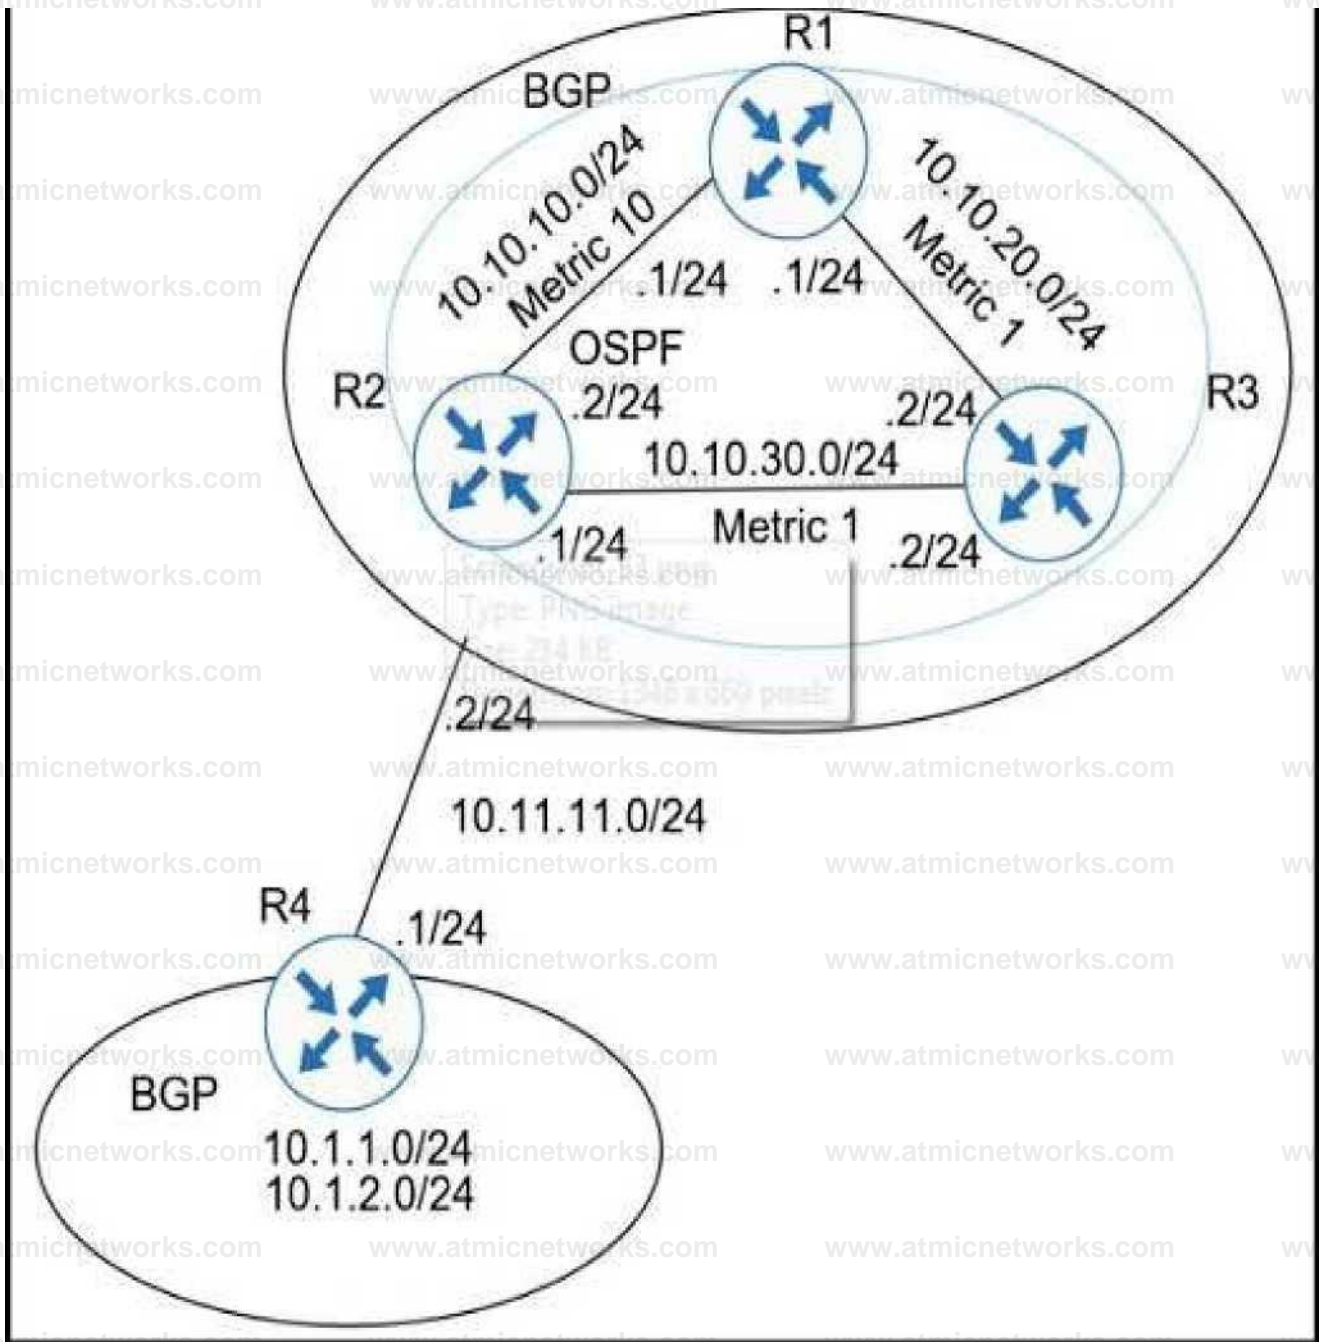
Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xr-3s/ip6f-xr-3s-book/ip6-snooping.pdf

Question: 137

Refer to the exhibit.

```
ip sla 10
tcp connect 10 1 1 1 80
ip sla schedule 10 life 30 start time now
```



A user has set up an IP SLA probe to test if a non-SLA host web server on IP address 10.1.1.1 accepts HTTP sessions prior to deployment. The probe is failing. Which action should the network administrator

recommend for the probe to succeed?

- A. Re-issue the ip sla schedule command.


- B. Add icmp-echo command for the host.
- C. Add the control disable option to the tcp connect.
- D. Modify the ip sla schedule frequency to forever.

Answer: C

Explanation:

Question: 138

Refer to the exhibit.



The network administrator must mutually redistribute routes at the Chicago router to the LA and NewYork routers. The configuration of the Chicago router is this:

```
router ospf 1
 redistribute eigrp 100
router eigrp 100
 redistribute ospf 1
```

After the configuration, the LA router receives all the NewYork routes, but NewYork router does not receive any LA routes. Which set of configurations fixes the problem on the Chicago router?

A)

router ospf 1 redistribute eigrp 100 metric 20

B)

router eigrp 100 redistribute ospf 1 metric 1010 101010

C)

router eigrp 100 redistribute ospf 1 subnets

D)

router ospf 1 redistribute eigrp 100 subnets

A. Option A

B. Option B

C. Option C

D. Option D

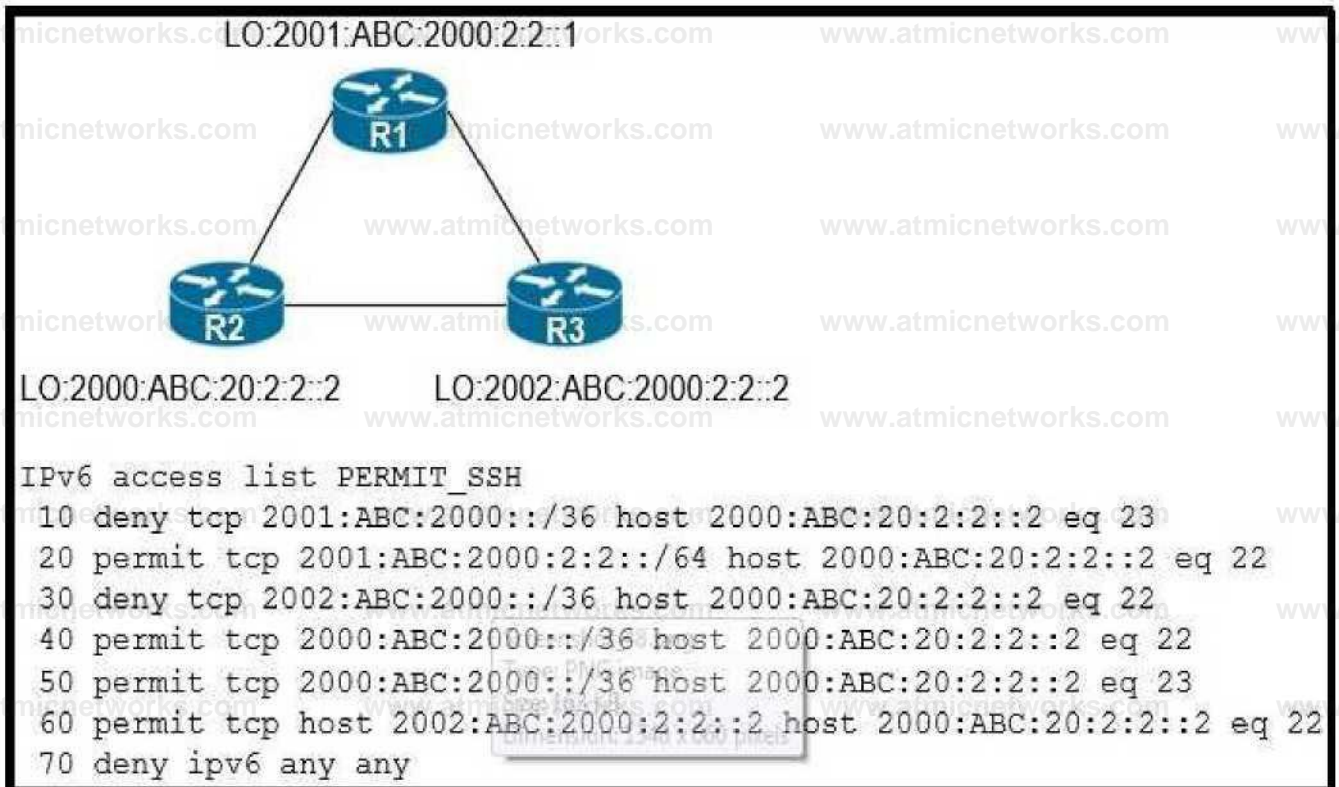
Answer: B

Explanation:

“LA router receives all the NewYork routes but it does not receive any LA routes” because when redistributing into EIGRP, we must configure the default metric.

Question: 139

Refer to the exhibit.



An IPv6 network was newly deployed in the environment and the help desk reports that R3 cannot SSH to the R2s Loopback interface. Which action resolves the issue?

- A. Modify line 10 of the access list to permit instead of deny.
- B. Remove line 60 from the access list.
- C. Modify line 30 of the access list to permit instead of deny.
- D. Remove line 70 from the access list.

Answer: C

Explanation:

Question: 140

An engineer configured SNMP notifications sent to the management server using authentication and encrypting data with DES. An error in the response PDU is received as "UNKNOWNUSERNAME.WORKS.COM WRONGDIGEST". Which action resolves the issue?

- A. Configure the correct authentication password using SNMPv3 authPriv .
- B. Configure the correct authentication password using SNMPv3 authNoPriv.
- C. Configure correct authentication and privacy passwords using SNMPv3 authNoPriv.
- D. Configure correct authentication and privacy passwords using SNMPv3 authPriv.


Answer: D

Explanation:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/xs-3se/3850/snmp-xe-3se-3850-book/nm-snmp-snmpv3.html>

Question: 141

Refer to the exhibits.



Chicago Router

```
ip route 192.168.1.0 255.255.255.0 10.1.1.2
ip route 192.168.2.0 255.255.255.0 10.1.1.2
!
```

```
router eigrp 100
 redistribute static
```

LA Router

```
ip route 0.0.0.0 0.0.0.0 10.1.1.1
```

A user on the 192.168.1.0/24 network can successfully ping 192.168.3.1, but the administrator cannot ping 192.168.3.1 from the LA router. Which set of configurations fixes the issue?

A)

Chicago Router

```
router eigrp WO
 redistribute static metric W W WWW
```

B)

Chicago Router

```
router eigrp WO
 redistribute connected
```

C)

Chicago Router**ip route 192.168.3.0 255.255.255.0 10.1.2.2****ip route 192.168.4.0 255.255.255.0 10.1.2.2**

D)

LA Router**ip route 192.168.3.0 255.2 55.255.0 10.1.1.1****io route 192.168.4.0 255.255.255.0 10.1.1.1**

A. Option A

B. Option B

C. Option C

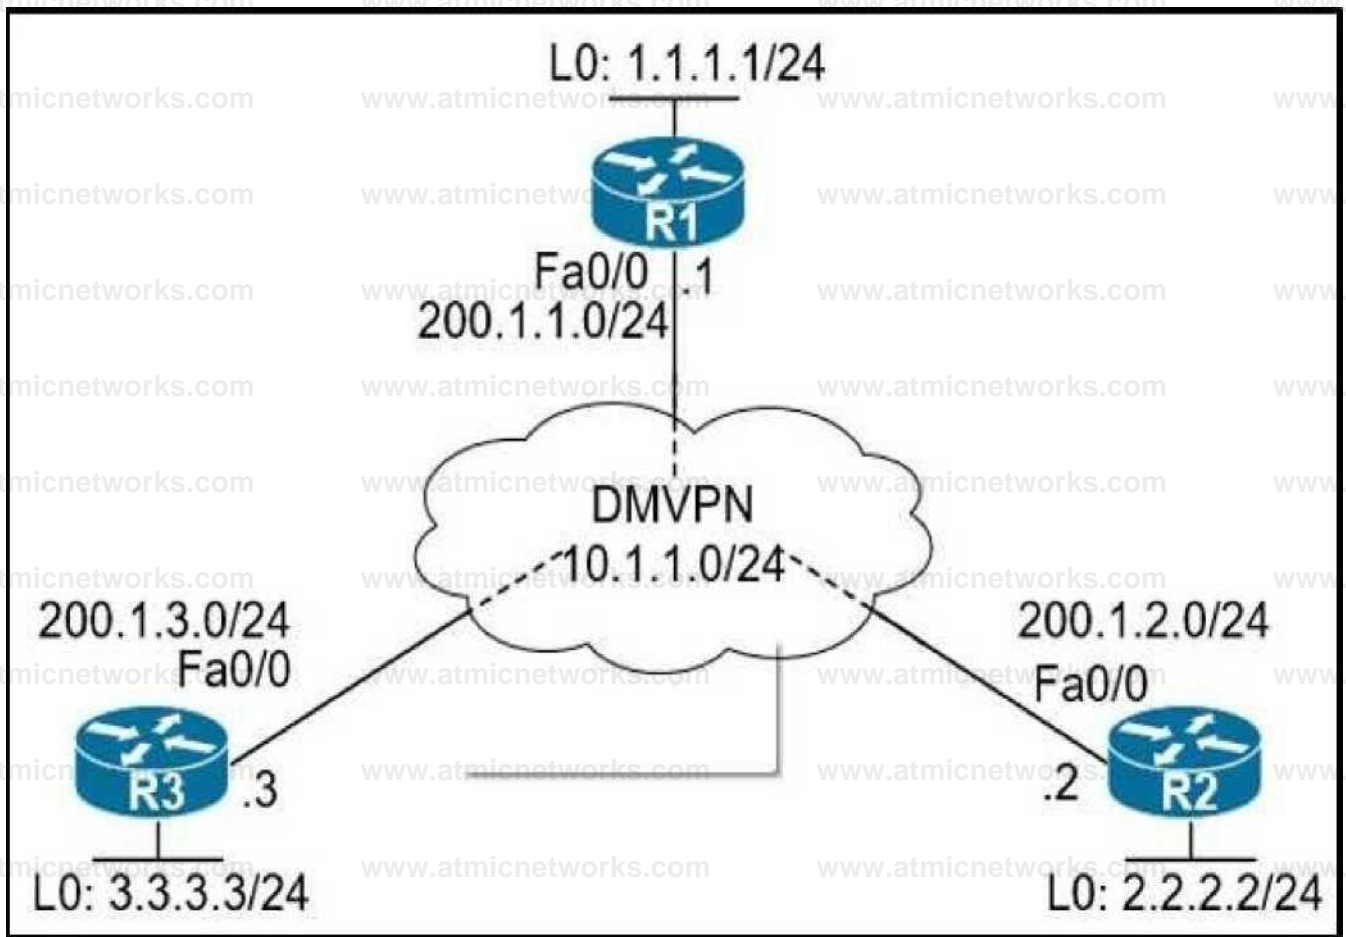
D. Option D

Answer: B

Explanation:

Question: 142

Refer to the exhibits.



```

R2(config)# crypto isakmp policy 10
R2(config-isakmp)# hash md5
R2(config-isakmp)# authentication pre-share
R2(config-isakmp)# group 2
R2(config-isakmp)# encryption 3des
R2(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
R2(cfg-crypto-trans)# mode transport
R2(config)# crypto ipsec profile TST
R2(ipsec-profile)# set transform-set TSET
R2(config)# interface tunnel 123
R2(config-if)# tunnel protection ipsec profile TST

```

When DMVPN is configured, which configuration allows spoke-to-spoke communication using loopback as a tunnel source?

- A. Configure crypto isakmp key cisco address 0.0.0.0 on the hub.
- B. Configure crypto isakmp key Cisco address 200.1.0.0 255.255.0.0 on the hub.

- C. Configure crypto isakmp key cisco address 200.1.0.0 255.255.0.0 on the spokes.
- D. Configure crypto isakmp key cisco address 0.0.0.0 on the spokes.

Answer: D

Explanation:

https://www.cisco.com/en/US/technologies/tk583/tk372/technologies_white_paper0900aecd802b_8f3c.html

Question: 143

What are two functions of IPv6 Source Guard? (Choose two.)

- A. It uses the populated binding table for allowing legitimate traffic.
- B. It works independent from IPv6 neighbor discovery.
- C. It denies traffic from unknown sources or unallocated addresses.
- D. It denies traffic by inspecting neighbor discovery packets for specific pattern.
- E. It blocks certain traffic by inspecting DHCP packets for specific sources.

Answer: A,C

Explanation:

IPv6 source guard is an interface feature between the **populated binding table** and **data traffic filtering**

IPv6 source guard can deny traffic from **unknown sources** or **unallocated addresses**,

Question: 144

An engineer configured access list NON-CISCO in a policy to influence routes

route-map PBR. deny, sequence 5

Match clauses:

ip address (access-list): NON-CISCO

Set clauses:

Policy routing matches: 0 packets, 0 bytes

route-map PBR, permit, sequence 10

Match clauses:

Set clauses:

ip next-hop 192.168.1.5

Policy routing matches: 38 3213827 packets, 2 220 096 85077 bytes

What are the two effects of this route map configuration? (Choose two.)

- A. Packets are not evaluated by sequence 10.
- B. Packets are evaluated by sequence 10.
- C. Packets are forwarded to the default gateway.
- D. Packets are forwarded using normal route lookup.
- E. Packets are dropped by the access list.

Answer: B,C

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html>

Question: 145

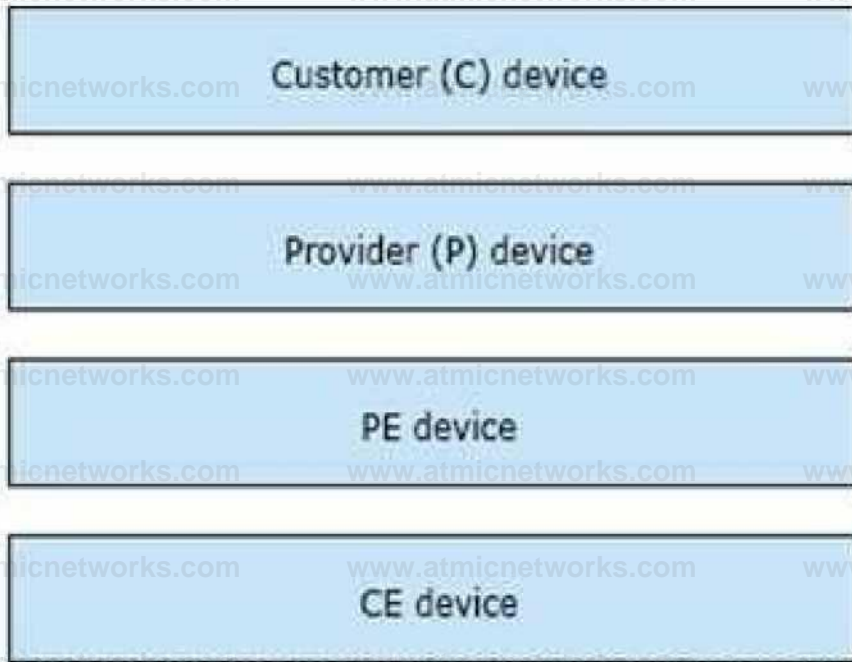
DRAG DROP

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

Customer (C) device	device in the core of the provider network that switches MPLS packets
CE device	device that attaches and detaches the VPN labels to the packets in the provider network
PE device	device in the enterprise network that connects to other customer devices
Provider (P) device	device at the edge of the enterprise network that connects to the SP network

Answer:

Explanation:



Graphical user interface,

application Description automatically generated

Question: 146

Refer to the exhibit.


```
R1#show policy-map control-plane
Control Plane
Service-policy input: CoPP
Class-map: PERMIT (match-all)
  50 packets, 3811 bytes
  5 minute offered rate 0000 bps
  Match: access-group 100
Class-map: ANY (match-all)
  210 packets, 19104 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 199
  drop
Class-map: class-default (match-any)
  348 packets, 48203 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any

R1#show access-list 100
Extended IP access list 100
  10 permit udp any any eq 23 (100 matches)
  20 permit tcp any any eq telnet (5 matches)
  30 permit tcp any eq telnet any (10 matches)

R1#show access-list 199
Extended IP access list 199
  10 deny tcp any eq telnet any (50 matches)
  50 permit ip any any (1 match)

R1#show running-config | section line vty
line vty 0 4
 login
 transport input telnet ssh
 transport output telnet ssh
```

Which two actions restrict access to router R1 by SSH? (Choose two.)

- A. Configure transport input ssh on line vty and remove sequence 30 from access list 100.
- B. Configure transport output ssh on line vty and remove sequence 20 from access list 100.
- C. Remove class-map ANY from service-policy CoPP
- D. Configure transport output ssh on line vty and remove sequence 10 from access list 199.
- E. Remove sequence 10 from access list 100 and add sequence 20 deny tcp any any eq telnet to access list 199

Answer: A,B

Explanation:

Question: 147

What is the minimum time gap required by the local system before putting a BFD control packet on the wire?

- A. Detect Mult
- B. Required Min Echo RX Interval
- C. Desired Min TX Interval
- D. Required Min RX Interval

Answer: C

Explanation:

Desired Min TX Interval: This is the minimum interval, in microseconds, that the local system would like to use when transmitting BFD Control packets, less any jitter applied. The value zero is reserved.

Required Min Echo RX Interval: This is the minimum interval, in microseconds, between received BFD Echo packets that this system is capable of supporting, less any jitter applied by the sender. If this value is zero, the transmitting system does not support the receipt of BFD Echo packets.

Reference: <https://tools.ietf.org/html/rfc5880>

Question: 148

Refer to the exhibit.

login block-for 15 attempts 10 within 120

login on-failure log

login on-success log

archive

log config

logging enable

logging size 300

notify syslog

snmp-server enable traps syslog

snmp-seiver host 17216.17.1 public syslog

The administrator can see the traps for the failed login attempts, but cannot see the traps of successful login attempts. What command is needed to resolve the issue?

- A. Configure logging history 2
- B. Configure logging history 3

C. Configure logging history 4

D. Configure logging history 5

Answer: D

Explanation:

By default, the maximum severity sent as a syslog trap is warning. That is why you see syslog traps for login failures. Since a login success is severity 5 (notifications), those syslog messages will not be converted to traps. To fix this, configure:

logging history 5

Syslog levels are listed below

Level	Keyword	Description
0	emergencies	System is unusable
1	alerts	Immediate action is needed
2	critical	Critical conditions exist
3	errors	Error conditions exist
4	warnings	Warning conditions exist
5	notification	Normal, but significant, conditions exist
6	informational	Informational messages
7	debugging	Debugging messages

Note:

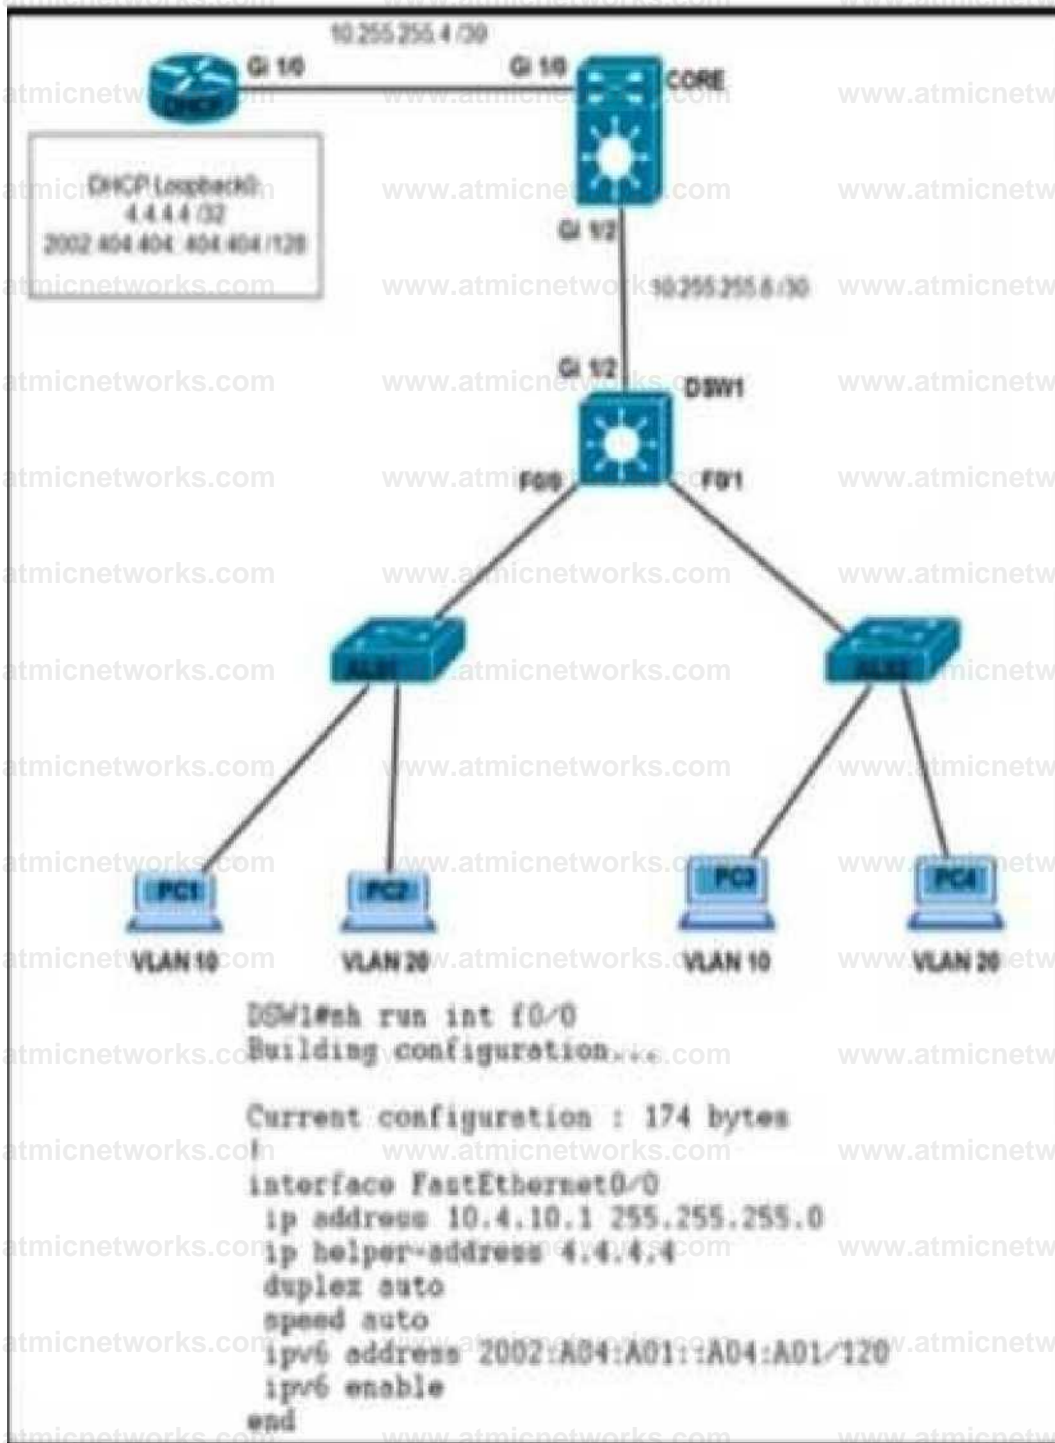
The syntax of login block is:

login block-for seconds attempts tries within seconds

Question: 149

Clients on ALS2 receive IPv4 and IPv6 addresses but clients on ALS1 receive only IPv4 addresses and not IPv6 addresses.

Which action on DSW1 allows clients on ALS1 to receive IPv6 addresses?



A. Configure DSW1(config-if)#ipv6 helper address 2002:404:404::404:404

- B. Configure DSW1(dhcp-config)#default-router 2002:A04:A01::A04:A01
- C. Configure DSW1(config)#ipv6 route 2002:404:404:404:404/128 FastEthernet1/0
- D. Configure DSW1(config-if)#ipv6 dhcp relay destination 2002:404:404::404:404 GigabitEthernet1/2
- E. Option A
- F. Option B
- G. Option C
- H. Option B

Answer: B

Explanation:

<https://community.cisco.com/t5/networking-documents/stateful-dhcpv6-relay-configuration-example/ta-p/3149338>

Question: 150

A network administrator is tasked to permit http and https traffic only toward the internet from the User1 laptop to adhere to company's security policy. The administrator can still ping to www.cisco.com Which interface should the access list 101 be applied to resolve this issue?

```
access list 101 permit tcp 192.168.10.0
0.0.0.255 any eq 80
access-hst 101 permit tcp 192.168.10.0
0.0.0.255 any eq 443
access-list 101 deny ip any any log
|
interface Serial 1/0
ip address 200.193.22.94 255.255.255.252
ip access-group 101 in
```

Internet



Internet Router



192 (>8 10.100

- A. Interface G0/48 in the incoming direction
- B. Interface G0/0 in the outgoing direction.
- C. Interface S1/0 in the outgoing direction.
- D. Interface G0/0 in the incoming direction.

Answer:

D

Explanation:

Question:

151

Refer to Exhibit.

```
HQ_R2 g0/0  
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.2 track 1  
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.6 5  
|  
BRANCH(config)# ip sla 1  
BRANCH(config-ip-sla)# icmp-echo 172.16.35.6  
BRANCH(config-ip-sla)# timeout 200  
BRANCH(config-ip-sla)# frequency 5  
|  
BRANCH(config)# ip sla schedule 1 life forever start-time now  
|  
BRANCH(config)# track 1 ip sla 1 reachability
```

Traffic from the branch network should route through HQ R1 unless the path is unavailable. An engineer tests this functionality by shutting down interface on the BRANCH router toward HQ_R1 router but 192.168.20.0/24 is no longer reachable from the branch router. Which set of configurations resolves the issue?

- A. HQ_R1(config)# ip sla responderHQ_R1(config)# ip sla responder icmp-echo 172.16.35.2
- B. BRANCH(config)# ip sla 1BRANCH(config-ip-sla)# icmp-echo 172.16.35.1
- C. HQ_R2(config)# ip sla responderHQ_R2(config)# ip sla responder icmp-echo 172.16.35.5
- D. BRANCH(config)# ip sla 1BRANCH(config-ip-sla)# icmp-echo 172.16.35.2

Answer: D

Explanation:

In the configuration above, the engineer has made a mistake as he was tracking 172.16.35.6 (the backup path) instead of tracking the main path (172.16.35.2). Therefore, when he shut down the main path, the track 1 was still up so traffic still went through the main path -> it failed.

To fix this issue, we just need to correct the tracking interface of the main path.

Question: 152

Refer to Exhibit.

```
ip dhcp excluded-address 172.16.16.1 172.16.16.2 1
ip dhcp pool 0 network 172.16.16.0 255.255.255.0
domain-name cisco.com dns-server 172.16.16.2
tease 30
```

```
interface Ethernet0/0 ip address 10.1.1.1 2
55.255.255.252 ip access-group 100 in
```

```
access-list 100 deny udp any any
access-list 100 permit ip any any
```

Which two configurations allow clients to get dynamic ip addresses assigned?

- A. Configure access-list100 permit udp any any eq 61 as the firstline
- B. Configure access-list100 permit udp any any eq 86 as the firstline
- C. Configure access-list100 permit udp any any eq 68 as the firstline
- D. Configure access-list100 permit udp any any eq 69 as the firstline
- E. Configure access-list100 permit udp any any eq 67 as the firstline

Answer: C,E

Explanation:

A DHCP server that receives a DHCPDISCOVER message may respond with a DHCPOFFER message on UDP port 68 (BootP client).

In the event that the DHCP server is not on the local subnet, the DHCP server will send the DHCPOFFER, as a unicast packet, on UDP port 67, back to the DHCP/BootP Relay Agent from which the DHCPDISCOVER came.

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html>

Question: 153

Which IPv6 first-hop security feature helps to minimize denial of service attacks?

- A. IPv6 Router Advertisement Guard
- B. IPv6 Destination Guard
- C. DHCPv6 Guard
- D. IPv6 MAC address filtering

Answer: B

Explanation:

The Destination Guard feature helps in minimizing denial-of-service (DoS) attacks. It performs address resolutions only for those addresses that are active on the link, and requires the FHS binding table to be populated with the help of the IPv6 snooping feature. The feature enables the filtering of IPv6 traffic based on

the destination address, and blocks the NDP resolution for destination addresses that are not found in the binding table. By default, the policy drops traffic coming for an unknown destination.

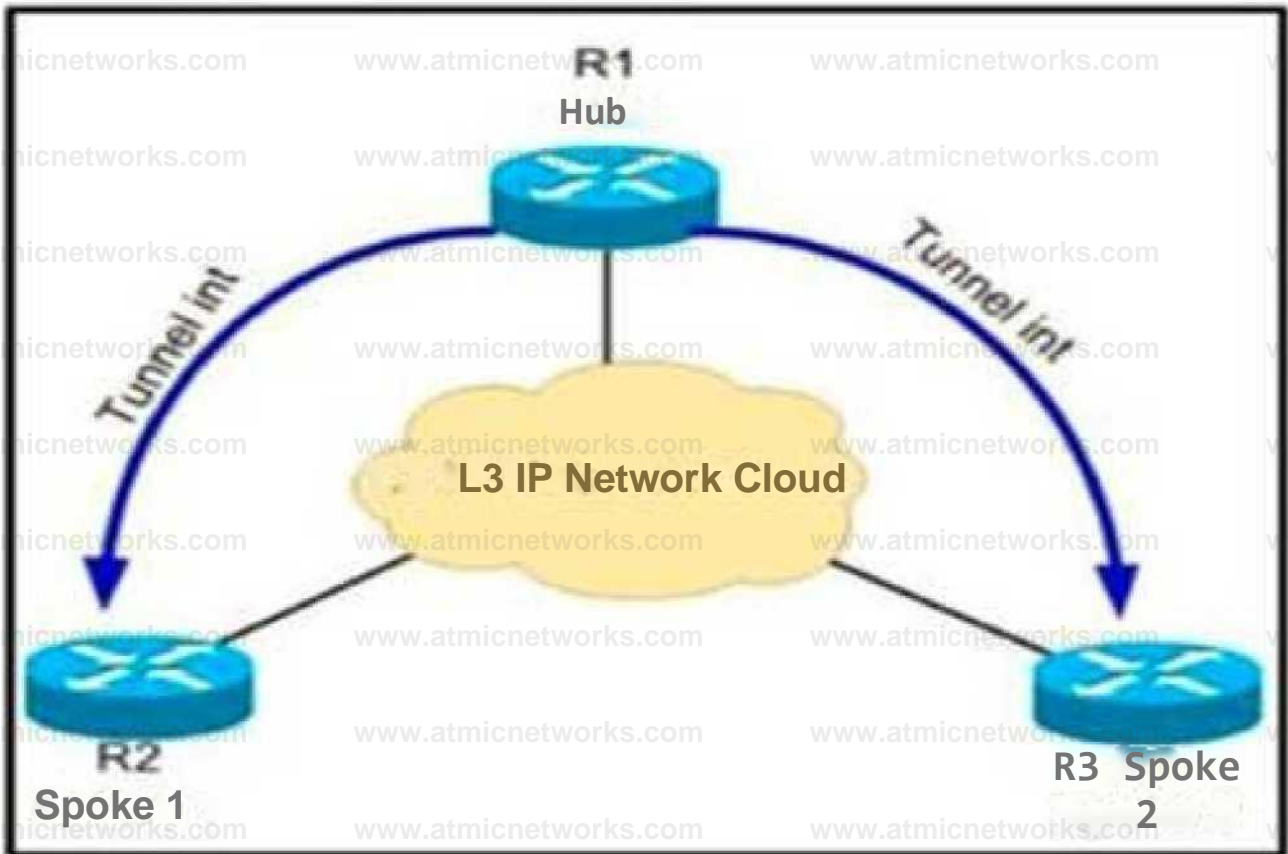
Reference:

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.pdf

Question: 154

Refer to Exhibit.

A network administrator has successfully configured DMVPN topology between a hub and two spoke routers. Which two configuration commands should establish direct communications between spoke 1 and spoke 2 without going through the hub? (Choose two).



- A. At the hub router, configure the ip nhrp shortcut command.
- B. At the spoke routers, configure the ip nhrp spoke-tunnel command.
- C. At the hub router, configure ip nhrp redirect the command
- D. At the spoke routers, configure the ip nhrp shortcut command.
- E. At the hub router, configure the ip nhrp spoke-tunnel command

Answer: C,D

Explanation:

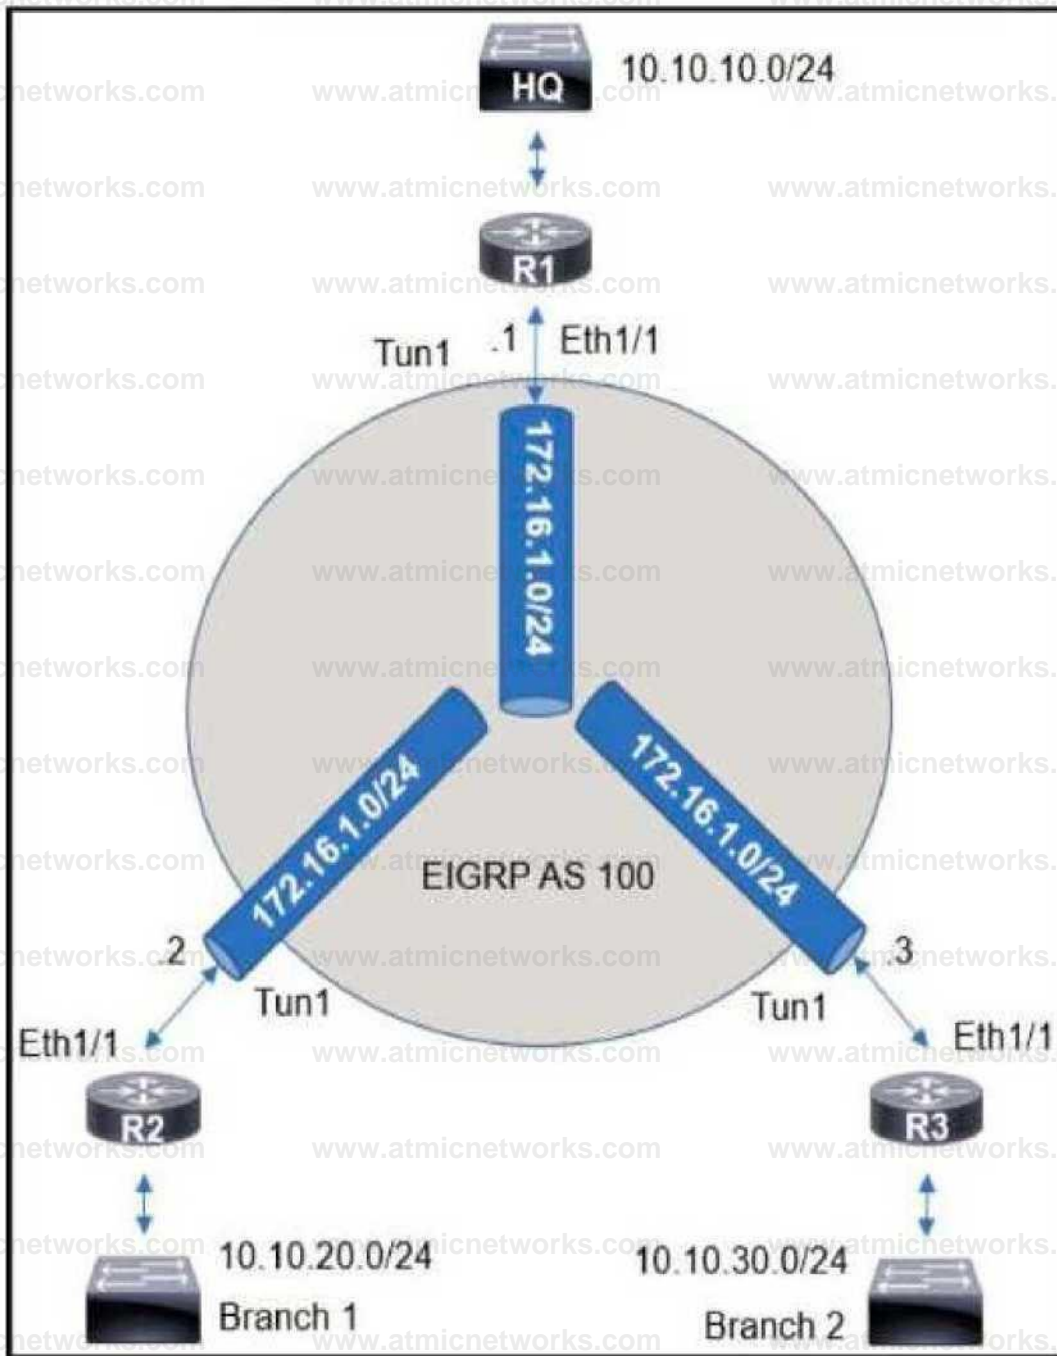
To configure Spoke to Spoke communication we can configure DMVPN Phase II or Phase III. But in Phase II, the first few packets would go through Hub. In order to totally ignore the hub, we have to

use DMVPN Phase III:

DMVPN Phase III is same as Phase 2 but removes some restrictions and complexities of Phase 2. Also allows greater variety of DMVPN network designs we use: + ip nhrp redirect in hub: tells the initiator spoke to look for a better path to the destination spoke than through the Hub. Upon receiving the NHRP redirect message the spokes communicate with each other over the hub and they have their NHRP replies for the NHRP Resolution Requests that they sent out. + ip nhrp shortcut in spokes: overwrite the CEF table on the spoke. It basically overrides the next-hop value for a remote spoke network from the default initial hubtunnel IP address to the NHRP resolved remote spoke tunnel IP address)

Question: 155

Refer to the exhibit.



An engineer sets up a DMVPN connection to connect branch 1 and branch 2 to HQ. Branch 1 and branch 2 cannot communicate with each other. Which change must be made to resolve this issue?

R2(config)# interface Tunnel1

R1(config)# ip split-horizon eigrp 100

R2(config)# router eigrp 100

R2(config)# neighbor 172.16.1.3

R3(config) "router eigrp 100

R3(config-if) router -neighbor 172.16.1.2

RKwnfigjeint tunnel 1

RUconfig.u «no Ip spllbhorizon eigrp 100

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

R1(config)#int tunnel 1

R1(config-if) no ip split-horizon eigrp 100

Question: 156

Refer to the exhibit.


```
access-list 1 permit 1.1.1.0 0.0.0.255
!
route-map FILTER1 deny 10
match ip address 1
!
router eigrp 1
distribute-list route-map FILTER1 in
```

Refer to the exhibit. Which action restores the routes from neighbors while still filtering 1.1.1.0/24?

- A. Add a second line in the access list to permit any.
- B. Modify the route map to permit the access list instead of deny it
- C. Modify the access list to deny instead of permit it.
- D. Add a second sequence in the route map permit 20

Answer: D

Explanation:

Question: 157

Which two components are needed for a service provider to utilize the LVPN MPLS application? (Choose two.)

- A. The P routers must be configured for MP-iBGP toward the PE routers
- B. The P routers must be configured with RSVP.
- C. The PE routers must be configured for MP-iBGP with other PE routers
- D. The PE routers must be configured for MP-eBGP to connect to CEs

E. The P and PE routers must be configured with LDP or RSVP

Answer: C,E

Explanation:

MPLS Network Protocols

+ IGP: OSPF, EIGRP, IS-IS on core facing and core links+ RSVP and/or LDP on core and/or core facing links ->

+ MP-iBGP on PE devices (for MPLS services), MP-BGP: Multiprotocol Border Gateway Protocol, used for MPLS L3 VPN -> .

Reference: [https://www.uio.no/studier/emner/matnat/ifi/IN3230/h19/kursmaterieell/mpls-](https://www.uio.no/studier/emner/matnat/ifi/IN3230/h19/kursmaterieell/mpls-lecture.pdf)

[lecture.pdf](https://www.uio.no/studier/emner/matnat/ifi/IN3230/h19/kursmaterieell/mpls-lecture.pdf)

Question: 158

What are two characteristics of VRF instance? (Choose two.)

- A. All VRFs share customers routing and CEF tables .
- B. An interface must be associated to one VRF.
- C. Each VRF has a different set of routing and CEF tables
- D. It is defined by the VPN membership of a customer site attached to a P device.
- E. A customer site can be associated to different VRFs

Answer: B,C

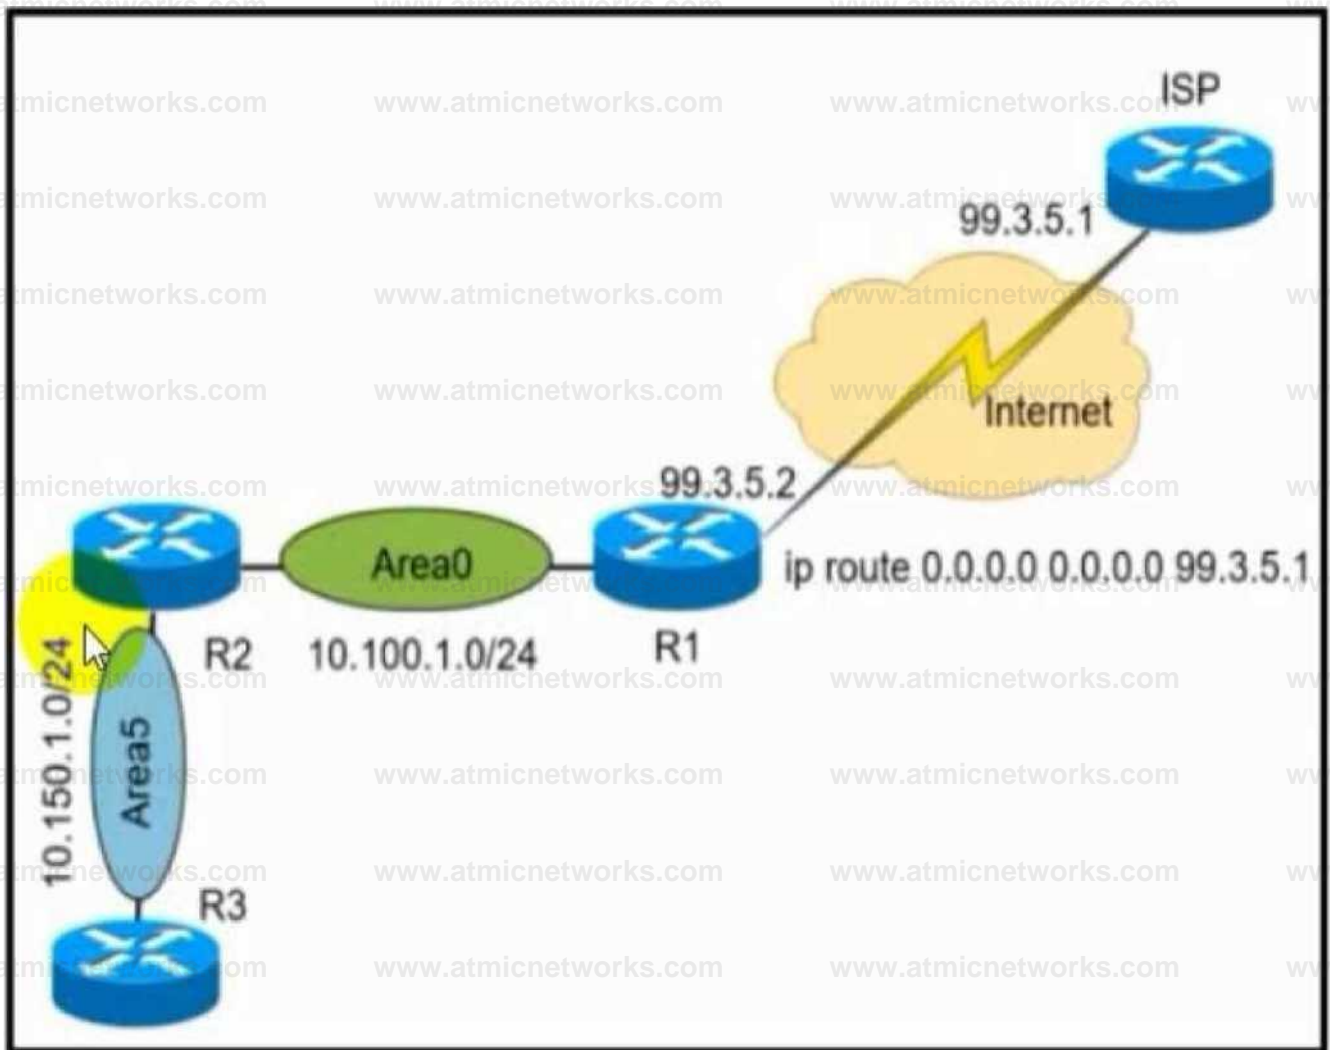
Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch_cef/configuration/xr-3s/isw-cef-xe-3s-book/isw-cef-basic-config.html

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-s/mp-l3-vpns-15-s-book/mp-bgp-mpls-vpn.pdf

Question: 159

Refer to the exhibit.



Refer to the exhibit. A network administrator redistributed the default static route into OSPF toward all internal routers to reach to Internet. Which set of commands restores reachability to the Internet by internal routers?

- A. `router ospf 1 default-information originate`
- B. `router ospf 1 network 0.0.0.0 0.0.0.0 area 0`
- C. `router ospf 1 redistribute connected 0.0.0.0`
- D. `router ospf 1 redistribute static subnets`

Answer: A

Explanation:

Question: 160

Refer to the exhibit.

```
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
  len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[10]
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
  len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[11]
%OSPF-5-ADJCHG: Process 1, Nbr 10.100.1.2 on GigabitEthernet0/1
from EXSTART to
DOWN, Neighbor Down: Too many retransmissions
```

Refer to the exhibit. The OSPF neighbor relationship is not coming up. What must be configured to restore OSPF neighbor adjacency?

- A. OSPF on the remote router
- B. matching hello timers
- C. use router ID
- D. matching MTU values

Answer: D

Explanation:

Question: 161

An engineer configured a DHCP server for Cisco IP phones to download its configuration from a TFTP server, but the IP phones failed to load the configuration. What must be configured to resolve the issue?

- A. BOOTP port 67
- B. DHCP option 66
- C. BOOTP port 68
- D. DHCP option 69

Answer: B

Explanation:

Command	Purpose
<code>dhcpd option 66 ascii server_name</code>	Provides the IP address or name of a TFTP server for option 66

Example:
hostname(config)# dhcpd option 66 ascii exampleserver

DHCP options 3, 66, and 150 are used to configure Cisco IP Phones. Cisco IP Phones download their configuration from a TFTP server. When a Cisco IP Phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information. + DHCP option 150 provides the IP addresses of a list of TFTP servers. + DHCP option 66 gives the IP address or the hostname of a single TFTP server.

Reference:

http://www.cisco.com/c/en/us/td/docs/security/asa/asa84/configuration/guide/asa_84_cli_config/basic_dhcp.pdf

Question: 162

Refer to the exhibit.

```
Ipv6 unicast-routing
```

```
!
```

```
Router ospfv3 4
```

```
Router-id 192.168.1.1
```

```
!
```

```
Interface E 0/0
```

```
Ipv6 enable
```

```
Ip address 10.1.1.1 255.255.255.0
```

```
Ospfv3 4 area 0 ipv4
```

```
No shut
```

```
!
```

```
Interface Loopback0
```

```
Ipv6 enable
```

```
Ipv4 172.16.1.1 255.255.255.0
```

```
Ospfv3 4 area 0 ipv4
```

Refer to the exhibit. The network administrator configured the branch router for IPv6 on the E 0/0 interface. The neighboring router is fully configured to meet requirements, but the neighbor relationship is not coming up. Which action fixes the problem on the branch router to bring the IPv6 neighbors up?

- A. Enable the IPv4 address family under the E 0/0 interface by using the address-family ipv4 unicast command
- B. Disable IPv6 on the E 0/0 interface using the no ipv6 enable command
- C. Enable the IPv4 address family under the router ospfv3 4 process by using the address-family ipv4 unicast command

D. Disable OSPF for IPv4 using the no ospfv3 4 area 0 ipv4 command under the E 0/0 interface.

Answer: C

Explanation:

Once again, Cisco changed the IOS configuration commands required for OSPFv3 configuration. The new OSPFv3 configuration uses the “ospfv3” keyword instead of the earlier “ipv6 router ospf” routing process command and “ipv6 ospf” interface commands.

The Open Shortest Path First version 3 (OSPFv3) address families feature enables both IPv4 and IPv6 unicast traffic to be supported. With this feature, users may have two processes per interface, but only one process per address family (AF).

Question: 163

An engineer configured two routers connected to two different service providers using BGP with default attributes. One of the links is presenting high delay, which causes slowness in the network. Which BGP attribute must the engineer configure to avoid using the high-delay ISP link if the second ISP link is up?

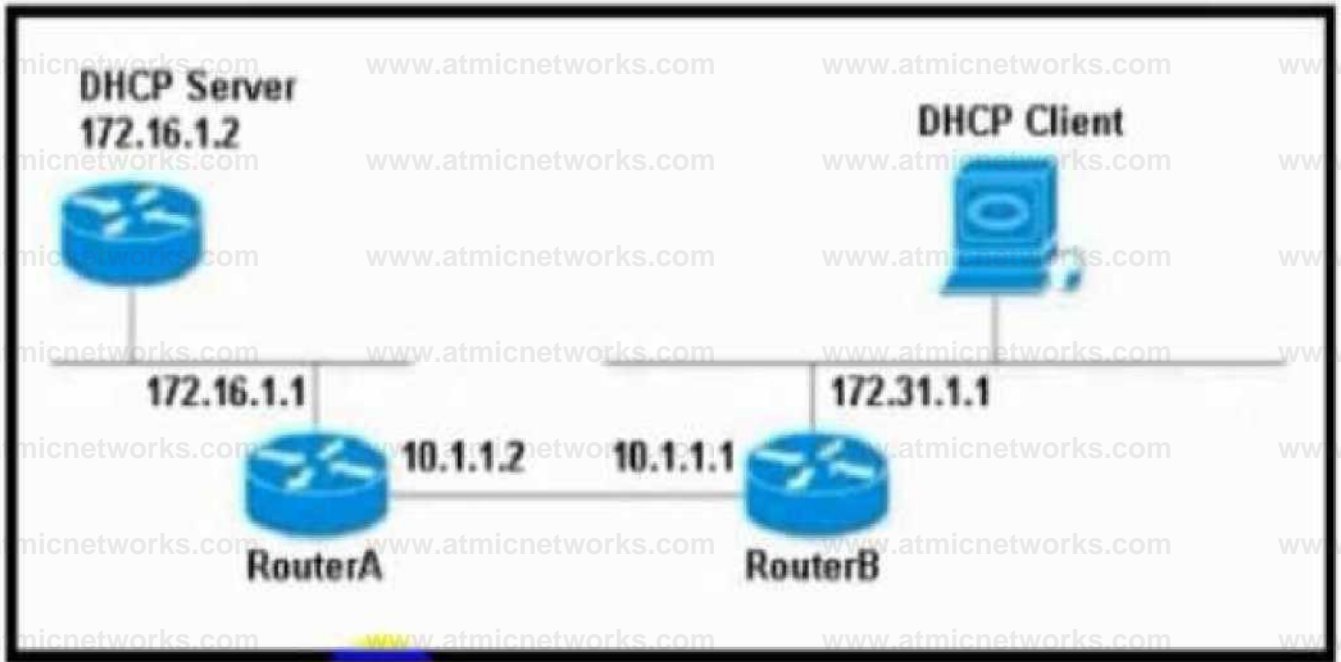
- A. LOCAL_PREF
- B. MED
- C. WEIGHT
- D. AS-PATH

Answer: A

Explanation:

Question: 164

Refer to the exhibit.



Refer to the exhibit. The DHCP client is unable to receive an IP address from the DHCP server RouterB is configured as follows:

```
Interface fastethernet 0/0
```

```
description Client DHCP ID 394482431
```

```
Ip address 172 31 11 255 255.255 0
```

```
!
```

```
ip route 172.16.1.0 255 255 255.0 10.1.1.2
```

Which command is required on the fastethernet 0/0 interface of RouterB to resolve this issue?

- A. RouterB(config-if)#ip helper-address 172.31.1.1
- B. RouterB(config-if)#ip helper-address 255.255.255.255
- C. RouterB(config-if)#ip helper-address 172.16.1.1
- D. RouterB(config-if)#ip helper-address 172.16.1.2

Answer: D

Explanation:

Question: 165

What are two purposes of using IPv4 and VPNv4 address-family configurations in a Layer 3 MPLS VPN?
(Choose two.)

- A. The VPNv4 address is used to advertise the MPLS VPN label.
- B. RD is prepended to the IPv4 route to make it unique.
- C. MP-BGP is used to allow overlapping IPv4 addresses between customers to advertise through the network.
- D. The IPv4 address is needed to tag the MPLS label.
- E. The VPNv4 address consists of a 64-bit route distinguisher that is prepended to the IPv4 prefix.


Answer: B,E

Explanation:

VPNv4 address consists of 64-bit Route Distinguisher (RD) prepended to IPv4 prefix. This is to make routes unique that are in different VRFs.

Question: 166

Refer to the exhibit.



Chicago

interface Tunnel 1

ip address 192.168.1.1 255.255.255.0

tunnel source E0/0

tunnel mode gre multipoint

ip nhrp network-id 1

ip nhrp map multicast dynamic

no ip next-hop-self eigrp 111

tunnel protection ipsec profile IPSec-PROFILE

!

router eigrp 111

network 192.168.1.0

network 10.0.0.0

Refer to the exhibit. The Los Angeles and New York routers are receiving routes from Chicago but not from each other. Which configuration fixes the issue?

- A. Interface Tunnel1no ip split-horizon eigrp 111

- B. Interface Tunnel1ip next-hop-self elgrp 111
- C. Interface Tunnel1tunnel mode Ipsec Ipv4
- D. Interface Tunnel1tunnel protection ipsec profile IPSec-PROFILE

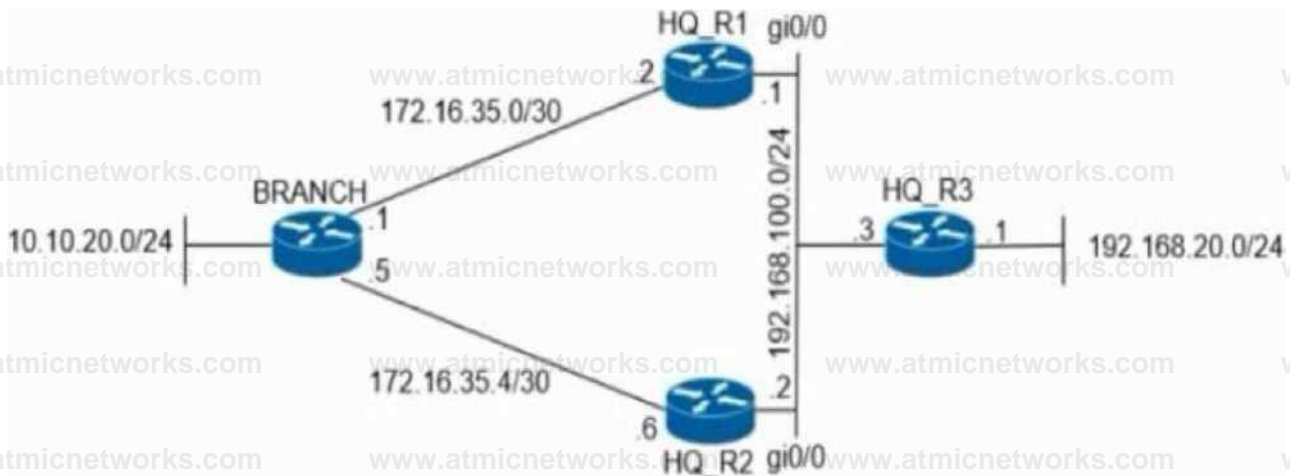
Answer: A

Explanation:

In this topology, Chicago router (Hub) will receive advertisements from Los Angeles (Spoke1) router on its tunnel interface. The problem here is that it also has a connection with New York (Spoke2) on that same tunnel interface. If we don't disable EIGRP split-horizon, then the Hub will not relay routes from Spoke1 to Spoke2 and the other way around. That is because it received those routes on interface Tunnel1 and therefore it cannot advertise back out that same interface (splithorizon rule). Therefore we must disable split-horizon on the Hub router to make sure the Spokes know about each other.

Question: 167

Refer to the exhibit.



```
BRANCH(config)* ip route 0.0.0.0 0.0.0.0 172.16.35.2 track 1
BRANCH(config)* ip route 0.0.0.0 0.0.0.0 172.16.35.6 5 |
BRANCH(config)* ip sla 1
BRANCH(config-lp-sla) i lamp-echo 172.16.35.2
BRANCH(config-ip-sla)* timeout 200
BRANCH(config-ip-sla)* frequency 5 I
BRANCH(config)* ip sla schedule 1 life forever start-time now I
BRANCH(config)* track 1 ip sla 1 reachability
```

Refer to the exhibit. An engineer has successfully set up a floating static route from the BRANCH router to the HQ network using HQ_R1 as the primary default gateway. When the g0/0 goes down on HQ_R1, the branch network cannot reach the HQ network 192.168.20.0/24. Which set of configurations resolves the issue?

- A. HQ_R3(config)# ip sla responderHQ_R3(config)# ip sla responder icmp-echo 172.16.35.1
- B. BRANCH(config)# ip sla 1BRANCH(config-ip-sla)# icmp-echo 192.168.100.2
- C. HQ_R3(config)# ip sla responderHQ_R3(config)# ip sla responder icmp-echo 172.16.35.5
- D. BRANCH(config)# ip sla 1BRANCH(config-ip-sla)# icmp-echo 192.168.100.1

Answer: D

Explanation:

Question: 168

What are two functions of MPLS Layer 3 VPNs? (Choose two.)

- A. LDP and BGP can be used for Pseudowire signaling.
- B. It is used for transparent point-to-multipoint connectivity between Ethernet links/sites.
- C. BGP is used for signaling customer VPNv4 routes between PE nodes.
- D. A packet with node segment ID is forwarded along with shortest path to destination.
- E. Customer traffic is encapsulated in a VPN label when it is forwarded in MPLS network.

Answer: C,E

Explanation:

MPLS Layer-3 VPNs provide IP connectivity among CE sites* MPLS VPNs enable full-mesh, hub- andspoke, and hybrid IP connectivity* CE sites connect to the MPLS network via IP peering across PE- CE links* MPLS Layer-3 VPNs are implemented via VRFs on PE edge nodes* VRFs providing customer routing and forwarding segmentation* BGP used for signaling customer VPN (VPNv4) routes between PE nodes* To ensure traffic separation, customer traffic is encapsulated in an additional VPN label when forwarded in MPLS network* Key applications are layer-3 business VPN services, enterprise network segmentation, and segmented layer-3 Data Center access

Reference: <https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2018/pdf/BRKMPL-1100.pdf>

Question: 169

Refer to the exhibit.

```
ip prefix-list DefaultRouteOnly seq 5 deny 0.0.0.0/0 le 32
ip prefix-list DefaultRouteOnly seq 10 permit 0.0.0.0/0

router eigrp ccnp
 address-family ipv4 unicast autonomous-system 1
 topology base
 distribute-list prefix DefaultRouteOnly out Tunnel0
```

Refer to the exhibit. The administrator configured route advertisement to a remote low resources router to use only the default route to reach any network but failed. Which action resolves this issue?

- A. Change the direction of the distribute-list command from out to in.
- B. Remove the line with the sequence number 5 from the prefix list.
- C. Remove the prefix keyword from the distribute-list command.
- D. Remove the line with the sequence number 10 from the prefix list.

Answer: B

Explanation:

Question: 170

Refer to the exhibit.

```
config t
flow record v4_r1
match ipv4 tos
match ipv4 protocol
match ipv4 source address
match ipv4 destination address
match transport source-port
match transport destination-port
collect counter bytes long
collect counter packets long
!
flow exporter EXPORTER-1
destination 172.16.10.2
transport udp 2055
exit
!
flow monitor FLOW-MONITOR-1
exporter EXPORTER-1
record v4_r1
exit
!
flow monitor v4_r1
```


Refer to the exhibit. The remote server is failing to receive the NetFlow data. Which action resolves the issue?

- A. Modify the flow transport command transport udp 2055 to move under flow monitor profile.
- B. Modify the interlace command to ip flow monitor FLOW-MONITOR-1 Input.
- C. Modify the udp port under flow exporter profile to ip transport udp 4739.
- D. Modify the flow record command record v4_r1 to move under flow exporter profile.

Answer: B

Explanation:

From the exhibit we see there are two flow monitors: the first one "FLOW-MONITOR-1" has been configured correctly but the second one "v4_r1" was left empty and interface E0/0.1 is using it. So the remote server does not receive any NetFlow data.

Question: 171

A DMVPN single hub topology is using IPsec + mGRE with OSPF. What should be configured on the hub to ensure it will be the designated router?

- A. tunnel interface of the hub with ip nhrp ospf dr
- B. OSPF priority to 0
- C. route map to set the metrics of learned routes to 110
- D. OSPF priority greater than 1

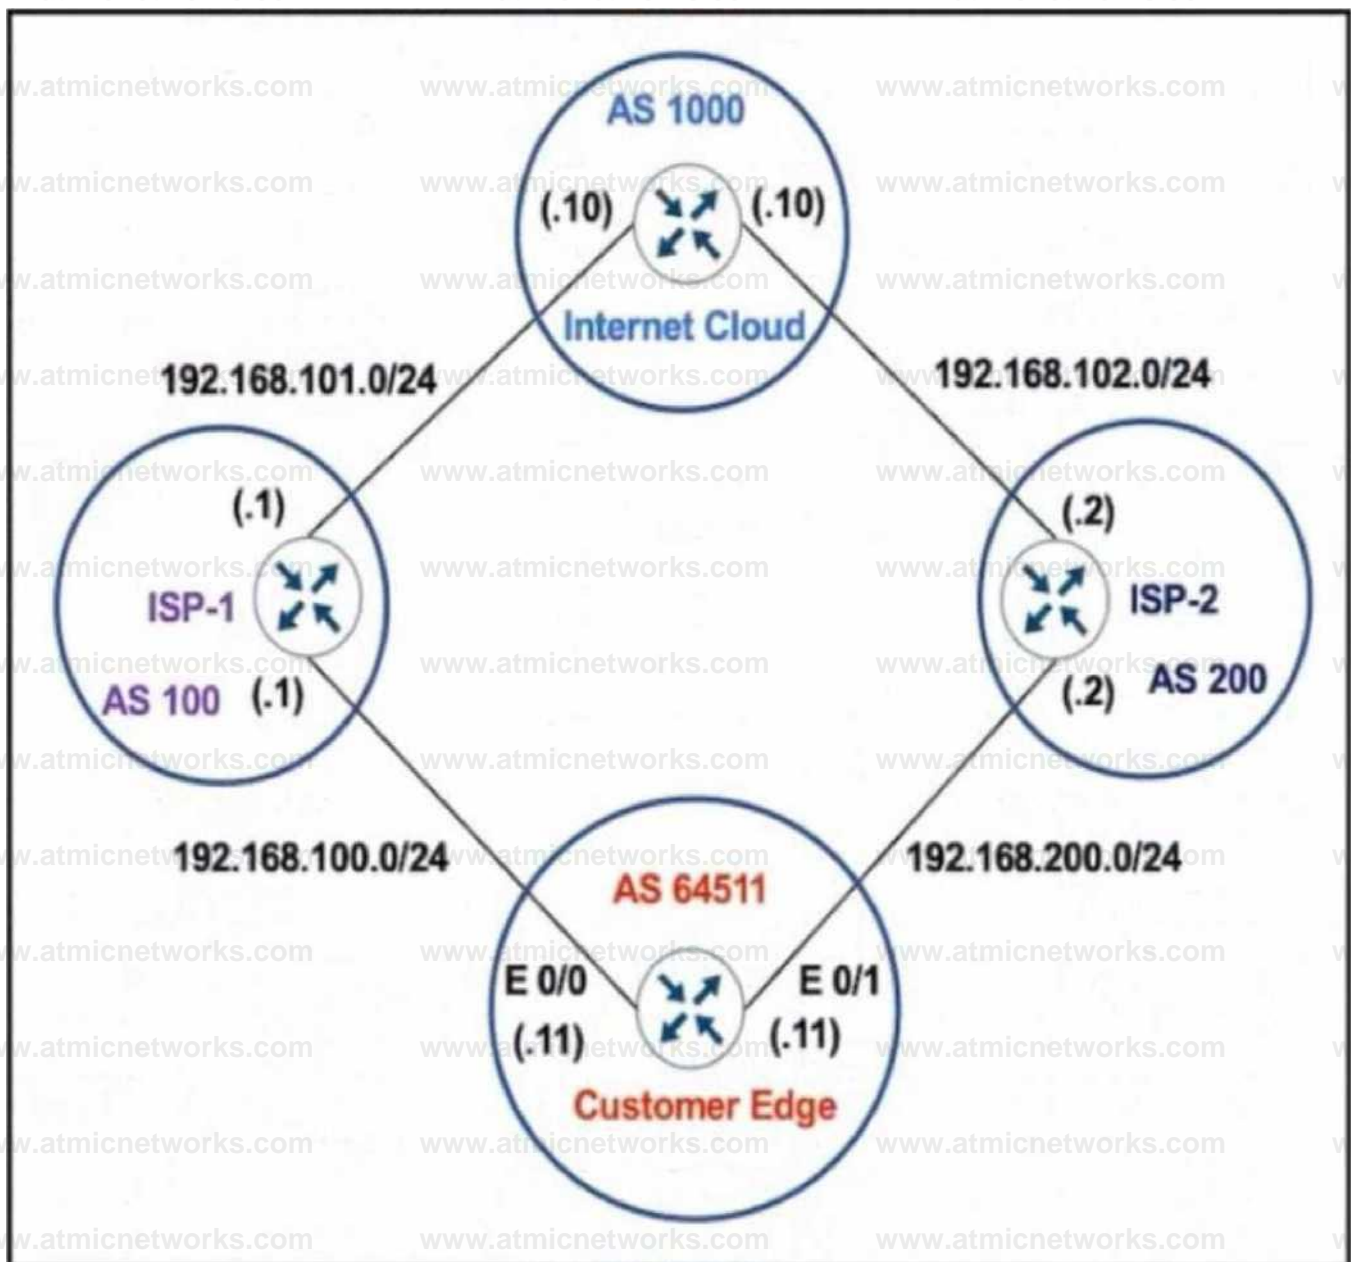
Answer: D

Explanation:

By default, the priority is 1 on all routers so we can set the OSPF priority of the hub to a value which is greater than 1 to make sure it would become the DR.

Question: 172

Refer to the exhibit.



Refer to the exhibit. The network administrator has configured the Customer Edge router (AS 64511) to send only summarized routes toward ISP-1 (AS 100) and ISP-2 (AS 200).

```
router bgp 64511
```

```
network 172.16.20.0 mask 255.255.255.0
```

```
network 172.16.21.0 mask 255.255.255.0
```

```
network 172.16.22.0 mask 255.255.255.0
```

```
network 172.16.23.0 mask 255.255.255.0
```

```
aggregate-address 172.16.20.0 255.255.252.0
```

After this configuration, ISP-1 and ISP-2 continue to receive the specific routes and the summary route. Which configuration resolves the issue?

- A. `router bgp 64511 aggregate-address 172.16.20.0 255.255.252.0 summary-only`
- B. `router bgp 64511 neighbor 192.168.100.1 summary-only neighbor 192.168.200.2 summary-only`
- C. `interface E 0/0 ip bgp suppress-map BLOCK_SPECIFIC ! interface E 0/1 ip bgp suppress-map BLOCK_SPECIFIC ! ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22 ge 24 ! route-map BLOCK_SPECIFIC permit 10 match ip address prefix-list PL_BLOCK_SPECIFIC`
- D. `ip prefix-list PL_BLOCK_SPECIFIC deny 172.16.20.0/22 ge 22 ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22 ! route-map BLOCK_SPECIFIC permit 10 match ip address prefix-list PL_BLOCK_SPECIFIC ! router bgp 64511 aggregate-address 172.16.20.0 255.255.252.0 suppress-map BLOCK_SPECIFIC`

Answer: A

Explanation:

When the `aggregate-address` command is used within BGP routing, the aggregated address is advertised, along with the more specific routes. The exception to this rule is through the use of the `summary-only` command. The “summary-only” keyword suppresses the more specific routes and announces only the summarized route.

Question: 173

What are two MPLS label characteristics? (Choose two.)

- A. The label edge router swaps labels on the received packets.
- B. Labels are imposed in packets after the Layer 3 header.
- C. LDP uses TCP for reliable delivery of information.
- D. An MPLS label is a short identifier that identifies a forwarding equivalence class.
- E. A maximum of two labels can be imposed on an MPLS packet.

Answer: C,D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html>

Question: 174

In which two ways does the IPv6 First-Hop Security Binding Table operate? (Choose two.)

- A. by IPv6 routing protocols to securely build neighborships without the need of authentication
- B. by the recovery mechanism to recover the binding table in the event of a device reboot
- C. by IPv6 HSRP to make sure neighbors are authenticated before being used as gateways
- D. by various IPv6 guard features to validate the data link layer address
- E. by storing hashed keys for IPsec tunnels for the built-in IPsec features

Answer: B,D**Explanation:****Overview of the IPv6 First-Hop Security Binding Table**

A database table of IPv6 neighbors connected to the device is created from information sources such as NDP snooping. This database, or binding table, is used by various IPv6 guard features to validate the link-layer address (LLA), the IPv4 or IPv6 address, and the prefix binding of the neighbors to prevent spoofing and redirect attacks.

IPv6 First-Hop Security Binding Table Recovery Mechanism The IPv6 first-hop security binding table recovery mechanism enables the binding table to recover in the event of a device reboot.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ipv6-fhs-bind-table.html

Question: 175

Refer to the exhibit.

```
ipv6 access-list inbound
permit tcp any any
deny ipv6 any any log
!
interface gi0/0
ipv6 traffic-filter inbound out
```

Refer to the exhibit. A network administrator configured an IPv6 access list to allow TCP return traffic only, but it is not working as expected. Which changes resolve this issue?

- A. ipv6 access-list inbound permit tcp any any deny ipv6 any any log ! interface gi0/0 ipv6 traffic-filter inbound out
- B. ipv6 access-list inbound permit tcp any any deny ipv6 any any log ! interface gi0/0 ipv6 traffic-filter inbound in
- C. ipv6 access-list inbound permit tcp any any established deny ipv6 any any log ! interface gi0/0 ipv6 traffic-filter inbound in
- D. ipv6 access-list inbound permit tcp any any established deny ipv6 any any log ! interface gi0/0 ipv6 traffic-filter inbound out

Answer: C

Explanation:

Question: 176

Refer to the exhibit.

```
Configuration output:
clock timezone PST -8
clock summer-time PDT recurring
service timestamps debug datetime
service timestamps log datetime
logging buffered 16000 debugging
ntp clock-period 17179272
ntp server 161.181.92.152

Debug output:
router#show clock
14:12:26.312 PDT Thu Apr 27 2019
router#config t
Enter configuration commands, one per line. End with CNTL/Z.
router(config)#exit
router#
Apr 27 21:12:28: %SYS-5-CONFIG_I: Configured from console by vty0
```

Refer to the exhibit. A network administrator configured NTP on a Cisco router to get synchronized time for system and logs from a unified time source. The configuration did not work as desired. Which service must be enabled to resolve the issue?

- A. Enter the service timestamps log datetime localtime global command.
- B. Enter the service timestamps log datetime synchronize global command.
- C. Enter the service timestamps log datetime console global command.
- D. Enter the service timestamps log datetime clock-period global command.

Answer: A

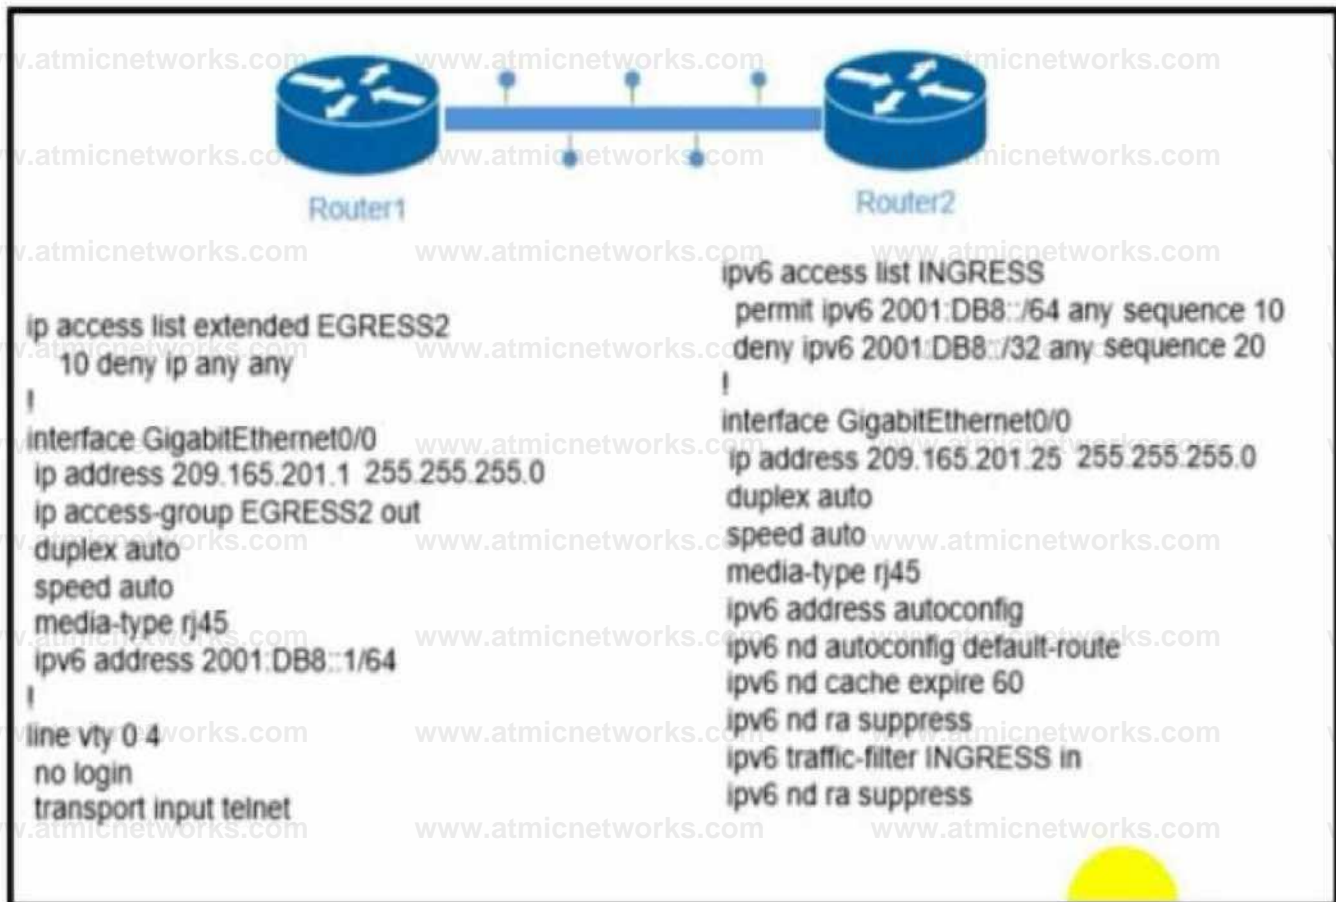
Explanation:

If a router is configured to get the time from a Network Time Protocol (NTP) server, the times in the router's log entries may be different from the time on the systemclock if the [localtime] option is not in the service timestamps log command. To solve this issue, add the [localtime] option to the service timestamps log command. The times should now be synchronized between the system clock and the log message timestamps.

Reference: <https://community.cisco.com/t5/networking-documents/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258>

Question: 177

Refer to the exhibit.



Refer to the exhibit. The engineer configured and connected Router2 to Router1. The link came up but could not establish a Telnet connection to Router1 IPv6 address of 2001:DB8::1. Which configuration allows Router2 to establish a Telnet connection to Router1?

- A. ipv6 unicast-routing
- B. permit ICMPv6 on access list INGRESS for Router2 to obtain IPv6 address
- C. permit ip any any on access list EGRESS2 on Router1
- D. IPv6 address on GigabitEthernet0/0

Answer: D**Explanation:**

R1

interface Ethernet0/0

ip address 209.165.201.1 255.255.255.0

ip access-group EGRESS2 out

ipv6 address 2001:DB8::1/64

end

R2

interface Ethernet0/0

ip address 209.165.201.25 255.255.255.0

ipv6 address 2001:DB8::2/64

ipv6 address autoconfig

ipv6 nd autoconfig default-route

ipv6 nd cache expire 60

ipv6 nd ra suppress

ipv6 traffic-filter INGRESS in

end

IOU_Router2#telnet 2001:DB8::1

Trying 2001:DB8::1 ... Open

IOU Router1>

Question: 178

Refer to the exhibit.

Filtered

00:00:46: %LINK-3-UPDOWN: Interface Port-channel 1, changed state to up

00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up

00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up

Desired

00:00:46: %LINK-3-UPDOWN: Interface Port-Channel 1, changed state to up

00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up

00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up

00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down

00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed state to down

00:00:48: %SYS-5-CONFIG_I: Configured from console by vty2

Refer to the exhibits. An engineer filtered messages based on severity to minimize log messages.

After applying the filter, the engineer noticed that it filtered required messages as well. Which action must the engineer take to resolve the issue?

- A. Configure syslog level 2.
- B. Configure syslog level 3.
- C. Configure syslog level 4.
- D. Configure syslog level 5.

Answer: D

Explanation:

Question: 179

An engineer configured policy-based routing for a destination IP address that does not exist in the routing table. How is the packet treated through the policy for configuring the set ip default next-hop command?

- A. Packets are not forwarded to the specific next hop.
- B. Packets are forwarded based on the routing table.
- C. Packets are forwarded based on a static route.

D. Packets are forwarded to the specific next hop.

Answer: D

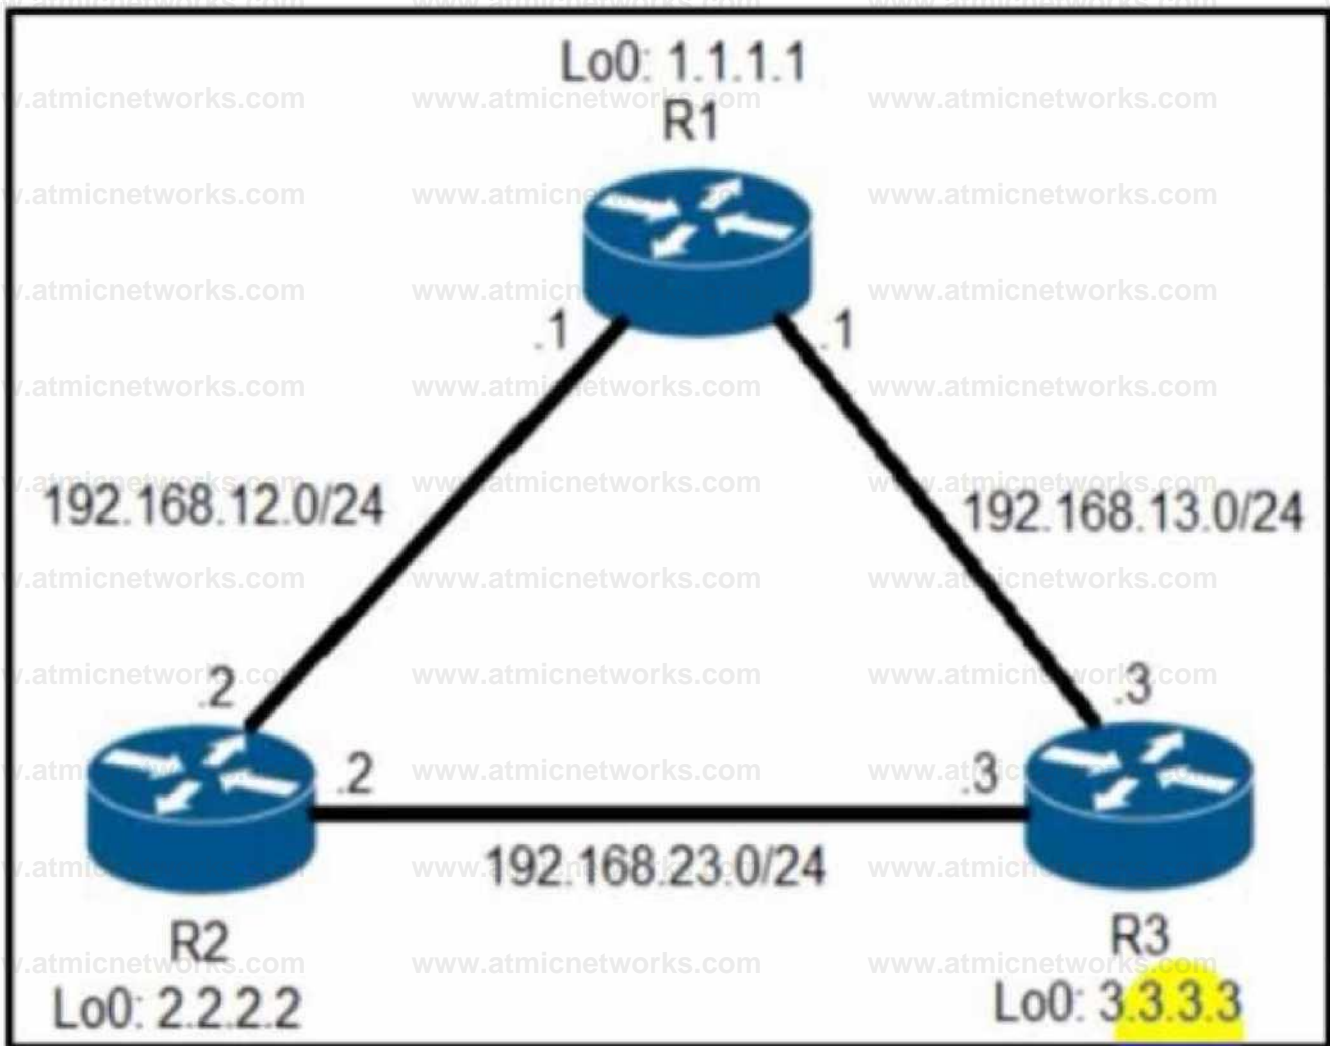
Explanation:

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...+ if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.+ if the destination IP address does not exist, the command policy routes the packet by sending it to the specified next hop.

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html>

Question: 180

Refer to the exhibit.



```
R2#show ip protocols | include eigrpfMaxxnum
```

```
Routing Protocol is "eigrp 1"
```

```
Maximum path 4
```

```
Maximum hopcount 100
```

```
Maximum metric variance 1
```

```
R2*show ip eigrp topology 192 168 13 0/24
```

```
EIGRP IPv4 Topology Entry for AS(1HD(2 2 2 2) for 192 168 13 0/24 State
```

```
is Passive Query origm flag is 1, 1 Successor(s) ED is 1075200 Descriptor
```

```
Blocks 192 168 23 3 (FastEthernetE1) from 192 168 23 3 Send flag is 0x0
```

```
Composite metric is (1075200 2 81600) route is Internal Vector metric
```

```
Minimum bandwidth is 2500 Kbit
```

```
Total delay ts 2000 microseconds
```

```
Reliability is 255255
```

```
Load is 255255
```

```
Minimum MTU is 1500
```

```
Hop count is 1
```

```
Originating router is 3 3 3 3
```

```
192 168 12 1 (FastEthernetE2) from 192 168 12 1 Send flag is 0x0
```

```
Composite metric is (2611200 281600) route is Internal Vector metric
```

```
Minimum bandwidth is 1000 Kbit
```

```
Total delay is 2000 microseconds
```

```
Reliability is 255255
```

```
Load is 1255
```

```
Minimum MTU is 1500
```

```
Hop count is 1
```

```
Originating router is 1 1 1 1
```

```
R2#show ip route 192 168 13 0
```

```
Routing entry for 192 168 13 0/24
```

```
Known via "eigrp" T distance 90 metric 1075200 type internal Redistributing  
via eigrp 1
```

```
Last update from 192 168 23 3 on FaMEthernetE1 00 00 57 ago Routing
```

```
Descriptor Blocks
```

```
• 192 168 23 3 to 192 168 23 3 0000 57 ago via 1 as 192 168 12 1 Route
```

```
metric is 1075200. Hattr share count is 1
```

Refer to the exhibit. R2 has two paths to reach 192.168.13.0/24. but traffic is sent only through R3.

Which action allows traffic to use both paths?

- A. Configure the bandwidth 2000 command under interface FastEthernet0/0 on R2.
- B. Configure the variance 4 command under the EIGRP process on R2.
- C. Configure the delay 1 command under interface FastEthernet0/0 on R2.
- D. Configure the variance 2 command under the EIGRP process on R2

Answer: B

Explanation:

From the output of the “show ip eigrp topology ...” command, we notice network 192.168.13.0/24 was learned via two routes:+ From 192.168.23.3 (R3) with FD = 1075200 and AD = 281600+ From 192.168.12.1 (R1) with FD = 2611200 and AD = 281600

From the output of the “show ip route ...” command, we learned that the best (and chosen) path is via 192.168.23.3 (R3).

To use both paths (called unequal cost load balancing) with EIGRP, the second path via R1 must satisfy the feasibility condition. The feasibility condition states that, the Advertised Distance (AD) of a route must be lower than the feasible distance of the current successor route.

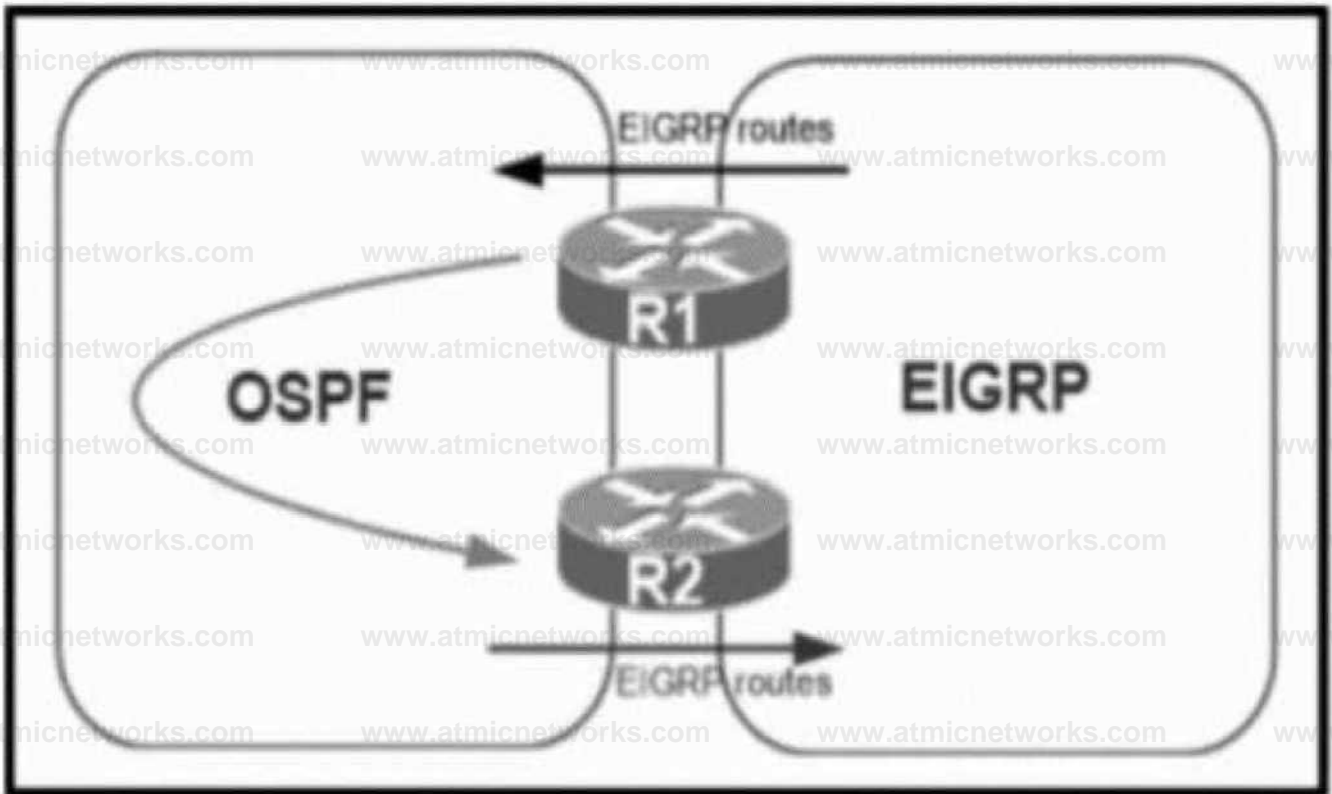
In this case, the second path satisfies the feasible condition as its AD (281600) is smaller than the FD (1075200) of the best path. Therefore we can configure loadbalancing with “variance” command.

In other words, EIGRP will install all paths with metric < variance * best_metric into the local routing table, provided that it meets the feasibility condition to prevent routing loop. Therefore we can calculate the variance > metric / best_metric = 2611200 / 1075200 = 2.4.

So with a variance greater than 2 (and must be an integer), we can load balance traffic to network 192.168.13.0/24.

Question: 181

Refer to the exhibit.



Refer to the exhibit. A network administrator configured mutual redistribution on R1 and R2 routers, which caused instability in the network. Which action resolves the issue?

- A. Set a tag in the route map when redistributing EIGRP into OSPF on R1. and match the same tag on R2 to deny when redistributing OSPF into EIGRP.
- B. Set a tag in the route map when redistributing EIGRP into OSPF on R1. and match the same tag on R2 to allow when redistributing OSPF into EIGRP.
- C. Advertise summary routes of EIGRP to OSPF and deny specific EIGRP routes when redistributing into OSPF.
- D. Apply a prefix list of EIGRP network routes in OSPF domain on R1 to propagate back into the EIGRP routing domain.

Answer: A

Explanation:

When doing mutual redistribution at multiple points (between OSPF and EIGRP on R1 & R2), we may create routing loops so we should use route-map to prevent redistributed routes from redistributing again into the original domain.

In the below example, the route-map "SET-TAG" is used to prevent any routes that have been redistributed into EIGRP from redistributed again into OSPF domain by tagging these routes with tag 1:

```
R3
route-map SET-TAG permit 10
set tag 1
```

These routes are prevented from redistributed again by route-map FILTER_TAG by denying any routes with tag 1 set:

```
R4
route-map FILTER-TAG deny 10 match tag 1
```

Question: 182

Refer to the exhibit.


```
R1
interface Loopback0
 ip address 172.16.1.1 255.255.255.255
interface FastEthernet0/0
 ip address 192.168.12.1 255.255.255.0
router eigrp 100
 no auto-summary
 network 192.168.12.0
 network 172.16.0.0
 neighbor 192.168.12.2 FastEthernet0/0

R2
interface Loopback0
 ip address 172.16.2.2 255.255.255.255
interface FastEthernet0/0
 ip address 192.168.12.2 255.255.255.0
router eigrp 100
 network 192.168.12.0
 network 172.16.0.0
 neighbor 192.168.12.1 FastEthernet0/0
 passive-interface FastEthernet0/0
```

Refer to the exhibit. R1 and R2 cannot establish an EIGRP adjacency. Which action establishes EIGRP adjacency?

- A. Remove the current autonomous system number on one of the routers and change to a different value.
- B. Remove the passive-interface command from the R2 configuration so that it matches the R1 configuration.

- C. Add the no auto-summary command to the R2 configuration so that it matches the R1 configuration.
- D. Add the passive-interface command to the R1 configuration so that it matches the R2 configuration.

Answer: B

Explanation:

Question: 183

When configuring Control Plane Policing on a router to protect it from malicious traffic, an engineer observes that the configured routing protocols start flapping on

that device. Which action in the Control Plane Policy prevents this problem in a production environment while achieving the security objective?

- A. Set the conform-action and exceed-action to transmit initially to test the ACLs and transmit rates and apply the Control Plane Policy in the output direction
- B. Set the conform-action and exceed-action to transmit initially to test the ACLs and transmit rates and apply the Control Plane Policy in the input direction
- C. Set the conform-action to transmit and exceed-action to drop to test the ACLs and transmit rates and apply the Control Plane Policy in the input direction
- D. Set the conform-action to transmit and exceed-action to drop to test the ACLs and transmit rates and apply the Control Plane Policy in the output direction

Answer: B

Explanation:

Question: 184

Refer to Exhibit.

```
Ipv6 unicast-routing
!
Router ospfv3 4
  Router-id 192.168.1.1
  !
  Interface E 0/0
    Ipv6 enable
    Ip address 10.1.1.1 255.255.255.0
    Ospf3 4 area 0 ipv4
    No shut
  !
  Interface Loopback0
    Ipv6 enable
    Ipv4 172.16.1.1 255.255.255.0
    Ospf3 4 area 0 ipv4
```

The network administrator configured the branch router for IPv6 on the E0/0 interface. The neighboring router is fully configured to meet requirements, but the neighbor relationship is not coming up. Which action fixes the problem on the branch router to bring the IPv6 neighbors up?

- A. Enable the IPv4 address family under the router ospfv3 4 process by using the address-family ipv4 unicast command
- B. Disable IPv6 on the E0/0 interface using the no ipv6 enable command
- C. Enable the IPv4 address family under the E0/0 interface by using the address-family ipv4 unicast command
- D. Disable OSPF for IPv4 using the no ospfv3 4 area 0 ipv4 command under the E0/0 interface

Answer: A

Explanation:

Explanation

Once again, Cisco changed the IOS configuration commands required for OSPFv3 configuration. The new OSPFv3 configuration uses the “ospfv3” keyword instead of the earlier “ipv6 router ospf” routing process command and “ipv6 ospf” interface commands.

The Open Shortest Path First version 3 (OSPFv3) address families feature enables both IPv4 and IPv6 unicast traffic to be supported. With this feature, users may

have two processes per interface, but only one process per address family (AF).

Question: 185

An engineer is troubleshooting on the console session of a router and turns on multiple debug commands. The console screen is filled with scrolling debug messages that none of the commands can be verified if entered correctly or display any output. Which action allows the engineer to see entered console commands while still continuing the analysis of the debug messages?

- A. Configure the logging synchronous command
- B. Configure the no logging console debugging command globally
- C. Configure the logging synchronous level all command
- D. Configure the term no mon command globally

Answer: A

Explanation:

Let's see how the “logging synchronous” command affect the typing command:

Without this command, a message may pop up and you may not know what you typed if that message is too long. When trying to erase (backspace) your command, you realize you are erasing the message instead.

XVboaZSll-llconf t Enter configuration contnands, one per line. End with CNTL/2. NVboa2811-

```
l(config)#*Z ^Vbos2811-im ban 18 16:38:02: ISY3-5-CCNFIG_I: Configured from console by admin on
vtyO (10.0.1.111)
```

With this command enabled, when a message pops up you will be put to a new line with your

typing command which is very

```
NVbos2811-l(config)fline con 0
MVbos28ii-i(config-line)llogging synch
NVbos281i-l(config-line)fline vty 0 4
NVbos2811-l (config-line)fllogging synchr
NVbos2811-1 (conf ig-line) flloggmg synchronous
NVbos2811-l(conf ig-lme) fl* 2
NVbos2811-lflsh ip
Jan 18 16:39:33: *SYS-5-C0NFIG I: Configured from console by admin
NVbos2811-lflsh ip |
```

Question: 186

An engineer must configure a Cisco router to initiate secure connections from the router to other devices in the network but kept failing. Which two actions resolve the issue? (Choose two.)

- A. Configure a source port for the SSH connection to initiate
- B. Configure a TACACS+ server and enable it
- C. Configure transport input ssh command on the console
- D. Configure a domain name
- E. Configure a crypto key to be generated

Answer: D,E

Explanation:

Follow these guidelines when configuring the switch as an SSH server or SSH client:

+ An RSA key pair generated by a SSHv1 server can be used by an SSHv2 server, and the reverse.+ If the SSH server is

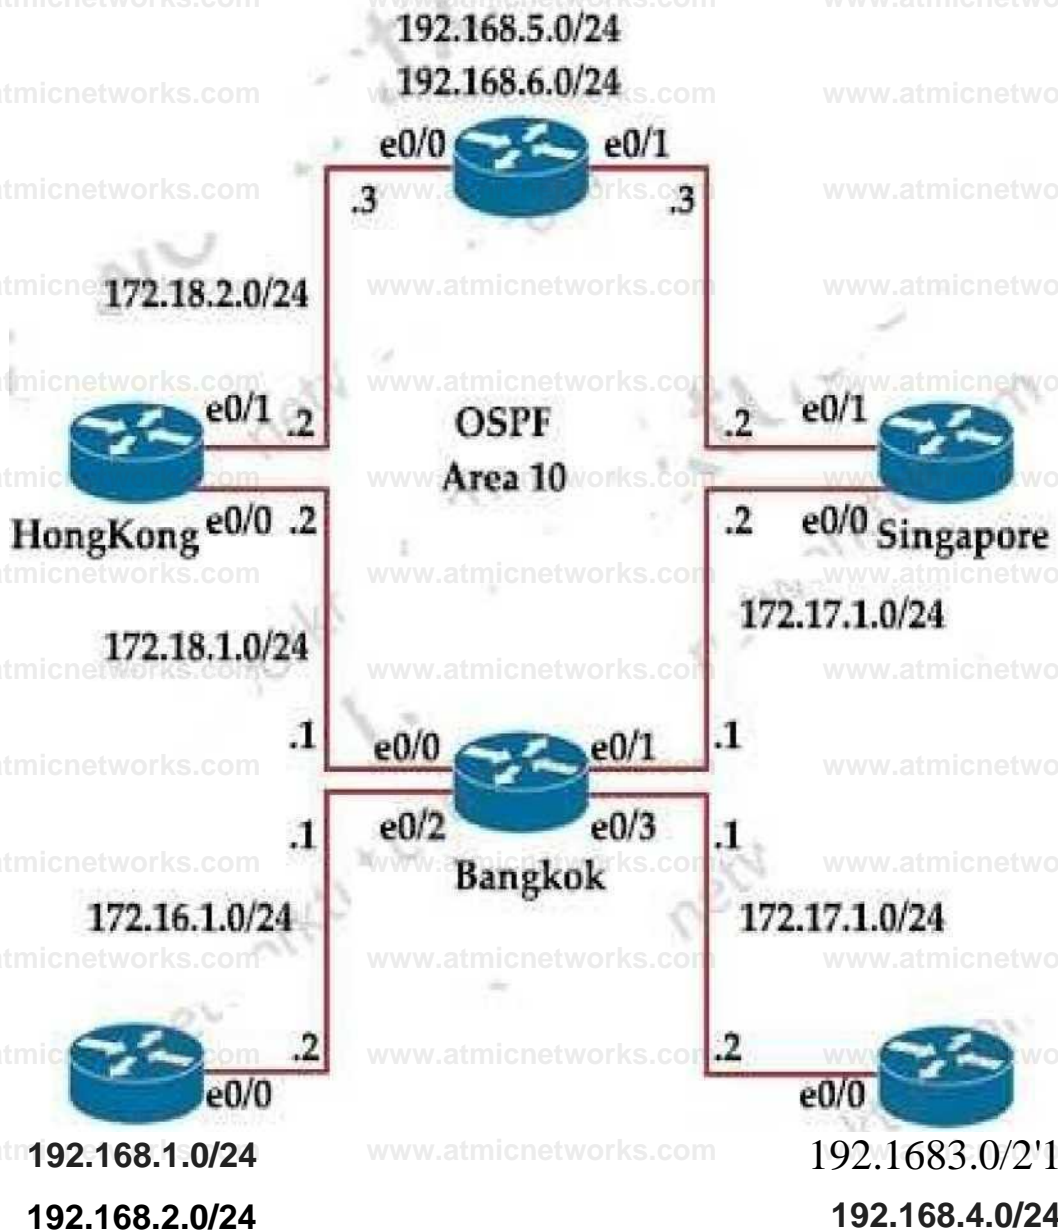
running on a stack master and the stack master fails, the new stack master uses the RSA key pair generated by the previous stack master

+ If you get CLI error messages after entering the crypto key generate rsa global configuration command, an RSA key pair has not been generated. Reconfigure the hostname and domain, and then enter the crypto key generate rsa command.+ When generating the RSA key pair, the message No host name specified might appear. If it does, you must configure a hostname by using the hostname global configuration command.+ When generating the RSA key pair, the message No domain specified might appear. If it does, you must configure an IP domain name by using the ip domain- name global configuration command.+ When configuring the local authentication and authorization authentication method, make sure that AAA is disabled on the console.

Reference:https://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3.2_0_s/e/multibook/configuration_guide/b_consolidated_config_guide_3850_chapter_01100

Question: 187

Exhibit:



Bangkok is using ECMP to reach to the 192.168.5.0/24 network. The administrator must configure Bangkok in such a way that Telnet traffic from 192.168.3.0/24 and 192.168.4.0/24 networks uses the HongKong router as the preferred router. Which set of configurations accomplishes this task?

- A. `access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255`
`access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255`
`route-map PBR1 permit 10`
`match ip address 101`
`set ip next-hop 172.18.1.2`
`interface Ethernet0/3`
`ip policy route-map PBR1`
- B. `access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23`
`access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23`
`route-map PBR1 permit 10`
`match ip address 101`
`set ip next-hop 172.18.1.2`
`interface Ethernet0/1`
`ip policy route-map PBR1`

- C. access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23
access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23!
route-map PBR1 permit 10 match ip address 101 set ip next-hop 172.18.1.2!
interface Ethernet0/3 ip policy route-map PBR1
- D. access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255
access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255!
route-map PBR1 permit 10 match ip address 101 set ip next-hop 172.18.1.2!
interface Ethernet0/1 ip policy route-map PBR1

Answer: C

Explanation:

We need to use Policy Based Routing (PBR) here on Bangkok router to match the traffic from 192.168.3.0/24 & 192.168.4.0/24 and "set ip next-hop" to HongKong router(172.18.1.2 in this case).

Note: Please notice that we have to apply the PBR on incoming interface e0/3 to receive traffic from 192.168.3.0/24 and 192.168.4.0/24.

Question: 188

Exhibit:

```
11:27:07.532: AAA/BIND (00000055): Bind 1/  
11:27:07.532: AAA/AU THEN/LOGIN (00000055): Pick method list 'default'  
11:27:07.532: TPLUS: Queuing AAA Authentication request 85 for processing  
11:27:07.537 TPLUS (00000055) login timer started 1020 sec timeout  
11:27:07.532:TPLUS: processing authentication start request id 8b  
11:27:07.5??-TPLUS: Authentication start packet created for 85(1  
11:27:07:53 2; TPLUS: Using server 10.106.60.182  
11:27:07.532 TPLUS (00000055)/0/NB_WAIT/225FE2DC: Started 5 sec timeout  
11:27:07,532: TPLUS (00000055)/0/NB WAIT: socket event 2  
11:27:07.532: TPLUS (000000551/0/NB WAIT: wrote entire 38 bytes request  
11:27:07.532: TPLUS (000000551/0/READ; socket event 1  
11:27:07.532 TPLUS (000000551/0/READ: Would block while reading  
11:27:07.532: TPLUS (00000055)/0/READ: socket event 1
```


1177:07.532: TPLUS (00000055)/0/RF AD- react entire 12 header bytes {expect 6 bytes data)
1377:07.532: TPLUS (000000551/0/READ: socket event 1
1177:07.532 TPLUS (0000005 5)/n/READ: read entire 18 bytes response
1177:07.532: TPLUS (00000055)/0/225FE2DC: Processing the reply packet
11:27:07.532: TPLUS: received bad AUTHEN packet: length = 6, expected 43974
1177:07.532; TPLUS: Invalid AUTHEN packet (check keys).

Which action resolves the authentication problem?

- A. Configure the user name on the TACACS+ server
- B. Configure the UDP port 1812 to be allowed on the TACACS+ server
- C. Configure the TCP port 49 to be reachable by the router
- D. Configure the same password between the TACACS+ server and router.

Answer: D

Explanation:

Explanation

From the last line of the output, we notice that the result was "Invalid AUTHEN packet". Therefore something went wrong with the username or password.

Reference: <https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html>

Question: 189

Refer to the exhibit.

LoO:
192.168.1.55
255.255.255.128



Admin PC
IP address
192.155.1.20D
255.255.255.128

```
aaa new-model
```

```
!
```

```
aaa authentication login default line enable
```

```
aaa authorization commands 15 default local
```

```
aaa authorization network default local
```

```
username admin privilege 15 password cisco 1231
```

```
ip ssh version 2
```

```
!
```

```
access-list 101 permit tcp 192.168.1.0 0.0.0.255 any eq 22
```

```
access-list 101 permit tcp 192.168.5.0 0.0.0.255 any range 22 smtp i
```

```
line vty 0 4 access-class 101 in password cisco transport input all J
```

```
line vty 5 15 access-class 101 in password cisco transport input all
```

The administrator successfully logs into R1 but cannot access privileged mode commands. What should be configured to resolve the issue?

- A. aaa authorization reverse-access
- B. secret cisco123! at the end of the username command instead of password cisco123!
- C. matching password on vty lines as cisco123!
- D. enable secret or enable password commands to enter into privileged mode

Answer:

D

Explanation:

Question:

190

DRAG DROP

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

label edge router

allows an LSR to remove the label
before forwarding the packet

label switch router

accepts unlabeled packets and imposes labels

forwarding equivalence class

group of packets that are forwarded
in the same manner

penultimate hop popping

receives labeled packets and swaps labels

Answer:**Explanation:**

- + allows an LSR to remove the label before forwarding the packet: penultimate hop popping
- + accepts unlabeled packets and imposes labels: label edge router
- + group of packets that are forwarded in the same manner: forwarding equivalence class
- + receives labeled packets and swaps labels: label switch router

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: 191

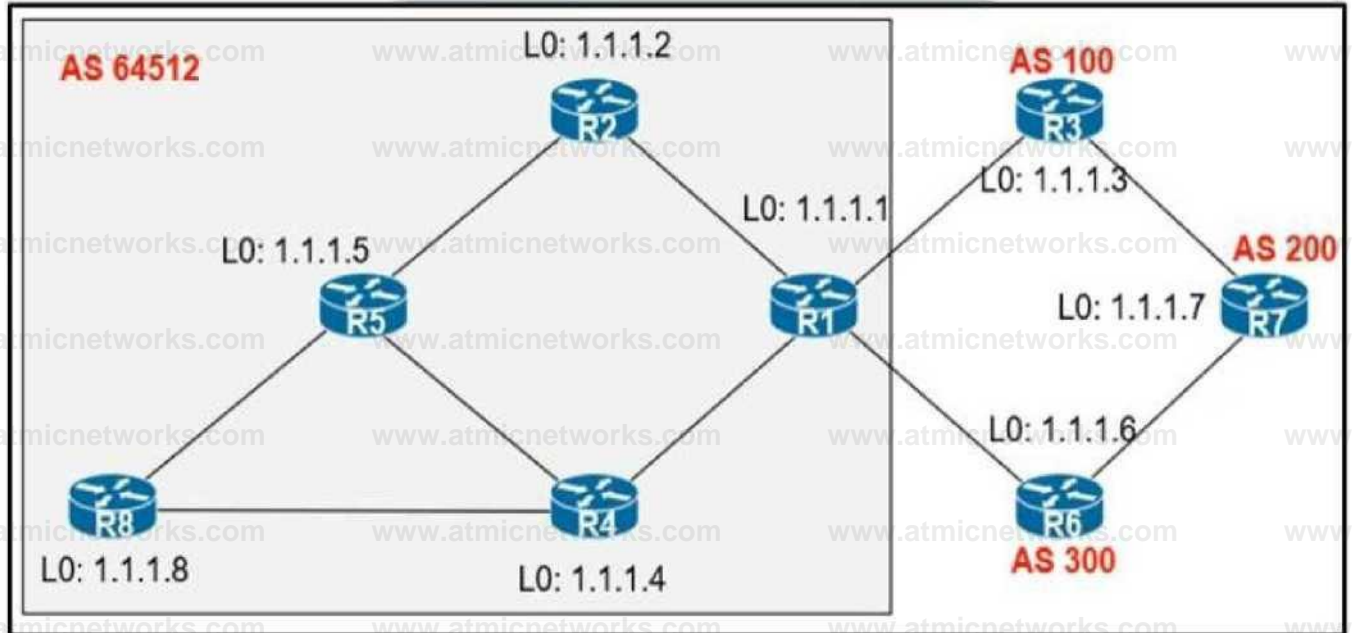
Which two protocols work in the control plane of P routers across the MPLS cloud? (choose two)

- A. LSP
- B. RSVP
- C. ECMP
- D. LDP
- E. MPLS OAM

Answer: B,D**Explanation:**

Question: 192

Exhibit:



An engineer configured R2 and R5 as route reflectors and noticed that not all routes are sent to R1 to advertise to the eBGP peers. Which iBGP routers must be configured as route reflectors to advertise all routes to restore reachability across all networks?

- A. R1 and R4
- B. R1 and R5
- C. R4 and R5
- D. R2 and R5

Answer: C

Explanation:

When R2 & R5 are route reflectors (RRs), routes from R4 & R8 are advertised to R5 and R5 advertises to R2. But R2

would drop them as R2 is also a RR. Therefore some routes are missing on R1 to advertise to eBGP peers.

Good reference:


<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2015/pdf/TECRST-2310.pdf>

Route reflectors (RR) must be fully iBGP meshed so we cannot configure RR on both R1 and R5.

We should choose routers at the center of the topology RRs, in this case R4 & R5.

Question: 193

Refer to exhibit.



Routing protocols are mutually redistributed on R3 and R1. Users report intermittent connectivity to services hosted on the 10.1.1.0/24 prefix. Significant routing update changes are noticed on R3 when the `show ip route profile` command is run. How must the services be stabilized?

- A. The issue with using BGP must be resolved by using another protocol and redistributing it into EIGRP on R3
- B. The routing loop must be fixed by reducing the admin distance of iBGP from 200 to 100 on R3
- C. The routing loop must be fixed by reducing the admin distance of OSPF from 110 to 80 on R3
- D. The issue with using iBGP must be fixed by running eBGP between R3 and R4

Answer: B**Explanation:**

After redistribution, R3 learns about network 10.1.1.0/24 via two paths:
+ Internal BGP (iBGP): advertised from R4 with AD of 200 (and metric of 0)
+ OSPF: advertised from R1 with AD of 110 (O E2) (and metric of 20)
Therefore R3 will choose the path with the lower AD via OSPF

But this is a looped path which is received from R3 -> R2 -> R1 -> R3. So when the advertised route from R4 is expired, the looped path is also expired soon and R3 will reinstall the main path from R4. This is the cause of intermittent connectivity. In order to solve this issue, we can lower the AD of iBGP to a value which is lower than 110 so that it is preferred over OSPF-advertised route.

Question: 194

Refer to Exhibit.



A network administrator added one router in the Cisco DNA Center and checked its discovery and health from the Network Health Dashboard. The network administrator observed that the router is still showing up as unmonitored. What must be configured on the router to mount it in the Cisco DNA Center?

- A. Configure router with NetFlow data
- B. Configure router with the telemetry data
- C. Configure router with routing to reach Cisco DNA Center
- D. Configure router with SNMPv2c or SNMPv3 traps

Answer: B

Explanation:

Unmonitored: Unmonitored devices are devices for which Assurance did not receive any

telemetry data during the specified time range.

Question: 195

Exhibit:

The screenshot shows the Cisco DNA Assurance interface. The main heading is "Excessive time lag between Cisco DNA Center and WLC *WLC-5520*". The status is "Open". The description reads: "The time on Cisco DNA Center and WLC *WLC-5520* has drifted too far apart. The drift between the two devices is *65.8 minutes*. Cisco DNA Center cannot process the wireless client data accurately if the time difference is more than 10 minutes." Under "Suggested Actions (3)", the actions are: 1. If NTP is enabled, check whether the NTP servers are reachable from Cisco DNA Center and the WLC. 2. If NTP servers are not configured, configure the NTP servers on Cisco DNA Center and WLC *WLC-5520*. 3. If NTP servers are not deployed, manually reset the time on Cisco DNA Center or WLC *WLC-5520* so that the time is synchronized.

NTP is configured across the network infrastructure and Cisco DNA Center. An NTP issue was reported on the Cisco DNA Center at 17:15. Which action resolves the issue?

- A. Check and resolve reachability between the WLC and the NTP server
- B. Reset the NTP server to resolve any synchronization issues for all devices
- C. Check and resolve reachability between Cisco DNA Center and the NTP server
- D. Check and configure NTP on the WLC and synchronize with Cisco DNA Center

Answer: D**Explanation:**

Excessive time lag between Cisco DNA Center and device: The time difference between Cisco DNA Center and the device IP Address has drifted too far apart. CiscoDNA Center cannot process the device data accurately if the time difference is more than 3 minutes.

Reference: [https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-2-](https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-2-10/b_cisco_dna_assurance_1_2_10_ug/b_cisco_dna_assurance_1_2_10_ug_chapter_01101.html)

[10/b_cisco_dna_assurance_1_2_10_ug/b_cisco_dna_assurance_1_2_10_ug_chapter_01101.html](https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-2-10/b_cisco_dna_assurance_1_2_10_ug/b_cisco_dna_assurance_1_2_10_ug_chapter_01101.html)

Question: 196

Refer to Exhibit.

```
Jan 9 15 29 29 713 DHCP SNOOPING process new DHCP packet, message type DHCPINFORM, input interface Po2, MAC da ffff ffff ffff. DHCP yiaddr 0 0 0 0, DHCP siaddr 0 0 0 0, DHCP giaddr 0000
```

```
Jan 9 15 29 29 713 DHCP_SNOOPING_SW bridge packet get invalid mat entry FFFF FFFF FFFF. packet is flooded to ingress VLAN (1)
```

```
Jan 9 152929 722 DHCP_SNOOPING_SW bridge packet send packet to cpu port Vlan1
```

```
Jan 9 152931 509 DHCP Snooping(hlrm_set_rfjinput) Setting ifjinput to Po2 for pak Was V11
```

```
Jan 9 152931 509 DHCP Snooping(hlrmsetifjinput) Setting ifjinput to V11 for pak Was Po2
```

```
Jan 9 15 29 31 509 DHCP Snooping(hlrm_set_ifInput): Setting if_input to Po2 for pakWas V11 Jan 9
```

```
15 29 31 517 DHCP_SNOOPING received new DHCP packet from input interface (Port-channel2)
```

A network administrator enables DHCP snooping on the Cisco Catalyst 3750-X switch and configures the uplink port (Port-channel2) as a trusted port. Clients are not receiving an IP address, but when DHCP snooping is disabled, clients start receiving IP addresses. Which global command resolves the issue?

- A. No ip dhcp snooping information option
- B. ip dhcp snooping
- C. ip dhcp relay information trust portchannel2
- D. ip dhcp snooping trust

Answer: A

Explanation:

Question: 197

Which configuration feature should be used to block rogue router advertisements instead of using the IPv6 Router Advertisement Guard feature?

- A. VACL blocking broadcast frames from nonauthorized hosts
- B. PVLANS with promiscuous ports associated to route advertisements and isolated ports for nodes
- C. PVLANS with community ports associated to route advertisements and isolated ports for nodes
- D. IPv4 ACL blocking route advertisements from nonauthorized hosts

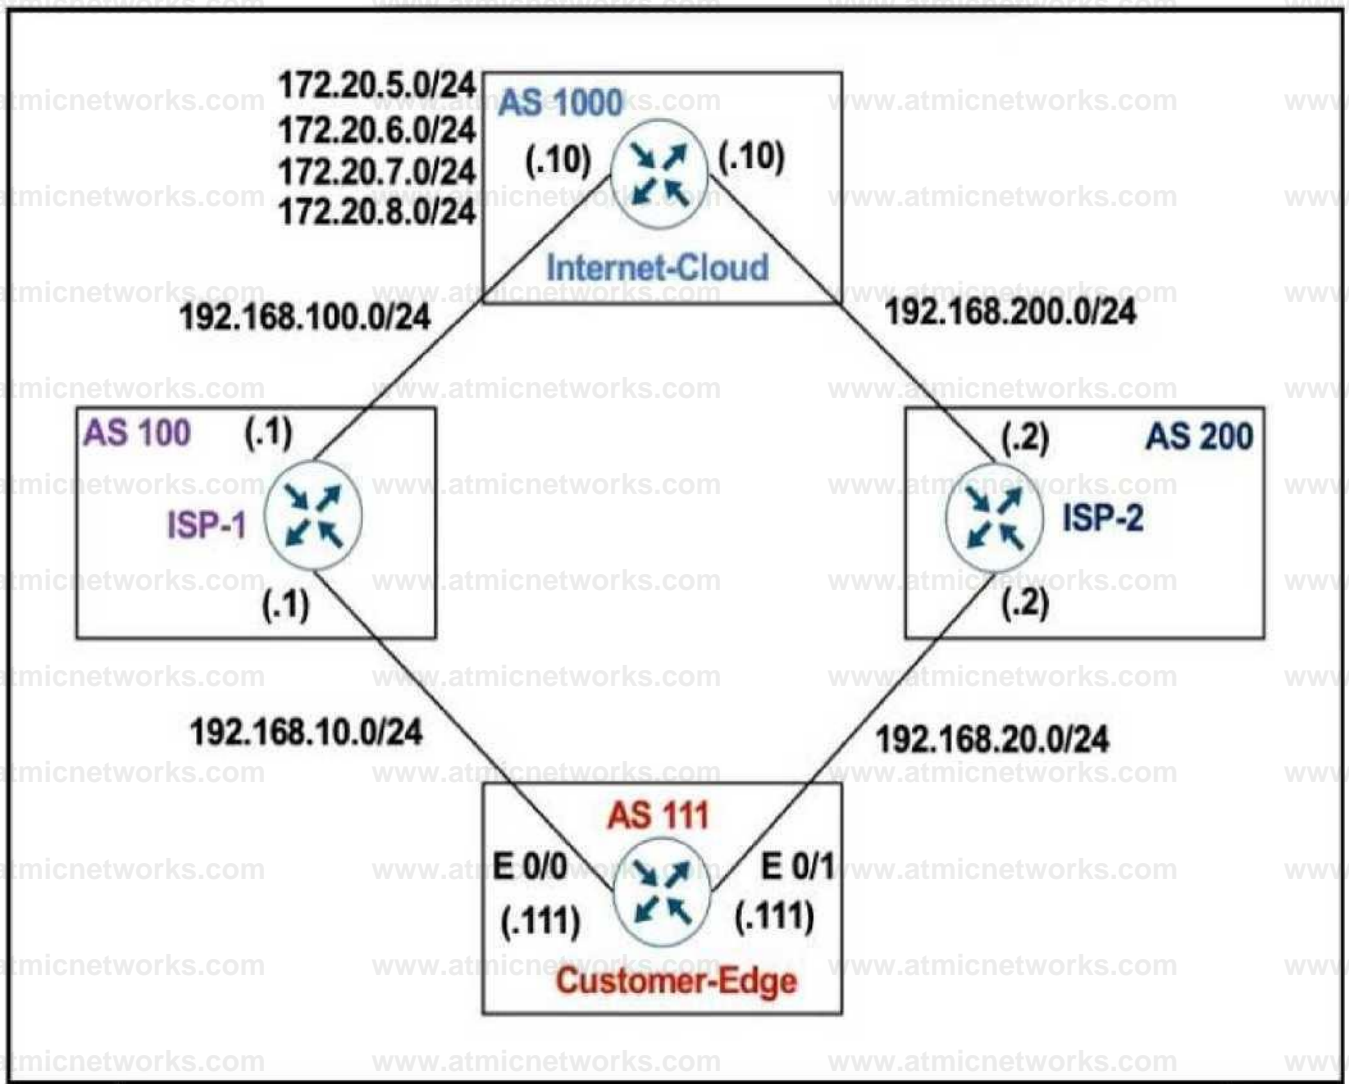
Answer: B

Explanation:

The IPv6 Router Advertisement Guard feature provides support for allowing the network administrator to block or reject unwanted or rogue router advertisement guard messages that arrive at the network device platform. Router Advertisements are used by devices to announce themselves on the link. The IPv6 Router Advertisement Guard feature analyzes these router advertisements and filters out router advertisements that are sent by unauthorized devices.

Certain switch platforms can already implement some level of rogue RA filtering by the administrator configuring Access Control Lists (ACLs) that block RA ICMP messages that might be inbound on "user" ports.

Reference: <https://datatracker.ietf.org/doc/html/rfc6104>



Customer-Edge

```
ip prefix-list PLIST1 permit 172.20.5.0/24
route-map SETLP permit 10
match ip address prefix-list PLIST1
set local-preference 90
router bgp 111
neighbor 192.168.10.1 remote-as 100
neighbor 192.168.10.1 route-map SETLP in
neighbor 192.168.20.2 remote-as 200
```

AS 111 wanted to use AS 200 as the preferred path for 172.20.5.0/24 and AS 100 as the backup. After the configuration, AS 100 is not used for any other routes. Which configuration resolves the issue?

- A. route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 99 route-map SETLP permit 20
- B. route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 110 route-map SETLP permit 20
- C. router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.10.1 route-map SETLP out
- D. router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.20.2 route-map SETLP in

Answer: A

Explanation:

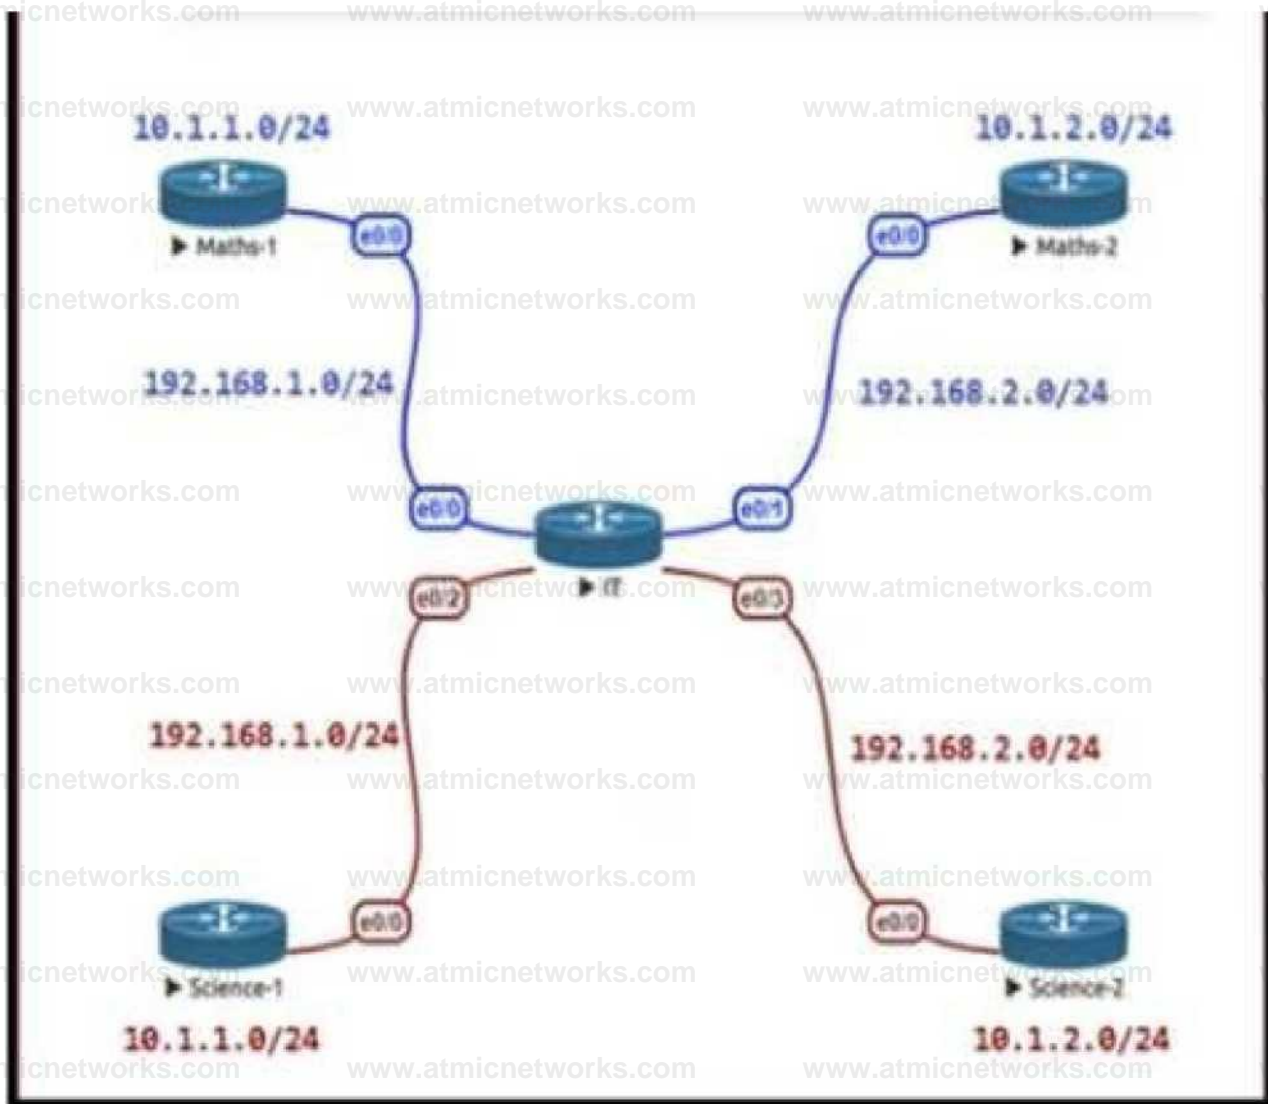
There is an implicit deny all at the end of any route-map so all other traffic that does not match 172.20.5.0/24 would be

dropped. Therefore we have to add a permitsequence at the end of the route-map to allow other traffic.

The default value of Local Preference is 100 and higher value is preferred so we have to set the local preference of AS100 lower than that of AS200.

Question: 199

Refer to the exhibit.



The Math and Science departments connect through the corporate IT router but users in the Math department must not be able to reach the Science department and vice versa. Which configuration accomplishes this task?

- A. vrf definition Science ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut
- B. vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 vrf forwarding Science no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 vrf forwarding Science no shut
- C. vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut
- D. vrf definition Science address-family ipv4 ! interface E 0/2 vrf forwarding Science ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 vrf forwarding Science ip address 192.168.2.1 255.255.255.0 no shut

Answer: D

Explanation:

Question: 200

An engineer configured Reverse Path Forwarding on an interface and noticed that the routes are dropped when a route lookup fails on that interface for a prefix that is available in the routing table. Which interface configuration resolves the issue?

- A. ip verify unicast source reachable-via rx
- B. ip verify unicast source reachable-via any
- C. ip verify unicast source reachable-via allow-default
- D. ip verify unicast source reachable-via 12-src

Answer: B


Explanation:

According to this question, uRPF is running in strict mode because packets are dropped even when that route exists in the routing table. Maybe packets are dropped because the receiving interface is different from the interface the local router uses to send packets to that destination. The ip verify unicast source reachable-via rx command enables Unicast RPF in strict mode.

To enable loose mode, administrators can use the any option (ip verify unicast source

reachable-via any). In loose mode, it doesn't matter if we use this interface to reach the source

or not.



The allow-default option allows the use of the default route in the source verification process.

Question: 201

Refer to the exhibit.

NY

```
router ospf 1 network 192.168.12.0 0.0.0.255
  area 0 network 172.16.2.0 0.0.0.255 area 0
```

```
interface E 0/0
```

```
  ip ospf authentication message-digest
```

```
  ip ospf message-digest-key 1 md5 Cisco123
```

The neighbor relationship is not coming up Which two configurations bring the adjacency up? (Choose two)

- A. NYrouter ospf 1area 0 authentication message-digest
- B. LAinterface E 0/0ip ospf message-digest-key 1 md5 Cisco123
- C. NYinterface E 0/0no ip ospf message-digest-key 1 md5 Cisco123ip ospf authentication-key Cisco123
- D. LAinterface E 0/0ip ospf authentication-key Cisco123
- E. LArouter ospf 1area 0 authentication message-digest

Answer: B,E

Explanation:

The configuration on NY router is good for OSPF authentication. So we must enable OSPF authentication on LA router with the following commands:

```
router ospf 1
```

```
area 0 authentication message-digest
```

```
interface E0/0
```

```
ip ospf message-digest-key 1 md5 Cisco123
```

Question: 202

Refer to the exhibit.

L 172.1.12.3/32 is directly connected. EthernetO/O

C 172.1.13.0/24 is directly connected. EthernetO/1

L 172.1.13.3/32 is directly connected. EthernetO/1

O 192.168.1.0/24 [110/2] via 172.1.12.1. 00:04:44 Ethernet

- O 192.168.2.0/24 [110/2 via 172.1.12.1,00:04:44, ElhernetO 0 192.168 3.0/24
[110/2 via 172 1 13 2. 00 04 44 ElhernetO/1
- O 192 168 4 0/24 [110/2 via 172.1 13 2,00 04 44 ElhernetO/1
192 168 5 0 '24 is variably subnelted 2 subnets. 2 masks
- C 192 168.5 0/24 is directly connected, LoopbackO
- L 192.168.5.1/32 is directly connected. LoopbackO 192.168.6.0/24 is variably
subnelted. 2 subnets. 2 masks
- C 192.168.6 0/24 is directly connected. Loopback 1
- L 192 168 6.1/32 is directly connected, Loopbackl

SanFrancisco and Boston routers are choosing slower links to reach each other despite the direct links being up Which configuration fixes the issue?

Boston Router

```
router ospf 1  
auto-cost reference-bandwidth 1000
```

SanFrancisco Router

```
router ospf 1  
auto-cost reference-bandwidth 1000
```

AH Routers

```
router ospf 1  
auto-cost reference-bandwidth 100
```

All Routers

```
router ospf 1  
auto-cost reference-bandwidth 1000
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 203

Refer to the exhibit.

```
Debug output:
username: USER55
password:
Aug 26 12:39:23.813: TPLUS: Queuing AAA Authentication request 4950 for processing
Aug 26 12:39:23.813: TPLUS(00001356) login timer started 1020 sec timeout
Aug 26 12:39:23.813: TPLUS: processing authentication continue request id 4950
Aug 26 12:39:23.813: TPLUS: Authentication continue packet generated for 4950
Aug 26 12:39:23.813: TPLUS(00001356)/0/WRITE/3A72C8D0: Started 5 sec timeout
!
|----- output omitted -----|
!
Aug 26 12:40:01.241: TAC+: using previously set server 192.168.1.3 from group tacacs+
Aug 26 12:40:01.241: TAC+: Opening TCP/IP to 192.168.1.3/49 timeout=5
Aug 26 12:40:01.249: TAC+: Opened TCP/IP handle 0x3BE31D1C to 192.168.1.3/49
Aug 26 12:40:01.249: TAC+: Opened 192.168.1.3 index=1
Aug 26 12:40:01.250: TAC+: 192.168.1.3 (3653537180) AUTHOR/START queued
Aug 26 12:40:01.449: TAC+: (3653537180) AUTHOR/START processed
Aug 26 12:40:01.449: TAC+: (-641430116): received author response status = FAIL
Aug 26 12:40:01.450: TAC+: Closing TCP/IP 0x3BE31D1C connection to 192.168.1.3/49
```

A network administrator logs into the router using TACACS+ username and password credentials, but the administrator cannot run any privileged commands. Which action resolves the issue?

- A. Configure TACACS+ synchronization with the Active Directory admin group
- B. Configure the username from a local database
- C. Configure full access for the username from TACACS+ server
- D. Configure an authorized IP address for this user to access this router

Answer: C

Explanation:

Question: 204

Refer to the exhibit.

```
ipv6 access-list INTERNET
permit ipv6 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA14::/64
permit tcp 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA13::/64 eq telnet
permit tcp 2001:DB8:AD59:BA21::/64 any eq http
permit ipv6 2001:DB8:AD59::/48 any
deny ipv6 any any log
```

When monitoring an IPv6 access list, an engineer notices that the ACL does not have any hits and is causing unnecessary traffic to pass through the interface. Which command must be configured to resolve the issue?

- A. access-class INTERNET in
- B. ipv6 traffic-filter INTERNET in
- C. ipv6 access-class INTERNET in
- D. ip access-group INTERNET in

Answer: C

Explanation:

Question: 205

Refer to the exhibit.


```
router ospf 1
 redistribute eigrp 1 subnets route-map EIGRP->OSPF
|
router eigrp 1
 network 10 0 106 0 0.0 0 255
```

```
route-map EIGRP->OSPF permit 10 match ip address WAN_PRE
FIXES route-map EIGRP->OSPF permit 20 match ip address
LOCAL_PREFIXES route-map EIGRP->OSPF permit 30 match ip
address VPN_PREFIXES
```

```
|
ip prefix-list LOCAL_PREFIXES seq 5 permit 172 16.0.0/12 le 24 ip
prefix-list VPN_PREFIXES seq 5 permit 192.168.0.0/16 le 24 ip
prefix-list WAN_PREFIXES seq 5 permit 10.0.0.0/8 le 24
```

The network administrator configured redistribution on an ASBR to reach to all WAN networks but failed Which action resolves the issue?

- A. The route map must have the keyword prefix-list to evaluate the prefix list entries
- B. The OSPF process must have a metric when redistributing prefixes from EIGRP.
- C. The route map EIGRP->OSPF must have the 10.0.106.0/24 entry to exist in one of the three prefix lists to pass
- D. EIGRP must redistribute the 10.0.106.0/24 route instead of using the network statement

Answer: A

Explanation:

In order to use a prefix-list in a route-map, we must use the keyword "prefix-list" in the "match" statement. . For

example:

match ip address prefix-list WAN_PREFIXES

Without this keyword, the router will try to find an access-list with the same name instead.

Question: 206

How does an MPLS Layer 3 VPN function?

- A. set of sites use multiprotocol BGP at the customer site for aggregation
- B. multiple customer sites interconnect through service provider network to create secure tunnels between customer edge devices
- C. set of sites interconnect privately over the Internet for security
- D. multiple customer sites interconnect through a service provider network using customer edge to provider edge connectivity

Answer: D

Explanation:

A Multiprotocol Label Switching(MPLS) Layer 3 Virtual Private Network (VPN) consists of a set of sites that are interconnected by means of an MPLS provider core network. At each customer site, one or more customer edge (CE)

routers attach to one or more provider edge (PE) routers. Reference:

https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r6-5/lxvpn/configuration/guide/b-l3vpn-cg-asr9000-65x/b-l3vpn-cg-asr9000-65x_chapter_010.pdf

Configuration Output:

```
aaa new-model
!
aaa authentication login default local
aaa authentication login VTY_AUTH local
aaa authorization exec default none
aaa authorization exec VTY_AUTH local
aaa accounting exec default start-stop group radius
!
```

```
password 7 K0AyUubDrfOgO4s
authorization exec VTY_AUTH
login authentication VTY_AUTH
```

Debug Output:

```
AAA/AUTHEN/LOGIN (000004B6): Pick method list 'default'
AAA/AUTHOR (0x4B6): Pick method list 'VTY_AUTH'
AAA/AUTHOR/EXEC(000004B6): Authorization FAILED
```

Which action resolves the failed authentication attempt to the router?

- A. Configure aaa authorization login command on line vty 0 4
- B. Configure aaa authorization login command on line console 0
- C. Configure aaa authorization console global command
- D. Configure aaa authorization console command on line vty 0 4

Answer: C

Explanation:

In the debug output, we see that the Authorization (not Authentication) failed so we need to correct the authorization.

In order to enable authorization, we must use the global command “aaa authorization console” first.

Reference:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/a1/sec-a1-cr-book/sec-cr-a1.html>

Question: 208

A customer reports to the support desk that they cannot print from their PC to the local printer id:401987778. Which tool must be used to diagnose the issue using Cisco DNA Center Assurance?

- A. application trace
- B. path trace
- C. ACL trace
- D. device trace

Answer: B

Explanation:

Question: 209

When determining if a system is capable of support, what is the minimum time spacing required for a BFD control packet to receive once a control packet is arrived?

- A. Desired Min TX Interval

- B. Detect Mult
- C. Required Min RX Interval
- D. Required Min Echo RX Interval

Answer: C

Explanation:

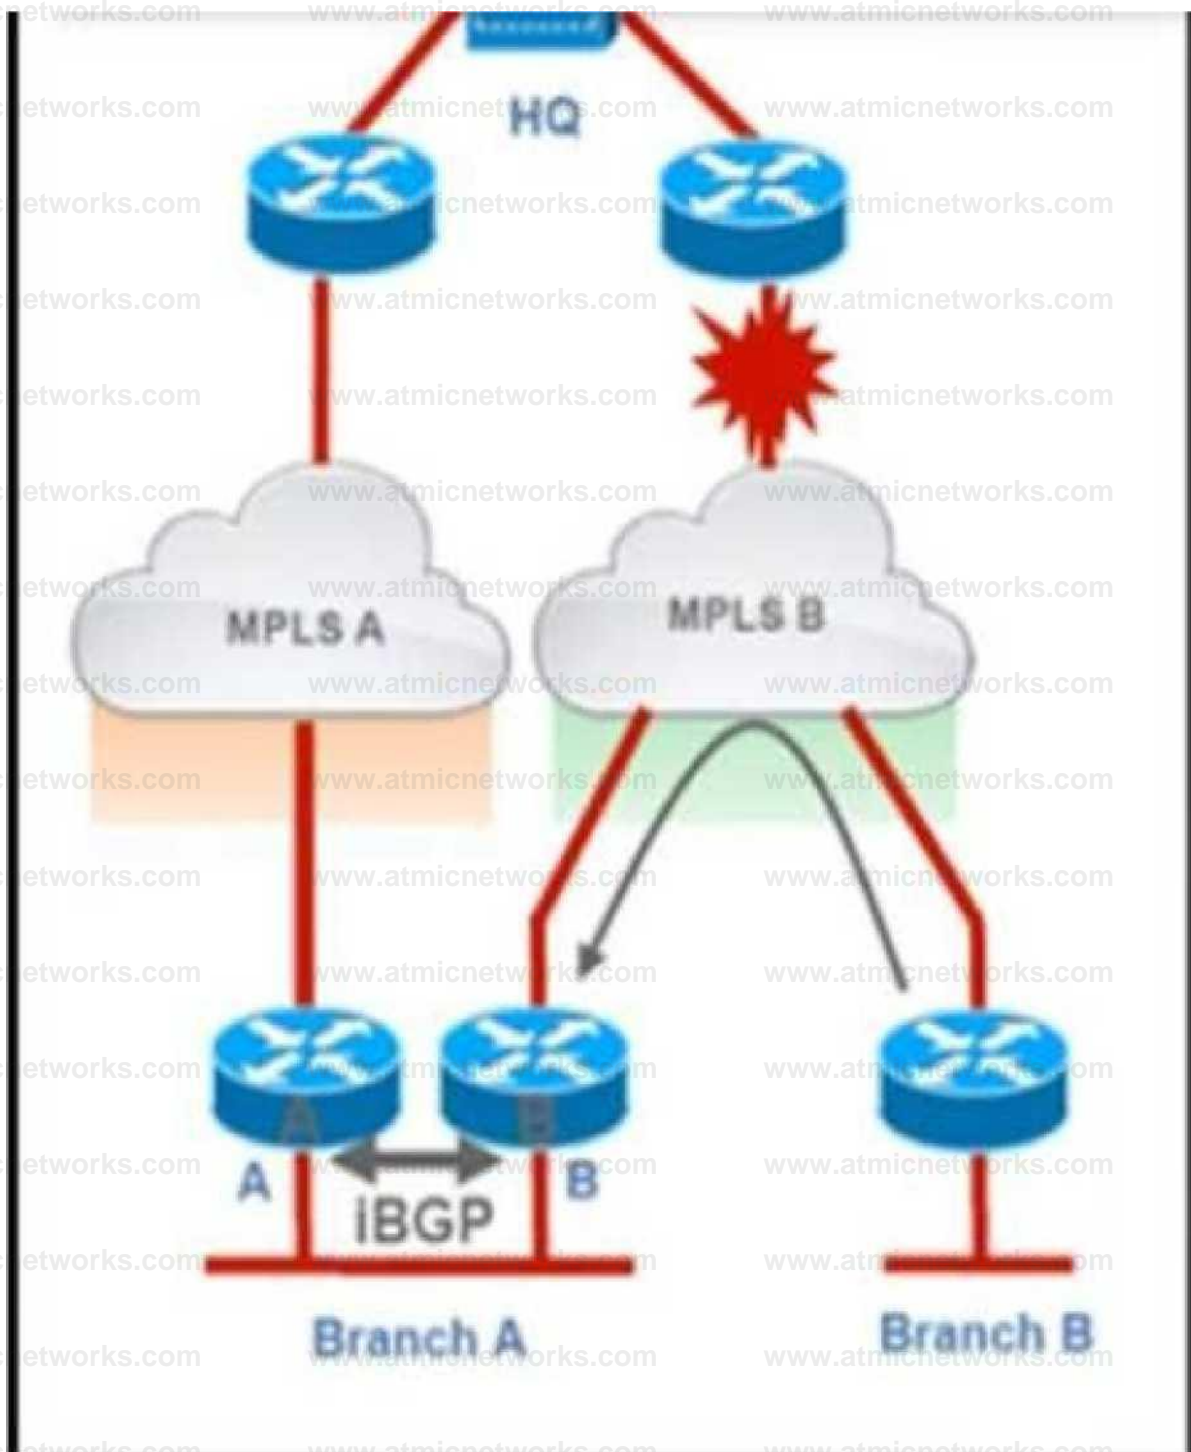
Required Min RX Interval: This is the minimum interval, in microseconds, between received BFD Control packets that this system is capable of supporting.

Reference:

https://www.cisco.com/en/US/technologies/tk648/tk365/tk480/technologies_white_paper0900aecd80244005.html

Question: 210

Refer to the exhibit.



Troubleshoot and ensure that branch B only ever uses the MPLS B network to reach HQ. Which action achieves this requirement?

A. Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP

B. Increase the local preference for all HQ prefixes received at branch B from the MPLS B network to be higher than the local preferences used on the MPLS A network

- C. Introduce AS path prepending on the branch A MPLS B network connection so that any HQ advertisements from branch A toward the MPLS B network are prepended three times
- D. Modify the weight of all HQ prefixes received at branch B from the MPLS B network to be higher than the weights used on the MPLS A network

Answer: A

Explanation:

If we modify the weight, increase local preference or use AS path prepending then we can only make MPLS B prefer over MPLS A. But when MPLS B is down then MPLS A will be used which does not meet the requirement of this question. Only with AS path filtering we can deny prefixes from certain AS and make sure branch B never uses MPLS A to reach HQ.

Question: 211

DRAG DROP

Drag and drop the LDP features from the left onto the descriptions on the right

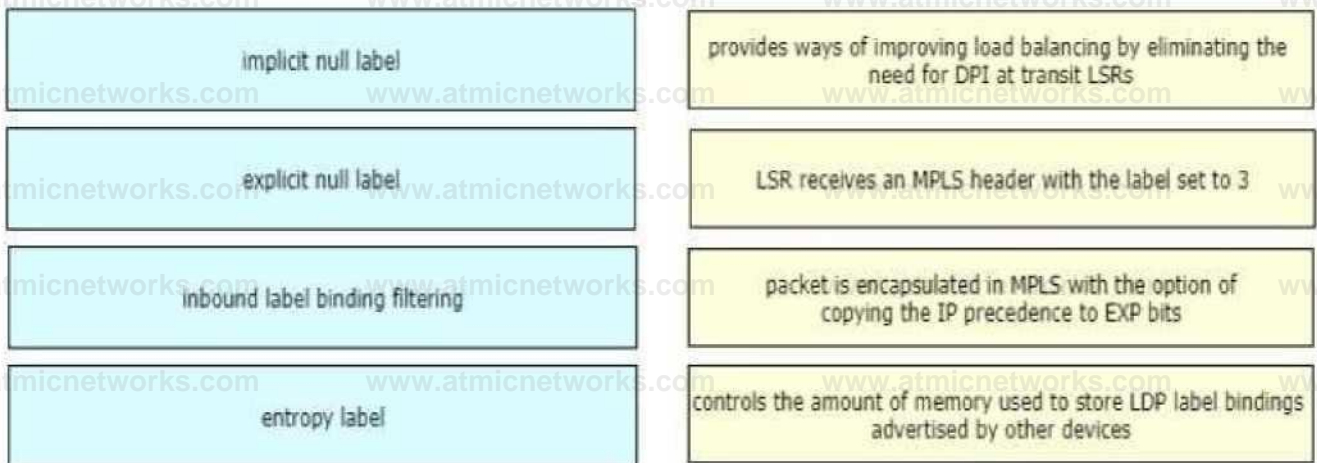
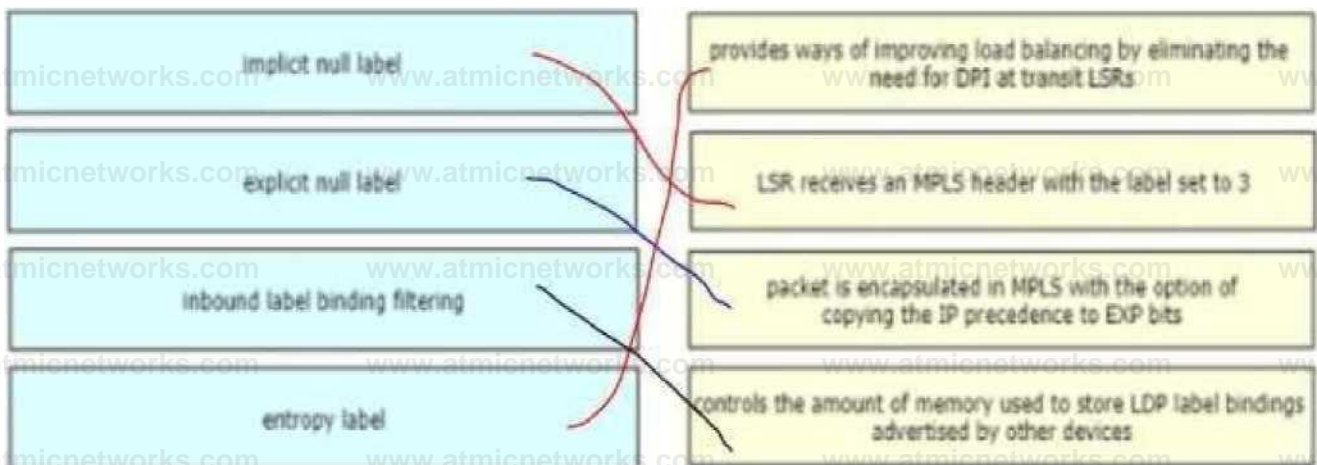
**Answer:****Explanation:**

Diagram Description automatically generated

The MPLS LDP Inbound Label Binding Filtering feature can be used to control the amount of memory used to store Label Distribution Protocol (LDP) label bindings advertised by other devices. For example, in a simple Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) environment, the VPN provider edge (PE) devices might require label switched paths (LSPs) only to their peer PE devices (that is, they do not need LSPs to core devices).

Inbound label binding filtering enables a PE device to accept labels only from other PE devices.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-sy/mp-ldp-15-sy-book/mp-ldp-inbound-filtr.html

Question: 212

Refer to the exhibit.

```
Router# show ip route
```

```
2.0.0.0/24 is subnetted, 1 subnets
```

```
C    2.2.2.0 is directly connected, Ethernet0/0
```

```
C    3.0.0.0/8 is directly connected, Serial1/0
```

```
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
```

```
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
```

```
131.108.0.0/24 is subnetted, 2 subnets
```

```
O    131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
```

```
O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
```

```
Router# show ip bgp
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.2.2.0/24	0.0.0.0	0	32768	?	
*> 131.108.1.0/24	2.2.2.2	84	32768	?	
*> 131.108.2.0/24	2.2.2.2	74	32768	?	

The OSPF routing protocol is redistributed into the BGP routing protocol, but not all the OSPF routes are distributed into BGP. Which action resolves the issue?

- A. Include the word external in the redistribute command
- B. Use a route-map command to redistribute OSPF external routes defined in an access list

- C. Include the word internal external in the redistribute command
- D. Use a route-map command to redistribute OSPF external routes defined in a prefix list.

Answer: C

Explanation:

If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default. You can use the internal keyword along with the redistribute command under router bgp to redistribute OSPF intra- and inter-area routes.

Use the external keyword along with the redistribute command under router bgp to redistribute OSPF external routes into BGP.

-> In order to redistribute all OSPF routes into BGP, we must use both internal and external keywords. The full command would be (suppose we are using OSPF 1): redistribute ospf 1 match internal external

Note: The configuration shows match internal external 1 external 2. This is normal because

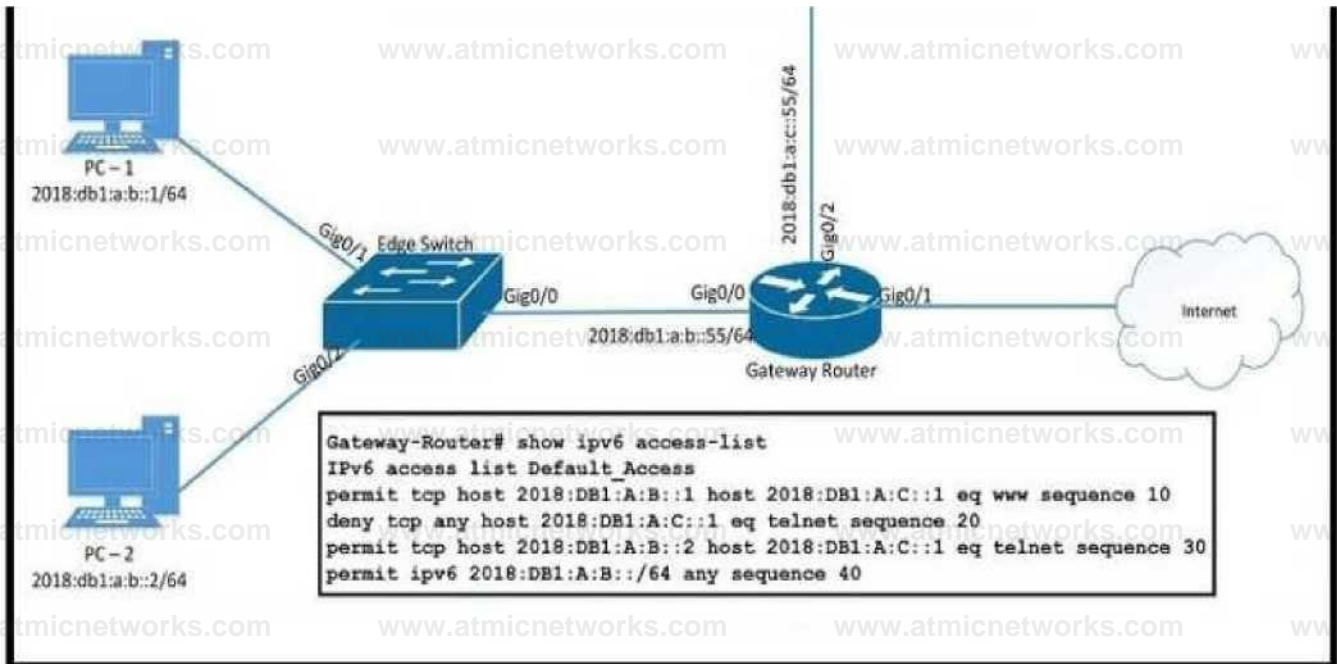
OSPF automatically appends "external 1 external 2" in the configuration. In other words, keyword external = external 1 external 2. External 1 = O E1 and External 2 = O E2.

Reference:

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redistribution.html>

Question: 213

Refer to the exhibit.



PC-2 failed to establish a Telnet connection to the terminal server. Which configuration resolves the issue?

```

Gateway-Router(config)=ipv6 access-list Default_Access
Gateway-Router(config-ipv6-acl-sequence 15 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet
Gateway-Router(config)#ipv6 access-list Default_Access
Gateway-Router(config-ipv6-acl-permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet
  
```

```

Gateway-Router(config)#ipv6 access-list Default_Access
Gateway-Router(config-ipv6-acl) :no sequence 20
Gateway-Router(config-ipv6-acl-sequence 5 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet
  
```

- Gateway-Router(config)#ipv6 access-list Default_Access


```
Gateway-Router(config-ipv6-acl-sequence 25 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

In fact in this question both answer A and answer C are correct but we believe answer A is the better choice as it only allows

PC-2 to telnet to terminal server. All other hosts are refused to telnet to terminal server via sequence 20.

Question: 214

What statement about route distinguishes in an MPLS network is true?

- A. Route distinguishes make a unique VPNv4 address across the MPLS network.
- B. Route distinguishers allow multiple instances of a routing table to coexist within the edge router.
- C. Route distinguishes are used for label bindings
- D. Route distinguishes define which prefixes are imported and exported on the edge router

Answer: A

Explanation:

Question: 215

Refer to the exhibit.

```
Router#show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(1)
Interface          Xmit Queue PeerQ      Mean Pacing Time Multicast F
Peers Un/Reliable Un/Reliable SRTT Un/Reliable Flow T
Lo0                0      0/0      0/0      0      0/0      0      0
Fa0/0              1      0/0      0/0      7      0/2      50     0

Router#show running-config | section eigrp
router eigrp 1
 network 172.16.0.0 0.0.0.255
 network 192.168.2.2 0.0.0.0
 network 192.168.12.2 0.0.0.0

Router#show running-config interface Fa0/3
Building configuration...

Current configuration : 93 bytes
!
interface FastEthernet0/3
 ip vrf forwarding CLIENT1
 ip address 172.16.0.1 255.255.255.0
```

While troubleshooting an EIGRP neighbor adjacency problem, the network engineer notices that the interface connected to the neighboring router is not participating in the EIGRP process. Which action resolves the issues?

- A. Configure the network command to network 172.16.0.1 0.0.0.0
- B. Configure the network command under EIGRP address family vrf CLIENT1
- C. Configure EIGRP metrics on interface FastEthernet0/3

D. Configure the network command under EIGRP address family ipv4

Answer: B

Explanation:

```
router eigrp 1
```

```
address-family ipv4 vrf CLIENT1 network 172.16.0.0 0.0.0.255 no auto-summary autonomous-system 1 exit-  
address-family
```

Question: 216

Refer to the exhibit.

```
admin@linux:~$ scp script.py admin@198.51.100.64:script.py
```

Password:

Administratively disabled.

```
admin@linux:~$ Connection to 198.51.100.64 closed by remote host.
```

A network administrator has developed a Python script on the local Linux machine and is trying to transfer it to the router. However, the transfer fails. Which action resolves this issue?

- A. The SSH service must be enabled with the crypto key generate rsa command.
- B. The SCP service must be enabled with the ip scp server enable command.
- C. The Python interpreter must first be enabled with the guestshell enable command.
- D. The SSH access must be allowed on the VTY lines using the transport input ssh command.


Answer: B

Explanation:

The error "Administratively disabled" means we need to enable SCP on the router with the command:
Router(config)#ip scp server enable

Question: 217

Refer to the exhibit.



The network administrator can see the DHCP discovery packet in R1. but R2 is not replying to the DHCP request. The R1 related interface is configured with the DHCP helper address. If the PC is directly connected to the FaO/1 interface on R2, the DHCP server assigns as IP address from the DHCP pool to the PC. Which two commands resolve this issue?

(Choose two.)

- A. service dhcp-relay command on R1
- B. ip dhcp option 82 command on R2
- C. service dhcp command on R1
- D. ip dhcp relay information enable command on R1
- E. ip dhcp relay information trust-all command on R2

Answer: C,E**Explanation:**

1. R1 received DHCP packet and its interface was configured with the DHCP helper address. But we are not sure if R1 forward DHCP packet to R2 or not. 2. If we connect PC directly to R2 then this problem will not appear - > DHCP Server function was configured on R2.

From these facts, the most likely problem is related to Option 82. Maybe R2 ignored DHCP request packets because it was receiving these packets with the giant field set to 0.0.0.0.

By default Cisco IOS devices reject packets with zero "giaddr" and by default Cisco Catalyst switches use "giaddr" of zero when configured for DHCP snooping! Reference:

<https://blog.ine.com/2009/07/22/understanding-dhcp-option-82>

If we can run the "debug ip dhcp server packet" on R2, we may see these messages:

```
*Feb 22 23:54:57.759: IP: s=0.0.0.0 (FastEthernet0/1), d=255.255.255.255, len 34 4, input feature, MCI Check(64), rtype 0,
forus FALSE, sendself FALSE, mtu 0, fw dchk FALSE *Feb 22 23:54:57.759: IP: s=0.0.0.0 (FastEthernet0/1),
d=255.255.255.255, len 34 4, rcvd 2 *Feb 22 23:54:57.759: IP: s=0.0.0.0 (FastEthernet0/1), d=255.255.255.255, len 34 4,
stop process pak for forus packet
```

```
*Feb 22 23:54:57.759: DHCPD: inconsistent relay information. *Feb 22 23:54:57.759: DHCPD: relay information option
exists, but giaddr is zero
```

We are receiving the DHCP packet from R1, source 0.0.0.0, and destination 255.255.255.255 broadcast, but if you notice from the debug output, R2, our DHCP Server, is complaining that the relay information is inconsistent. Option 82, Information Option, is contained in the packet but the GIADDR is zero. The GIADDR stands for Gateway IP Address, which is the IP Address of the relaying agent. The Option 82, Information Option, would then contain the receiving port and hostname of the Relaying Agent by default.

R2 sees the Option 82 information, signalling that the DHCP packet might have been relayed, BUT there is no relaying IP Address. This is the behavior of DHCP Snooping when enabling it on a switch, and since the switchport does not contain an IP Address, since it's Layer 2, no GIADDR will be added.

Instead, just the Option 82 Information is added and this is the problem we have, but there are options:

J. You could trust all on R2 the DHCP Server, which will cause the server to not be so suspicious: – ip dhcp relay information trust-all – ip dhcp relay information trusted 2. Disable the addition of Option 82 information on SW: – no ip dhcp snooping information option 3. Trust the port that is receiving the DHCP Discover: – ip dhcp snooping trust

Any of these options will fix our predicament.


Reference: <https://evilttl.com/wiki/DHCP-Snooping>

But in the answer choices, we only have 1 correct answer which is the command “ip dhcp relay information trust-all”. We checked if we need any “service dhcp...” command on both IOS version 12.4 and 15.1:

Therefore we only have the “service dhcp” command, we don’t have any “service dhcp-relay” command available. But the description of the “service dhcp” command says that it enables both DHCP server and relay agent so this is the best answer left.

Question: 218

Refer to the exhibit.



```
R1#sh bgp ipv6 sum
BGP router identifier 1.1.1.1, local AS number 6501
BGP table version is 1, main routing table version 1

Neighbor          V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down State/PfxRcd
AB01:2011:7:100::3 4 6502    0      0        1     0    0   never      Idle
```

```
R1#debug ip bgp all
* Nov 8 17:22:11.223: BGP: AB01:2011:7:100::3 active went from Idle to Active
* Nov 8 17:22:11.223: BGP: AB01:2011:7:100::3 open active, local address AB01:2011:7:100::1
* Nov 8 17:22:11.224: BGP: AB01:2011:7:100::3 open failed: Connection refused by remote host
* Nov 8 17:22:11.224: BGP: AB01:2011:7:100::3 Active open failed - tcb is not available, open
active delayed 11264 ms (35000ms max, 60% jitter)
* Nov 8 17:22:11.224: BGP: ses global AB01:2011:7:100::3 (0xC3F49FF0:0) act Reset (Active open failed)
* Nov 8 17:22:11.232: BGP: AB01:2011:7:100::3 active went from Active to Idle
* Nov 8 17:22:11.232: BGP: nrb global AB01:2011:7:100::3 Active open failed - open timer running
```

```
R1#ping ipv6 AB01:2011:7:100::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to AB01:2011:7:100::3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
! ! ! ! !
```

```
St'nii-'ig 5. IOC :?;ir C^iP Ecto 10 AB01 ^OikT-tnt); ,< l'nrr.jt 'iJ seconds; IREI
```

```
Success fate is 100 percent |5/5|, round-trip rrmn/avg/max: - 1/1/1 ms
```

An engineer configured BGP between routers R1 and R3. The BGP peers cannot establish neighbor adjacency to be able to exchange routes. Which configuration resolves this issue?

- A. R3router bgp 6502address-family ipv6neighbor AB01:2011:7:100::1 activate
- B. R1router bgp 6501address-family ipv6neighbor AB01:2011:7:100::3 activate
- C. R3router bgp 6502neighbor AB01:2011:7:100::1 ebgp-multihop 255
- D. R1router bgp 6501 neighborAB01:2011:7:100::3ebgp-multihop255


Answer: A**Explanation:**

From the output, we learned that R1 was trying to establish BGP neighbor relationship with R3 but failed. Both of them were using physical interface to establish neighbor relationship so we don't need the "... ebgp-multihop" command here.

The only reasonable answer is R3 has not been configured to activate BGP neighbor relationship with R1.

Question: 219

Refer to the exhibit.



AS65510 iBGP is configured for directly connected neighbors. R4 cannot ping or traceroute network 192.168.100.0/24

Which action resolves this issue?

- A. Configure R4 as a route reflector server and configure R1 as a route reflector client

- B. Configure R1 as a route reflector server and configure R2 and R3 as route reflector clients
- C. Configure R4 as a route reflector server and configure R2 and R3 as route reflector clients.
- D. Configure R1 as a route reflector server and configure R4 as a route reflector client


Answer: D

Explanation:

A route received from one iBGP peer will NOT be advertised to another iBGP peer. Therefore R4 could not receive advertisement for network 192.168.100.0/24. We can overcome this BGP limitation by configuring R1 as a route reflector server and R4 as a route reflector client so that R1 sends advertisements for R4.

Question: 220

Refer to the exhibit.



BGP and EIGRP are mutually redistributed on R3, and EIGRP and OSPF are mutually redistributed on R1. Users report packet loss and interruption of service to applications hosted on the 10.1.1.0/24 prefix. An engineer tested the link from R3 to R4 with no packet loss present but has noticed frequent routing changes on R3 when running the debug ip route command.

Which action stabilizes the service?

- A. Tag the 10.1.1.0/24 prefix and deny the prefix from being redistributed into OSPF on R1.
- B. Repeat the test from R4 using ICMP ping on the local 10.1.1.0/24 prefix, and fix any Layer 2 errors on the host or switch side of the subnet. ^ C. Place an OSPF distribute-list outbound on R3 to block the 10.1.10/24 prefix from being advertised back to R3.
- C. Reduce frequent OSPF SPF calculations on R3 that cause a high CPU and packet loss on traffic traversing R3.

Answer: A**Explanation:**

After redistribution, R3 learns about network 10.1.1.0/24 via two paths:

+ Internal BGP (IBGP): advertised from R4 with AD of 200 (and metric of 0)

+ OSPF: advertised from R1 with AD of 110 (O E2) (and metric of 20)

Therefore R3 will choose the path with the lower AD via OSPF

But this is a looped path which is received from R3 -> R2 -> R1 -> R3. So when the advertised route from R4 is expired, the

looped path is also expired soon and R3 will reinstall the main path from R4. This is the cause of intermittent

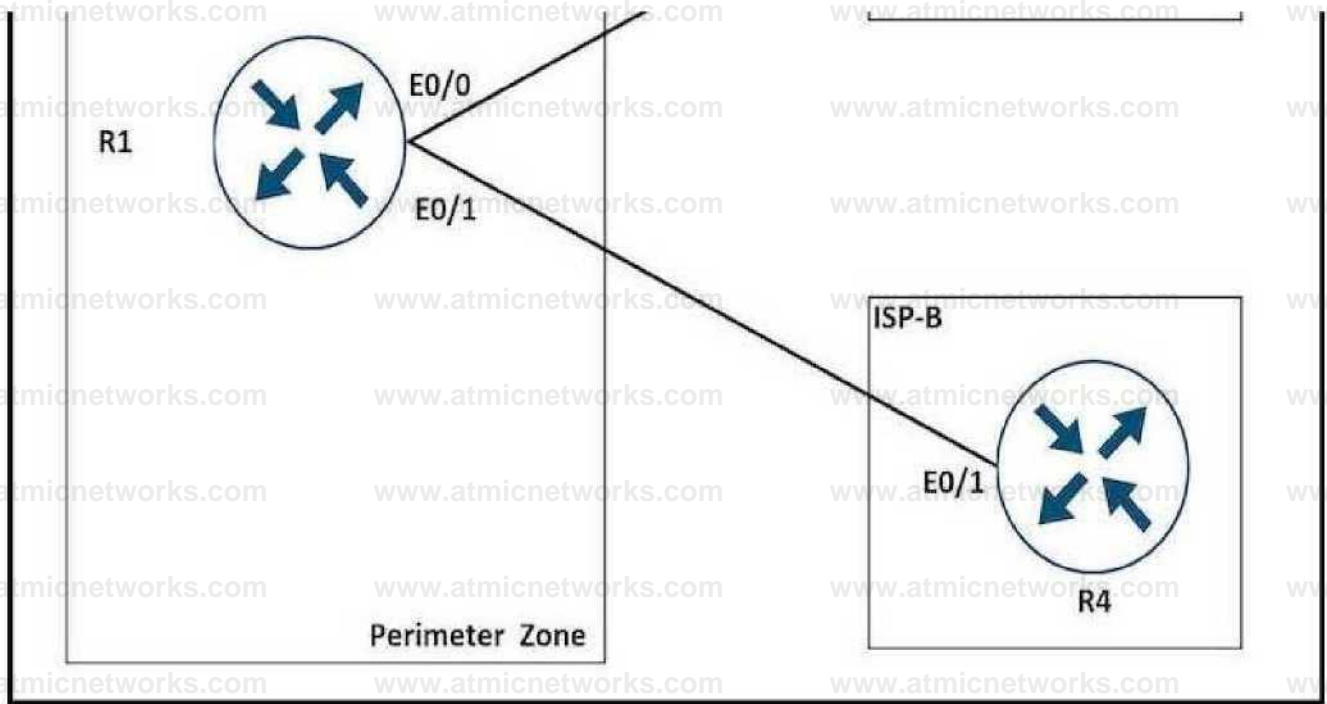
connectivity.

We can solve this problem by denying the 10.1.1.0/24 prefix from being redistributed into OSPF on R1. So R3 will not learn this prefix from R1.

Or another solution is to place an OSPF distribute-list inbound on R3 to block the 10.1.1.0/24 prefix from being advertised back to R3.

Question: 221

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. no terminal monitor
- B. (config)#terminal no monitor
- C. #terminal no monitor
- D. (config)#no terminal monitor

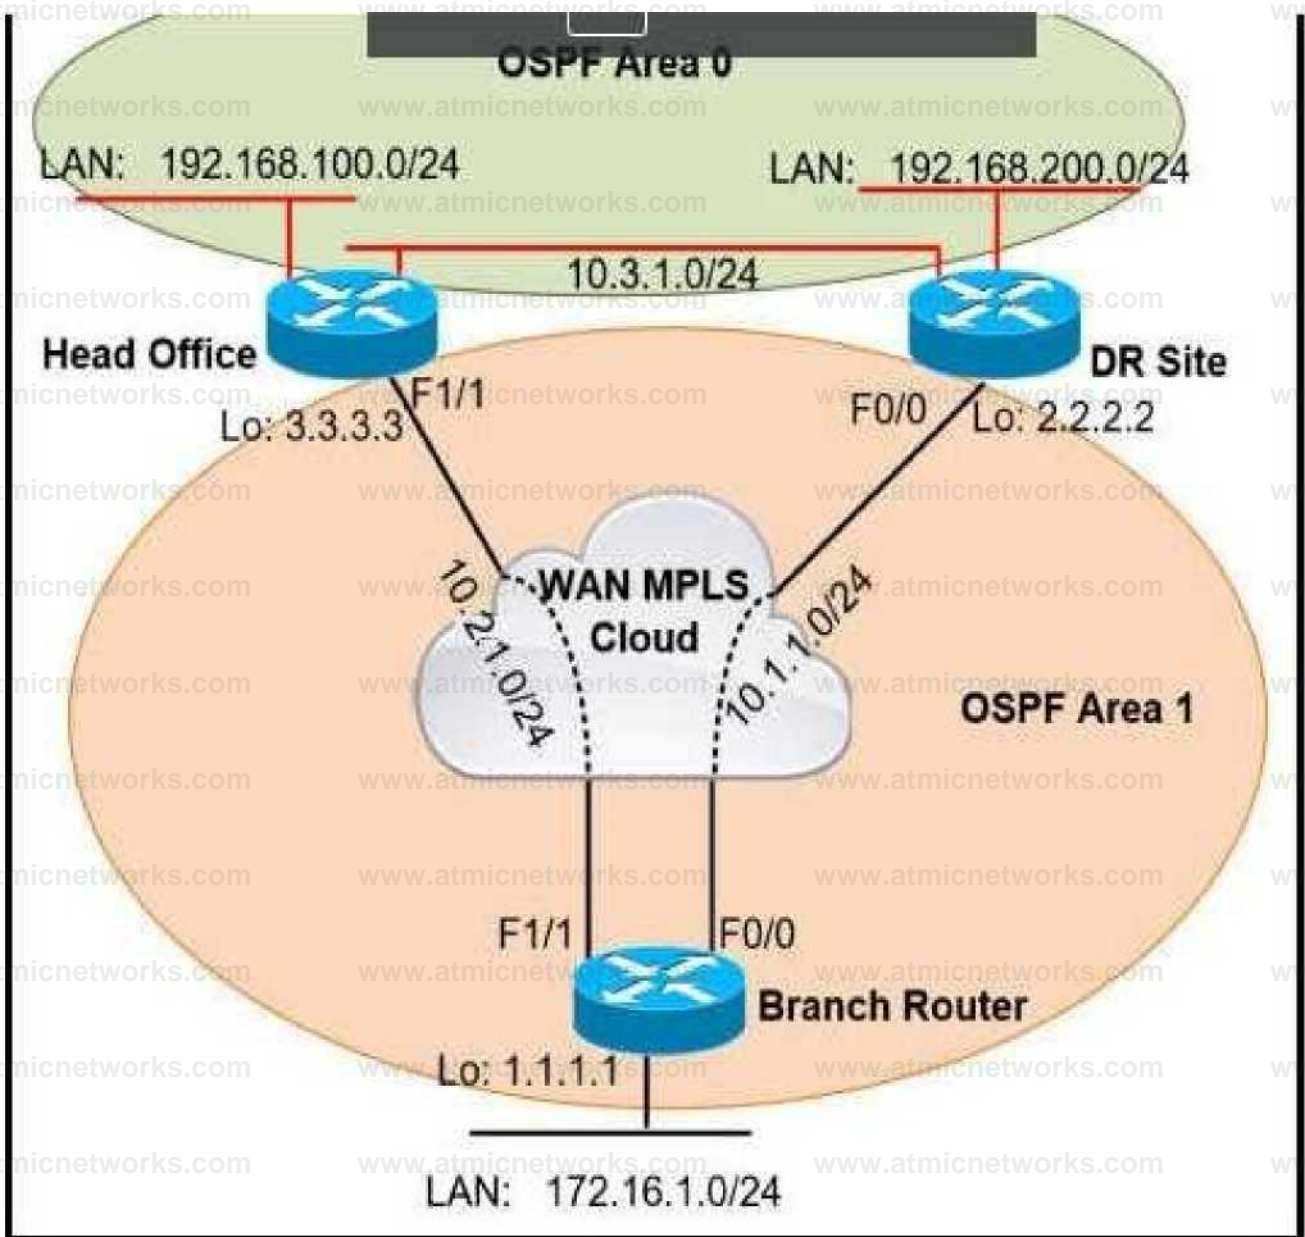
Answer: C

Explanation:

To turn off terminal monitor, use "terminal no monitor" in the enable mode

Question: 222

Refer to the exhibit.



A network administrator reviews the branch router console log to troubleshoot the OSPF adjacency issue with the DR router. Which action resolves this issue?

- A. Advertise the branch WAN interface matching subnet for the DR site.
- B. Configure matching hello and dead intervals between sites.
- C. Configure the WAN interface for DR site in the related OSPF area.

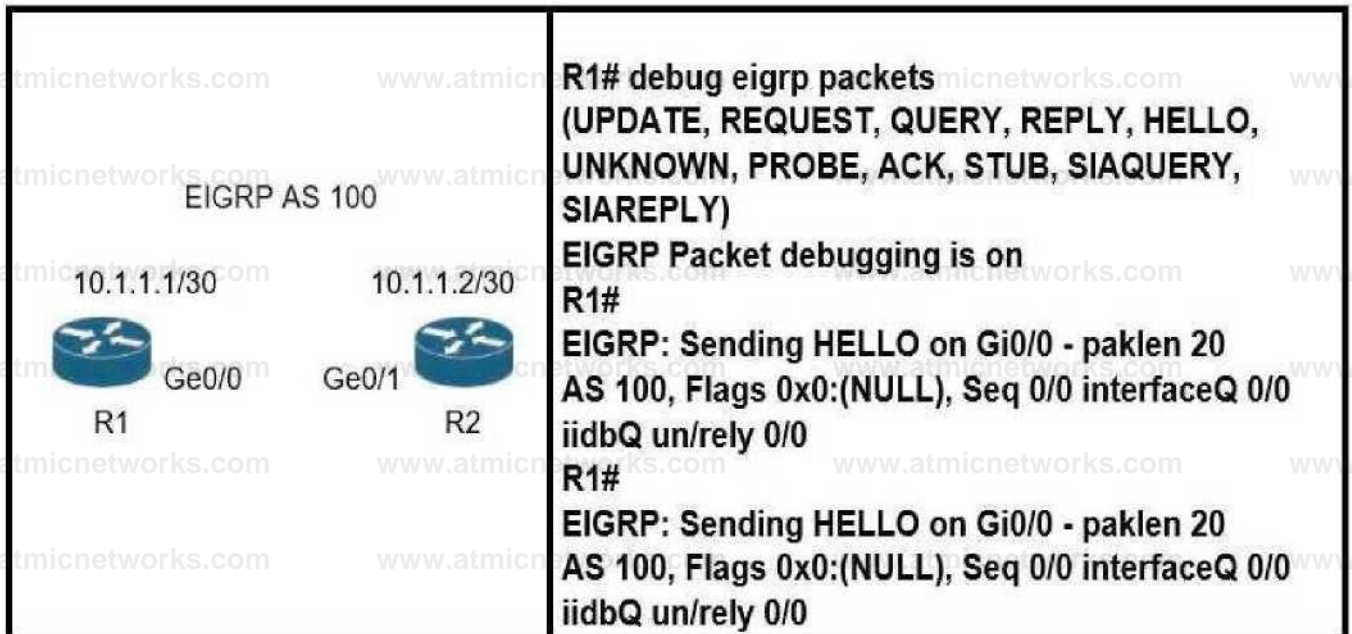
D. Stabilize the DR site flapping link to establish OSPF adjacency.

Answer: A

Explanation:

Question: 223

Refer to the exhibit.



Which action resolves the adjacency issue?

- A. Match the hello interval timers.
- B. Configure the same EIGRP process IDs.
- C. Match the authentication keys.
- D. Configure the same autonomous system numbers.

Answer: D

Explanation:

EIGRP does not have process ID as it uses Autonomous System (AS) numbers only.

This is not an authentication problem or we would see this error from the debug:


EIGRP: Ethernet0/0: ignored packet from 10.1.1.3, opcode = 1 (missing authentication or key-chain missing)

If the AS numbers between two routers are different then the neighbor relationship cannot be formed.

Topic 3, Exam Pool C

Question: 224

Refer to the exhibit.



A network administrator is trying to access a branch router using TACACS+ username and password credentials, but the administrator cannot log in to the router because the WAN connectivity is down. The branch router has following

AAA configuration:

aaa new-model

aaa authorization commands 15 default group tacacs+

aaa accounting commands 1 default stop-only group tacacs+

aaa accounting commands 15 default stop-only group tacacs+ tacacs-server host 10.100.50.99

tacacs-server key CiSco123

Which command will resolve this problem when WAN connectivity is down?

A. `aaa authentication login default group tacacs+ local`

B. `aaa authentication login default group tacacs+ enable`

- C. aaa authentication login default group tacacs+ console
- D. aaa authentication login console group tacacs+ enable

Answer: A

Explanation:

With the “aaa authentication login default group tacacs+ local” command configured, when logging in, the password supplied will be attempted to be verified by the TACACS+ server before access is granted. If the server is unavailable/unreachable, then the switch will fall back to using the local authentication database.

Question: 225

Users report issues with reachability between areas as soon as an engineer configured summary routes between areas in a multiple area OSPF autonomous system. Which action resolves the issue?

- A. Configure the summary-address command on the ASBR.
- B. Configure the summary-address command on the ABR.
- C. Configure the area range command on the ABR.
- D. Configure the area range command on the ASBR.

Answer: C

Explanation:

For OSPF, we can only summary at the ABR with the command “area range” or at the ASBR with the command “summary-address” -> Therefore answer A and answer B are not correct.

In this question, the most likely problem is that when doing summarization, the network mask is configured wrong and summarization doesn't work because of the misconfiguration. When configuring the area range command, make sure that the summarization mask is in the form of a prefix mask rather than a wildcard mask (that is, 255.255.255.0 instead of 0.0.0.255).

Good reference: <https://www.configrouter.com/troubleshooting-route-summarization-ospf-14082/>

Question: 226

A network administrator is troubleshooting a high utilization issue on the route processor of a router that was reported by NMS. The administrator logged into the router to check the control plane policing and observed that the BGP process is dropping a high number of routing packets and causing thousands of routes to recalculate frequently. Which solution resolves this issue?

- A. Police the cir for BGP, conform-action transmit, and exceed action transmit.
- B. Shape the pir for BGP, conform-action set-prec-transmit, and exceed action set-frde-transmit.
- C. Shape the cir for BGP, conform-action transmit, and exceed action transmit.
- D. Police the pir for BGP, conform-action set-prec-transmit, and exceed action set-clp-transmit.

Answer: D


Explanation:

CIR (Committed Information Rate) is the minimum guaranteed traffic delivered in the network.

PIR (Peak Information Rate) is the top bandwidth point of allowed traffic in a non-busy time without any guarantee.

Two Rates & Three Colors


Information Rate (bps)



+ Policing: is used to control the rate of traffic flowing across an interface. During a bandwidth exceed (crossed the maximum configured rate), the excess traffic is generally dropped or remarked. The result of traffic policing is an output rate that appears as a saw-tooth with crests and troughs. Traffic policing can be applied to inbound and outbound interfaces. Unlike traffic shaping, QoS policing avoids delays due to queuing. Policing is configured in bytes.

+ Shaping: retains excess packets in a queue and then schedules the excess for later transmission over increments of time. When traffic reaches the maximum configured rate, additional packets are

queued instead of being dropped to proceed later. Traffic shaping is applicable only on outbound interfaces as buffering and queuing happens only on outbound interfaces. Shaping is configured in bits per second.



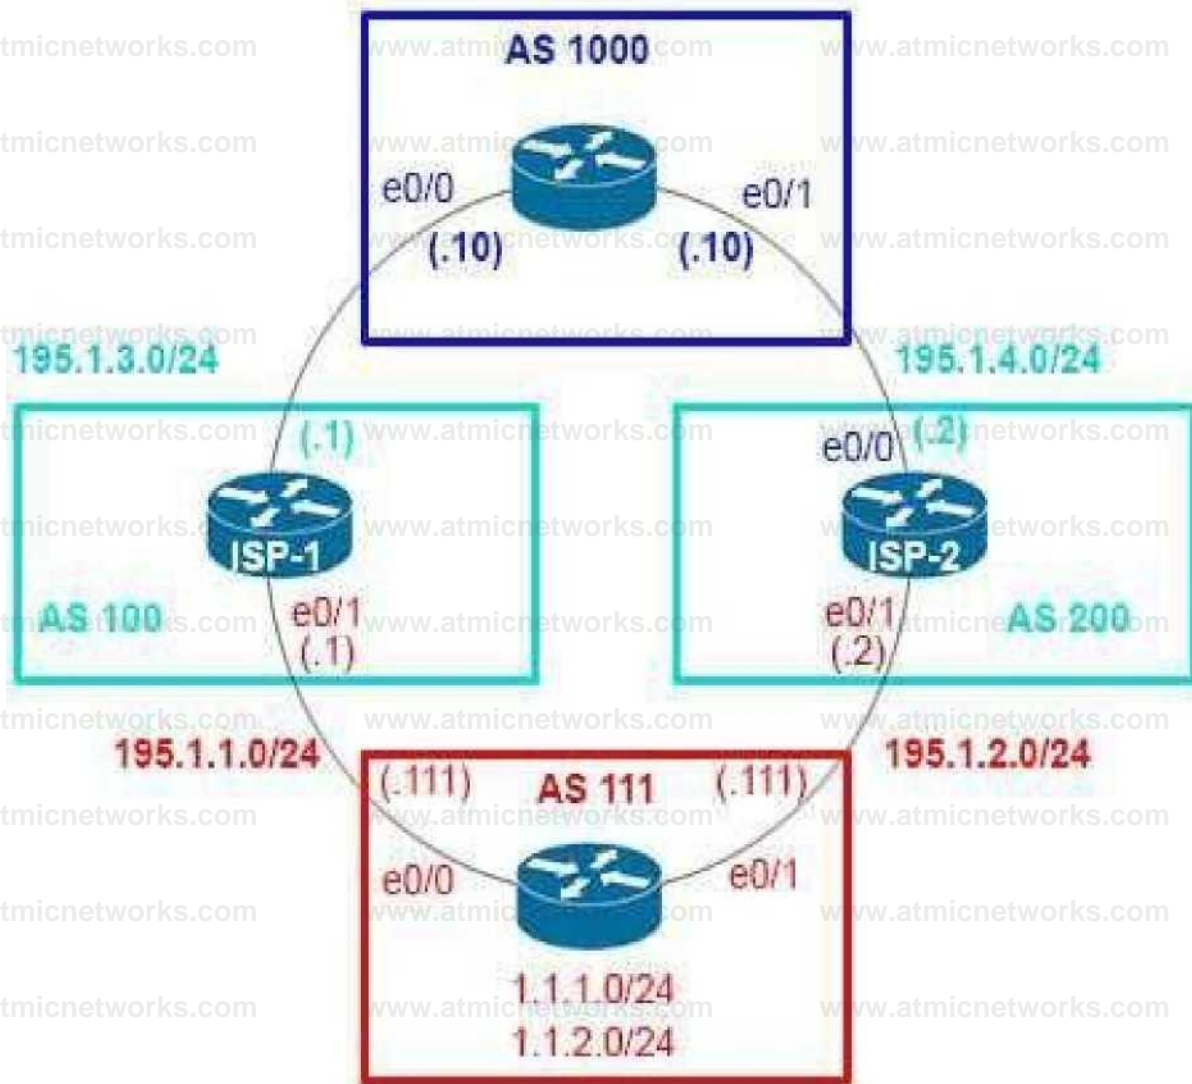
Therefore in this case we can only policing, not shaping as traffic shaping is applicable only on outbound interfaces as buffering and queuing happens only on outbound interfaces. Moreover, BGP traffic is not important so we can drop the excess packets without any problems.

And we only policing the PIR traffic so that the route processor is not overwhelmed by BGP calculation.

Note: The "set-prec-transmit" is the same as "transmit" command except it sets the IP Precedence level as well. The "set-clp-transmit" sets the ATM Cell Loss Priority (CLP) bit from 0 to 1 on the ATM cell and transmits the packet.

Question: 227

Refer to the exhibit.



AS111

Router bgp 111

Neighbor 195.1.1.1 remote-as 100

Neighbor 195.1.1.1 allowas-in

Neighbor 195.1.2.2 remote-as 200

Neighbor 195.1.2.2 allowas-in

AS111 is receiving its own routes from AS200 causing a loop in the network. Which configuration provides loop prevention?

A)

```
router bgp 111
neighbor 195.1.1.1 as-override
neighbor 195.1.2.2 as-override
```

B)

```
router bgp 111
neighbor 195.1.1.1 as-override
no neighbor 195.1.2.2 allowas-in
```

C)

```
router bgp 111
no neighbor 195.1.1.1 allowas-in
no neighbor 195.1.2.2 allowas-in
```

D)

```
router bgp 111
neighbor 195.1.2.2 as-override
no neighbor 195.1.1.1 allowas-in
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

A router discards BGP network prefixes if it sees its ASN in AS-Path as a loop prevention mechanism. The “allowas-in” feature allows routes to be received and processed even if router detects its own ASN in AS-Path.

Question: 228

Refer to the exhibit.

```
ip address 10.0.0.1 255.255.255.0
```

```
interface FastEthernet W Description ***** WAN link ***** ip address  
10.0.0.1 255.255.255.0
```

```
interface FastEthernet!/1  
Description ***** LAN Network ***** ip address 192.168.1.1  
255.255.255.0
```

```
router ospf 1  
router-id 4.4.4.4  
log-adjacency-changes  
network 4.4.4.4 0.0.0.0 area 0  
network 10.0.0.1 0.0.0.0 area 0  
network 192.168.1.1 0.0.0.0 area 10
```

A)

```
interface loopback0  
ip address 4.4.4.4 255.255.255.0  
ip ospf network broadcast
```

B)

```
interface loopback0  
ip address 4.4.4.4 255.255.255.0  
ip ospf interface type network
```

C)

```
interface loopback0  
ip address 4.4.4.4 255.255.255.0  
ip ospf network point-to-point
```

D)

```
interface loopback0
ip address 4.4.4 4 255.255.255.0
ip ospf interface area 10
```

- A. Option
- B. Option
- C. Option
- D. Option

Answer: A

Explanation:

Question: 229

Refer to the exhibit.

```
P 172.29.0.0/16, 1 successors, FD is 307200, serno 2
via 192.168.254.2 (307200/281600), FastEthernet0/1
via 192.168.253.2 (410200/352300), FastEthernet0/0
```

When the FastEthernet0/1 goes down, the route to 172.29.0.0/16 via 192.168.253.2 is not installed in the RIB. Which action resolves the issue?

- A. Configure reported distance greater than the feasible distance
- B. Configure feasible distance greater than the successor's feasible distance.
- C. Configure reported distance greater than the successor's feasible distance.
- D. Configure feasible distance greater than the reported distance

Answer: D**Explanation:**

From the exhibit, we notice network 172.29.0.0/16 was learned via two routes:

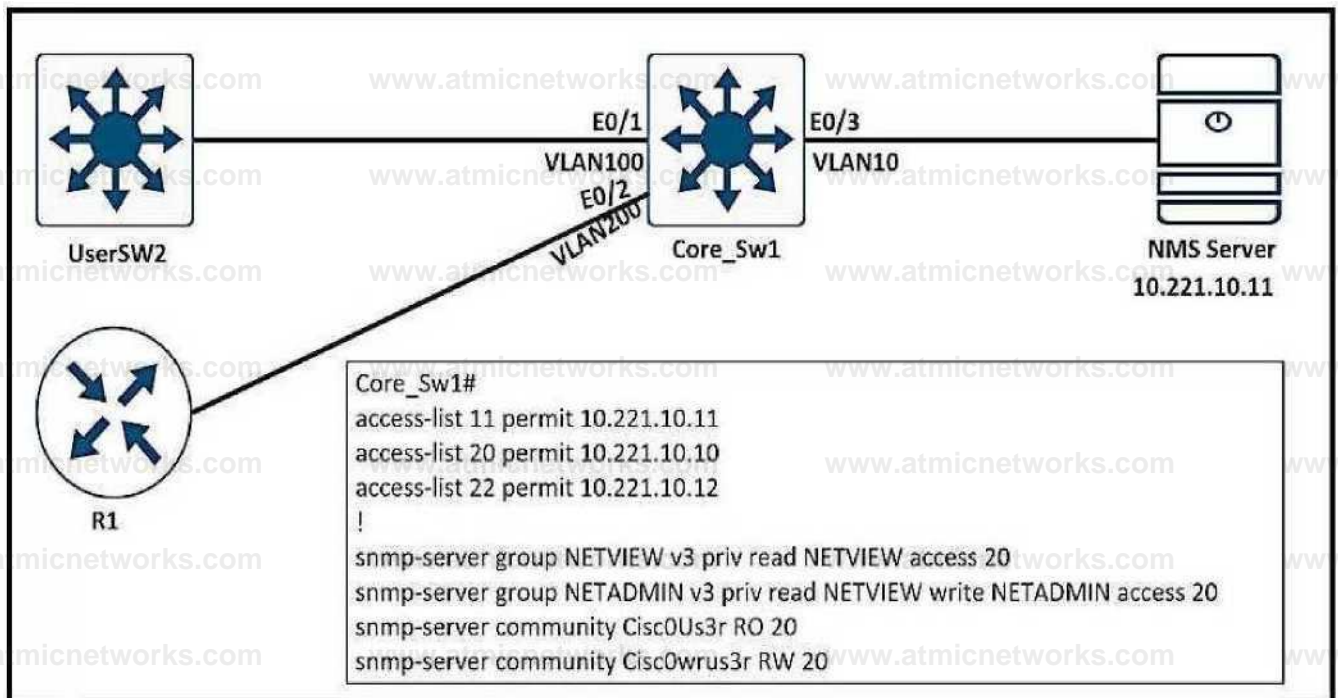
- + From 192.168.254.2 with FD = 307200 and AD = 281600
- + From 192.168.253.2 with FD = 410200 and AD = 352300

The first route is installed into the RIB as the successor route because of lower FD.

When the first route fails, router will not use the second route as it does not satisfy the feasibility condition. The feasibility condition states that, the Advertised Distance (AD, also called the reported distance) of a route must be lower than the feasible distance of the current successor route.

Question: 230

Refer to the exhibit.



- A. access-list 20 permit 10.221.10.12
- B. snmp-server group NETVIEW v2c priv read NETVIEW access 20

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

D. access-list 20 permit 10.221.10.11

Answer: D

Explanation:

Question: 231

IPv6 is enabled in the infrastructure to support customers with an IPv6 network over WAN and to connect the head office to branch offices in the local network. One of the customers is already running IPv6 and wants to enable IPv6 over the DMVPN network infrastructure between the headend and branch sites. Which configuration command must be applied to establish an mGRE IPv6 tunnel neighborship?

A. tunnel protection mode ipv6

B. ipv6 unicast-routing

C. ipv6 nhrp holdtime 30

D. tunnel mode gre multipoint ipv6

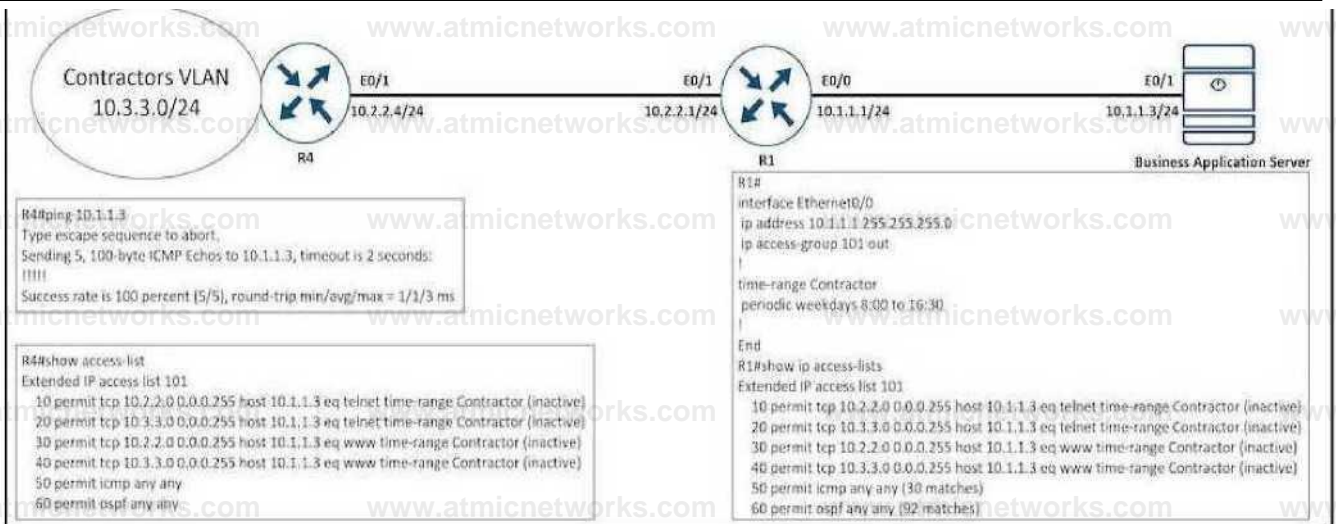
Answer: D

Explanation:

The command "tunnel mode gre multipoint ipv6" sets the encapsulation mode of the tunnel to mGRE IPv6.

Question: 232

Refer to the exhibit.



An engineer is troubleshooting failed access by contractors to the business application server via Telnet or HTTP during the weekend. Which configuration resolves the issue?

A)

R1
time-range Contractor
no periodic weekdays 8:00 to 16:30
periodic daily 8:00 to 16:30

B)

R4
time-range Contractor
no periodic weekdays 17:00 to 23:59
periodic daily 8:00 to 16:30

C)

R4
no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

D)

R1
no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

A. Option

B. Option

c. Option

d. Option

Answer: A

Explanation:

Question: 233

Refer to the exhibit.

```
Route-map PBR, permit, sequence 10
Match clauses:
  ip address (access-lists): FILTER_ACL
Set clauses:
  ip next-hop verify-availability 209.165.202.129 1 track 100 [down]
  ip next-hop verify-availability 209.165.202.131 2 track 200 [up]
Policy routing matches: 0 packets, 0 bytes
route-map PBR, deny, sequence 20
Match clauses:
Set clauses:
  ip next-hop 209.165.201.30
Policy routing matches: 275364861 packets, 12200235037 bytes
```

An engineer has configured policy-based routing and applied the configured to the correct interface. How is the configuration applied to the traffic that matches the access list?

- A. It is sent to 209.165.202.131.
- B. It is sent to 209.165.202.129.
- C. It is dropped.
- D. It is forwarded using the routing table lookup.

Answer: A

Explanation:

The set ip next-hop verify-availability command in route-map configuration mode to configure policy routing to verify the reachability of the next hop of a route map before the router performs policy routing to that next hop. In this question we see track 100 is down so the PBR will not use this next-hop, it will use the second next-hop with track 200 (up).

Question: 234

How is VPN routing information distributed in an MPLS network?

- A. The top level of the customer data packet directs it to the correct CE device
- B. It is established using VPN IPsec peers.
- C. It is controlled using of VPN target communities.
- D. It is controlled through the use of RD.

Answer: C

Explanation:

The distribution of virtual private network (VPN) routing information is controlled through the use of VPN route target communities, implemented by Border Gateway Protocol (BGP) extended communities.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-cfg-layer3-vpn.html

Question: 235

Which mechanism must be chosen to optimize the reconvergence time for OSPF at company location 407173257 that is less CPU-intensive than reducing the hello and dead timers?

- A. BFD
- B. Dead Peer Detection keepalives
- C. SSO
- D. OSPF demand circuit

Answer: A

Explanation:

Question: 236

A network administrator performed a Compact Flash Memory upgrade on a Cisco Catalyst 6509 Switch. Everything is functioning normally except SNMP, which was configured to monitor the bandwidth of key interfaces but the interface indexes are changed. Which global configuration resolves the issue?

- A. snmp-server ifindex permanent
- B. snmp ifindex permanent
- C. snmp-server ifindex persist
- D. snmp ifindex persist

Answer: C

Explanation:

The SNMP ifIndex persistence feature provides an interface index (ifIndex) value that is retained and used when the router reboots. The ifIndex value is a unique identifying number associated with a physical or logical interface. In the following example, SNMP ifIndex persistence is enabled for all interfaces:

```
router(config)# snmp-server ifindex persist
```

Question: 237

Refer to the exhibit.

- * Sep 26 19:50:43.504: SNMP: Packet received via UDP from 192.168.1.2 on

GigabitEthernetO/1SrParseV3SnmpMessage: No matching Engine ID.

SrParseV3SnmpMessage: Failed.

SrDoSnmp: authentication failure, Unknown Engine ID

* Sep 26 19:50:43.504: SNMP: Report, reqid 29548, errstat 0, erridx 0
internet.6.3.15.1.1.4.0 = 3

* Sep 26 19:50:43.508: SNMP: Packet sent via UDP to 192.168.1.2
process_mgmt_req_int: UDP packet being de-queued

Which two commands provide the administrator with the information needed to resolve the issue? (Choose two.)

- A. Show snmp user
- B. debug snmp engine-id
- C. debug snmpv3 engine-id
- D. debug snmp packet
- E. showsnmpv3 user

Answer: A,D

Explanation:

There are 3 values in the SNMPv3 header that must match for the communication to take place: snmpEngineID, snmpEngineTime, snmpEngineBoots. The error received indicates a problem with the EngineID value: "authentication failure, Unknown Engine ID"

To specify the Engine ID, we can use the command "show snmp user". The following example specifies the username as abcd with Engine ID: 0000000902000000C025808:

```
Router#show snmp user abed
User name: abed
Engine ID: 00000009020000000C025808
storage-type: nonvolatile active access-list: 10
Rowstatus: active
Authentication Protocol: MD5
Privacy protocol: 3DES
Group name: VacmGroupName
Group name: VacmGroupName
```

The “debug snmp packet” command displays all SNMP packets that are arriving and being replied to.

Question: 238

Refer to the exhibit.

```
*Sep 26 19:50:43.504: SNMP: Packet received via UDP from 192.168.1.2 on
GigabitEthernetO/1SrParseV3SnmpMessage: No matching Engine ID.
```

```
SrParseV3SnmpMessage: Failed.
```

```
SrDoSnmp: authentication failure, Unknown Engine ID
```

```
*Sep 26 19:50:43.504: SNMP: Report, reqid 29548, errstat 0, erridx 0
internet.6.3.15.1.1.4.0 = 3
```

```
*Sep 26 19:50:43.508: SNMP: Packet sent via UDP to 192.168.1.2
process_mgmt_req_int: UDP packet being de-queued
```

Which two commands provide the administrator with the information needed to resolve the issue? (Choose two.)

- A. snmp user
- B. debug snmp engine-id

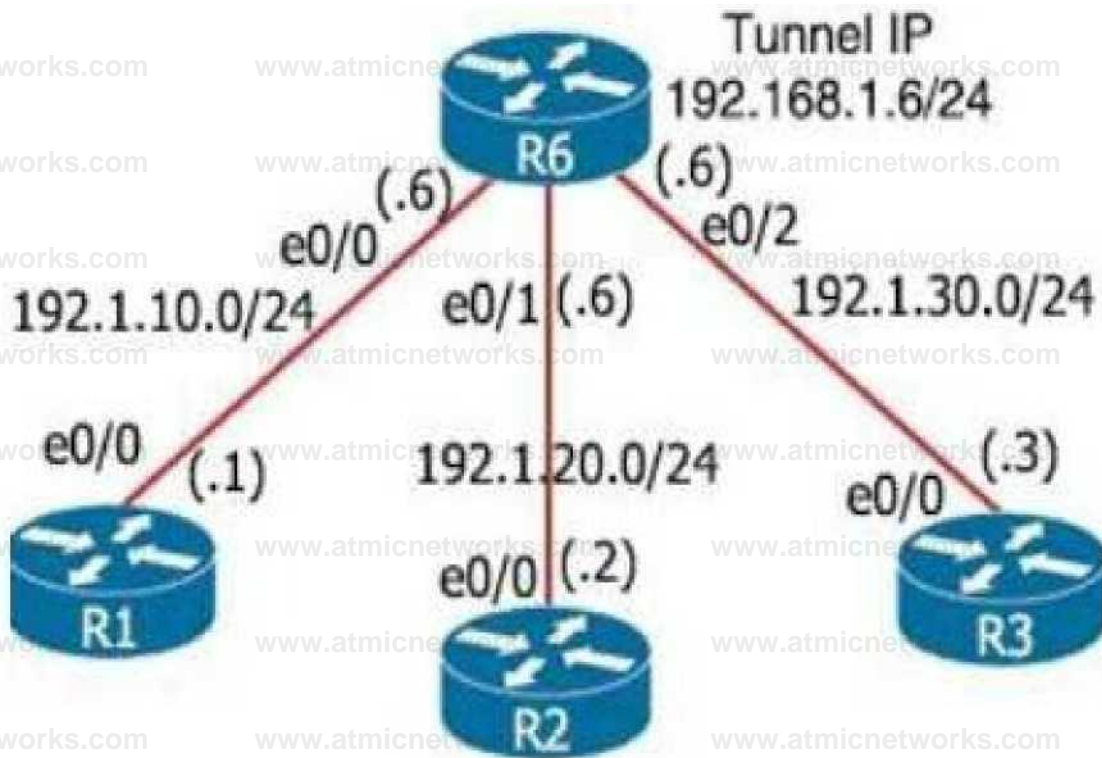
- C. debug snmpv3 engine-id
- D. debug snmp packet
- E. showsnmpv3 user

Answer: A,E

Explanation:

Question: 239

Refer to the exhibit.



An engineer must establish multipoint GRE tunnels between hub router R6 and branch routers R1, R2, and R3. Which configuration accomplishes this task on R1?

A)

```
interface Tunnel 1
ip address 192.168.1.1 255.2 55.255.0
tunnel source e0/1
tunnel mode gre multipoint
ip nhrp nhs 192.168.1.6
ip nhrp map 192.168.1.6 192.1.10.6
```

B)

```
interface Tunnel 1
ip address 192.168.1.1 255.255.255.0
tunnel source e0/1
tunnel mode gre multipoint
ip nhrp network-id 1
ip nhrp nhs 192.168.1.6
ip nhrp map 192.168.1.6 192.1.10.1
```



```
ip nhrp map 192.168.1.2 192.1.20.2  
ip nhrp map 192.168.1.3 192.1.30.3
```

C)

```
interface Tunnel 1  
ipaddress 192.168.1.1 255.2 55.255.0  
tunnel source e0/0  
tunnel mode gre multipoint  
ip nhrp nhs 192.168.1.6  
ip nhrp map 192.168.1.6 192.1.10.1  
ip nhrp map 192.168.1.2 192.1.20.2  
ip nhrp map 192.168.1.3 192.1.30.3
```

D)

```
interface Tunnel 1  
ipaddress 19 2.168.1.1 255.2 55.255.0  
tunnel source e0/0  
tunnel mode gre multipoint  
ip nhrp network-id 1  
ip nhrp nhs 192.168.1.6  
ip nhrp map 192.168.1.6 192.1.10.6
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

We have an example of how to configure DMVPN Phase II and we show the configuration here for your reference:

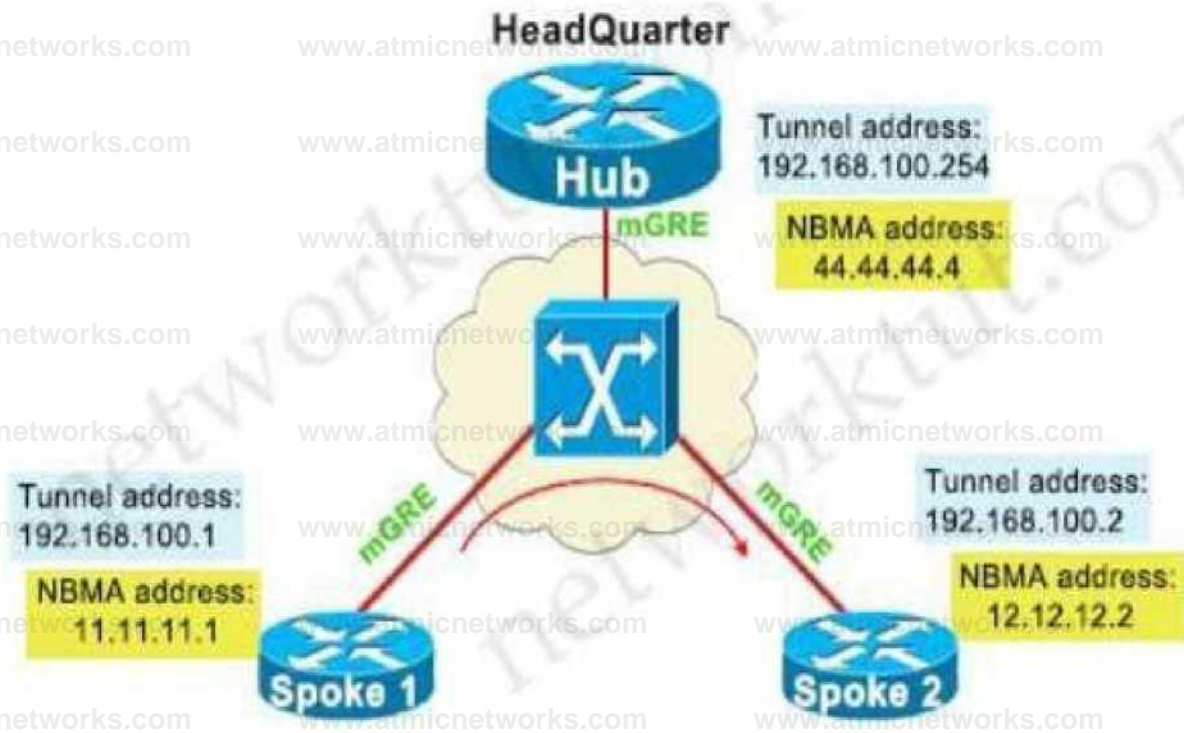


Diagram Description automatically generated

DMVPN Phase II – Dynamic Mapping

Hub

```
interface tunnel 1
ip address 192.168.100.254
255.255.255.0
tunnel source 44.44.44.4
tunnel mode gre multipoint ip
nhrp network 10
```

Spoke 1

```
interface tunnel 1
ip address 192.168.100.1
255.255.255.0
tunnel source 11.11.11.1
tunnel mode gre multipoint
ip nhrp network 10
ip nhrp map 192.168.100.254
44.44.44.4
ip nhrp nhs 192.168.100.254
```

Spoke 2

```
interface tunnel 1
ip address 192.168.100.2
255.255.255.0
tunnel source 12.12.12.2
tunnel mode gre multipoint
ip nhrp network 10
ip nhrp map 192.168.100.254
44.44.44.4
ip nhrp nhs 192.168.100.254
```

Text Description automatically generated

Note: Although Phase II – Dynamic Mapping is “dynamic” but we still need to add a static entry for the hub because without that entry, the NHRP registration cannot be sent.

Question: 240

Refer to the exhibit.

```
interface loopback0 ip address 4.4.4.4 255.255.255.0
```

```
interface FastEthernet1/0
```

```
Description "WAN link" ip address 10.0.0.1 255.255.255.0
```

```
interface FastEthernet0/1
```

```
Description "LAN Network"
```

```
ip address 192.168.1.1 255.255.255.0
```

```
router ospf 1
```

```
router-id 44.4.4
```

```
log-adjacency-changes
```

```
network 4.4.4.0 0.0.0.0 area 0
```

```
network 10.0.0.1 0.0.0.0 area 0 network 192.168.1.0 0.0.0.0 area 10
```

Which set of commands restore reachability to loopback0?

A)

```
interface loopback0 ip address 4.4.4.4 255.255.255.0 (no ospf redistribute ip in HO-p0)
```

B)

```
interface loopback0 ip address 4.4.4.4 255.255.255.0 no ospf redistribute ip in HO-p0
```

C)

```
interface loopback0
ip address 4.4.4.4 255.255.255.0
ip ospf interface area 10
```

D)

```
interface loopback0
ip address 4.4.4.4 255.255.255.0
ip ospf interface type network
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

We tested this config in GNS3 (except the LAN interface) but R1 loopback0 was advertised normally on R2 and R2 could reach this loopback0.

```
R1#sh run [ b interface interface Loopback0
ip address 4.4.4.4 255.255.255,0
interface FastEthernet0/0
```

```
ip address 10.0.0.1 255.255.255.0
duplex auto
speed auto

router ospf 1
 log-adj acency-change s
 network 4.4.4.4 0.0.0.0 area 0
 network 10.0.0.1 0.0.0.0 area 0
```

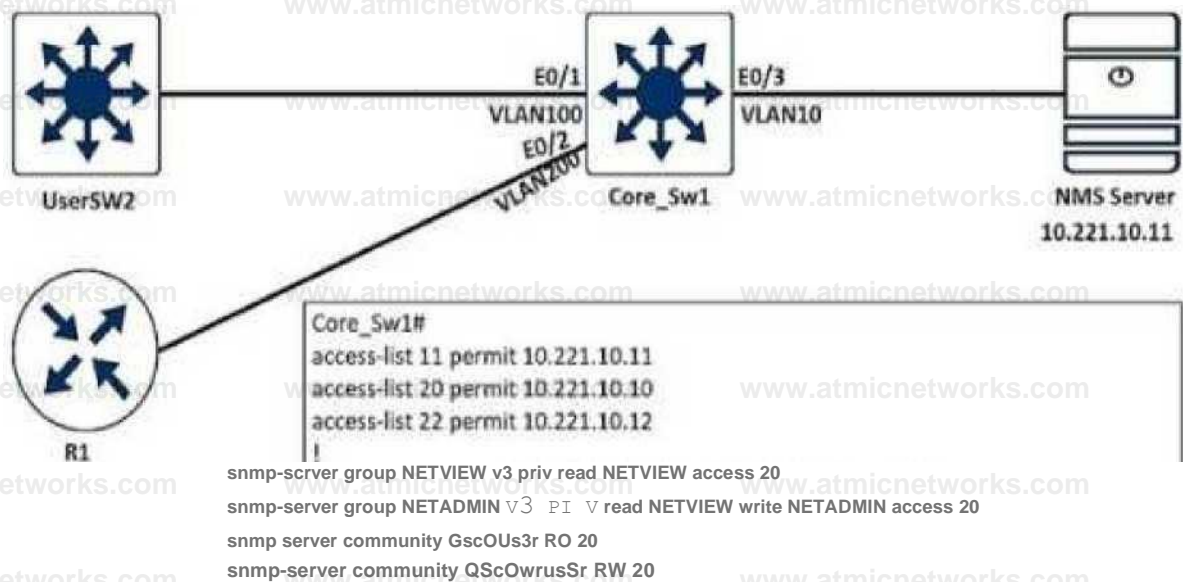
```
R2#sh ip route ospf
 4.0.0.0/32 is subnetted, 1 subnets
 0      4.4.4.4 [110/2] via 10.0.0.1, 00:41:03, FastEthernet0/0
R2#ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds: ! I!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 24/34/56 ms
```

Note: Although the configured loopback address is 4.4.4.4/24 but by default OSPF will advertise this route to loopback0 as 4.4.4.4/32 (most specific route to that loopback). In order to override this, we have to change the network type to point-to-point. After this OSPF will advertise the address to loopback as 4.4.4.0/24.

Question: 241

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 NETVIEW v2c priv read NETVIEW access 20
- B. access-list 20 permit 10.221.10.11
- C. access-list 20 permit 10.221.10.12
- D. snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 22

Answer: B

Explanation:

Question: 242

What is a characteristic of Layer 3 MPLS VPNs?

- A. LSP signaling requires the use of unnumbered IP links for traffic engineering.
- B. Traffic engineering supports multiple IGP instances

- C. Traffic engineering capabilities provide QoS and SLAs.
- D. Authentication is performed by using digital certificates or preshared keys.

Answer: C

Explanation:

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/15-mt/mp-te-diffserv-15-mt-book/mp-te-diffserv-aw.html

MPLS traffic engineering supports only a single IGP process/instance

The MPLS traffic engineering feature does not support routing and signaling of LSPs over unnumbered

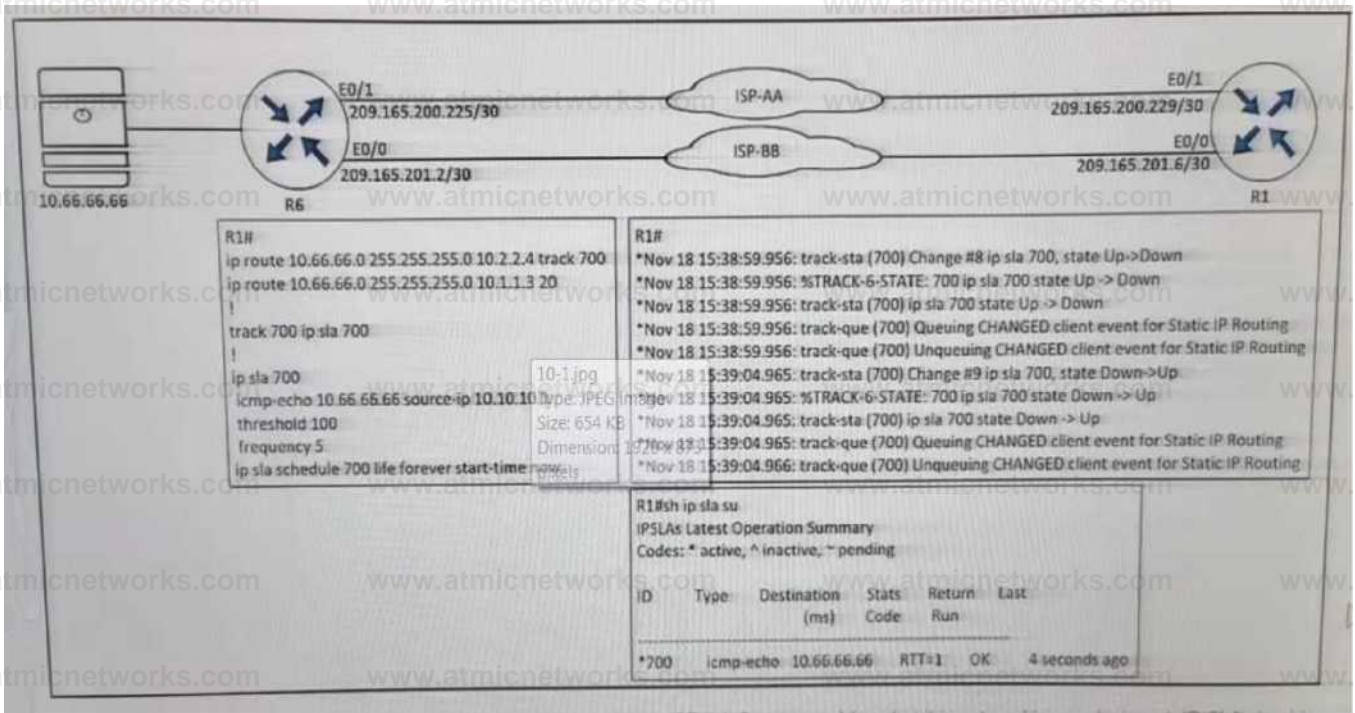
IP links.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_path_setup/configuration/xen-3s/mp-te-path-setup-xe-3s-book/mp-te-enhance-xe.html

3s/mp-te-path-setup-xe-3s-book/mp-te-enhance-xe.html

Question: 243

Refer to the exhibit.



An engineer configured IP SLA on R1 to avoid the ISP link flapping problem. but it is not working as designed IP SLA should wait 30 seconds before switching traffic to a secondary connection and then revert to the primary link after waiting 20 seconds, when the primary link is available and stabilized. Which configuration resolves the issue?

- A. R1(config)#ip sla 700R1(config-ip-sla)#delay down 30 up 20
- B. R1(config)#ip sla 700R1(config-ip-sla)#delay down 20 up 30
- C. R1(config)#track 700 ip sla 700R1(config-track)#delay down 30 up 20
- D. R1(config)#track 700 ip sla 700R1(config-track)#delay down 20 up 30

Answer: C

Explanation:


“wait 30 seconds before switching traffic to a secondary connection” -> delay down 30

“then revert to the primary link after waiting 20 seconds” -> up 20

Under the track object, you can specify delays so we have to configure delay under “track 700 ip sla 700” (not under “ip sla 700”).

Question: 244

Refer to the exhibit.



- Mar 1 17 19 04 051 %OSPF-5ADJCHG Process 100 Nbr 1 1 1 1 on TunneHOO from LOADING to FULL, Loading Done
- Mar 1 17 19 06 375 %OSPF-5 ADJCHG Process 100 Nbr 1111 on Tunnel100 from FULL to DOWN Neighbor Down Adjacency forced to reset
- Mar 1 17 19 06 627 %OSPF-5ADJCHG Process 100 Nbr 2 2 2 2 on Tunnel100 from LOADING to FULL, Loading Done
- Mar 1 17 19 10 123 %OSPF-5-ADJCHG. Process 100 Nbr 2 2 2 2 on TunneHOO from FULL to DOWN Neighbor Down Adjacency forced to reset
- Mar 1 17 19 14 499 %OSPF-5 ADJCHG Process 100 Nbr 101010 10 on TunneHOO from LOADING to FULL. Loading Done
- Mar 1 17 19 19 139 %OSPF-5 ADJCHG Process 100 Nbr 10 10 10 10 on TunneHOO from EXSTART to DOWN Neighbor Down Interface down or detached
- Mar 1 17 01 51 975 %OSPF 4 NONCIGHDOR Received database description from unknown neighbor 192168 1 1
- Mar 1 17 01 57 783 OSPF Rev LS UPD from 192 168 1 1 on TunneHOO length 88 LSA count 1
- Mar 1 17 01 57 155 OSPF Send UPD to 10 255 253 1 on TunneHOO length 100 LSA count 2

A network administrator sets up an OSPF routing protocol for a DMVPN network on the hub router.

Which configuration required to establish a DMVPN tunnel with multiple spokes?

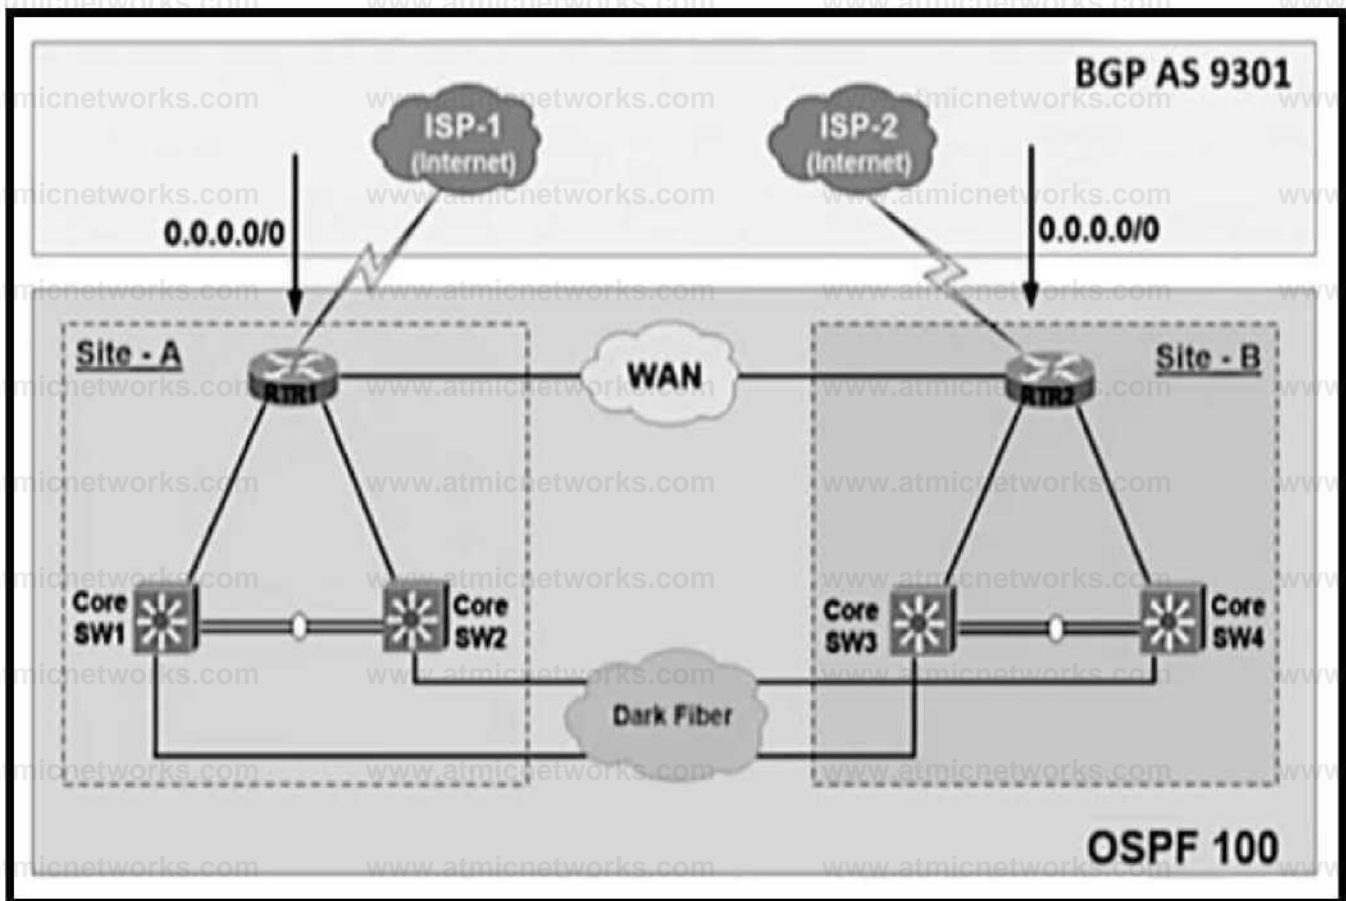
- A. ip ospf network point-to-multipoint on both spoke routers
- B. ip ospf network point-to-point on the hub router
- C. ip ospf network point-to-multipoint on One spoke router
- D. ip ospf network point-to-point on both spoke routers

Answer: A

Explanation:

Question: 245

Refer to the exhibit.



The Internet traffic should always prefer Site-A ISP-1 if the link and BGP connection are up; otherwise, all Internet traffic should go to ISP-2. Redistribution is configured between BGP and OSPF routing protocols and it is not working as expected.

What action resolves the issue?

- A. Set metric-type 2 at Site-A RTR1, and set metric-type 1 at Site-B RTR2
- B. Set OSPF cost 100 at Site-A RTR1, and set OSPF Cost 200 at Site-B RTR2
- C. Set OSPF cost 200 at Site: A RTR1 and set OSPF Cost 100 at Site-B RTR2
- D. Set metric-type 1 at Site-A RTR1, and set metric-type 2 at Site-B RTR2

Answer: D

Explanation:

OSPF type 1 route is always preferred over a type 2 route for the same destination so we can set metric-type 1 at Site-A

RTR1 so that it is preferred over Site-B RTR2.

Note:

Routes are redistributed in OSPF as either type 1 (E1) routes or type 2 (E2) routes, with type 2 being the default.

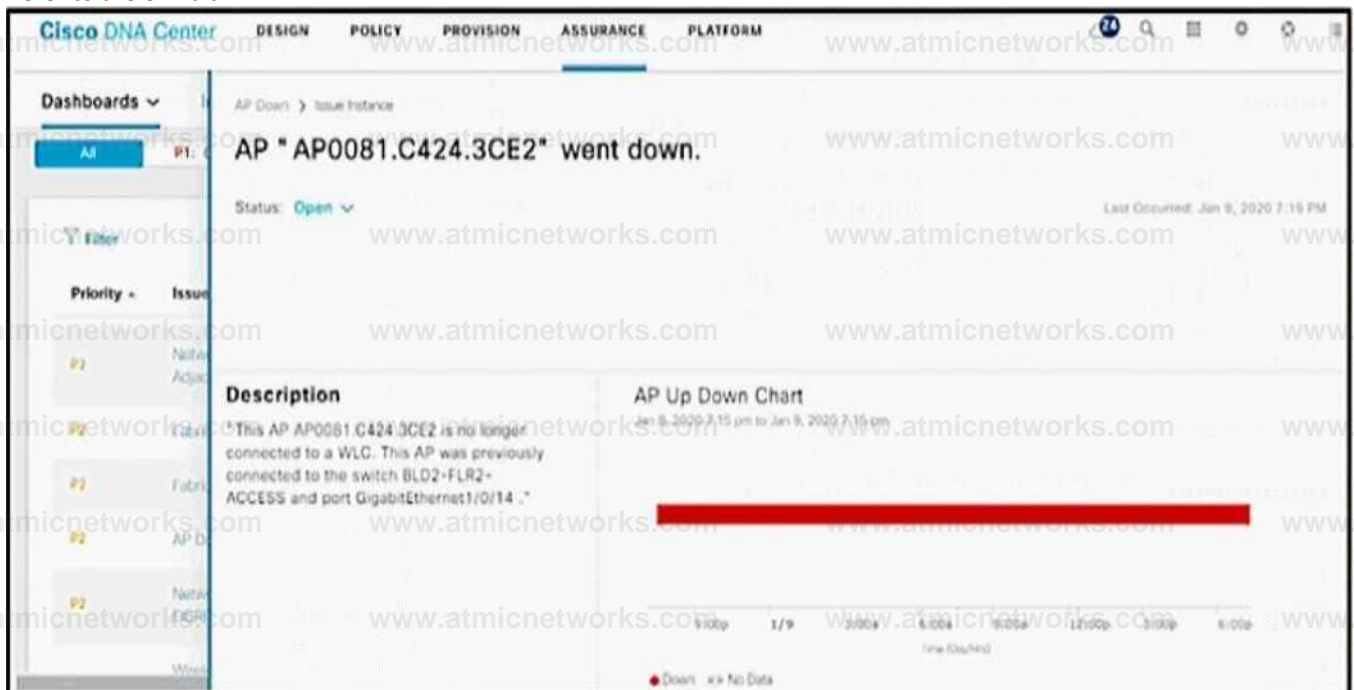
- A type 1 route has a metric that is the sum of the internal OSPF cost and the external redistributed cost.
- A type 2 route has a metric equal only to the redistributed cost.
- If routes are redistributed into OSPF as type 2 then every router in the OSPF domain will see the same cost to reach

the external networks.

- If routes are redistributed into OSPF as type 1, then the cost to reach the external networks could vary from router to router.

Question: 246

Refer to the exhibit.



The AP status from Cisco DNA Center Assurance Dashboard shows some physical connectivity issues from access switch interface G1/0/14. Which command generates the diagnostic data to resolve the physical connectivity issues?

- A. test cable-diagnostics tdr interface GigabitEthernet1/0/14
- B. Check cable-diagnostics tdr interface GigabitEthernet1/0/14
- C. show cable-diagnostics tdr interface GigabitEthernet1/0/14
- D. Verify cable-diagnostics tdr interface GigabitEthernet1/0/14

Answer: A

Explanation:

The Time Domain Reflectometer (TDR) feature allows you to determine if a cable is OPEN or SHORT when it is at fault.

To start the TDR test, perform this task:

Step 1 (Starts the TDR test): test cable-diagnostics tdr {interface {interface-number}}

Step 2 (Displays the TDR test counter information): show cable-diagnostics

tdr {interface interface-number}

<https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/16->

[11/configuration_guide/int_hw/b_1611_int_and_hw_9600_cg/checking_port_status_and_connectivity.pdf](https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/16-11/configuration_guide/int_hw/b_1611_int_and_hw_9600_cg/checking_port_status_and_connectivity.pdf)

TDR test stalled on interface Gi 1/0/14

A TDR test can take a few seconds to inn on an interface

Use 'show cable-diagnostics tdf to read the TDR results.

Wait 10 seconds and then issue the command to show the cable diagnostics result:

```
TDR test last run on: December 05 18:50:53
```

```
Interface Speed Local pair Pair length Remote pair Pair status
```

```
Gil1/0/14s1000M Pair A 19 +/- 10 meters Pair B Normal
```

```
Pair B 19 +/- 10 meters Pair A Normal
```

Pair C 19 +/- 10 meters Pair D Normal

Pair D 19 +/- 10 meters Pair C Normal

Notice that the results are "Normal" in the above example. Other results can be: + Open: Open circuit. This means that one (or more) pair has "no pin contact". + Short: Short circuit.
+ Impedance Mismatched: Bad cable.)

Text, table Description automatically generated

Question: 247

An engineer creates a Cisco DNA Center cluster with three nodes, but all the services are running on **one host node**. Which action resolves this issue?

- A. Restore the link on the switch interface that is connected to a cluster link on the Cisco DNA Center
- B. Click the master host node with all the services and select services to be moved to other hosts
- C. Enable service distribution from the Systems 360 page.
- D. Click system updates, and upgrade to the latest version of Cisco DNA Center.

Answer: C

Explanation:

To deploy Cisco DNA Center on a three-node cluster with High Availability (HA) enabled, complete the following procedure:

Step 1: Configure Cisco DNA Center on the first node in your cluster...

Step 2: Configure Cisco DNA Center on the second node in your cluster...

Step 3: Configure Cisco DNA Center on the third node in your cluster...

Step 4: Enable high availability on your cluster:

- a. In the Cisco DNA Center GUI, click and choose System Settings. The System 360 tab is displayed by default.
- b. In the Hosts area, click Enable Service Distribution.

After you click Enable Service Distribution, Cisco DNA Center enters into maintenance mode. In this mode, Cisco DNA Center is unavailable until the redistribution of services is completed. You should take this into account when scheduling an HA deployment.

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3-3-0/ha_guide/b_cisco_dna_center_ha_guide_1_3_3_0.html

Therefore we can choose "Enable Service Distribution" to distribute services to other host nodes.

Question: 248

R1 and R2 are configured as eBGP neighbor, R1 is in AS100 and R2 is in AS200. R2 is advertising these networks to R1:



The network administrator on R1 must improve convergence by blocking all subnets of 172.16.0.0/16 major network with a mask lower than 23 from coming in, Which set of configurations accomplishes the task on R1?

- A. `ip prefix-list PL-1 deny 172.16.0.0/16 le 23`
`ip prefix-list PL-1 permit 0.0.0.0/0 le 32`
`router bgp 100`
`neighbor 192.168.100.2 remote-as 200`
`neighbor 192.168.100.2 prefix-list PL-1 in`
- B. `ip prefix-list PL-1 deny 172.16.0.0/16 ge 23`
`ip prefix-list PL-1 permit 0.0.0.0/0 le 32`
`router bgp 100`
`neighbor 192.168.100.2 remote-as 200`
`neighbor 192.168.100.2 prefix-list PL-1 in`
- C. `access-list 1 deny 172.16.0.0 0.0.254.255`
`access-list 1 permit any`
`router bgp 100`
`neighbor 192.168.100.2 remote-as 200`
`neighbor 192.168.100.2 distribute-list 1 in`
- D. `ip prefix-list PL-1 deny 172.16.0.0/16`
`ip prefix-list PL-1 permit 0.0.0.0/0`
`router bgp 100`
`neighbor 192.168.100.2 remote-as 200`
`neighbor 192.168.100.2 prefix-list PL-1 in`

Answer: A**Explanation:**

“Blocking all subnets of 172.16.0.0/16 major network with a mask lower than 23 from coming in”

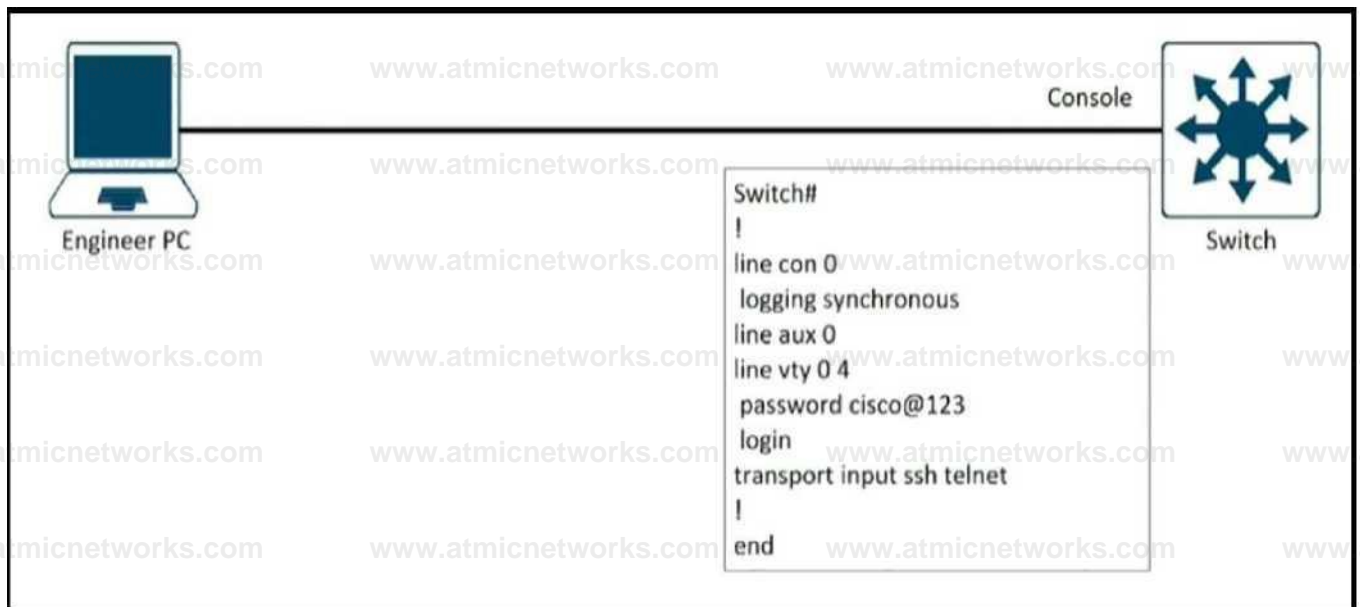
would block 172.16.16.0/20.

The first prefix-list “ip prefix-list PL-1 deny 172.16.0.0/16 le 23” means “all networks that fall within the 172.16.0.0/16 range AND that have a subnet mask of /23 or less” are denied.

The second prefix-list “ip prefix-list PL-1 permit 0.0.0.0/0 le 32” means allows all other prefixes.

Question: 249

Refer to the exhibit.



An engineer must block access to the console ports for all corporate remote Cisco devices based on the recent corporate security policy but the security team still can connect through the console port. Which configuration on the console port resolves the issue?

- A. transport input telnet
- B. login and password
- C. no exec
- D. exec 0.0

Answer: C

Explanation:

“no exec” will disable access to a line. It is used if we want to allow only outgoing session (and disable incoming session) so this command will block all console port access.

There is no “exec 0 0” command. We can only find the “exec prompt” command in IOS Version 15.4(2)T4.

```
Router(config-line)#exec ?
  prompt EXEC prompt
  <cr>

RouterCconfig-line^exec pro
Router(config-line)^exec prompt ?
  timestamp Print rimestamps for show commands

Router(config-line)#exec prompt | _____
```

The most similar command is “exec-timeout 0 0” command, which is used to prevent Telnet/SSH sessions from timing out.

Question: 250

The network administrator configured R1 to authenticate Telnet connections based on Cisco ISE using TACACS+. ISE has been configured with an IP address of 192.168.1.5 and with a network device pointing toward R1(192.168.1.1) with a shared secret password of Cisco123.

```
aaa new-model
```

```
tacacs server tSE1
```

```
address ipv4 192/168.1.5
```

```
key Cisco123
```

```
group server tacacs* TAC-SERV server name ISE1
```

```
aaa authentication login telnet group TAC-SERV
```

The administrator cannot authenticate to R1 based on ISE. Which configuration fixes the issue?

A. ip tacacs-server host 192.168.1.5 key Cisco123

B. line vty 0 4login authentication TAC-SERV

C. line vty 0 4login authentication telnet

D. tacacs-server host 192.168.1.5 key Cisco123

Answer: C

Explanation:

The last command “aaa authentication login telnet group TAC-SERV” created the method list name telnet so we need to assign it to line vty.

Reference: [https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/200208-](https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/200208-Configure-ISE-2-0-IOI-TACACS-Authentic.html)

[Configure-ISE-2-0-IOI-TACACS-Authentic.html](https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/200208-Configure-ISE-2-0-IOI-TACACS-Authentic.html)

Question: 251

Refer to the exhibit.

```
aaa new-model
aaa group server radius RADIUS-SERVERS
aaa authentication login default group RADIUS-SERVERS local
aaa authentication enable default group RADIUS-SERVERS enable
aaa authorization exec default group RADIUS-SERVERS if-authenticated
aaa authorization network default group RADIUS-SERVERS if-authenticated
aaa accounting send stop-record authentication failure
aaa session-id common
!
line con 0
logging synchronous
stopbits 1
line vty 0 4
logging synchronous
transport input ssh
```

A network administrator successfully logs in to a switch using SSH from a RADIUS server. When the network administrator uses a console port to access the switch, the RADIUS server returns shell:priv-lvl=15 and the switch asks to enter the enable command. The command is entered, it gets rejected. Which command set is used to troubleshoot and resolve this issue?

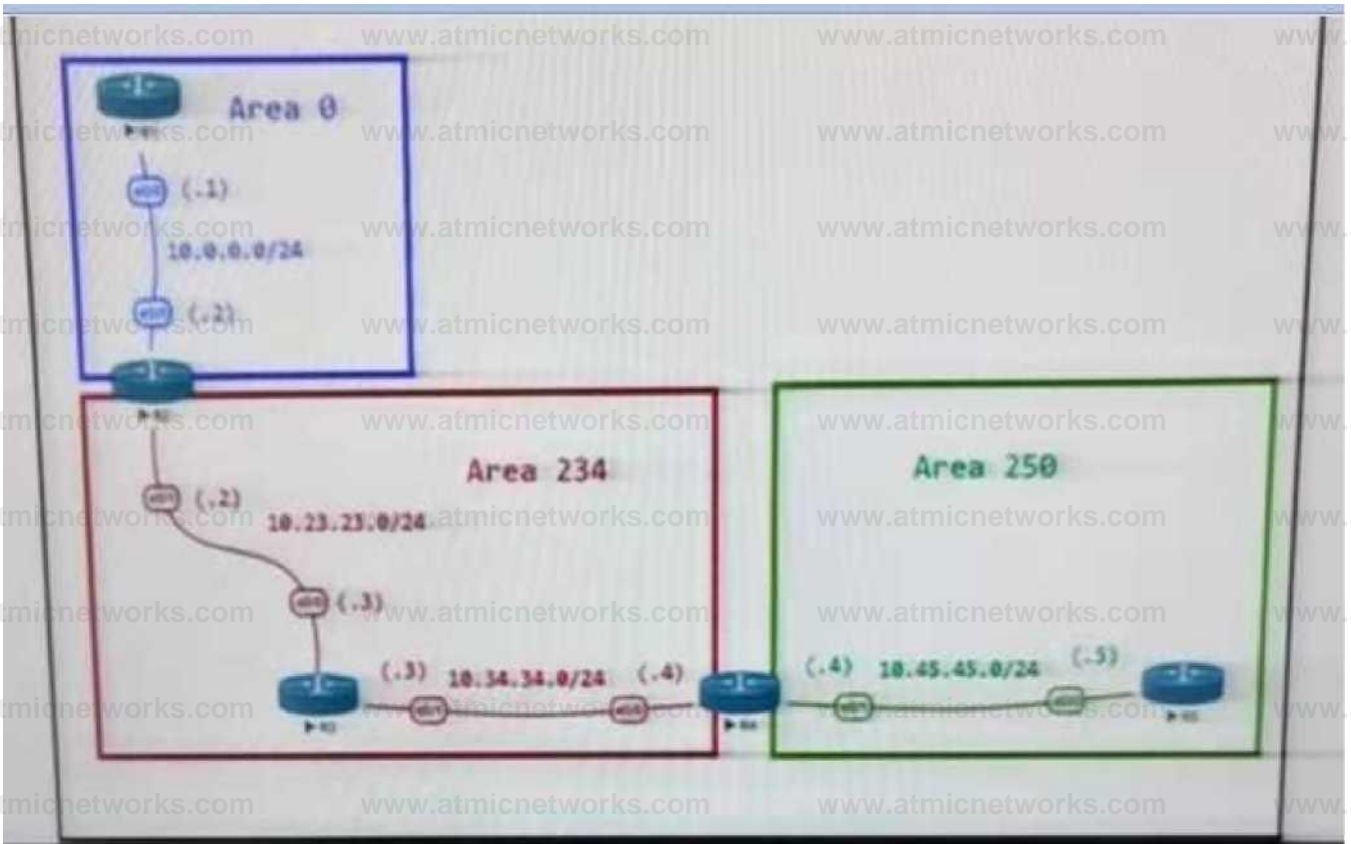
- A. line con 0aaa authorization console!line vty 0 4transport input ssh
- B. line con 0aaa authorization console!line vty 0 4authorization exec
- C. line con 0aaa authorization console!line vty 0 4authorization exec
- D. line con 0aaa authorization console!line vty 0 4transport input ssh

**Answer:
A**

Explanation:

**Question:
252**

Refer to the exhibit.



```

ABR Configurations
R2
router ospf 1
router-id 0.0.0.22
area 234 virtual-link 10.34.34.4
network 10.0.0.0 0.0.0.255 area 0
network 10.2.2.0 0.0.0.255 area 0
network 10.22.22.0 0.0.0.255 area 234
network 10.23.23.0 0.0.0.255 area 234

R4
router ospf 1
router-id 0.0.0.44
area 234 virtual-link 10.23.23.2
network 10.34.34.0 0.0.0.255 area 234
network 10.44.44.0 0.0.0.255 area 234
network 10.45.45.0 0.0.0.255 area 250

Virtual Link Status
R2 -> sh ip ospf virtual-links
Virtual Link OSPF_VL0 to router 10.34.34.4 is down
Run as demand circuit
DoNotAge LSA allowed
Transit area 234
Topology-MTID Cost Disabled Shutdown Topology Name
0 65535 no no Base
Transmit Delay is 1 sec, State DOWN,

```

The network administrator configured the network to connect two disjointed networks and all the connectivity is up except the virtual link which causes area 250 to be unreachable. Which two configurations resolve this issue? (Choose two.)

- A. R2router ospf 1router-id 10.23.23.2
- B. R2router ospf 1no area area 234 virtual-link 10.34.34.4area 0 virtual-link 0.0.0.44
- C. R4router ospf 1no area 234 virtual-link 10.23.23.2area 234 virtual-link 0.0.0.22
- D. R2router ospf 1no area 234 virtual-link 10.34.34.4area 234 virtual-link 0.0.0.44
- E. R4router ospf 1no area area 234 virtual-link 10.23.23.2area 0 virtual-link 0.0.0.22

Answer: C,D

Explanation:

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html>

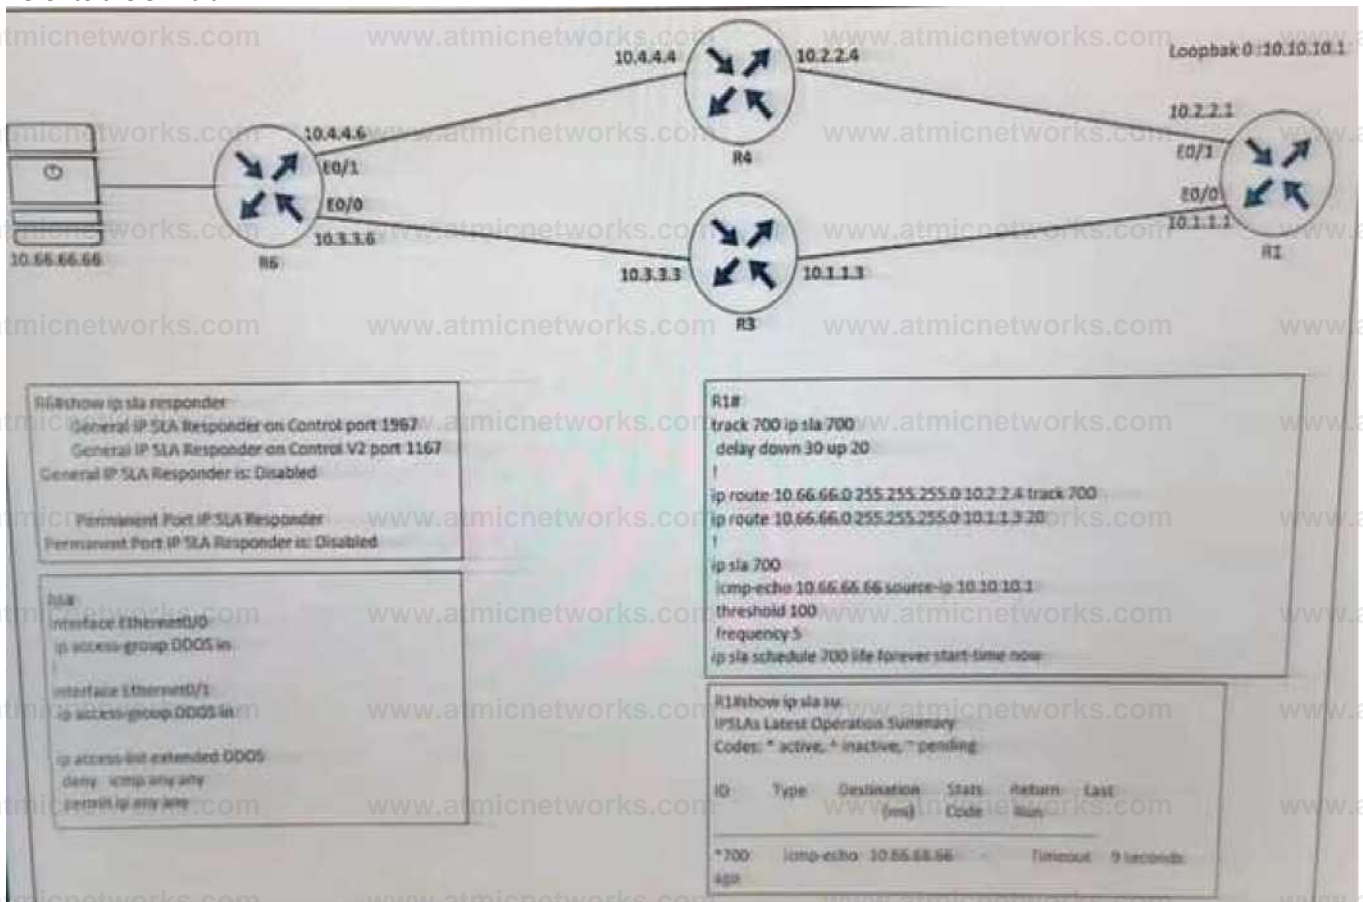
An important thing to remember when configuring virtual-link is we need to configure the OSPF Router

ID and NOT the IP address of the ABR. Therefore in this question we have to use the command "area

234 virtual-link 0.0.0.44" on R2 and "area 234 virtual-link 0.0.0.22" on R4.

Question: 253

Refer to the exhibit.



R1 is configured with IP SLA to check the availability of the server behind R6 but it kept failing. Which configuration resolves the issue?

A. R6(config)# ip sla responder

- B. R6(config)# ip sla responder udp-echo ip address 10.10.10.1 port 5000
- C. R6(config)# ip access-list extended DDOSR6(config ext-nac)# 5 permit icmp host 10.66.66.66 host 10.10.10.1
- D. R6(config)# ip access-list extended DDOSR6(config ext-nac)# 5 permit icmp host 10.10.10.1 host 10.66.66.66

Answer: D

Explanation:

In this IP SLA tracking, we don't need a IP SLA Responder so the command "ip sla responder" on R6 is not necessary.

We also notice that the ACL is blocking ICMP packets on both interfaces E0/0 & E0/1 of R6 so we need to allow ICMP from source 10.10.10.1 to destination 10.66.66.66.

Question: 254

Which mechanism provides traffic segmentation within a DMVPN network?

- A. RSVP
- B. BGP
- C. MPLS
- D. iPsec

Answer: C

Explanation:

To use the DMPVN – Traffic Segmentation Within DMVPN feature you must configure Multiprotocol

Label Switching (MPLS) by using the mpls ip command.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/xen-16/sec-conn-dmvpn-xe-16-book/sec-conn-dmvpn-dmvpn.html

Question: 255

What are two characteristics of IPv6 Source Guard? (Choose two.)

- A. requires IPv6 snooping on Layer 2 access or trunk ports
- B. used in service provider deployments to protect DDoS attacks
- C. requires the user to configure a static binding
- D. requires that validate prefix be enabled
- E. recovers missing binding table entries

Answer: D,E

Explanation:

IPv6 Source Guard uses the IPv6 First-Hop Security Binding Table to drop traffic from unknown sources or bogus IPv6 addresses not in the binding table. The switch also tries to recover from lost address information, querying DHCPv6 server or using IPv6 neighbor discovery to verify the source IPv6 address after dropping the offending packet(s).

Reference: <https://blog.ip-space.net/2013/07/first-hop-ipv6-security-features-in.html>

Question: 256

How does an MPLS Layer 3 VPN differentiate the IP address space used between each VPN?

- A. by RD
- B. by address family
- C. by MP-BGP
- D. by RT

Answer: A

Explanation:

Question: 257

Refer to the exhibit.

```
R1#show ip interface GigabitEthernet0/0 | include drops
0 verification drops
0 suppressed verification drops

R1#show ip interface GigabitEthernet0/1 | include drops
5 verification drops
0 suppressed verification drops
```

R1 is configured with uRPF, and ping to R1 is failing from a source present in the R1 routing table via the GigabitEthernet 0/0 interface. Which action resolves the issue?

- A. Remove the access list from the interface GigabitEthernet 0/0
- B. Modify the uRPF mode from strict to loose
- C. Enable Cisco Express Forwarding to ensure that uRPF is functioning correctly
- D. Add a floating static route to the source on R1 to the GigabitEthernet 0/1 interface

Answer: B

Explanation:

Question: 258

Which OSI model is used to insert an MPLS label?

- A. between Layer 5 and Layer 6
- B. between Layer 1 and Layer 2

- C. between Layer 3 and Layer 4
- D. between Layer 2 and Layer 3

Answer: D

Explanation:

Question: 259

Which function does LDP provide in an MPLS topology?

- A. It enables a MPLS topology to connect multiple VPNs to P routers.
- B. It provides hop-by-hop forwarding in an MPLS topology for LSRs.
- C. It exchanges routes for MPLS VPNs across different VRFs.
- D. It provides a means for LSRs to exchange IP routes.

Answer: B

Explanation:

LDP provides a standard methodology for hop-by-hop, or dynamic label, distribution in an MPLS network by assigning labels to routes that have been chosen by the underlying Interior Gateway Protocol (IGP) routing protocols. The resulting labeled paths, called label switch paths (LSPs), forward label traffic across an MPLS backbone to particular destinations.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/12-4t/mp-ldp-12-4t-book.pdf

Question: 260

An engineer is implementing a coordinated change with a server team. As part of the change, the engineer must configure interface GigabitEthernet2 in an existing VRF "RED" then move the interface to an existing VRF "BLUE" when the server team is ready. The engineer configured interface GigabitEthernet2 in VRF "RED"

```
interface GigabitEthernet2
```

```
description Migration ID: S410A60D0806G06 vrf forwarding RED
```

```
ip address 10.0.0.0 255.255.255.254 negotiation auto
```

Which configuration completes the change?

- A. interface GigabitEthernet2no ip addressvrf forwarding BLUE
- B. interface GigabitEthernet2no vrf forwarding REDvrf forwarding BLUEip address 10.0.0.0 255.255.255.254
- C. interface GigabitEthernet2no vrf forwarding REDvrf forwarding BLUE
- D. interface GigabitEthernet2no ip addressip address 10.0.0.0 255.255.255.254vrf forwarding BLUE

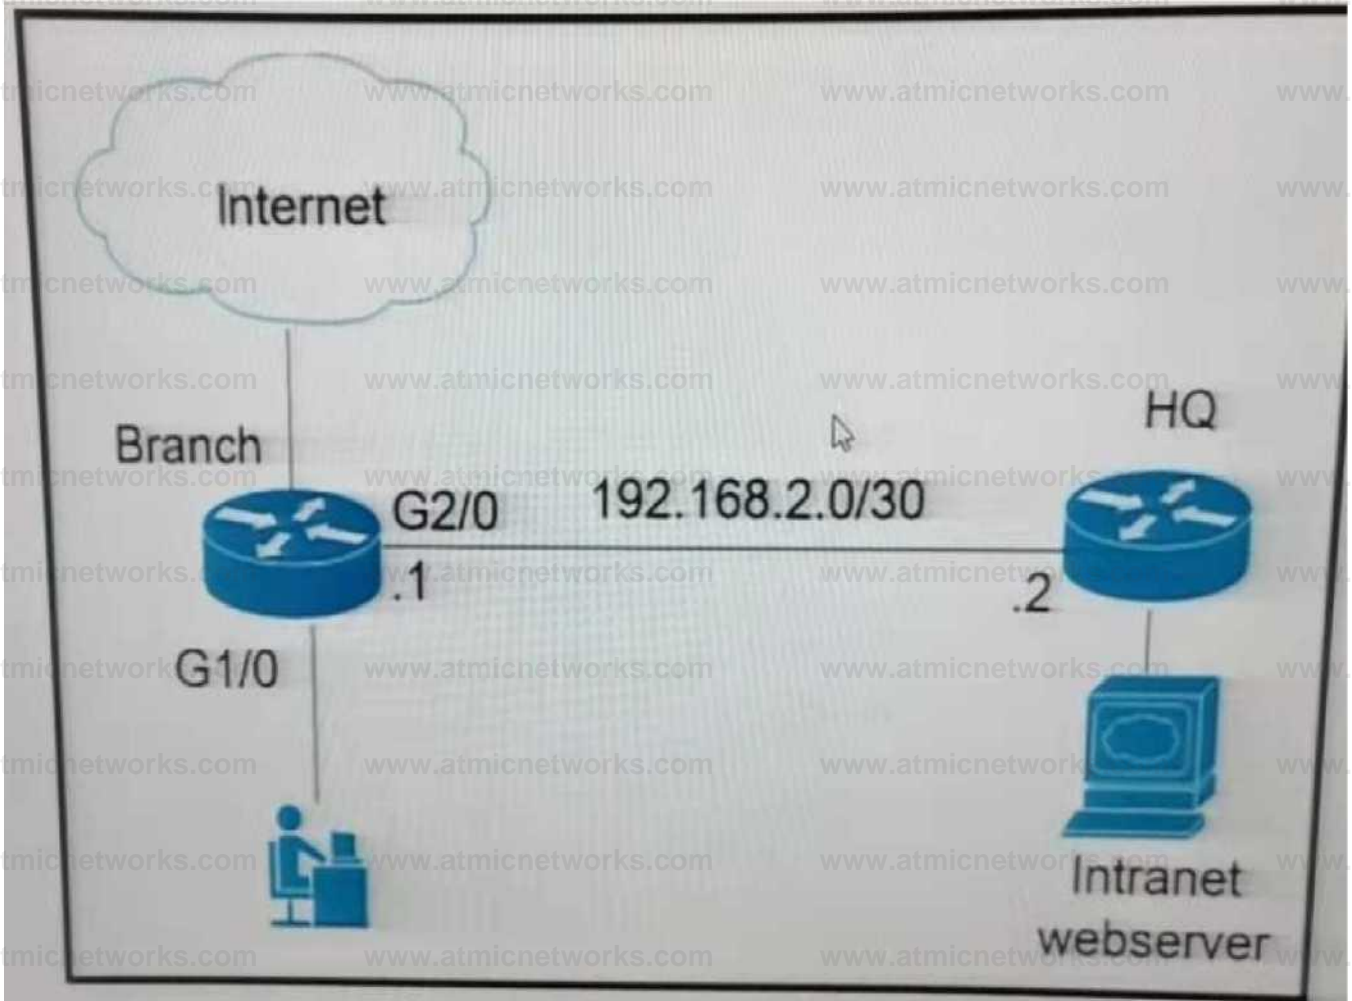
Answer: B

Explanation:

When assigning an interface to a VRF, the IP address will be removed so we have to reassign the IP address to that interface.

Question: 261

Refer to the exhibit.



The branch router is configured with a default route toward the internet and has no routes configured for the HQ site that is connected through interface G2/0. The HQ router is fully configured and does not require changes. Which configuration on the branch router makes the intranet website (TCP port 80) available to the branch office users?

A)

```
access-list 100 permit tcp any host intranet-webserver-ip eq 80
```

```
route-map pbr permit 10  
match ip address 100 set ip next-hop 191168.12
```

interface G10 ip policy route-map pbr

B)

access-list 101 permit tcp any any eq 80

access-list 102 permit tcp any host 192.168.1.2

route-map pbr permit 10

match ip address 101 102 set ip next-hop 192.168.1.2

interface G1 ip policy route-map pbr

C)

access-list 101 permit tcp any any eq 80

access-list 102 permit tcp any host 192.168.1.2

route-map pbr permit 10

match ip address 101

set ip next-hop 192.168.1.2

route-map pbr permit 20

match ip address 102 set ip next-hop 192.168.1.2

interface G10

ip policy route-map pbr

D)

access-list 100 permit tcp host 192.168.1.2 eq 80 any

route-map pbr permit 10

match ip address 100

set ip next-hop 192.168.2.2

interface G1/0

ip policy route-map pbr

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

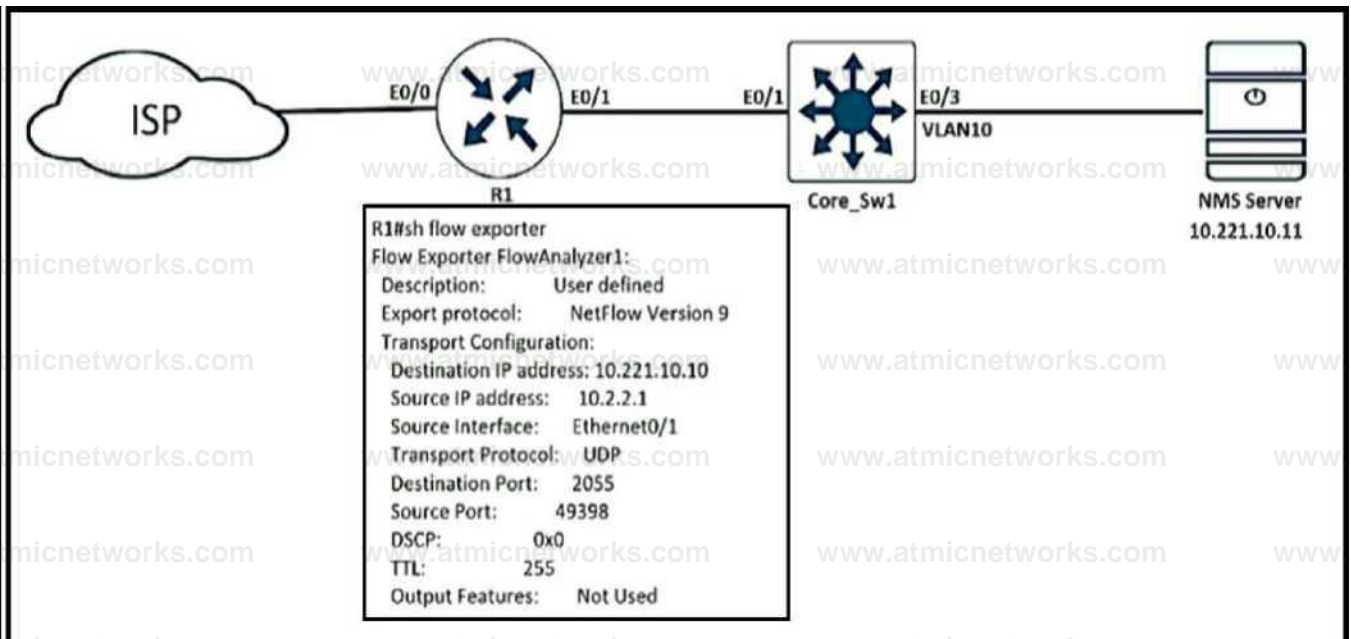
the ACL 101 matches all HTTP packets while the ACL 102 matches TCP packets destined

to Intranet webserver. These packets will be sent to HQ router.

If a match command refers to several objects in one command, either of them should match (the logical OR algorithm is applied). For example, in the match ip address 101 102 command, a route is permitted if it is permitted by access list 101 or access list 102.

Question: 262

Refer to the exhibit.



An engineer configured NetFlow on R1, but the NMS server cannot see the flow from R1. Which configuration resolves the issue?

- A. flow monitor Flowmonitor1destination 10.221.10.11
- B. flow exporter FlowAnalyzer1destination 10.221.10.11
- C. interface Ethernet0/1flow-destination 10.221.10.11
- D. interface Ethernet0/0flow-destination 10.221.10.11

Answer: B

Explanation:

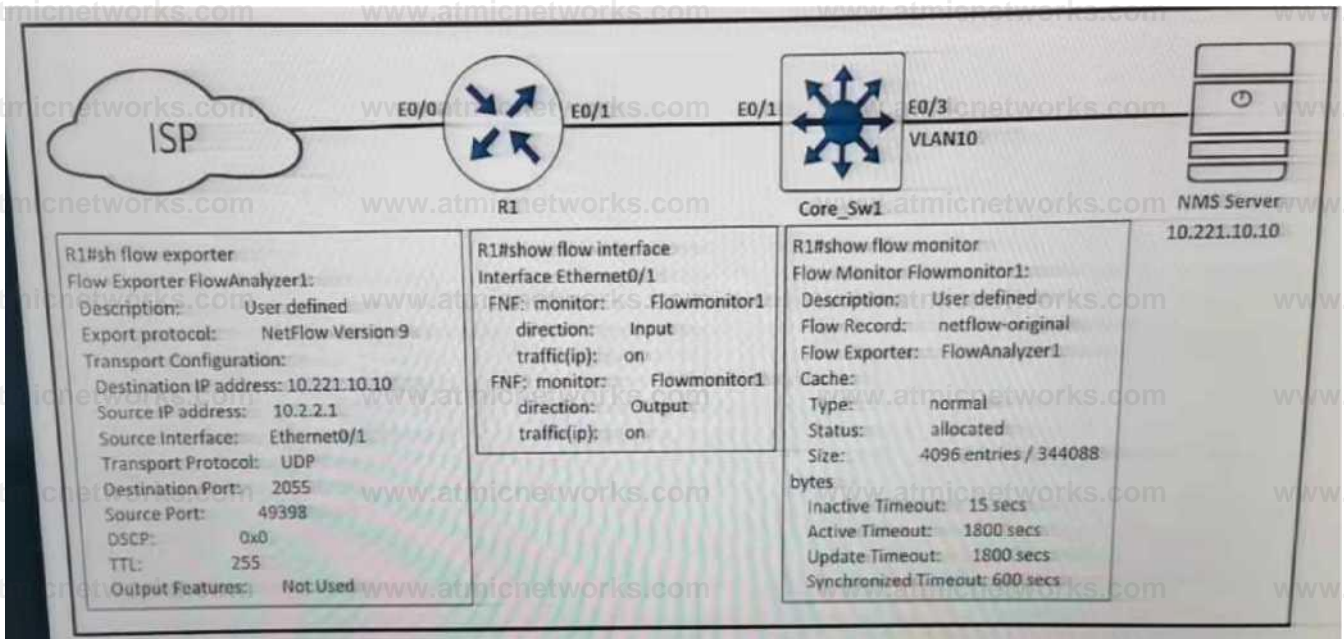
From the output we notice that the destination IP address is not correct. The NMS server IP address should be 10.221.10.11, not 10.221.10.10. Therefore we have to change this information under “flow exporter ...” configuration.

NetFlow configuration reference: <https://www.cisco.com/c/en/us/td/docs/iosxml/>

[ios/fnetflow/configuration/15-mt/fnf-15-mt-book/cfg-de-fnflow-exprts.html](https://www.cisco.com/c/en/us/td/docs/ios/fnetflow/configuration/15-mt/fnf-15-mt-book/cfg-de-fnflow-exprts.html)

Question: 263

Refer to the exhibit.



An engineer configured NetFlow on R1, but the NMS server cannot see the flow from ethernet 0/0 of R1. Which configuration resolves the issue?

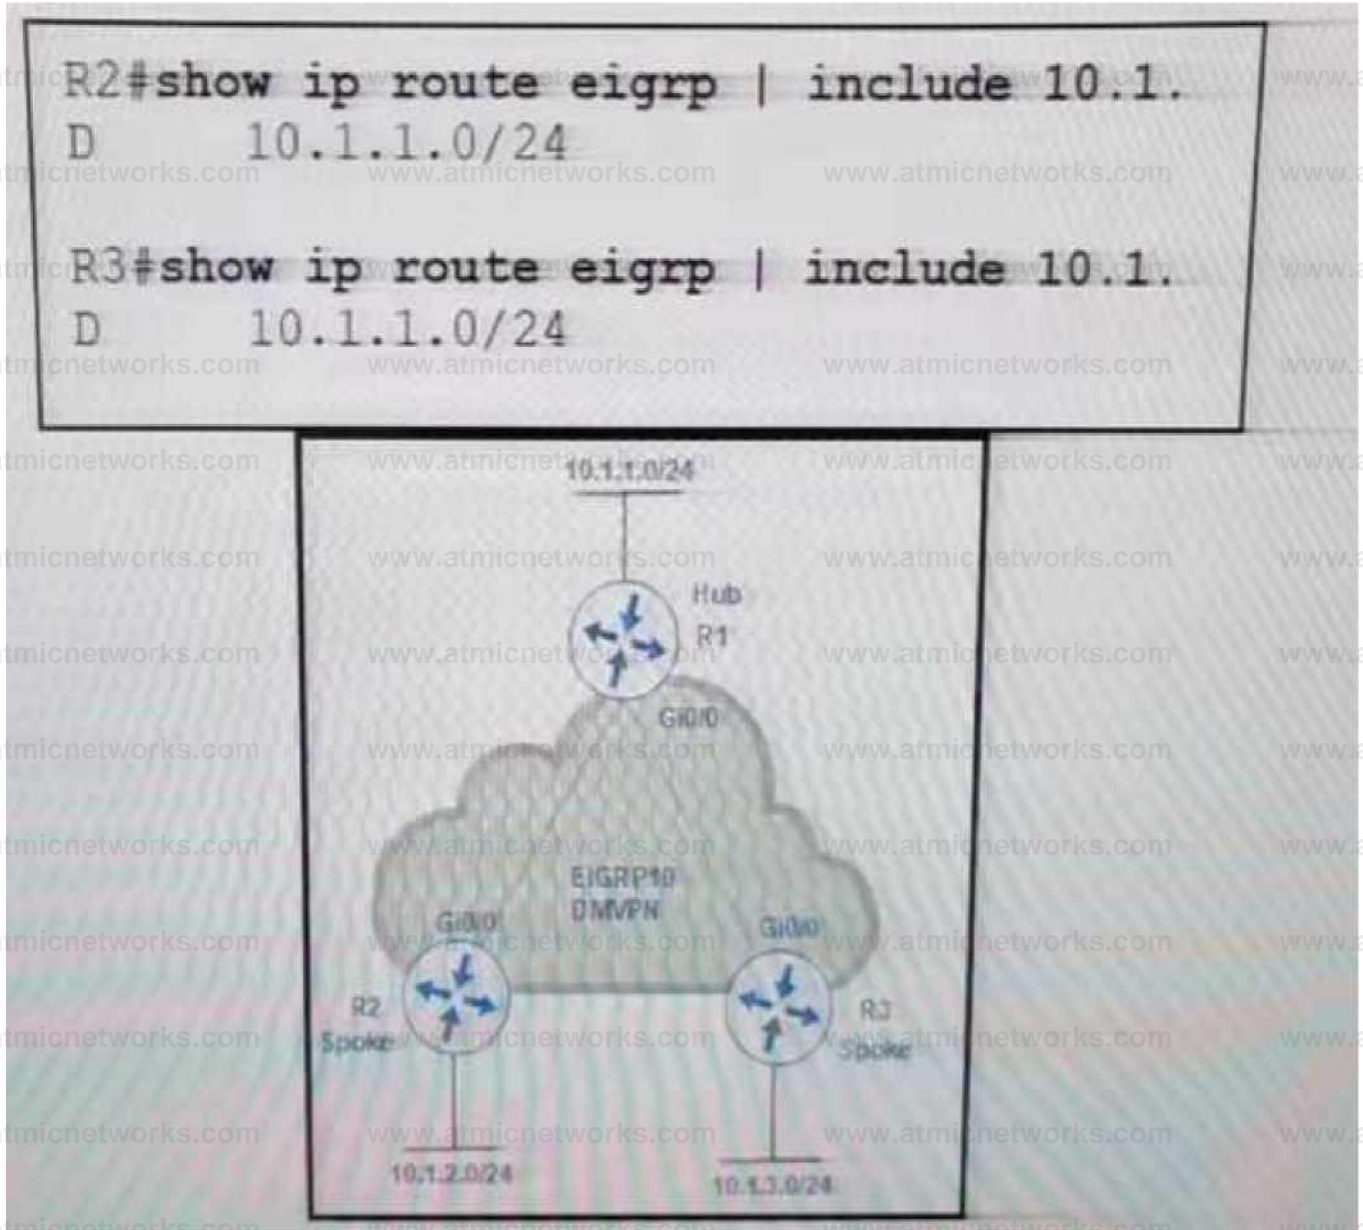
- A. flow monitor Flowmonitor1 source Ethernet0/0
- B. interface Ethernet0/1 ip flow monitor Flowmonitor1 input ip flow monitor Flowmonitor1 output
- C. interface Ethernet0/0 ip flow monitor Flowmonitor1 input ip flow monitor Flowmonitor1 output
- D. flow exporter FlowAnalyzer1 source Ethernet0/0

Answer: C

Explanation:

Question: 264

Refer to the exhibit.



An engineer configures DMVPN and receives the hub location prefix of 10.1.1.0/24 on R2 and R3. The R3 prefix of 10.1.3.0/24 is not received on R2, and the R2 prefix 10.1.2.0/24 is not received on R3. Which action resolves the issue?

- A. Split horizon prevents the routes from being advertised between spoke routers; it should be disabled with the command `no ip split-horizon eigrp 10` on the tunnel interface of R1.
- B. There is no spoke-to-spoke connection; DMVPN configuration should be modified to enable a tunnel connection between R2 and R3 and neighbor relationship confirmed by use of the `show ip eigrp neighbor` command.

- C. Split horizon prevents the routes from being advertised between spoke routers it should be disabled with the `no ip split-horizon eigrp 10` command on the Gi0/0 interface of R1.
- D. There is no spoke-to-spoke connection DMVPN configuration should be modified with a manual neighbor relationship configured between R2 and R3 and confirmed by use of the `show ip eigrp neighbor` command.

Answer: A

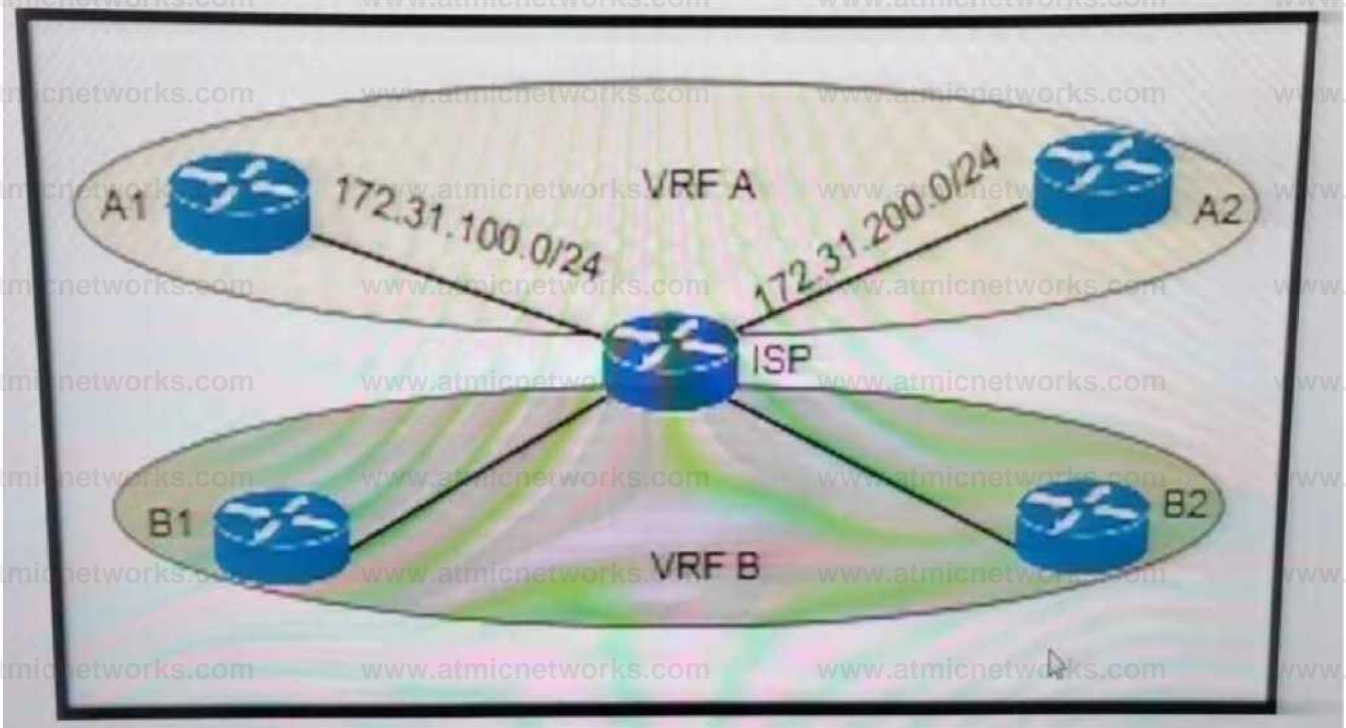
Explanation:

In this topology, the Hub router will receive advertisements from R2 Spoke router on its tunnel interface. The problem here is that it also has a connection with R3 Spoke on that same tunnel interface. If we don't disable split-horizon, then the Hub will not relay routes from R2 to R3 and the other way around. That is because it received those routes on the same interface tunnel and therefore

it cannot advertise back out that same interface (split-horizon rule). Therefore we must disable `split-horizon` on the Hub router to make sure the Spokes know about each other.

Question: 265

Refer to the exhibit. The ISP router is fully configured for customer A and customer B using the VRF- Lite feature. What is the minimum configuration required for customer A to communicate between routers A1 and A2?



A. A1interface fa0/0description To->ISPip add 172.31.100.1 255.255.255.0no shut!router ospf 100net 172.31.100.1 0.0.0.255 area 0A2interface fa0/0description To->ISPip add 172.31.200.1 255.255.255.0no shut!router ospf 100net 172.31.200.1 0.0.0.255 area 0

B. A1interface fa0/0description To->ISPip vrf forwarding Aip add 172.31.100.1 255.255.255.0no shut!router ospf 100net 172.31.100.1 0.0.0.255 area 0A2interface fa0/0description To->ISPip vrf forwarding Aip add 172.31.200.1 255.255.255.0no shut!router ospf 100net 172.31.200.1 0.0.0.255 area 0

C. A1interface fa0/0description To->ISPip add 172.31.200.1 255.255.255.0no shut!router ospf 100net 172.31.200.1 0.0.0.255 area 0A2interface fa0/0description To->ISPip add 172.31.100.1 255.255.255.0no shut!router ospf 100net 172.31.100.1 0.0.0.255 area 0

D. A1interface fa0/0description To->ISPip vrf forwarding Aip add 172.31.100.1 255.255.255.0no shut!router ospf 100 vrf Anet 172.31.200.1 0.0.0.255 area 0A2interface fa0/0description To->ISPip vrf forwarding Aip add 172.31.100.1 255.255.255.0no shut!router ospf 100 vrf Anet 172.31.200.1 0.0.0.255 area 0

Answer: C

Explanation:

A1 and A2 routers do not know they belong to VRF A. The two

interfaces of ISP (which are connected to A1 & A2) should be configured like this (we only show the configure of one

interface):

ISP router:

```
interface g0/0
```

```
description ISP->To_CustomerA
```

```
ip vrf forwarding A
```

```
ip address 172.31.100.2 255.255.255.0
```

```
router ospf 100 vrf A
```

```
network 172.31.200.2 0.0.0.255 area 0
```

Question: 266

The network administrator configured R1 for Control Plane Policing so that the inbound Telnet traffic is policed to 100 kbps.

This policy must not apply to traffic coming in from 10.1.1.1/32 and 172.16.1.1/32. The administrator has configured this:

```
access-list 101 permit tcp host 10.1.1.1 any eq 23
```

```
access-list 101 permit tcp host 172.16.1.1 any eq 23 [
```

```
class-map CoPP-TELNET match access-group 101 !
```

```
policy-map PM-CoPP class CoPP-TELNET police 100000 conform transmit  
exceed drop !
```

```
control-plane
```

```
service-policy input PM-CoPP
```

The network administrator is not getting the desired results. Which set of configurations resolves this issue?

- A. control-plane no service-policy input PM-CoPP
interface Ethernet 0/0 service-policy input PM-CoPP
- B. control-plane no service-policy input PM-CoPP
service-policy input PM-CoPP
- C. no access-list 101
access-list 101 deny tcp host 10.1.1.1 any eq 23
access-list 101 deny tcp host 172.16.1.1 any eq 23
access-list 101 permit ip any any
- D. no access-list 101
access-list 101 deny tcp host 10.1.1.1 any eq 23
access-list 101 deny tcp host 172.16.1.1 any eq 23
access-list 101 permit ip any any
interface E0/0 service-policy input PM-CoPP

Answer: C

Explanation:

Packets that match a deny rule are excluded from that class and cascade to the next class (if one exists) for classification.

Therefore if we don't want to CoPP traffic from 10.1.1.1/32 and 172.16.1.1/32, we must "deny" them in the ACL.

Question: 267

Refer to the exhibit.

```
R2U show Ip ospf neighbor Neighbor ID      Pri      State      Dead Timo
      Address      Interface
192.168.99.2      1  EXCHANGE/  00:00:36  192.168.99.1  Serlalo/1
router-611
```

```
R3I 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
```



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. Restore the Layer 2/Layer 3 connectivity issue in the ISP network.
- B. Match MTU on both router interfaces or ignore MTU.
- C. Administrative "shut then no shut" both router interfaces.

D. Enable OSPF on the interface, which is required.

Answer: B

Explanation:

After two OSPF neighboring routers establish bi-directional communication and complete DR/BDR election (on multi-access networks), the routers transition to the exstart state. In this state, the neighboring routers establish a master/slave relationship and determine the initial database descriptor (DBD) sequence number to use while exchanging DBD packets.

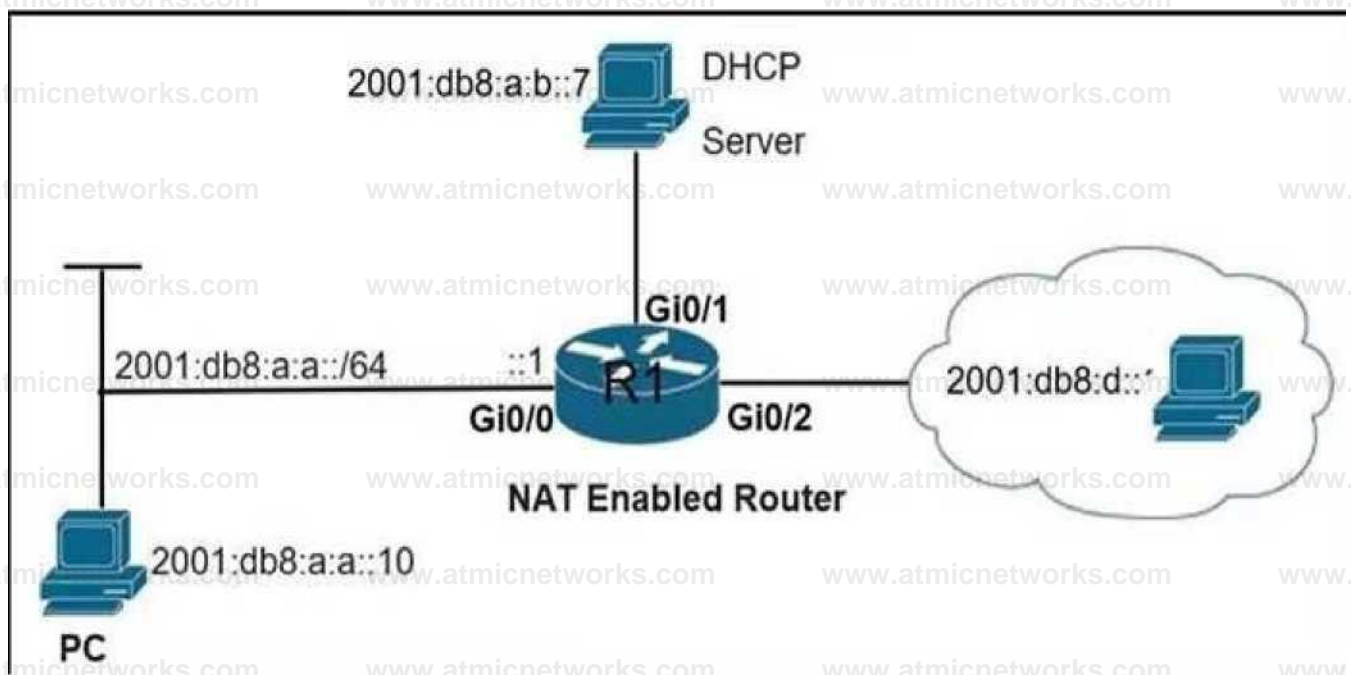
Neighbors Stuck in Exstart/Exchange State

The problem occurs most frequently when attempting to run OSPF between a Cisco router and another vendor's router. The problem occurs when the maximum transmission unit (MTU) settings for neighboring router interfaces don't match. If the router with the higher MTU sends a packet larger than

the MTU set on the neighboring router, the neighboring router ignores the packet.

Question:
268

Refer to the exhibit.



C:\PO ping 2001:db8:a:b::7

Pinging 2001 :db8 :a:b : :7 with 32 bytes of data:

Reply from 2001:db8:a:b::7: time=46ms

Reply from 2001:db8:a:b::7: time=40ms

Reply from 2001:db8:a:b::7: time=40ms

Reply from 2001:db8;a:b::7: time=40ms

Ping statistics for 2001:db8:a:b::7:

Packets: Sent = 4, Received = 4. Lost = 0 (0% loss).

Approximate round trip times in milli-seconds:

Minimum = 40ms. Maximum = 46ms. Average = 41 ms

RI* telnet 2001:db8:a:b::7

Trying 2001:DB8:A:B::7... Open

User Access Verification

Password:

Ri- show ipv6 access-list TSHOOT

IPv6 access list TSHOOT

deny tcp any host 2001:DB8:A:B::7 eq telnet (6 matches) sequence 10

permit tcp host 2001:DB8:A:A::10 host 2001:DB8:A:B::7 eq telnet sequence 20

permit tcp host 2001:DB8:A:A::10 host 2001:DB8:D::1 eq www sequence 30

permit ipv6 2001:DB8:A:A::64 any (67 matches) sequence 40

An engineer is troubleshooting a failed Telnet session from PC to the DHCP server. Which action resolves the issue?

- A. Remove sequence 30 and add it back to the IPv6 traffic filter as sequence 5.
- B. Remove sequence 20 and add it back to the IPv6 traffic filter as sequence 5.
- C. Remove sequence 10 to add the PC source IP address and add it back as sequence 10.
- D. Remove sequence 20 for sequence 40 in the access list to allow Telnet.

Answer: B

Explanation:

Question: 269

Refer to the exhibit.

```
ip sla 1 icmp-echo 8.8.8.8 threshold 1000 timeout 2000
  frequency 5
ip sla schedule 1 life forever start-time now i
track 1 ip sla 1 i
ip route 0.0.0.0 0.0.0.0 203.0.113.1 name ISP1 track 1
ip route 0.0.0.0 0.0.0.0 198.51.100.1 2 name ISP2
```

The administrator noticed that the connection was flapping between the two ISPs instead of switching to ISP2 when the ISP1 failed. Which action resolves the issue?

- A. Include a valid source-interface keyword in the icmp-echo statement.
- B. Reference the track object 1 on the default route through ISP2 instead of ISP1.
- C. Modify the static routes to refer both to the next hop and the outgoing interface.
- D. Modify the threshold to match the administrative distance of the ISP2 route.

Answer: A

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/ip-routing/200785-ISP-Failover-withdefault-routes-using-l.html>


```
RR# show running-config |
interface EthernetO/1 no ip address
  ipv6 address 2001:DB8:1:12::2/€4
  ipv6 traffic-filter ACL in
  ipv6 access-list ACL
sequence 10 pennis tcp any any eq 22
sequence 20 permit tcp any eq 22 any
sequence 30 permit tcp any any eq bgp
sequence 40 permit tcp any eq bgp any
sequence 50 permit udp any any eq ntp
sequence 60 permit udp any eq ntp any
sequence 70 permit udp any any eq snmp
sequence 80 deny ipv6 any any log
```

```
RR# show ipv6 cef ::/0
```

```
::/0
```

```
  nexthop 2001:DB8:1:12::1 EthernetO/i
```

```
*Feb 23 00:23:17.211: %IPV6 ACL-6-ACCESSLOGDP: list ACL/80 denied
icmpv6 2001:DB8:1:12::1 -> FF02::1:FF00:2 (135/0), 7321 packets
```

After a security audit, the administrator implemented an ACL in the route reflector. The RR became unreachable from any router in the network. Which two actions resolve the issue? (Choose two.)

- A. Enable the ND proxy feature on the default gateway.
- B. Configure a link-local address on the Ethernet0/1 interface.
- C. Permit ICMPv6 neighbor discovery traffic in the ACL.
- D. Remove the ACL entry 80.
- E. Change the next hop of the default route to the link-local address of the default gateway.

Answer: C,D

Explanation:

Question: 271

Refer to the exhibit.

```
R1 (config)# ip vrf CCNP
R1 (config-vrf)# rd 1:100
R1 (config-vrf)# exit
R1 (config)# interface Loopback0
R1 (config-if)# ip address 10.1.1.1 255.255.255.0
R1 (config-if)# ip vrf forwarding CCNP
R1 (config-if)# exit
R1 (config)# exit
R1#ping vrf CCNP 10.1.1.1
% Unrecognized host or address, or protocol not running
```

Which command must be configured to make VRF CCNP work?

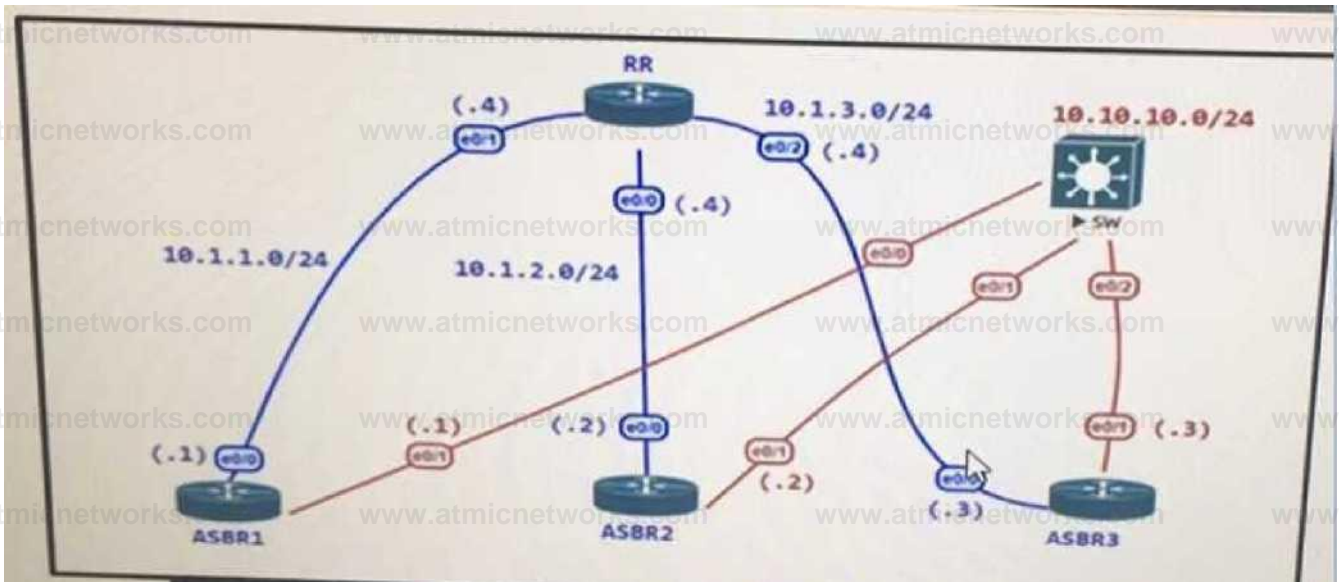
- A. interface Loopback0vrf forwarding CCNP
- B. interface Loopback0ip address 10.1.1.1 255.255.255.0
- C. interface Loopback0ip address 10.1.1.1 255.255.255.0vrf forwarding CCNP
- D. interface Loopback0ip address 10.1.1.1 255.255.255.0ip vrf forwarding CCNP

Answer: B**Explanation:**

From the exhibit, we learn that the command "ip address 10.1.1.1 255.255.255.0" has been issued before the command "ip vrf forwarding CCNP". But the second command removed the IP address configured in the first command so we have to retype the IP address command.

Question: 272

Refer to the exhibit.



RR

```
router bgp 100
  neighbor 10.1.1.1 remote-as 100
  neighbor 10.1.2.2 remote-as 100
  neighbor 10.1.3.3 remote-as 100
```

ASBR2

```
router bgp 100
  neighbor 10.1.1.4 remote-as 100
```

ASBR3

```
router bgp 100
  neighbor 10.1.2.4 remote-as 100
```

ASBR4

```
router bgp 100
  neighbor 10.1.3.4 remote-as 100
```

The administrator configured the network device for end-to-end reachability, but the ASBRs are not propagating routes to each other. Which set of configuration resolves this issue?

A)

```
router bgp 100
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.1.2.2 route-reflector-client
  neighbor 10.1.3.3 route-reflector-client
```

B)

```
router bgp 100
neighbor 10.1.1.1
neighbor 10.1.2.2 next-hop-self
neighbor 10.1.3.3 next-hop-self
```

C)

```
router bgp 100
neighbor 10.1.1.1 update-source Loopback0
neighbor 10.1.2.2 update-source Loopback0
neighbor 10.1.3.3 update-source Loopback0
```

D)

```
router bgp 100
neighbor 10.1.1.1 ebgp-multihop
neighbor 10.1.2.2 ebgp-multihop
neighbor 10.1.3.3 ebgp-multihop
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer:**A**

Explanation:

Question: 273

A company is expanding business by opening 35 branches over the Internet. A network engineer must configure DMVPN at the branch routers to connect with the hub router and allow NHRP to add spoke routers securely to the multicast NHRP mappings automatically. Which configuration meets this requirement at the hub router?

A)

interface Tunnel0**ip address 10 C.0.1 253.255.255.0****ip nhrp authentication KEY1 ip nhrp nhs dynamic ip nhrp network-id
10 tunnel mode mgre auto**

B)

interface Tunnel0**ip address 10 0.0 1 2 55.255.2 55.0****ip nhrp authentication KEY1 ip nhrp registration no-unique ip nhrp
networked 10 tunnel mode gre nmba**

C)

interface Tunnel0**ip address 10 0 0.1 255.255.255 0****ip nhrp authentication KEY1 ip nhrp map multicast dynamic ip nhrp
network-id 10 tunnel mode gre multipoint**

D)

interface Tunnel0**ip address 10 0 0.1 255.255.255 0****ip nhrp authentication KEY 1****ip nhrp map multicast 224.0 0 0****ip nhrp network-id 10****tunnel mode gre ipv4**

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

The command "ip nhrp map multicast dynamic" allows NHRP to automatically add spoke routers to the multicast NHRP mappings.

Question: 274

What is an advantage of implementing BFD?

- A. BFD provides faster updates for any flapping route.
- B. BFD provides millisecond failure detection.
- C. BFD is deployed without the need to run any routing protocol.
- D. BFD provides better capabilities to maintain the routing table.

Answer: B

Explanation:

Question: 275

What is a function of IPv6 Source Guard?

- A. It works with address glean or ND to find existing addresses.
- B. It inspects ND and DHCP packets to build an address binding table.

- C. It denies traffic from known sources and allocated addresses.
- D. It notifies the ND protocol to inform hosts if the traffic is denied by it.

Answer: A

Explanation:

IPv6 source guard is an interface feature between the populated binding table and data traffic filtering. This feature enables the device to deny traffic when it is originated from an address that is not stored in the binding table. IPv6 source guard does not inspect ND or DHCP packets; rather, it works in conjunction with IPv6 neighbor discovery (ND) inspection or IPv6 address glean, both of which detect existing addresses on the link and store them into the binding table.

Question: 276

What is the purpose of the DHCPv6 Guard?

- A. It messages between a DHCPv6 server and a DHCPv6 client (or relay agent).
- B. It shows that clients of a DHCPv5 server are affected.
- C. It block DHCPv6 messages from relay agents to a DHCPv6 server.
- D. It allows DHCPv6 replay and advertisements from (rouge) DHCPv6 servers.

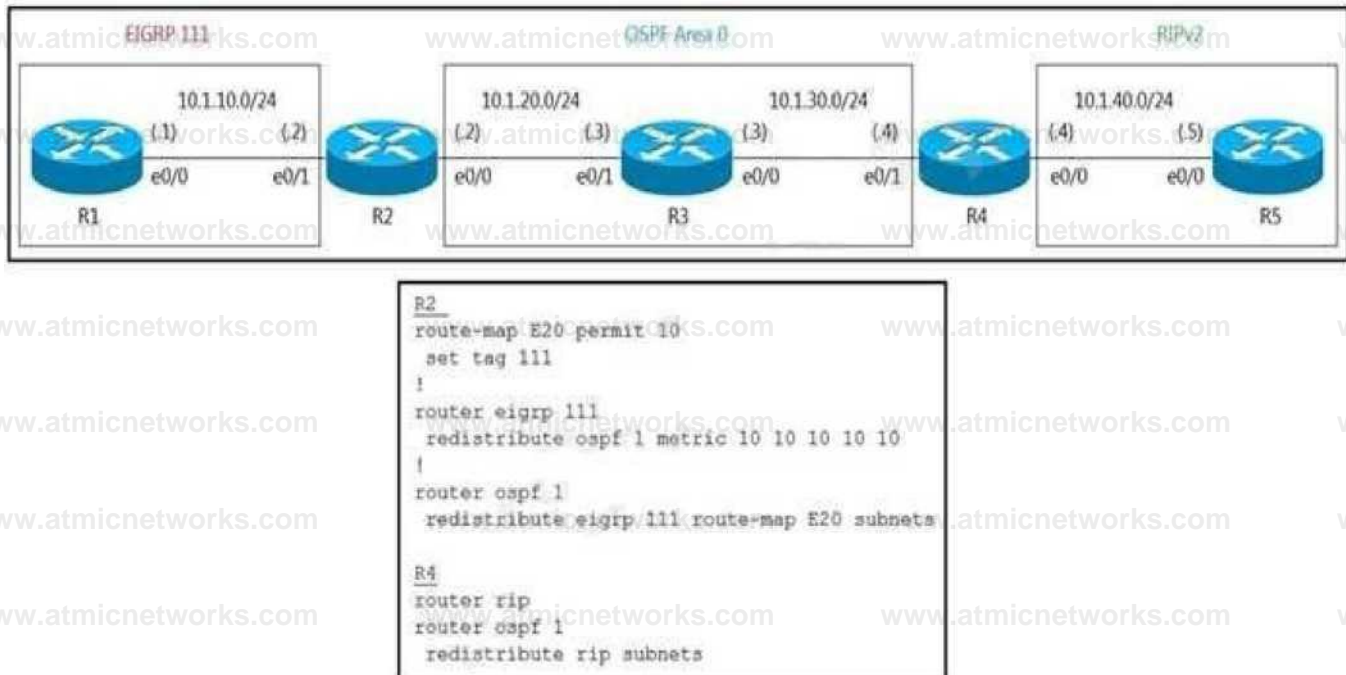
Answer: A

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xr-16/ip6fxe-16-book/ip6-dhcpv6-guard.html

Question: 277

Refer to the exhibit.



R5 should not receive any routes originated in the EIGRP domain. Which set of configuration changes removes the EIGRP routes from the R5 routing table to fix the issue?

- A. R4route-map O2R deny 10match tag 111route-map O2R permit 20!router ripredistribute ospf 1 route-map O2R metric 1
- B. R2route-map E20 deny 20R4route-map O2R deny 10match tag 111!router ripredistribute ospf 1 route-map O2R metric 1
- C. R4route-map O2R permit 10match tag 111route-map O2R deny 20!router ripredistribute ospf 1 route-map O2R metric 1
- D. R4route-map O2R deny 10match tag 111!router ripredistribute ospf 1 route-map O2R metric 1

Answer: A**Explanation:**

In this question, routes from EIGRP domain are redistributed into OSPF (with tag 111) then RIPv2 but without any filtering so


R5 learns all routes from both EIGRP and OSPF domain. If we only want R5 to learn routes from OSPF domain then we must filter out routes with tag 111 and permit other routes. The line "route-map O2R permit 20" is important to allow other

routes because of the implicit deny all

at the end of each route-map.

Question: 278

Refer to the exhibit.



An engineer has configured R1 as EIGRP stub router. After the configuration, router R3 failed to reach to R2 loopback address.

Which action advertises R2 loopback back into the R3 routing table?

- A. Add a static route for R2 loopback address in R1 and redistribute it to advertise to R3.
- B. Use a leak map on R1 that matches the required prefix and apply it with the distribute list command toward R3.
- C. Use a leak map on R3 that matches the required prefix and apply it with the EIGRP stub feature.

D. Add a static null route for R2 loopback address in R1 and redistribute it to advertise to R3.

Answer: B

Explanation:

The EIGRP stub feature is useful to prevent unnecessary EIGRP queries and to filter some routes that you advertise. What if you want to configure your router as a stub router but still make an exception to some routes that it advertises? That is possible with the leak-map feature. This is how to configure leakmap in this

question:

R1

```
(config)#ip access-list standard R2_L0
```

```
R1(config-std-nacl)#permit host 2.2.2.2
```

```
R1(config)#route-map R2_L0_LEAK
```

```
R2(config-route-map)#match ip address R2_L0
```

```
R1(config)#router eigrp 1
```

```
R1(config-router)#eigrp stub leak-map R2_L0_LEAK
```

Question: 279

Refer to the exhibit.

```
R1#sh ip route
10.0.0.0/8 is variably subnetted, 3 subnets, 1 masks
D    10.1.2. 0/24    [90/409600] via 10.1.100.10, 00:08:45,
```

```
FastEthernet0/0
D      10.1.1.0/24 [90/409600] via 10.1.100.10, 00:08:45,
FastEthernet0/0
C      10.1.100.0/24 is directly connected, FastEthernet0/0
```

An engineer configures the router 10.1.100.10 for EIGRP autosummarization so that R1 should receive the summary route of 10.0.0.0/8. However, R1 receives more specific /24 routes.

Which action resolves this issue?

- A. Router R1 should configure ip summary address eigrp (AS number) 10.0.0.0 255.0.0.0 for the R1 Fast Ethernet 0/0 connected interface.
- B. Router R1 should configure ip route 10.0.0.0 255.0.0.0 null 0 for the routes that are received on R1.
- C. Router 10.1.100.10 should configure ip route 10.0.0.0 255.0.0.0 null 0 for the routes that are summarized toward R1.
- D. Router 10.1.100.10 should configure ip summary address eigrp (AS number) 10.0.0.0 255.0.0.0 for the R1 Fast Ethernet 0/0 connected interface.

Answer: D

Explanation:

Question: 280

DRAG DROP

Drag and drop the IPv6 first hop security device roles from the left onto the corresponding descriptions on the right.

- host
- router
- monitor
- switch

Receives router advertisements from valid routers, and no router solicitation are received.

Receives router solicitation and sends router advertisements.

Receives valid and rogue router advertisements and all router solicitation.

Received router advertisements are trusted and are flooded to synchronize states.

Answer:

Explanation:

router

host

switch

monitor

Graphical user interface, text, application, email Description automatically generated

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_011011.pdf

Question: 281

Refer to the exhibit.

*17:40:07.826: AAA/BIND(00000055): Bind i/f

*17:40:07.826: AAA/AUTHEN/LOGIN (00000055): Pick method list 'default'

* 17:40:07.826: TPLUS: Queuing AAA Authentication request 85 for processing

* 17:40:07.826: TPLUS: TPLUS(OOOOOO55) login timer started 1020 sec timeout

```
* 17:40:07.826: TPLUS: processing authentication start request id 85
* 17:40:07.826: TPLUS: Authentication start packet created for 850
* 17:40:07.826: Using server 10.106.60.182
* 17:40:07.826: TPLUS(00000055)/O/NB_WAIT/225FE2DC: Started 5 sec timeout
* 17:40:07.830: TPLUS(00000055)/O/NB_WAIT: socket event 2
* 17:40:07.830: TPLUS(OOOOOO55)/O/NB_WAIT: wrote entire 38 bytes request
* 17:40:07.830: TPLUS(OOOOOO55)/O/READ: socket event 1
*17:40:07.830: TPLUS(OOOOOO55)/O/READ: Would block while reading
*17:40:07.886: TPLUS(OOOOOO55)/O/READ: socket event 1
*17:40:07.886: TPLUS(OOOOOO55)/O/READ: read entire 12 header bytes (expect 6 bytes data)
* 17:40:07.886: TPLUS(OOOOOO55)/O/READ: socket event 1
* 17:40:07.886: TPLUS(OOOOOO55)/O/READ: read entire 18 bytes response
* 17:40:07.886: TPLUS(00000055)/O/225FE2DC: Processing the reply packet
* 17:40:07.886: TPLUS: received bad AUTHEN packet: length = 6, expected 43974
* 17:40:07.886: TPLUS: Invalid AUTHEN packet (check keys).
```

An engineer is troubleshooting a TACACS problem.

Which action resolves the issue?

- A. Configure a matching TACACS server IP.
- B. Configure a matching preshared key.
- C. Generate authentication from a relative source interface.
- D. Apply a configured AAA profile to the VTY.

Answer: B

Explanation:

Reference:

<https://community.cisco.com/t5/network-access-control/issues-with-tacacs-authentication/td-p/3412001>

The last line shows us the reason, which is "Invalid AUTHEN packet (check keys)" so the most likely cause of this problem is key mismatch.

Question: 282

The network administrator configured CoPP so that all HTTP and HTTPS traffic from the administrator device located at 172.16.1.99 toward the router CPU is limited to 500 kbps. Any traffic that exceeds this limit must be dropped.

```
access-list 100 permit ip host 172.16.1.99 any
```

```
!
```

```
class-map CM-ADMIN
```

```
match access-group 100
```

```
!
```

```
policy-map PM-COPP
```

```
class CM-ADMIN
```

```
police 500000 conform-action transmit
```

```
!
```

```
interface E0/0
```

```
service-policy input PM-COPP
```

CoPP failed to capture the desired traffic and the CPU load is getting higher.

Which two configurations resolve the issue? (Choose two.)

A. interface E0/0 no service-policy input PM-COPP
control-plane service-policy input PM-COPP

- B. policy-map PM-COPPclass CM-ADMINno police 500000 conform-action transmitpolice 500 conform-action transmit!control-planeservice-policy input PM-COPP
- C. no access-list 100access-list 100 permit tcp host 172.16.1.99 any eq 80
- D. no access-list 100access-list 100 permit tcp host 172.16.1.99 any eq 80access-list 100 permit tcp host 172.16.1.99 any eq 443
- E. policy-map PM-COPPclass CM-ADMINno police 500000 conform-action transmitpolice 500 conform-action transmit

Answer: A

Explanation:

Question: 283

Refer to the exhibit.

```
ipv6 access-list INTERNET
permit ipv6 2001:DBS:AD59:BA21::/64 2001:DBS:COAB:BAI4::/64
permit tcp 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA13::/64 eq telnet
permit tcp 2001:DB8:AD59:BA21::/64 any eq http
permit ipvc 2001:DBS:AD59::/4S any
deny ipv6 any any log
```

While monitoring VTY access to a router, an engineer notices that the router does not have any filter and anyone can access the router with username and password even though an ACL is configured.

Which command resolves this issue?

- A. access-class INTERNET in

B. ip access-group INTERNET in

C. ipv6 traffic-filter INTERNET in

D. ipv6 access-class INTERNET in

Answer: D

Explanation:

Question: 284

Refer to the exhibit.

ipv6 dhcp server:

ipv6 unicast-routing

int e0/1

ipv6 enable

ipv6 add 200111::1/64

ipv6 nd other-config-flag no shut

ipv6 dhcp server !Pv6Pool i

ipv6 dhcp pool IPv6Pool dns-server

2002:555"! domain-name my.net

ipv6 dhcp client:

interface Ethernet0/1 no ip address ipv6
address dhcp ipv6 enable no shut

i-----1

A network administrator is troubleshooting IPv6 address assignment for a DHCP client that is not getting an IPv6 address from the server.

Which configuration retrieves the client IPv6 address from the DHCP server?

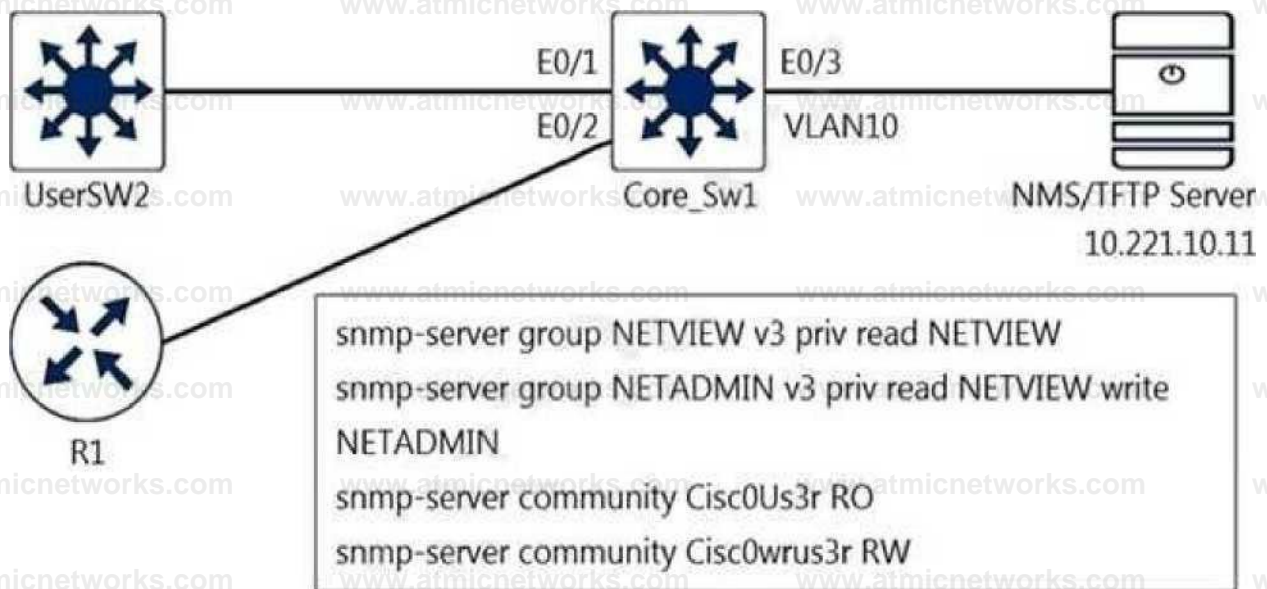
- A. ipv6 address autoconfig command on the interface
- B. ipv6 dhcp server automatic command on DHCP server
- C. ipv6 dhcp relay-agent command on the interface
- D. service dhcp command on DHCP server

Answer: A

Explanation:

Question: 285

Refer to the exhibit.



A junior engineer configured SNMP to network devices. Malicious users have uploaded different configurations to the network devices using SNMP and TFTP servers.

Which configuration prevents changes from unauthorized NMS and TFTP servers?

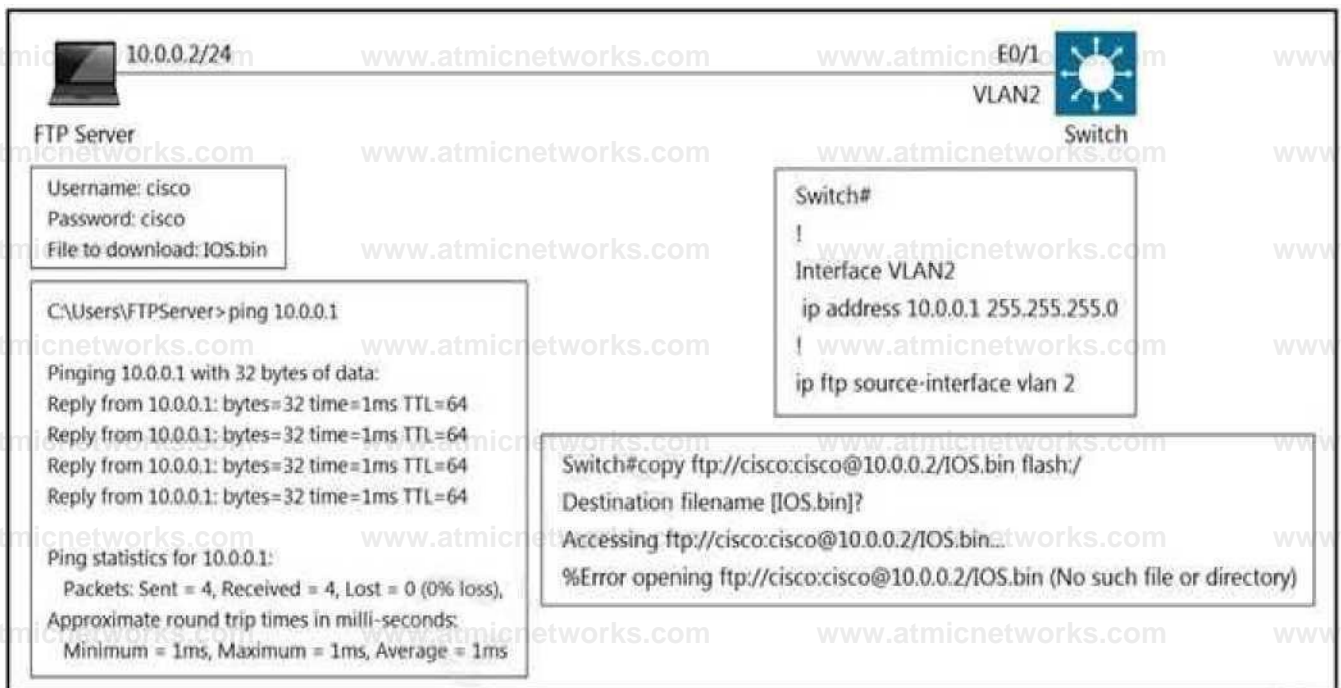
- A. `access-list 20 permit 10.221.10.11``access-list 20 deny any log!``snmp-server group NETVIEW v3 priv read NETVIEW``access 20snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN``access 20snmp-server community Cisc0Us3r RO``20snmp-server community Cisc0wrus3r RW``20snmp-server tftp-server-list 20`
- B. `access-list 20 permit 10.221.10.11``access-list 20 deny any log!``snmp-server group NETVIEW v3 priv read NETVIEW``access 20snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN``access 20snmp-server community Cisc0wrus3r RO``20snmp-server community Cisc0Us3r RW``20snmp-server tftp-server-list 20`
- C. `access-list 20 permit 10.221.10.11``access-list 20 deny any log`
- D. `access-list 20 permit 10.221.10.11`

Answer:**A**

Explanation:

Question:**286**

Refer to the exhibit.



An engineer cannot copy the IOS.bin file from the FTP server to the switch.

Which action resolves the issue?

A. Allow file permissions to download the file from the FTP server.

B. Add the IOS.bin file, which does not exist on FTP server.

- C. Make memory space on the switch flash or USB drive to download the file.
- D. Use the copy flash:/ ftp://cisco@10.0.0.2/IOS.bin command.

Answer: B

Explanation:

Question: 287

What does the MP-BGP OPEN message contain?

- A. MPLS labels and the IP address of the router that receives the message
- B. the version number and the AS number to which the router belongs
- C. IP routing information and the AS number to which the router belongs
- D. NLRI, path attributes, and IP addresses of the sending and receiving routers

Answer: B

Explanation:

Question: 288

Refer to the exhibit.

```
Router1#sh ip route
      10.0.0.0/8 is variably subnetted, 3 subnets, 1 masks
D 10.1.2.0/24 [90/409600] via 10.1.100.10, 00:08:45, FastEthernet0/0
D 10.1.1.0/24 [90/409600] via 10.1.100.10, 00:08:45,
```

FastEthernet0/0

C 10.1.100.0/24 is directly connected, FastEthernet0/0

Although summarization is configured for R1 to receive 10.0.0.0/8, more specific routes are received by R1. How should the 10.0.0.0/8 summary route be received from the neighbor, attached to R1 via Fast Ethernet0/0 interface?

- A. R1 should configure the ip summary-address eigrp <AS number> 10.0.0.0.255.0.0.0 command under the Fast Ethernet 0/0 interface.
- B. The summarization condition is not met Router 10 1 100.10 requires a route for 10 0.0.0/8 that points to null 0
- C. The summarization condition is not met. The network 10.1.100.0/24 should be changed to 172.16.0.0/24.
- D. R1 should configure the ip summary-address eigrp <AS number> 10.0.0.0 0.0.0.255 command under the Fast Ethernet 0/0 interface.

Answer: D

Explanation:

Question: 289

Refer to the exhibit.

```
Ri (config)#ip prefix-list EIGRP seq 10 deny 0.0.0.0/0 le 32
RI(config)#ip prefix-list EIGRP seq 20 permit 10.0.0.0/8
RI(config)#router eigrp 10
P1(config-router)#distribute-list prefix EIGRP in Ethernet0/0
RI#show ip route eigrp
```

A prefix list is created to filter routes inbound to an EIGRP process except for network 10 prefixes. After the prefix list is applied, no network 10 prefixes are visible in the routing table from EIGRP. Which configuration resolves the issue?

- A. ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9.
- B. ip prefix-list EIGRP seq 10 permit 0.0.0.0/0 le 32
- C. ip prefix-list EIGRP seq 5 permit 10.0.0.0/8 ge 9 no ip prefix-list EIGRP seq 20 permit 10.0.0.0/8
- D. ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9 ip prefix-list EIGRP seq 10 permit 0.0.0.0/0 le 32

Answer: C

Explanation:

Question: 290

Refer to the exhibit.

Tunnel source 195.1.1.1, destination 200.1.1.3 Tunnel protocol/transport GRE/IP

Key disabled, sequencing disabled Checksumming of packets disabled Tunnel TTL 255, Fast tunneling enabled Tunnel transport MTU 1476 bytes Tunnel transmit bandwidth 8000 (kbps) Tunnel receive bandwidth 8000 (kbps)

An engineer must establish a point-to-point GRE VPN between R1 and the remote site. Which configuration accomplishes the task for the remote site?

- A. Interface Tunnel1 tunnel source 199.1.1.1 tunnel destination 200.1.1.3 ip address 192.168.1.3 255.255.255.0
- B. Interface Tunnel1 tunnel source 200.1.1.3 tunnel destination 199.1.1.1 ip address 192.168.1.1 255.255.255.0
- C. Interface Tunnel1 tunnel source 200.1.1.3 tunnel destination 199.1.1.1 ip address 192.168.1.3 255.255.255.0
- D. Interface Tunnel1 tunnel source 199.1.1.1 tunnel destination 200.1.1.3 ip address 192.168.1.1 255.255.255.0

Answer: C

Explanation:

Question: 291

What are the two prerequisites to enable BFD on Cisco routers? (Choose two)

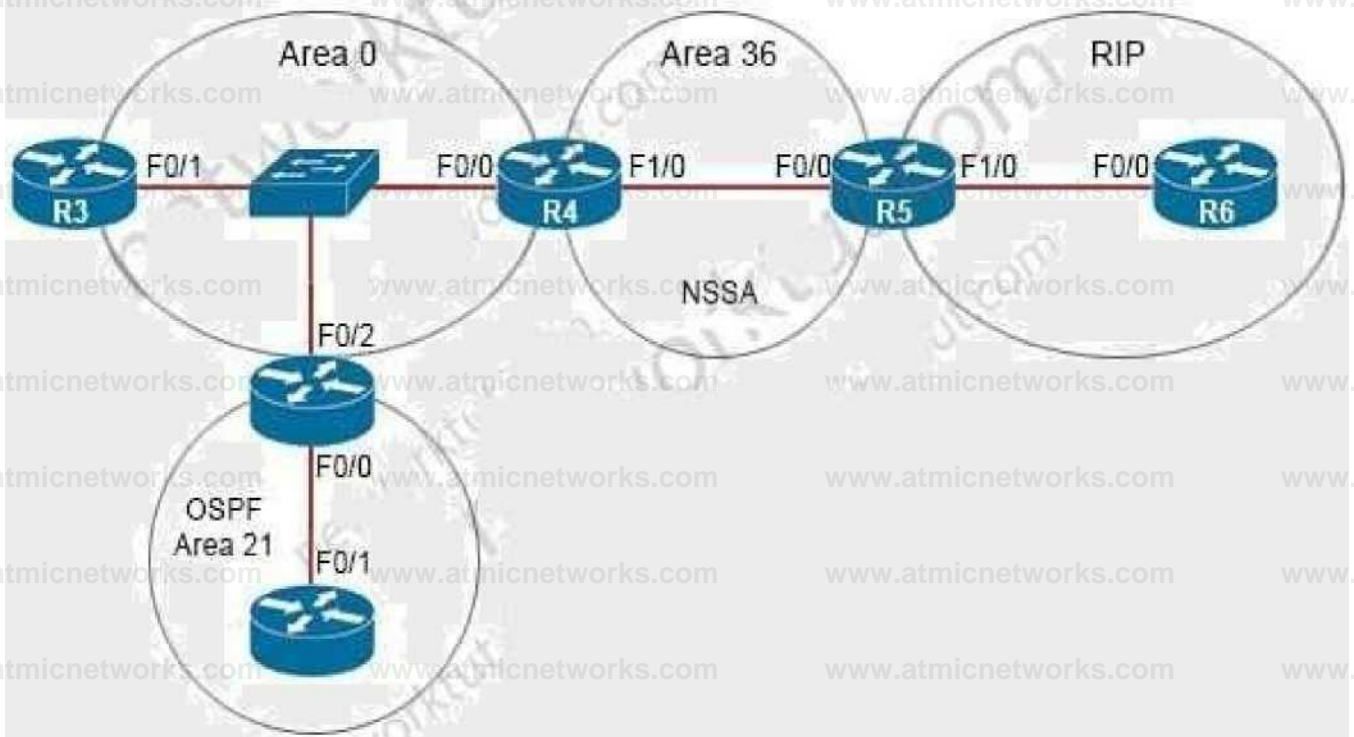
- A. A supported IP routing protocol must be configured on the participating routers.
- B. OSPF Demand Circuit must run BFD on all participating routers.
- C. ICMP must be allowed on all participating routers.
- D. UDP port 1985 must be allowed on all participating routers.
- E. Cisco Express Forwarding and IP Routing must be enabled on all participating routers.

Answer: C,E

Explanation:

Question: 292

Refer to the exhibit.



```
R5# show ip ospf 1 | begin Area 36
Area 36
Number of interfaces in this area is 2
It is a NSSA area
Area has no authentication
SPF algorithm last executed 00:32.46.376 ago
SFF algorithm executed 13 times
Area ranges are
 172.16.0.0/16 Passive Advertise
```

The network engineer configured the summarization of the RIP routes into the OSPF domain on R5 but

still sees four different 172.16.0.0/24 networks on R4. Which action resolves the issue?

- A. R5(config)#router ospf 1R5(config-router)#no areaR5(config-router)#summary-address 172.16.0.0 255.255.252.0
- B. R4(config)#router ospf 99R4(config-router)#network 172.16.0.0 0.255.255.255 area 56R4(config-router)#area 56 range 172.16.0.0 255.255.255.0
- C. R4(config)#router ospf 1R4(config-router)#no areaR4(config-router)#summary-address 172.16.0.0 255.255.252.0
- D. R5(config)#router ospf 99R5(config-router)#network 172.16.0.0 0.255.255.255 area 56R5(config-router)#area 56 range 172.16.0.0 255.255.255.0

Answer: A**Explanation:****Explanation**

Area 36 is a NSSA so R5 is an ASBR so we can summarize external routes using the “summaryaddress”

command. The command “area area-id range” can only be used on ABR so it is not correct.

The summarization must be done on the ASBR which is R5, not R4 so the correct answer must be

started with “R5(config)#router ospf 1”.

Note: The “no area” command is used to remove any existing “area ...” command (maybe “area 56 range ...” command).

Question: 293

The network administrator configured the router for Control Plane Policing to limit OSPF traffic to be policed to 1 Mbps.

Any traffic that exceeds this limit must also be allowed at this point for traffic analysis. The router configuration is:

```
access-list 100 permit ospf any any
```

```
!
```

```
class-map CM-OSPF
```

```
match access-group 100
```

```
!
```

```
policy-map PM-COPP
```

```
class CM-OSPF
```

police 1000000 conform-action transmit !

control-plane

service-policy output PM-COPP

The Control Plane Policing failed to monitor and police OSPF traffic. Which configuration resolves this issue?

no access-list 100

ip access-list 100 permit tcp any any eq 179 access-list 100 permit tcp any any range 22 23

class-map CM-MGMT

no match access-group 100

match access-group 101

central-policy

no service-policy output PM-COPP service-policy input PM-COPP

no access-list 100

access-list 100 permit tcp any any eq 179 access-list 100 permit tcp any any range eq 52 access-list 100 permit tcp any any range eq 23 access-list 100 permit ospf any any

central-policy

no service-policy output PM-COPP service-policy input PM-COPP

no access-list 100

access-list 100 permit tcp any any eq 179 access-list 100 permit ospf any any access-list 101 permit tcp any any range 22 23

class-map CM-MGMT

no match access-group 100

match access-group 101

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 294

Which feature minimizes DoS attacks on an IPv6 network?

- A. IPv6 Binding Security Table
- B. IPv6 Router Advertisement Guard
- C. IPv6 Prefix Guard
- D. IPv6 Destination Guard

Answer: D

Explanation:

The Destination Guard feature helps in minimizing denial-of-service (DoS) attacks. It performs address resolutions only for those addresses that are active on the link, and requires the FHS binding table to be populated with the help of the IPv6 snooping feature. The feature enables the filtering of IPv6 traffic based on the destination address, and blocks the NDP resolution for destination addresses that are not found in the binding table. By default, the policy drops traffic coming for an unknown destination.

Reference:

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_1_5_0s_book/IPv6_Security.pdf

Question: 295

Refer to Exhibit.

```
RI#sh ipv6 access-list GUARD
IPv6 access list GUARD
deny tcp any host 2001 :DB8:A:B:: 10 eq telnet (6 matches) sequence 10
pennit tcp host 2001 :DB8:A:A::20 host 2001 :DB8:A:B: 10 eq tehiet sequence
20
pennit tcp host 2001:DB8:A:A::2 host 2001:DB8:D:: 1 eq www sequence 3 0
pennit ipv6 2001:DB8:A:A::/64 any (67 matches) sequence 40
```

PC2 is directly connected to R1. A user at PC2 cannot Telnet to 2001:db8:a:b::10. The user can ping 2001:db8:a:b::10 and receive DHCP-related information from the DHCP server. Which action resolves the issue?

- A. Remove sequence 10 and put it back as sequence 25.
- B. Remove sequence 20 and put it back as sequence 45.
- C. Remove sequence 30 and put it back as sequence 5.
- D. Remove sequence 40 and put it back as sequence 15.

Answer: A

Explanation:

Question: 296

A CoPP policy is applied for receiving SSH traffic from the WAN interface on a Cisco ISR4321 router.

However, the SSH response from the router is abnormal and stuck during the high link utilization. The problem is identified as SSH traffic does not match in the ACL. Which action resolves the issue?

- A. Rate-limit SSH traffic to ensure dedicated bandwidth.
- B. Apply CoPP on the control plane interface.
- C. Increase the IP precedence value of SSH traffic to 6.
- D. Apply CoPP on the WAN interface inbound direction.

Answer: B

Explanation:

Explanation

The problem is "SSH traffic does not match in the ACL" and "CoPP policy is applied for receiving SSH traffic from the WAN interface" so we should apply CoPP on the control plane interface instead.

Question: 297

Refer to the exhibit.

```
Router#show ip bgp vpnv4 rd 1100:1001 10.30.116.0/23
```

```
BGP routing table entry Tor 3100:1001:10 30.116.0/23, version 2676527 5
```

```
Paths: (9 available, best #6, no table)
```

```
Advertised to update-groups:
```

```
12      3      '
```

```
(65001 64955 65003) 650 89, (Received from a RR-dient)
```

```
172.16.254.226 (metric 20645) from 172.16.224.236(172.16.224.2 36) Origin IGP
```

```
metric 0, localpref 100, valid, confed-internal
```

```
Extended Community: RT: 1100:1001 mpls labels in/out nolabel/362
```

```
(65008 64955 65003) 65 0 89
```

```
172.16.254.226 (metric 20645) from 10.131 123.71 (10.131.123.71) Origin IGP,
```

```
metric 0, localpref 100. valid, confed-external Extended Community: RT: 1100:1001
```

```
mpls labels in/out nolabel/362
```

```
(65001 64955 65003) 650 89
```

```
172.16.254.226 (metric 20645) from 172.16.216.253(172.16.216.253) Origin IGP,
```

```
metric 0. localpref 100. valid, confed-external Extended Community: RT: 1100:1001
```

```
mpls labels in/out nolabel/362
```

```
(65001 64955 65003) 650 89
```

```
172.16.254.226 (metric 20645) from 172.16.216.252 (172.16.216.252) Origin IGP,
```

```
metric 0, localpref 100, valid, confed-external Extended Community: RT: 1100:1001
```

```
mpls labels in/out nolabel/362
```

```
[64955 65003) 650 89
```

```
172.16.254.226 (metric 20645) from 10.77.255.57 (10 77.255.57) Origin IGP, metric
```

```
0, localpref 100. valid, confed-external Extended Community RT: 1100:1001 mpls
```

```
labels.in/out nolabel/362
```

(64955 65003) 650 89

172.16.254.226 (metric 20645) from 10.57.255.11 (10.57.255.11) Origin IGP, metric 0, localpref 100, valid, confed-external, best Extended Community RTTI 00:1001 mpls labels in/out nolabel/362

(64955 65003) 650 89

172.16.254.226 (metric 20645) from i 72.16.224.253 (172.16.224.253) Origin IGP. metric 0: localpref 100, valid, confed-external Extended Community RT: 1100:1001 mpls labels in/out nolabel/362
165003)65089

An engineer configured BGP and wants to select the path from 10.77.255.57 as the best path instead

of current best path. Which action resolves the issue?

- A. Configure AS_PATH prepend for the current best path
- B. Configure higher MED to select as the best path
- C. Configure AS_PATH prepend for the desired best path
- D. Configure lower LOCAL_PREF to select as the best path

Answer: D

Explanation:

Explanation

From the output, we learn that the current best path is from 10.57.255.11 (which includes "...valid, confed-external, best")

and this path is 2 ASes away (64955 65003). Although there are some paths with only 1 AS away (path from

172.16.254.234 for example) but they were not chosen the best path so AS_PATH was not used to determine the best path

-> Answers A and answer C are not correct.

All the paths in the output have metric of 0 and this is the lowest (best) value for this attribute. If we configure higher MED

then it is less preferred over other paths -> Answer B is not correct.

Only answer D is left but LOCAL_PREF attribute should be configured with higher value to be preferred

so we hope "lower LOCAL_PREF" here means higher value. But this is the best answer.

Question: 298

Refer to the exhibit.

```
Ft2(config)# int tun0
```

```
Mun 23 00:42:06.179: %LINEPROTO-5-UPDOWN: Line protocol on  
Interface Tunnel0, changed state to down
```

```
R2(config-if)# ip address 192.168.12.2 255.255.0
```

```
R2(config-if)# tunnel source lo0
```

```
R2(config-if)# tunnel destination 10.255.255.1
```

```
Mun 23 00:42:15.845: %LINEPROTO-5-UPDOWN: Line protocol on  
Interface Tunnel0, changed state to up
```

```
R2(config-if)# router eigrp 1
```

```
R2(config-router)# address-family ipv4 autonomous-system 1  
R2(config-router-af)# net 192.168.12.2 0.0.0.0
```

```
Mun 23 00:43:05.730: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1:  
Neighbor 192.168.2.1 (Tunnel0) is up: new adjacency
```

```
* Jun 23 00:43:05.993: %ADJ-5-PARENT Midchain parent maintenance  
for IP midchain out of Tunnel0 - looped chain attempting to stack  
Mun 23 00:43:15.193: %TUN-5-RECURDOWN: Tunnel0 temporarily disabled  
due to recursive routing
```

```
'Jun 23 00:43:15.193: %LINEPROTO-5-UPDOWN: Line protocol on  
Interface Tunnel0. changed state to down
```

An administrator is configuring a GRE tunnel to establish an EIGRP neighbor to a remote router. The other tunnel endpoint is already configured. After applying the configuration as shown, the tunnel started flapping. Which action resolves the issue?

- A. Modify the network command to use the Tunnel0 interface netmask
- B. Advertise the Loopback0 interface from R2 across the tunnel
- C. Stop sending a route matching the tunnel destination across the tunnel
- D. Readdress the IP network on the Tunnel0 on both routers using the /31 netmask

Answer: C

Explanation:

Explanation

In this question we are advertising the tunnel IP address 192.168.12.2 to the other side. When other end receives the EIGRP advertisement, it realizes it can reach the other side of the tunnel via EIGRP.

In other words, it reaches the tunnel destination through the tunnel itself -> This causes "recursive routing" error.

Note: In order to avoid this error, do not advertise the tunnel destination IP address on the tunnel interface to other side.

Good recursive routing reference: <https://networklessons.com/cisco/ccie-routing-switching/gretunnel-recursive-routing-error>

Question: 299

Which two solutions are used to overcome a flapping link that causes a frequent label binding exchange between MPLS routers? (Choose two)

- A. Create link dampening on links to protect the session.
- B. Increase input queue on links to protect the session.

- C. Create targeted hellos to protect the session.
- D. Increase a hold-timer to protect the session.
- E. Increase a session delay to protect the session.

Answer: A,C

Explanation:

Explanation

To avoid having to rebuild the LDP session altogether, you can protect it. When the LDP session between two directly connected LSRs is protected, a targeted LDP session is built between the two LSRs. When the directly connected link does

go down between the two LSRs, the targeted LDP session

is kept up as long as an alternative path exists between the two LSRs.

For the protection to work, you need to enable it on both the LSRs. If this is not possible, you can enable it on one LSR,

and the other LSR can accept the targeted LDP Hellos by configuring the command `mpls ldp discovery`

`targeted-hello accept`.

Reference: <https://www.ccexpert.us/mps-network/mps-ldp-session-protection.html>

Or from the reference

at <https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdf>

Troubleshooting LDP Issues

Problem:

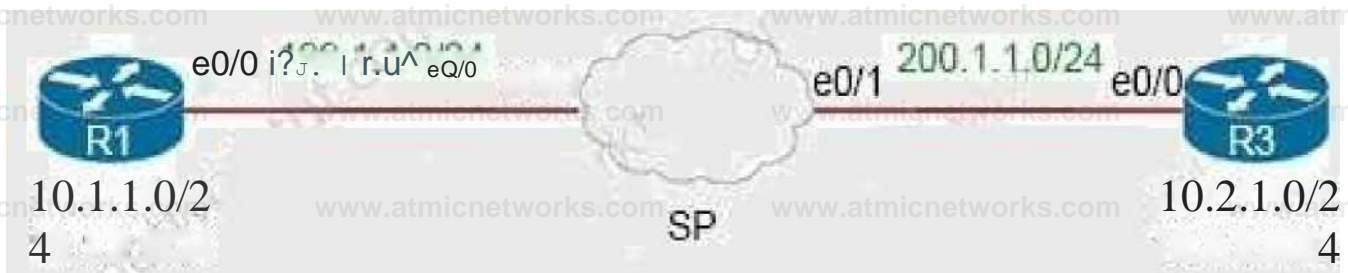
I. When a link flaps (for a short time),

Solution:

+ When LDP session supported by link hello is setup, create a targeted hello to protect the session.

Question: 300

Refer to the exhibit.



An engineer must configure a LAN-to-LAN IPsec VPN between R1 and the remote router. Which IPsec

Phase 1 configuration must the engineer use for the local router?

- A. `crypto isakmp policy 5 authentication pre-share encryption 3 des hash sha group 2 ! crypto isakmp key cisco123 address 200.1.1.3`
- B. `crypto isakmp policy 5 authentication pre-share encryption 3 des hash md5 group 2 ! crypto isakmp key cisco123 address 200.1.1.3`
- C. `crypto isakmp policy 5 authentication pre-share encryption 3 des hash md5 group 2 ! crypto isakmp key cisco123 address 199.1.1.1`
- D. `crypto isakmp policy 5 authentication pre-share encryption 3 des hash md5 group 2 ! crypto isakmp key cisco123! address 199.1.1.1`

Answer: A

Explanation:

Explanation

In the "crypto isakmp key ... address" command, the address must be of the IP address of the other end (which is 200.1.1.3 in this case) so Option A and Option B are correct. The difference between these two options are in the hash SHA or MD5 method but both of them can be used although SHA is better than MD5 so we choose Option A the best answer.

Note: Cisco no longer recommends using 3DES, MD5 and DH groups 1, 2 and 5.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_imgmt/configuration/xr-16-5/sec-ipsec-management-xe-16-5-book/sec-ipsec-usability-enhance.html

Question: 301

What is a function of an end device configured with DHCPv6 guard?

- A. If it is configured as a server, only prefix assignments are permitted.
- B. If it is configured as a relay agent, only prefix assignments are permitted.
- C. If it is configured as a client, messages are switched regardless of the assigned role.
- D. If it is configured as a client, only DHCP requests are permitted.

Answer: C

Explanation:
Explanation

The DHCPv6 Guard feature blocks reply and advertisement messages that come from unauthorized DHCP servers and relay agents.

Packets are classified into one of the three DHCP type messages. All client messages are always

switched regardless of device role. DHCP server messages are only processed further if the device role

is set to server. Further processing of server messages includes DHCP server advertisements (for

source validation and server preference) and DHCP server replies (for permitted prefixes).

If the device is configured as a DHCP server, all the messages need to be switched, regardless of the device role configuration.

Question: 302

A customer requested a GRE tunnel through the provider network between two customer sites using loopback to hide internal networks. Which configuration on R2 establishes the tunnel with R1?

- A. R2(config)# interface Tunnel 1R2(config-if)# ip address 172.20.1.2 255.255.255.0R2(config-if)# ip mtu 1400R2(config-if)# ip tcp adjust-mss 1360R2(config-if)# tunnel source 192.168.20.1R2(config-if)# tunnel destination 192.168.10.1
- B. R2(config)# interface Tunnel 1R2(config-if)# ip address 172.20.1.2 255.255.255.0R2(config-if)# ip mtu 1400R2(config-if)# ip tcp adjust-mss 1360R2(config-if)# tunnel source 10.10.2.2R2(config-if)# tunnel destination 10.10.1.1
- C. R2(config)# interface Tunnel 1R2(config-if)# ip address 172.20.1.2 255.255.255.0R2(config-if)# ip mtu 1500R2(config-if)# ip tcp adjust-mss 1360R2(config-if)# tunnel source 192.168.20.1R2(config-if)# tunnel destination 10.10.1.1
- D. R2(config)# interface Tunnel 1R2(config-if)# ip address 172.20.1.2 255.255.255.0R2(config-if)# ip mtu 1500R2(config-if)# ip tcp adjust-mss 1360R2(config-if)# tunnel source 10.10.2.2R2(config-if)# tunnel destination 10.10.1.1

Answer: D

Explanation:

Question: 303

A network administrator added a new spoke site with dynamic IP on the DMVPN network. Which configuration command passes traffic on the DMVPN tunnel from the spoke router?

- A. ip nhrp registration ignore
- B. ip nhrp registration no-registration
- C. ip nhrp registration dynamic
- D. ip nhrp registration no-unique

Answer: D

Explanation:

Question: 304

Which IPv6 feature enables a device to reject traffic when it is originated from an address that is not stored in the device binding table?

- A. IPv6 Snooping
- B. IPv6 Source Guard
- C. IPv6 DAD Proxy
- D. IPv6 RA Guard

Answer: B

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xr-3s/ipv6-xr-3s-book/ipv6-src-guard.html

R1#sh ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, 0 - 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, LI - 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, H - NHRP, 1 - LISP a - application route

+ - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

D 10.0.0.0/8 [90/409600] via 172.16.1.200, 00:00:28, Ethernet0/0
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks

C 172.16.1.0/24 is directly connected, Ethernet0/0

L 172.16.1.100/32 is directly connected, Ethernet0/0

172.17.0.0/24 is variably subnetted, 2 subnets, 2 masks

C 192.168.1.0/24 is directly connected, Loopback0

L 192.168.1.100/32 is directly connected, Loopback0

RI#

The R2 loopback interface is advertised with RIP and EIGRP using default values. Which configuration changes make R1 reach the R2 loopback using RIP?

- A. R1(config)# router ripR1(config-router)# distance 90
- B. R1(config)# router ripR1(config-router)# distance 100
- C. R1(config)# router eigrp 1R1(config-router)# distance eigrp 130 120
- D. R1(config)# router eigrp 1R1(config-router)# distance eigrp 120 120

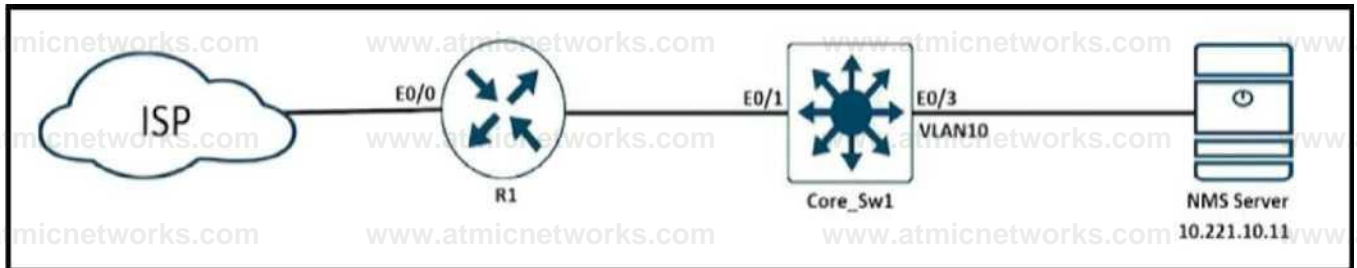
Answer: C

Explanation:

distance (AD Number u want to change to) (neighbor IP) (Wildcard Mask) (access-list number)

Question: 306

Refer to the exhibit.



During ISP router maintenance, the network produced many alerts because of the flapping interface. Which configuration on R1 resolves the issue?

- A. no snmp trap link-status
- B. snmp trap link-status down
- C. snmp trap ip verify drop-rate
- D. ip verify drop-rate notify hold-down 60

Answer: D

Explanation:

Question: 307

Refer to the exhibit.

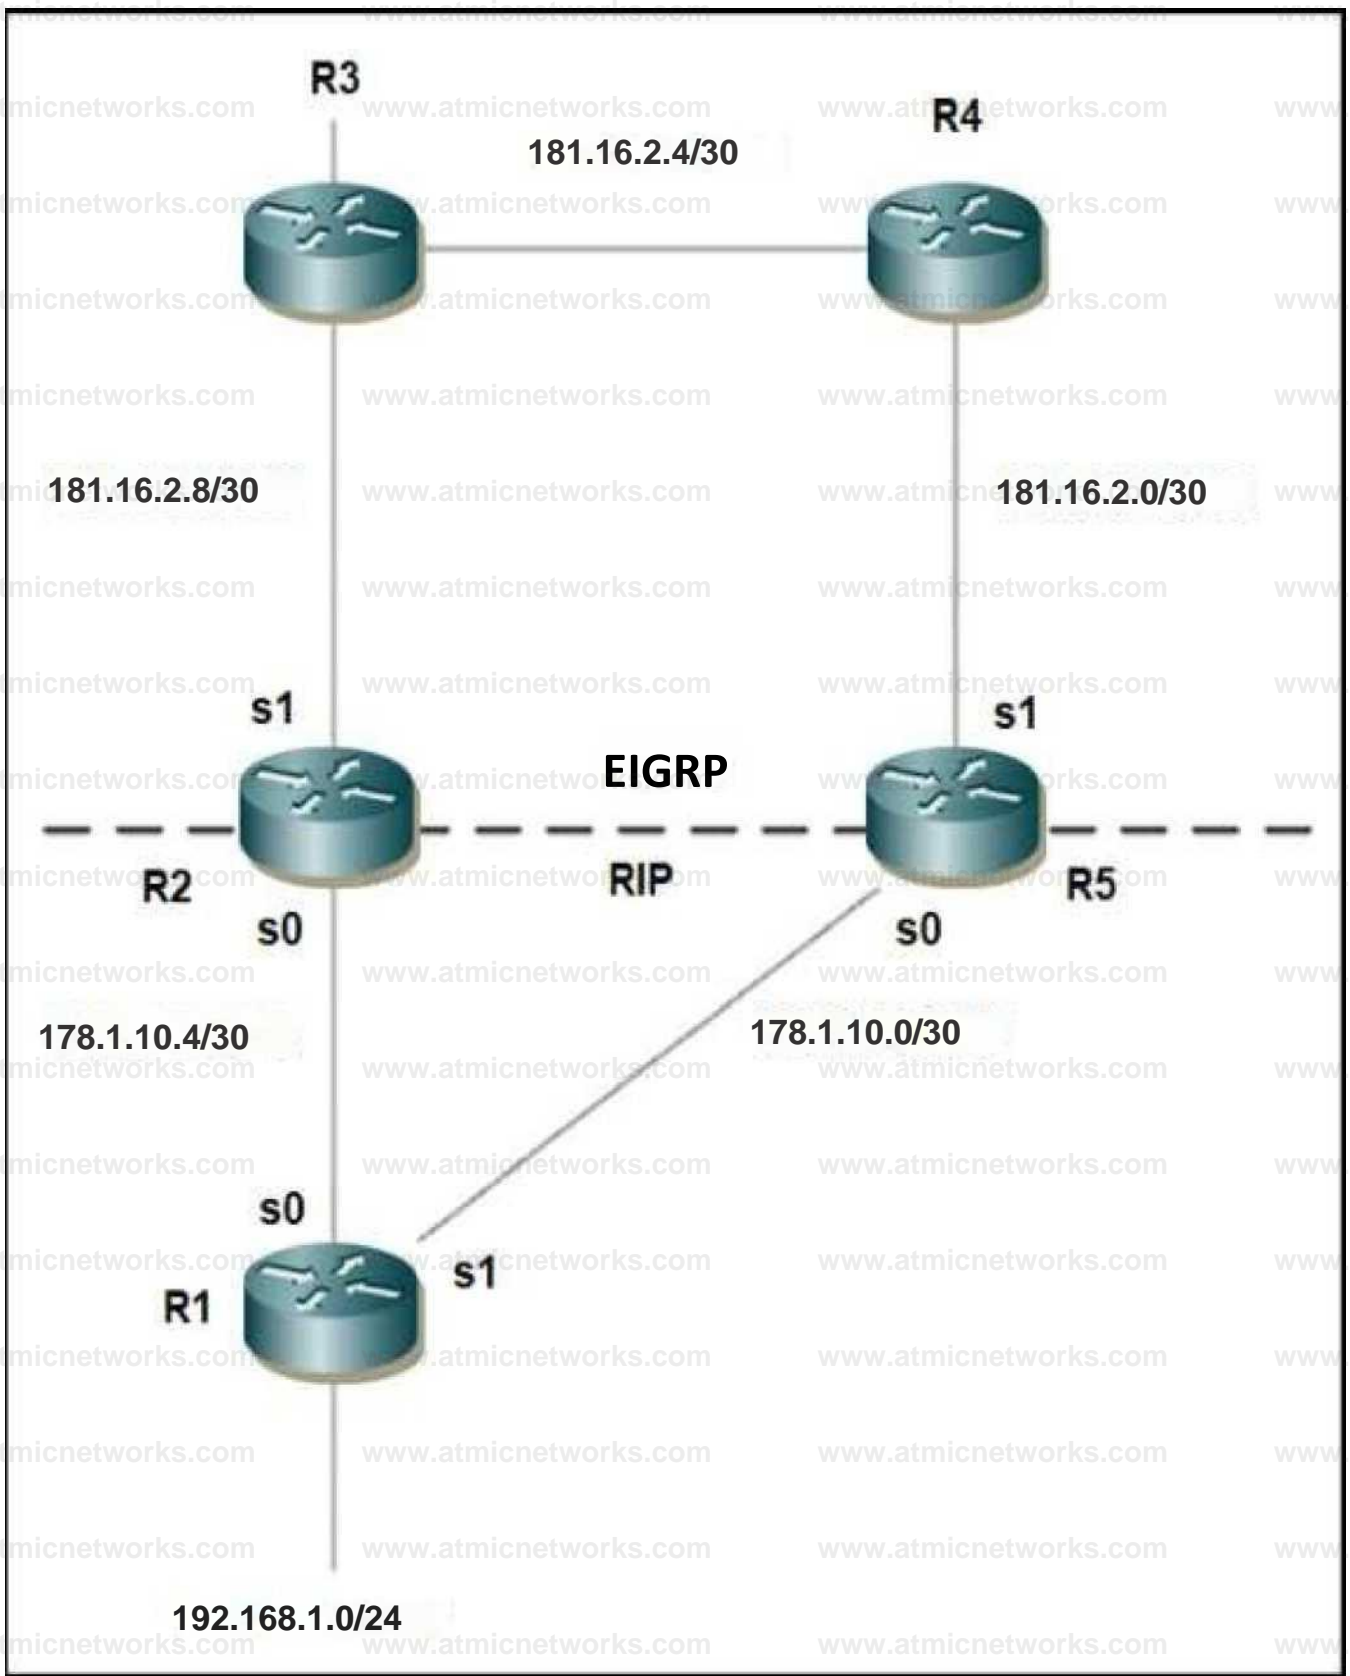
```
ip vrf CCNP
rd 1:1
interface Ethernet1
ip vrf forwarding CCNP
ip address 10.1.1.1 255.255.255.252
!
interface Ethernet2
ip vrf forwarding CCNP
ip address 10.2.2.2 255.255.255.252
```

Which configuration enables OSPF for area 0 interfaces to adjacency with a neighboring router with the same VRF?

- A. router ospf 1 vrf CCNPinterface Ethernet1ip ospf 1 area 0.0.0.0interface Ethernet2ip ospf 1 area 0.0.0.0
- B. router ospf 1interface Ethernet1ip ospf 1 area 0.0.0.0interface Ethernet2ip ospf 1 area 0.0.0.0
- C. router ospf 1 vrf CCNPnetwork 10.1.1.1 0.0.0.0 area 0network 10.2.2.2 0.0.0.0 area 0
- D. router ospf 1 vrf CCNPnetwork 10.0.0.0 0.0.255.255 area 0

Answer: C

Explanation:



Mutual redistribution is enabled between RIP and EIGRP on R2 and R5. Which configuration resolves the routing loop for the 192.168.1.0/24 network?

- A. R2:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s1!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit
anyR5:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s0!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit any
- B. R2:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s0!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit
anyR5:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s0!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit any
- C. R2:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s0!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit
anyR5:router eigrp 10network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s1!router
ripnetwork 178.1.0.0redistribute eigrp 10 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit any
- D. R2:router eigrp 7network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s1!router
ripnetwork 178.1.0.0redistribute eigrp 7 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit
anyR5:router eigrp 7network 181.16.0.0redistribute rip metric 1 1 1 1distribute-list 1 in s1!router
ripnetwork 178.1.0.0redistribute eigrp 7 metric 2!access-list 1 deny 192.168.1.0access-list 1 permit any

Answer: D

Explanation:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html>

Question: 309

Refer to the exhibit.



An engineer must advertise routes into IPv6 MP-BGP and failed. Which configuration resolves the issue on R1?

- A. `router bgp 65000 no bgp default ipv4-unicast address-family ipv6 multicast network 2001:DB8::/64`
- B. `router bgp 65000 no bgp default ipv4-unicast address-family ipv6 unicast network 2001:DB8::/64`
- C. `router bgp 64900 no bgp default ipv4-unicast address-family ipv6 unicast network 2001:DB8::/64`
- D. `router bgp 64900 no bgp default ipv4-unicast address-family ipv6 multicast neighbor 2001:DB8:7000::2 translate-update ipv6 multicast`

Answer: B

Explanation:

Question: 310

An engineer failed to run diagnostic commands on devices using Cisco DNA Center. Which action in Cisco DNA Center resolves the issue?

- A. Enable Command Runner

B. Enable APIs

C. Enable CDP

D. Enable Secure Shell

Answer:

A

Explanation:

Question:

311

Which two components are required for MPLS Layer 3 VPN configuration? (Choose two)

A. Use pseudowire for Layer 2 routes

B. Use MP-BGP for customer routes

C. Use OSPF between PE and CE

D. Use a unique RD per customer VRF

E. Use LDP for customer routes

Answer:

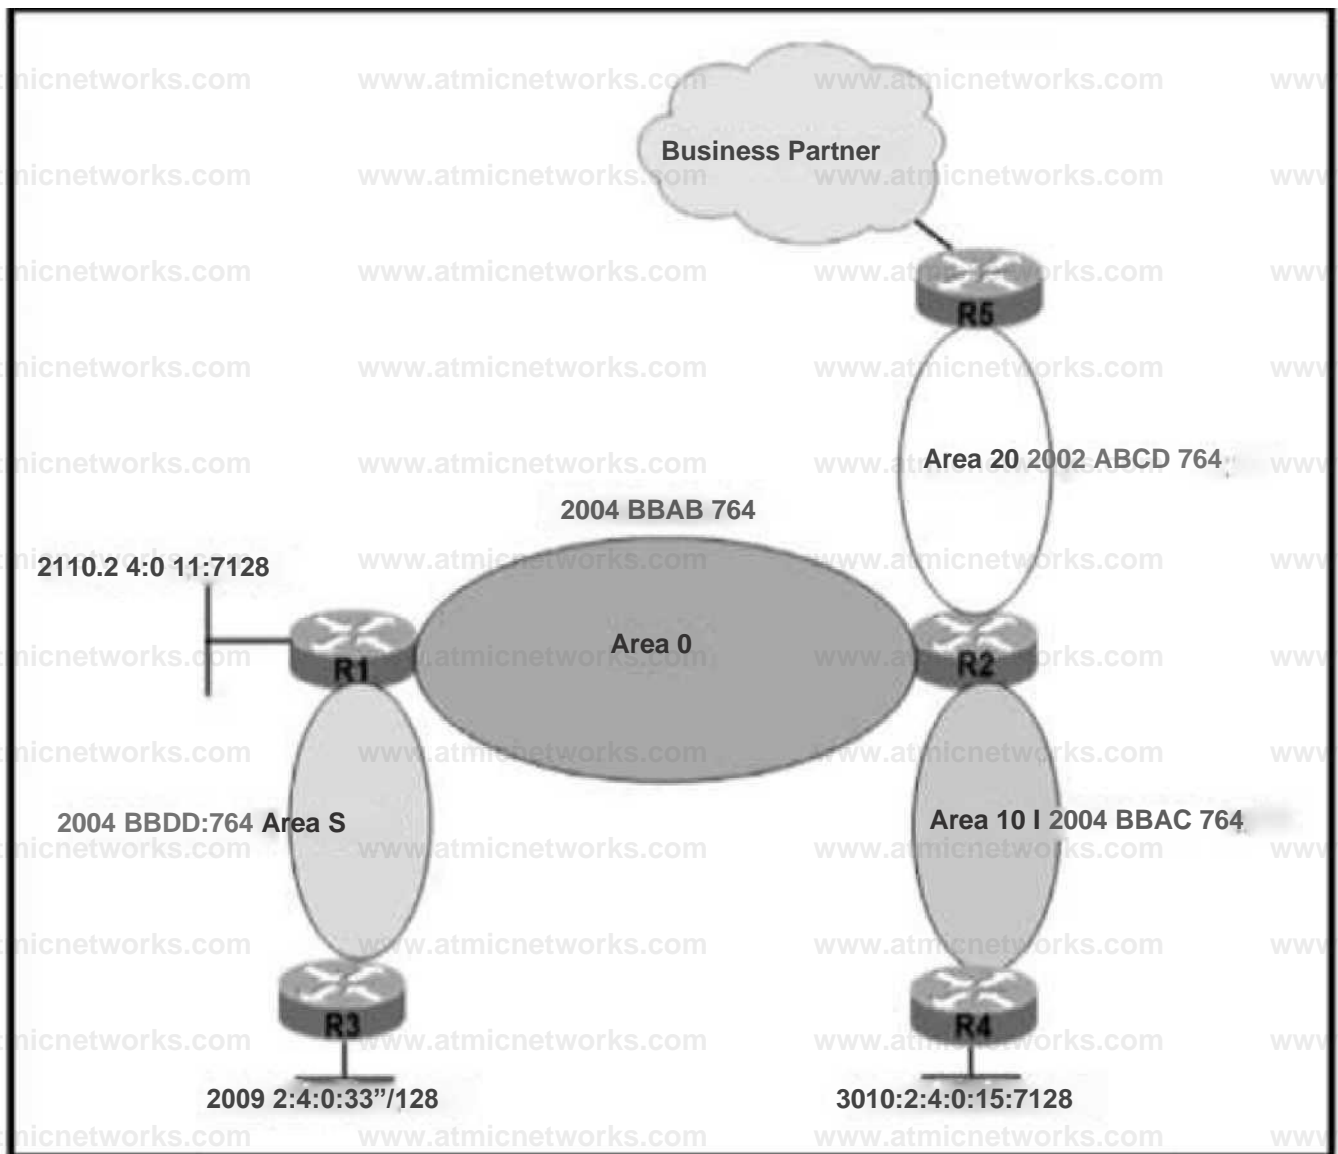
C,D

Explanation:

Question:

312

Refer to the exhibit.



```

R28wh lpv6 mute oapf
O 2002.ABCD;;/64 [110/11
  via FaetKthorMtO/1, directly connected
O 2004:BBAB::/64 1110/11
  via Fast Ethernet 0/0, directly connected O
2004:BBAC: :/M 1110/11
  via FastEthernet 1/0, directly connected O 3010:2.4
0.15:;/12# [110/1]
  via FE80::C804:IDFF:FE20:8r FaetEthernetotO/O

```


A network engineer applied a filter for LSA traffic on OSPFv3 interarea routes on the area 5 ABR to protect advertising the internal routes of area 5 to the business partner network. All other areas should receive the area 5 internal routes. After the respective route filtering configuration is applied on the ABR, area 5 routes are not visible on any of the areas. How must the filter list be applied on the ABR to resolve this issue?

- A. in the "in" direction for area 5 on router R1
- B. in the "out" direction for area 5 on router R1
- C. in the "in" direction for area 20 on router R2
- D. in the "out" direction for area 20 on router R2

Answer: D

Explanation:

Question: 313

Refer to the exhibit.

```
ipv6 dhcp pool DHCPPOOL
address prefix 2001:0:1:4::/64 lifetime infinite infinite
```

```
interface FastEthernet0/0
ipaddress 10.0.0.1 255.255.255.240
duplex auto
speed auto
ipv6 address 2001:0:1:4:: 1/64
ipv6 enable
ipv6 nd ra suppress
ipv6 ospf 1 area 1
ipv6 dhcp server DHCPPOOL
```

Reachability between servers in a network deployed with DHCPv6 is unstable. Which command must be removed from the configuration to make DHCPv6 function?

- A. ipv6 dhcp server DHCPPOOL
- B. ipv6 address 2001:0:1:4::/64
- C. ipv6 nd ra suppress
- D. address prefix 2001:0:1:4::/64 lifetime infinite infinite

Answer: C

Explanation:

Question: 314

Refer to the exhibit.

```
ip prefix-list DMZ-STATIC seq 5 permit 10.1.1.0/24
i
route-map DMZ permit 10
    match ip address prefix-list DMZ-STATIC i
router ospf 1
network 0.0.0.0 0.0.0.0 area 0
redistribute static route-map DMZ i
ip route 10.1.1.0 255.255.255.0 10.20.20.1
```

The static route is not present in the routing table of an adjacent OSPF neighbor router. Which action resolves the issue?

- A. Configure the next hop of 10.20.20.1 in the prefix list DMZ-STATIC
- B. Configure the next-hop interface at the end of the static router for it to get redistributed
- C. Configure a permit 20 statement to the route map to redistribute the static route
- D. Configure the subnets keyword in the redistribution command

Answer: D

Explanation:

Question: 315

Refer to the exhibit.

ACL for CoPP Routing class-map

```
access-list 120 permit tcp any gt 1024 eq bgp log
access-list 120 permit tcp any bgp gt 1024 established
access-list 120 permit tcp any gt 1024 eq 639
access-list 120 permit tcp any eq 639 gt 1024 established
access-list 120 permit tcp any eq 646
access-list 120 permit udp any eq 646
access-list 120 permit ospf any
access-list 120 permit ospf any host 224.0.0.5
access-list 120 permit ospf any host 224.0.0.6
access-list 120 permit eigrp any
access-list 120 permit eigrp any host 224.0.0.10
access-list 120 permit udp any any eq pim-auto-rp
```

The control plane is heavily impacted after the CoPP configuration is applied to the router. Which command removal lessens the impact on the control plane?

- A. access-list 120 permit udp any any eq pim-auto-rp
- B. access-list 120 permit eigrp any host 224.0.0.10
- C. access-list 120 permit ospf any
- D. access-list 120 permit tcp any gt 1024 eq bgp log

Answer: A

Explanation:

Question: 316

Refer to the exhibit.

```
snmp-server community Public RO 90
snmp-server community Private RW 90
```

R1#show access-list 90

```
Standard IP access list 90
  permit 10.11.110.11
  permit 10.11.111.12
```

```
Nov 6 06:45:11: %SNMP-3-AUTHFAIL: Authentication failure for SNMP req from host
10.11.110.12
```

```
Nov 6 06:45:12: %SNMP-3-AUTHFAIL: Authentication failure for SNMP req from host
10.11.110.12
```

A network administrator notices these console messages from host 10.11.110.12 originating from interface E1/0. The administrator considers this an unauthorized attempt to access SNMP on R1. Which action prevents the attempts to reach R1 E1/0?

- A. Configure IOS control plane protection using ACL 90 on interface E1/0
- B. Configure IOS management plane protection using ACL 90 on interface E1/0
- C. Create an inbound ACL on interface E1/0 to deny SNMP from host 10.11.110.12
- D. Add a permit statement including the host 10.11.110.12 into ACL 90

Answer: C

Explanation:

Question: 317

Refer to the exhibit.

```
CPE# ping 10.0.2.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.2.4, timeout is 2
seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
1/1/1 ms
CPE# copy flash:/packages.conf tftp://10.0.2.4/
Address or name of remote host [10.0.2.4]?
Destination filename [packages.conf]?
%Error opening tftp://10.0.2.4/packages.conf (Undefined error)
```

The administrator is trying to overwrite an existing file on the TFTP server that was previously uploaded by another router. However, the attempt to update the file fails. Which action resolves this issue?

- A. Make the packages.conf file executable by all on the TFTP server
- B. Make the packages.conf file writable by all on the TFTP server
- C. Make sure to run the TFTP service on the TFTP server
- D. Make the TFTP folder writable by all on the TFTP server

Answer: B

Explanation:

Question: 318

Refer to the exhibit.

```
R2#show ip route
```

```
Gateway of last resort is not set
10.0.0.0/8 is variably subnetted, 12 subnets, 3 masks
C    10.1.3.0/30 is directly connected, FastEthernet0/1
C    10.1.2.0/30 is directly connected, FastEthernet0/0
C    10.1.1.0/30 is directly connected, FastEthernet1/0
O E2 10.19.0.0/24 [110/20] via 10.1.3.2, 00:02:04, FastEthernet0/1
D    10.55.13.0/24 (90/4096001 via 10.1.2.2. 00:01:00. FastEthernet0/0
D    10.37.100.0/24 (90/4096001 via 10.1.2.2. 00:01:00. FastEthernet0/0
C    10.100.10.0/29 is directly connected, FastEthernet2/0.10
D    10.55.72.0/24 (90/409600] via 10.1.2.2. 00:01:01. FastEthernet0/0
C    10.100.20.0/29 is directly connected. FastEthernet2/0.20
O E2 10.144.1.0/24 /110/201 via 10.1.3.2. 00:12:51. FastEthernet0/1
D    10.55.144.0/24 (90/4096001 via 10.1.2.2. 00:01:01. FastEthernet0/0
O E2 10.123.187.0/24 (110/20] via 10.1.3.2. 00:12:51, FastEthernet0/1
```

```
R2#show ip eigrp topology
```

```
IP-EIGRP Topology Table for AS(100)/TD(10.100.20.2)
```

```
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply, N - Reply Status, S - aia Status
P 10.1.3.0/30, 1 successors, FD is 281600 via Connected, Fast Ethernet 0/1
P 10.1.7.0/30, 1 successors, FD is 281600 via Connected, FastEthernet0/0
P 10.1.1.0/30, 1 successors, to is 28160 via Connected, FastEthernet1/0
P 10.55.13.0/24, 1 successors, FD is 409600 via 10.1.2.2 (409600/1282661, FastEthernet0/0
P 10.37.100.0/24, 1 successors, FD In 409600 via 10.1.2.2 (409600/128256). FastEthernet0/0
P 10.55.72.0/24, 1 successors. ED is 409600 via 10.1.2.2 (409600/128256), FastEthernet0/0
P 10.55.144.0/24, 1 successors, FD is 409600 via 10.1.2.2 (409600/128256), FastEthernet0/0
P 10.123.187.0/24, 0 successors, FD is Inaccessible via 10.1.2.2 (409600/128256), FastEthernet0/0
```

Router R2 should be learning the route for 10.123.187.0/24 via EIGRP. Which action resolves the issue without introducing more issues?

- A. Use distribute-list to modify the route as an internal EIGRP route
- B. Redistribute the route in EIGRP with metric, delay, and reliability

- C. Use distribute-list to filter the external router in OSPF
- D. Remove route redistribution in R2 for this route in OSPF

Answer: C

Explanation:

Question: 319

Refer to the exhibit.


```

R2#show ip eigrp neighbors
1P-EIGRR neighbors for process 100
H   Address          Interface           Hold Uptime       SRTT  ATO  0  Seq
1   142.164.10.1     Ser/0              12 00'00:30      1  5000  2  0
•Jan 1 10:40:21.200: 0DUSL-O-NBRCHANGE: IP-EIGRP(O) 100: Neighbor 102.161.10.1 (teriall/0) is down: retry Unit exceeded
•Jan 1 15:40:51.561: tDUAL-5-NBRCHANGE: IF-EIGRP10) 100: Neighbor 102.140.10.1 (Serlall/0) is up: new adjacency
•Jan 1 15:42:11.101: tOUAL-5-NBRCHANGE: IF-EIGRP0) 100: Neighbor 102.140.10.1 tSorlall/0) la down: retry Unit exceeded
•Jan 1 15:42:14.(10:4DUAL-5-NBRCHANGE: IP-EIGRP(O) 100: Neighbor 102.140.10.1 tSorlall/0) la up: now adjacency

```

RI(show ip eigrp neighbor! IP-BtGRP neighbor! for process 100

RI Configuration:

```

key chain ciaco
key 2
  key-string abc 1
interface Loopback0
ip address 10.10.1.1 255.255.255.0 1
interface Serial/0 ip address 192.168.10.1
255.255.255.0 ip authentication mode eigrp 100
md5

```

```

ip authentication key-chain eigrp 100 ciaco ip address 192.168.10.2 255.255.255.0
serial restart-delay 0 f
router eigrp 100
network 10.10.1.0 0.0.0.255
network 192.168.10.0
no auto-summary

```

R2 configuration:

```

key chain cisco
key 1
  key-string 123
key 2
  key-string abc
1
interface Loopback0 ip address 10.10.2.2
255.255.255.0

```

```

interface Serial/0
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 cisco no fair-queue
1

router eigrp 100
network 10.10.2.0 0.0.0.255
network 192.168.10.0

```

R1 and R2 are configured for EIGRP peering using authentication and the neighbors failed to come up. Which action resolves the issue?

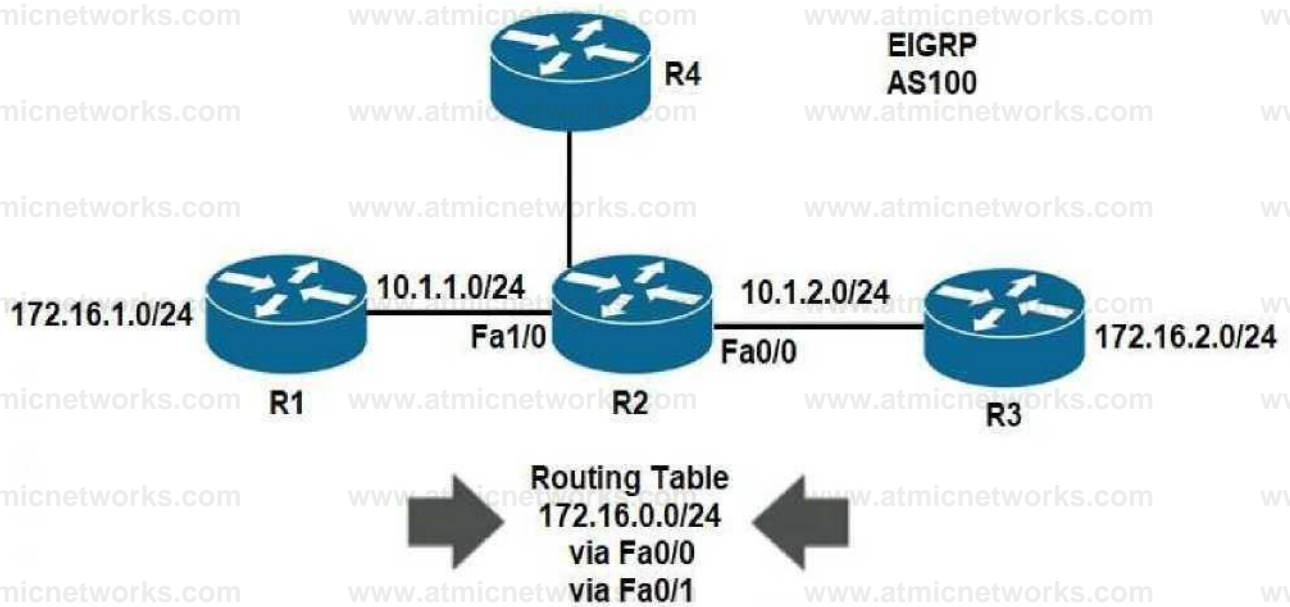
- A. Configure a matching key-id number on both routers
- B. Configure a matching lowest key-id on both routers
- C. Configure a matching key-chain name on both routers
- D. Configure a matching authentication type on both router

Answer: A

Explanation:

Question: 320

Refer to the exhibit.



R4 is experiencing packet drop when trying to reach 172.16.2.7 behind R2. Which action resolves the issue?

- A. Insert a /16 floating static route on R2 toward R3 with metric 254
- B. Insert a /24 floating static route on R2 toward R3 with metric 254
- C. Enable auto summarization on all three routers R1, R2, and R3
- D. Disable auto summarization on R2

Answer: D

Explanation:

Question: 321

Refer to the exhibit.

```
access-list 1 permit 209.165.200.215
access-list 2 permit 209.165.200.216

interface ethernet 1
ip policy route-map Texas

route-map Texas permit 10
match ip address 1
set ip precedence priority
set ip next-hop 209.165.200.217

route-map Texas permit 20
match ip address 2
set ip next-hop 209.165.200.218
```

Packets arriving from source 209.165.200.215 must be sent with the precedence bit set to 1, and packets arriving from source 209.165.200.216 must be sent with the precedence bit set to 5. Which action resolves the issue?

- A. set ip precedence critical in route-map Texas permit 10
- B. set ip precedence critical in route-map Texas permit 20
- C. set ip precedence immediate in route-map Texas permit 10
- D. set ip precedence priority in route-map Texas permit 20


Answer:
B

Explanation:

Question:

322

Refer to the exhibit.



R6 should reach R1 via R5>R2>R1. Which action resolves the issue?

- A. Increase the cost to 61 between R2-R3-R1
- B. Increase the cost to 61 between R2 and R3
- C. Decrease the cost to 2 between R6-R5-R2
- D. Decrease the cost to 41 between R2 and R1

Answer: B

Explanation:

Question: 323

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

Explanation:

Question: 324

Refer to the exhibit.



```
R1(config)#username Admin password 7 Cisco@123
Invalid encrypted password: Cisco@123
```

An engineer is trying to add an encrypted user password that should not be visible in the router configuration. Which two configuration commands resolve the issue? (Choose two)

- A. password encryption aes
- B. username Admin password Cisco@maedeh motamedi
- C. username Admin password 5 Cisco@maedeh motamedi
- D. username Admin secret Cisco@maedeh motamedi
- E. no service password-encryption
- F. service password-encryption

**Answer:
D,F**

Explanation:

Question: 325

Refer to the exhibit.

```

R2#show ip ospf interface brief
interface          PtO   Area   PAMnuMMK
Cost   State   Mts
Lo0    1        1    10002/32
Fa0/0  1        1    101010 1/30
R2#show ip ospf interface l1
interface l1
Buildup configuration.

```

Current configuration 116 bytes

```

interface FastEthernet0/0
ip address 10.10.10.255 255.255.252
ip mtu 1100
ip ospf 1 area 1 duplex mi end
R2#show ip ospf neighbor

```

```

Neighbor to Pri State      Dead Time
10.0.0.1      1 EXSTARDBDR 00:00:37

```

```

R1#show ip ospf interface brief
interface          PtO   Area   PAMnuMMK
Cost   State   Mts
Lo0    1        1    1000102
Fa0/0  1        1    101010 240
R1#show ip ospf interface l1
interface l1
Buildup configuration.

```

Current configuration 115 bytes

```

interface FastEthernet0/0
ip address 10.10.10.255 255.252.0 ospf
1 area 1 duplex auto speed auto
end
R1#show ip ospf neighbor

```

```

Neighbor ID Pri state      Dead Time Address interface
10.0.0.1      1 EXSTARDBDR 00:00:37

```

Which action restores OSPF adjacency between R1 and R2?

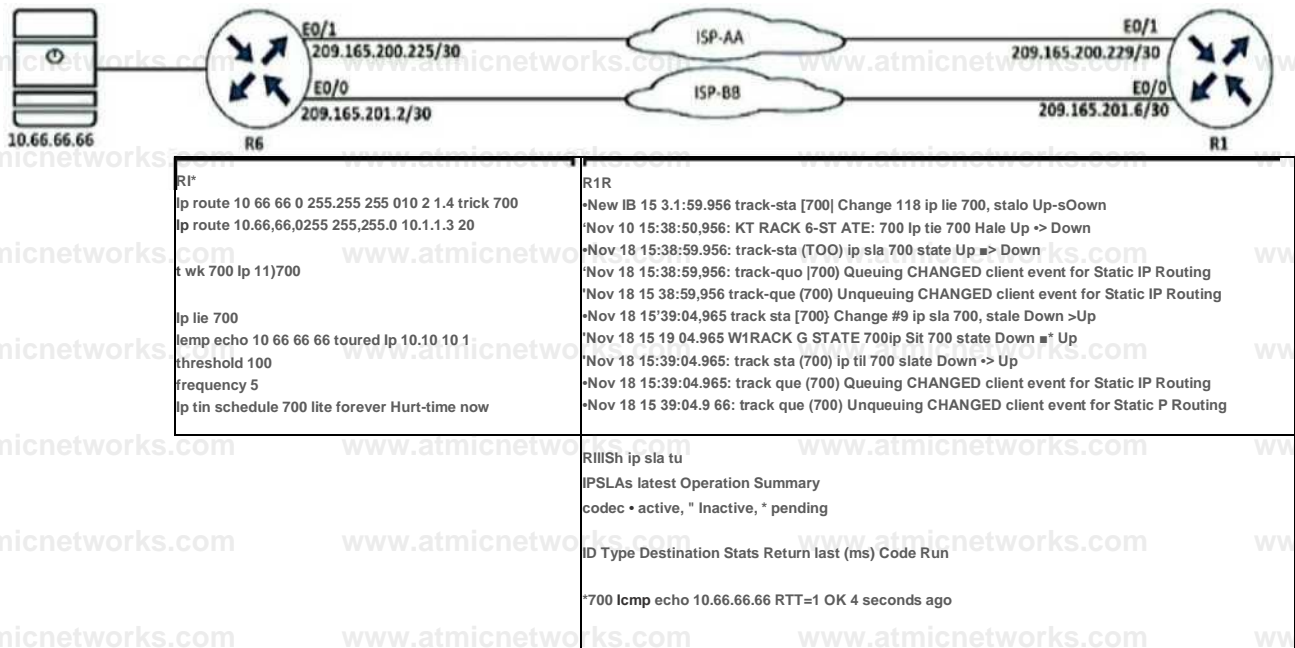
- A. Change the IPMTU of R1 Fa1/0 to 1300
- B. Change the IPMTU of R2 Fa0/0 to 1300
- C. Change the IPMTU of R1 Fa1/0 to 1500
- D. Change the IPMTU of R2 Fa0/0 to 1500

**Answer:
D**

Explanation:

**Question:
326**

Refer to the exhibit.



R1 is configured with IP SLA to check the availability of the server behind R6 but it kept failing. Which configuration resolves the issue?

A. R1(config)# ip sla 700R1(config-track)# delay down 30 up 20

B. R1(config)# ip sla 700R1(config-track)# delay down 20 up 30

C. R1(config)# track 700 ip sla 700R1(config-track)# delay down 30 up 20

D. R1(config)# track 700 ip sla 700R1(config-track)# delay down 20 up 30

Answer:

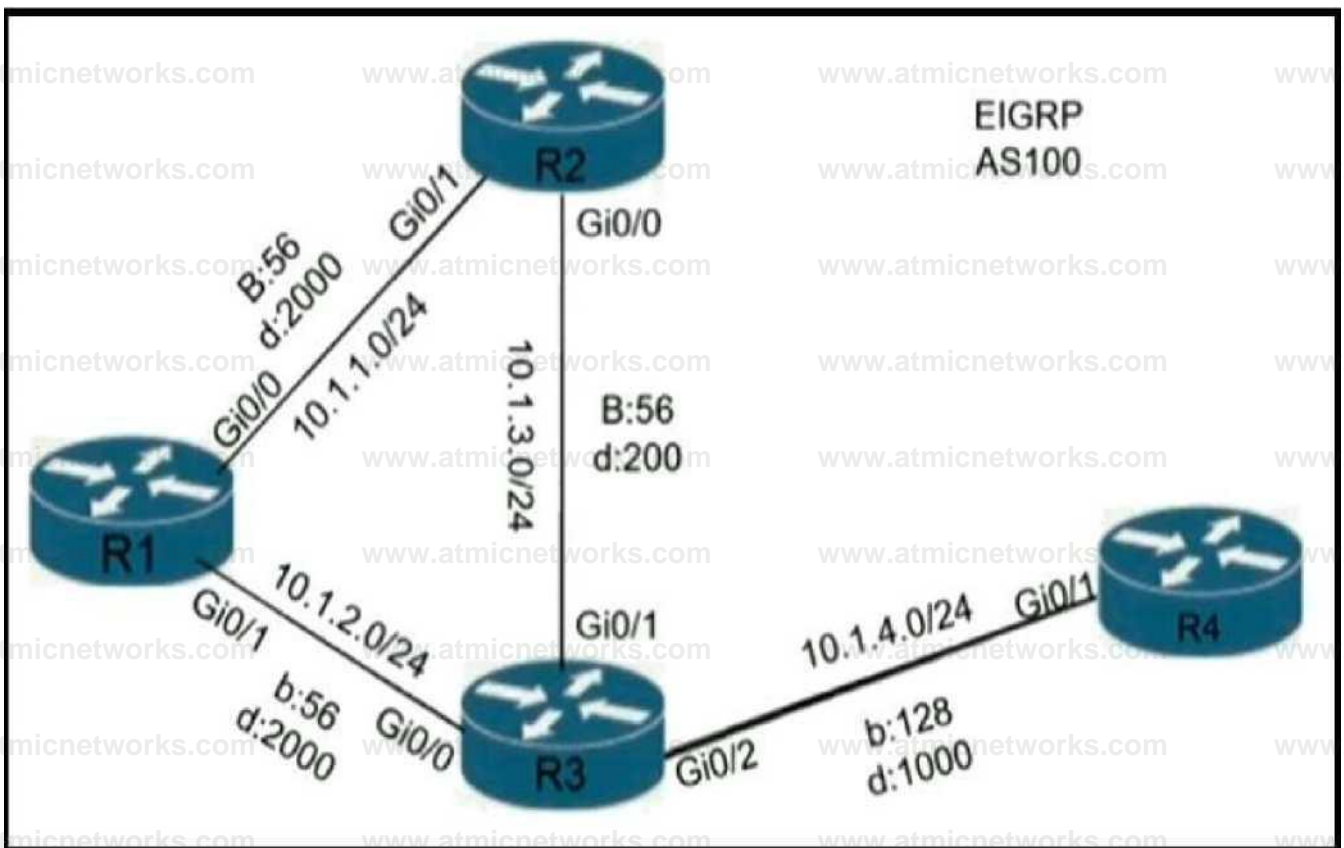
C

Explanation:

Question:

327

Refer to the exhibit.



A loop occurs between R1, R2, and R3 while EIGRP is run with poison reverse enabled. Which action prevents the loop between R1, R2, and R3?

- A. Configure route tagging
- B. Enable split horizon
- C. Configure R2 as stub receive-only
- D. Configure route filtering

Answer: B

Explanation:

Question: 328

A customer reports that traffic is not passing on an EIGRP enabled multipoint interface on a router configured as below:

```
interface Serial0/0
```

```
no ip address
```

```
interface Server0/0/0.9 multipoint ip address 10.1.1.1 255.255.255.248
```

```
ip split-horizon eigrp 1
```

Which action resolves the issue?

- A. Enable poison reverse
- B. Enable split horizon

- C. Disable poison reverse
- D. Disable split horizon

Answer: D

Explanation:

Question: 329

A newly installed spoke router is configured for DMVPN with the ip mtu 1400 command. Which configuration allows the spoke to use fragmentation with the maximum negotiated TCP MTU over GRE?

- A. ip tcp adjust-mss 1360 crypto ipsec fragmentation after-encryption
- B. ip tcp adjust-mtu 1360 crypto ipsec fragmentation after-encryption
- C. ip tcp adjust-mss 1360 crypto ipsec fragmentation mtu-discovery
- D. ip tcp adjust-mtu 1360 crypto ipsec fragmentation mtu-discovery

Answer: A

Explanation:

<https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/111976-dmvpn-troubleshoot-00.html>

Question: 330

What are the two goals of micro BFD sessions? (Choose two.)

- A. The high bandwidth member link of a link aggregation group must run BFD
- B. Run the BFD session with 3x3 ms hello timer
- C. Continuity for each member link of a link aggregation group must be verified
- D. Eny member link on a link aggregation group must run BFD
- E. Each member link of a link aggregation group must run BFD.

Answer: C,E

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-16-8/irb-xe-16-8-book/irb-micro-bfd.html

Question: 331

An engineer configured a router with this configuration

```
ip access-hst DENY TELNET
```

```
10 deny tcp any any eq 23 log-input
```

The router console starts receiving log message :%SEC-6-IPACCESSLOGP: list DENY_TELNET denied tcp

```
192.168.1.10(1022)(FastEthernet1/0 D508.89gb.003f) ->192.168.2.20(23), 1 packet"
```

Which action stops messages on the console while still denying Telnet?

- A. Configure a 20 permit ip any any command
- B. Remove log-Input keyword from the access list.
- C. Replace log-input keyword with the log keyword in the access list.
- D. Configure a 20 permit ip any any log-input command.

Answer: B

Explanation:

Question: 332

Refer to the exhibit.

```
R1#sh run | s bgp
router bgp 65001 no
synchronization bgp router-id
10.100.1.50 bgp log-neighbor-
changes
i network 10.1.1.0 mask
255.255.255.252
I network 10.1.1.12 mask
255.255.255,252
```

```
network 10.100.1.50 mask
255.255.255.255
 timers bgp 20 60
 neighbor R2 peer-group
 neighbor R4 peer-group
 neighbor 10.1.1.2 remote-as
65001 neighbor 10.1.1.2 peer-
group R2
 neighbor 10.1.1.14 remote-as
65001
 neighbor 10.1.1.14 peer-group
R4
 no auto-summary
```

While troubleshooting a BGP route reflector configuration, an engineer notices that reflected routes are missing from neighboring routers. Which two BGP configurations are needed to resolve the issue? (Choose two)

- A. neighbor 10.1.1.14 route-reflector-client
- B. neighbor R2 route-reflector-client
- C. neighbor 10.1.1.2 allowas-in
- D. neighbor R4 route-reflector-client
- E. neighbor 10.1.1.2 route-reflector-client

Answer: A,E

Explanation:

Question: 333

Which IPv6 first hop security feature controls the traffic necessary for proper discovery of neighbor device operation and performance?

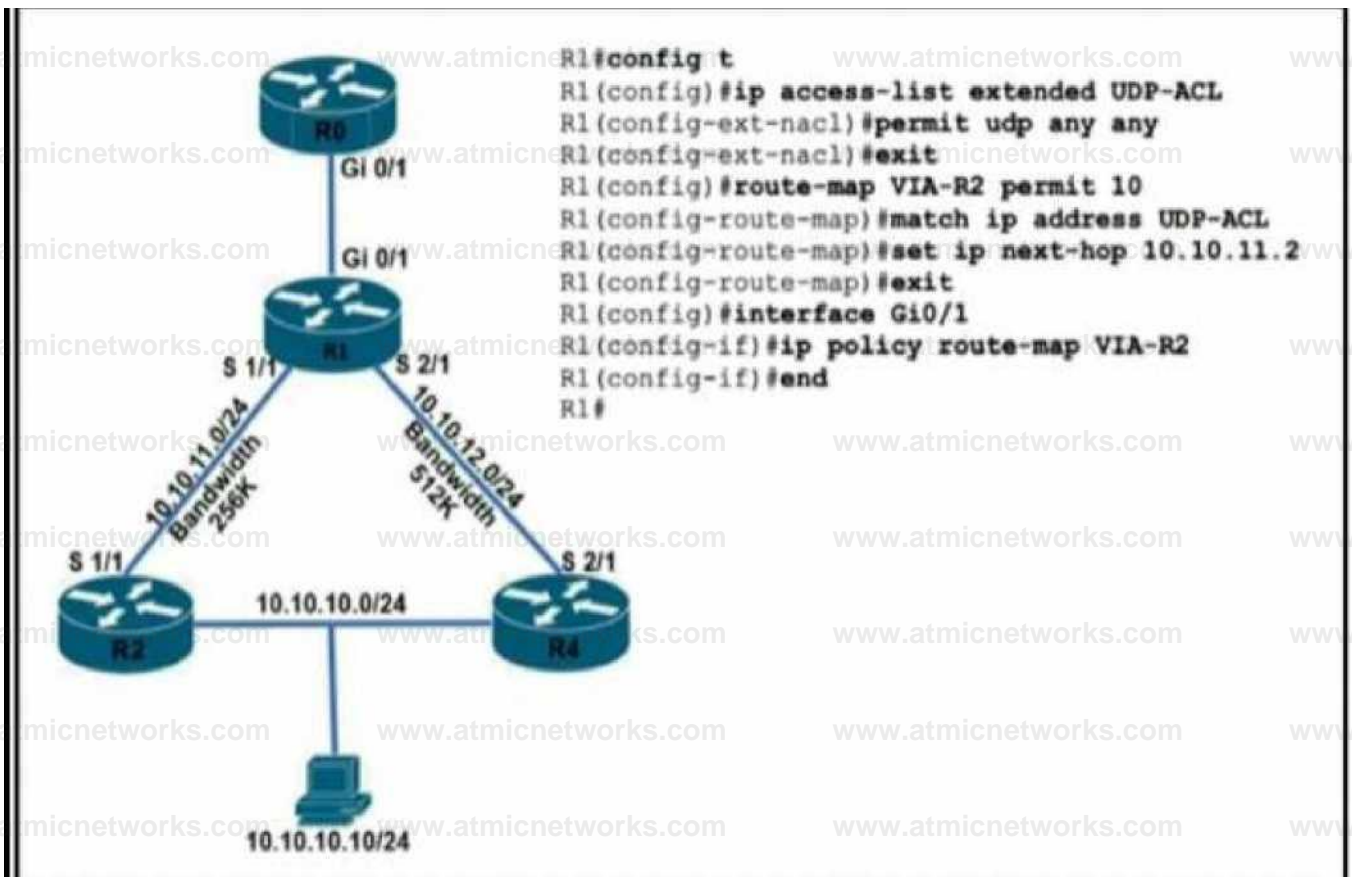
- A. RA Throttling
- B. Source or Destination Guard
- C. ND Multicast Suppression
- D. IPv6 Snooping

Answer: D

Explanation:

Question: 334

Refer to the exhibit.



TCP traffic should be reaching host 10.10.10.10/24 via R2. Which action resolves the issue?

- A. TCP traffic will reach the destination via R2 without any changes
- B. Add a permit 20 statement in the route map to allow TCP traffic
- C. Allow TCP in the access list with no changes to the route map
- D. Set IP next-hop to 10.10.12.2 under the route-map permit 10 to allow TCP traffic.

Answer: C

Explanation:

Question: 335

A network administrator must optimize the segment size of the TCP packet on the DMVPN IPsec protected tunnel interface, which carries application traffic from the head office to a designated branch. The TCP segment size must not overwhelm the MTU of the outbound link. Which configuration must be applied to the router to improve the application performance?

```
interface tunne 130
ip mtu 1400
ip tcp packet-size 1360
```

crypto ipsec fragmentation after-encryption

```
interface tunne 130
ip mtu 1400
ip tcp payload-size 1360
```

crypto ipsec fragmentation before-encryption

```
interface tunneB0
ip mtu 1400
ip tcp adjust-mss 1360
```

crypto ipsec fragmentation after-encryption

```
interface tunne B0
ip mtu 1400
ip tcp max-segment 1360
```

crypto ipsec fragmentation before-encryption

A. Option A

B. Option B

C. Option C

D. Option D

Answer:

C

Explanation:

Question:

336

Refer to the exhibit.

PI* show ip ospf database self-originate

```

toft K?bt<f yW U <10,255,255,11 rp™ec=3 ip 1)
Router Link Stat+c
IM^- o>
Link lb      ADV Rout*?   Age          Seq#         Checksum
Link Hunt 10+255,255,1
            10+255,255,1   «           O<<9W003BD 0*001AD9
3
Smwry Het Link States tArw 0)
Link ID      ADV Router   Age          Seq#         Checksum
14,0.34,0   10.255.255.1 3404        MOT Ct      X&0i:m^
10+255.^,4  10,255.255.1 3004        &XCW#
Type-S AS External
Link States
Link ID      ADV Sauter   Age          Seq#         Checksum
Tig 0.0.0.0  J0,25&.2&5.1 3<<0<
0           OxMOO01DO  Ox001CBC
•Feb 2 2 22: Mt 39.5231 W-'PF-4-FIWP_WAP;Process 1 flushes LSA
ID 0.0.0.0 type-5 adv-rtr 10.255.255.1 in area n

```

After configuring OSPF in R1, some external destinations in the network became unreachable.

Which action resolves the issue?

- A. Clear the OSPF process on R1 to flush stale LSAs sent by other routers.
- B. Change the R1 router ID from 10.255.255.1 to a unique value and clear the process.
- C. Increase the SPF delay interval on R1 to synchronize routes.
- D. Disconnect the router with the OSPF router ID 0.0.0.0 from the network.

Answer:

B

Explanation:

Question:

337

What is the function of BFD?

- A. It provides uniform failure detection regardless of media type.
- B. It creates high CPU utilization on hardware deployments.
- C. It negotiates to the highest version if the neighbor version differs.
- D. It provides uniform failure detection on the same media type.

Answer:


A

Explanation:

Question:

338

Refer to the exhibit.



A network engineer must establish communication between three different customer sites with these requirements:

Site-A: must be restricted to access to any users at Site-B or Site-C.

Site-B and Site-C must be able to communicate between sites and share routes using OSPF.

PE interface configuration.

```
interface FastEthernet0/0
ip vrf forwarding Site-A
```

```
interface FastEthernet0/1
ip vrf forwarding SharedSites
```

```
interface FastEthernet0/2 ip vrf forwarding SharedSites
```

Which configuration meets the requirements?

```
PE(config)#router ospf 10 vrf Site-A
```

```
PE(Config-router)#network 0.0.0.0 255.255.255.255 area 0
```

```
PE(config)#router ospf 10 vrf SharedSites
```

```
PE(config-router)#network 0.0.0.0 255.255.255.255 area 1
```

```
PE(config)#router ospf 10 vrf Site-A
PE(config-router)#network 0.0.0.0 255.255.255.255 area 0
PE(config)#router ospf 10 vrf SharedSites
PE(config-router)#network 0.0.0.0 255.255.255.255 area 0
```

```
PE(config)#router ospf 10 vrf Site-A
PE(config-router)#network 0.0.0.0 255.255.255.255 area 0
PE(config)#router ospf 20 vrf SharedSites
PE(config-router)#network 0.0.0.0 255.255.255.255 area 0
```

```
PE(config)#router ospf 10 vrf Site-A
PE(config-router)#network 0.0.0.0 255.255.255.255 area 0
PE(config)#router ospf 20 vrf SharedSites
PE(config-router)#network 0.0.0.0 255.255.255.255 area 1
```


- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 339

Refer to the exhibit.



I» Mt 1 M24

Site1 Show ip route

```

GaUrvwi) of inV^1 vxt n nc4 w4
IS? th# 1WM I* wnWy '.<nrrtra.? <unrw>ti 2 mm**
C ' <i? IM 1 fl 74 H <Mfci9W*a<C i OOptecM)
I 192 W01102 • <#<>¥ ctt0tedM'ioc*iDariO
D 142 W8 3 0 74 (9i '2>>16< ^ 73 tmw-hteOO
D 192 IM 4 004 |W2S1600J w 192 IM 11 2 00 DO 23 EBWUO 0 192
1M%6'M|W WJBOOI MM 192IM 12 3,00 H 71 fr^wn i (90'4 35 700] via
192 IM 11 2, 00 00 23 UntefwOU
D 1W W# 6 074 100&tfoTy MM 1'12 1M 12 3. 00 DO 2 3 H' ^1 J. i
|1>0MW21>>1 *» 19? <<• 11 ? 00 00 2 3 EthrrwrWD 192 163 11 Gi24
a v*ubv wbrwMd_3 wb *<* 2 num
C ' 92 168 11 0 24 n <lr*Cty OMMctod, Eff rrrnutO'Q
L 192 160 11 1/32 A diiecby COTeteCM, ElOrnHAO
5 13 2 lte no24 ;DM6 320Gj .41 192 166 12 3.60 00 2 3 EFK'rctO 1 MMWi&vW res n 2,
watt), Ejhvwpo
    
```

5h.iT Show Ip myip ktpol^y

```

P i&2 in H - 4 t VJCLI'&w? F DB 2X400
VW 1W1M 'I2(24OI2«2W^ Etwttciao
vw 192 1 SB 12 3 (691200 2048001. Em<ne«yi
P )&? IM 12 171 1 UJCCMW4 Hirst iffIOC
wiCornriri FUwjtrfDJt
p T&2 i6j? :J 024 2 ^uccwm FDH X72P0
    
```

VW 192 IM 12 3 (56 320)/7«M0k Ellwfnetyi
 VM it' <^11? 3WM0^1600) EthMWio0
P is? IM 'CH 1 SKJ -<^< I' D M i?fi25fi
 VW Cvntcw J Lc :i Jbji KQ
 P 192 168 6 024.2 sXCMM»« FQa4 35200
 vw 19 2 168 12 3 (665600 123756? EttwnwtO 1
wi 192 188 11 2 <4 35X« 4096001 EttMffWfOO
P I&ff 168.4 074 I idOO^'v f D!*7.tf4D»
 vw 192 IM 112(2*1600 -2tJ2X ENWHMCI-0
 VW 19 2 tea 12 3 ;C912OC 2044D0ju EthwrwtU I
P 192 168 S0 24 2«UCCMC0ff FDW43MQ0
 vw 192 168 12 3 16«560a 12*25«^ EUxwK) 1
 vw 192 168 11 2 <4 35200 409600>. EthMnrtJU
 P t&J 16 8 11 0 21 1 w#CMUM^ FOW 1536 30
 mConnMtal EtwnwtiM)

Sh*t Sbuw mn | UKUCIII iQut*< vigip route* ««n tOD vwtene« 2
 nvtwotk 192 1€\$ 1 0 hvf^K^X 192 tt| 2 0 r^iwrfk 192 i«4 11 Q fMMCK 192 164 120

Refer to the exhibit. Site1 must perform unequal cost load balancing toward the segments behind Site2 and Site3. Some of the routes are getting load balanced but others are not. Which configuration allows Site1 to load balance toward all the LAN segments of the remote routers?

SiteZ

router eigrp 100 variance 3

Site2

router eigrp 100 variance 2

Site3

**router eigrp 100
 variance 2**

Site1

router eigrp 100 variance 3

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 340

Refer to the exhibit.

```
interface Loopback0 no ip address 100*
```

```
Or OK : 1/10 ip*4 enable ip ospf 10 area 0
```

```
interface Loopback0 no ip address ip* address
```

```
400*:0:400C::1/44 ip ospf 10 area 0
```

```
interface Serial 1/0 no ip address ip address
```

```
10.10.10.10 : 10.10.10.10 ip address enable ip ospf network
```

```
point-to-point ip ospf 10 area 0
```

```
ip traffic-filter DDW TtLkT U4 in
```

```
serial restart-delay 0 clock rate UNO
```

```
ip router ospf 10 router-id 11.1.1.1 ip adjacency-  
change
```

```
ip access-list 1 deny ip 10.10.10.10
```

```
sequence 20 deny tcp host 100.10.10.10 host end
```

```
BL
```

```
interface Loopback0 no ip address ipvt address
```

```
1001 AM:2011:7::1/44 ipv6 enable ipv6 ospf 10
```

```
area 0
```

```
interface Serial 1/0 no ip address
```

```
ipvt address 10.10.10.10 : 10.10.10.10 aux 44 ipv6
```

```
enable
```

```
ip ospf network point-to-point ipvt ospf 10
```

```
area 0 serial restart-delay 0
```

```
ip router ospf 10 router-id 2.2.2.2 log-
```

```
adjacency-changes
```

```
end
```

```
0 40K . 1 eq telnet permit ip any any
```

```

R1#
interface Loopback4
no ip address
ip address 100.1.1.1/24
ipvt enable
ip address 100.1.1.1/24

interface Loopback4
no ip address
ip address 100.1.1.1/24
interface Serial1/3
no ip address
ip address 100.1.1.1/24
ipvt enable
ip address 100.1.1.1/24
ipvt traffic-filter OBIT TEJBT M in serial restart-
delay 0 "
clock rate 44000

ipvt router ospf 10 router-id 1.1.1.1 log-adjacency-changes

ipvt access list LINY Tiua? L04 sequence 20 deny tcp host 100.1.1.1 eq telnet permit
any any
end

ipvt access-list LINDY TtWtTjM
sequence 20 deny tcp host 100.1.1.1 eq telnet permit any any and

```

Refer to the exhibit. An engineer implemented an access list on R1 to allow anyone to Telnet except R2 Loopback0 to R1 Loopback4 How must sequence 20 be replaced on the R1 access list to resolve the issue?

- A. sequence 20 permit tcp host 100.1.1.1/24 host 100.1.1.1 eq telnet
- B. sequence 20 deny tcp host 100.1.1.1/24 host 100.1.1.1 eq telnet
- C. sequence 20 deny tcp host 100.1.1.1/24 host 100.1.1.1 eq telnet
- D. sequence 20 permit tcp host 100.1.1.1/24 host 100.1.1.1 eq telnet

Answer: C

Explanation:

Question: 341

An engineer notices that R1 does not hold enough log messages to identify the root cause during troubleshooting. Which command resolves this issue?

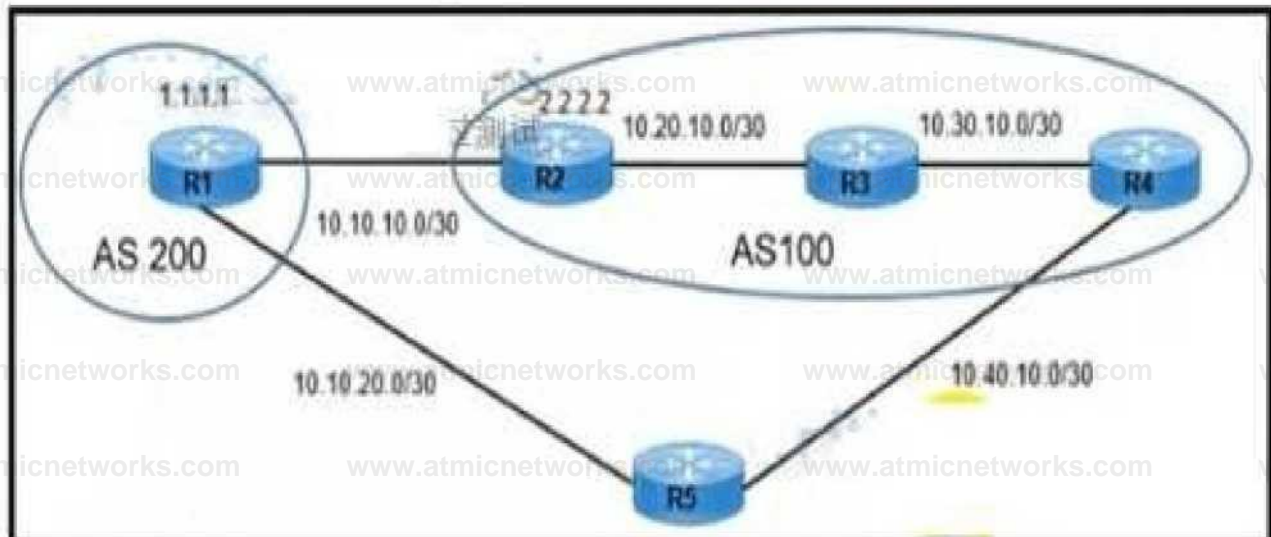
- A. #logging buffered 4096 critical
- B. (config)#logging buffered 16000 informational
- C. #logging buffered 16000 critical
- D. (config)#logging buffered 4096 informational

Answer: B

Explanation:

Question: 342

Refer to the Exhibit.



```

R2#
router eigrp 100
network 10.10.10.0 0.0.0.3
network 10.20.10.0 0.0.0.3
!
router ospf 100
network 10.10.10.0 0.0.0.3 area 0
network 10.20.10.0 0.0.0.3 area 0
!
!
router bgp 100
distance 100 10.20.10.0 0.0.0.3
distance 100 10.10.10.0 0.0.0.3
neighbor 1.1.1.1 remote-as 200
neighbor 10.10.10.1 remote-as 200
network 10.20.10.0 mask 255.255.255.252

R1#
router eigrp 100
network 10.10.10.0 0.0.0.3
network 10.10.20.0 0.0.0.3
network 1.1.1.1 0.0.0.0
!
router ospf 100
network 10.10.10.0 0.0.0.3 area 0
network 10.10.20.0 0.0.0.3 area 0
!
router bgp 200
distance 100 10.10.10.0 0.0.0.3
distance 100 10.20.10.0 0.0.0.3
neighbor 2.2.2.2 remote-as 100
neighbor 10.10.10.2 remote-as 100
network 10.10.10.0 mask 255.255.255.252
network 10.20.10.0 mask 255.255.255.252

```

R1 and R2 use IGP protocol to route traffic between AS 100 and AS 200 despite being configured to use BGP. Which action resolves the issue and ensures the use of BGP?

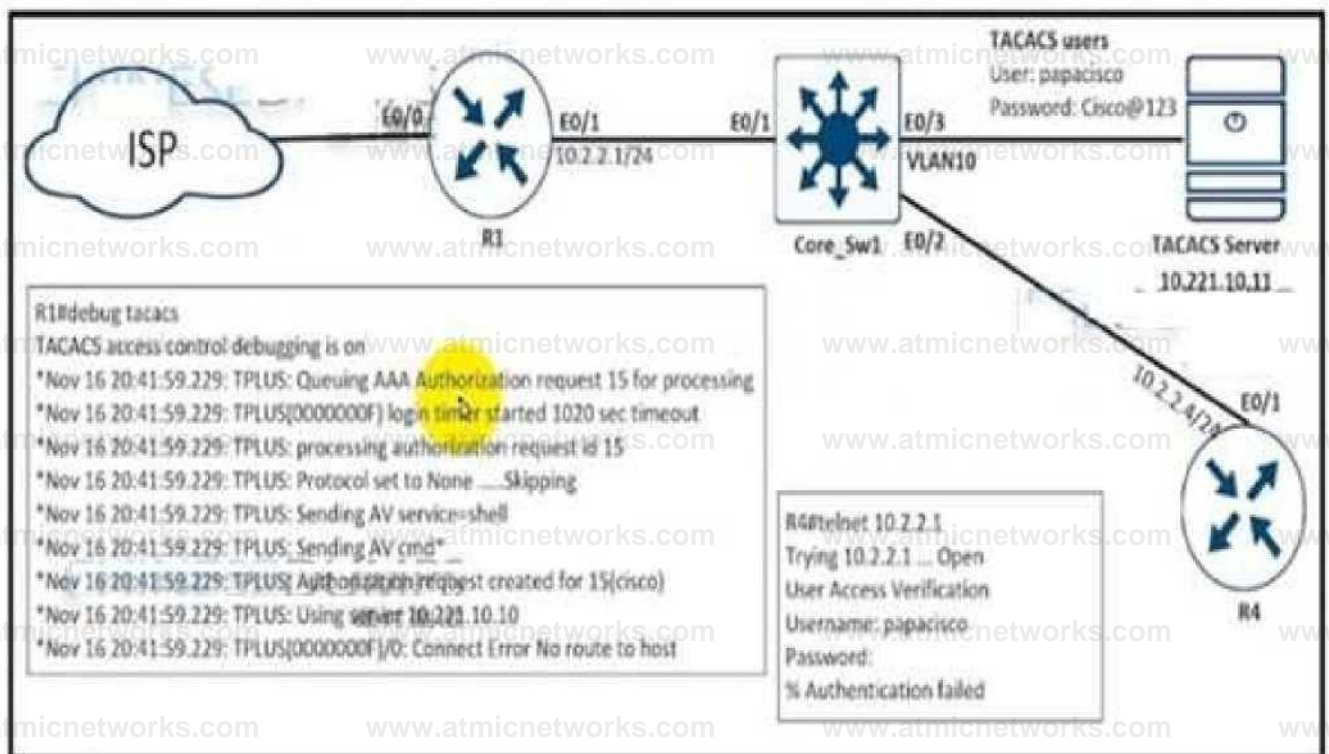
- A. Configure distance to 100 under the EIGRP process of R1 and R2.
- B. Remove distance commands under BGP AS 100 and AS 200.
- C. Remove distance commands under BGP AS 100.
- D. Configure distance to 100 under the OSPF process of R1 and R2

Answer: B

Explanation:

Question: 343

Refer to the exhibit.



An engineer is trying to connect to R1 via Telnet with no success. Which configuration resolves the issue?

```

tacacs server prod
address ipv4 10.221.10.10 exit
  
```

```

ip route 10.221.10.10 255.255.255.255 ether not 0/1
  
```

```

taoc* server prod
address ipv4 10.221.10.11
8 Kit
  
```

ip route 10.221.0.11 2 55.255.2 55.25 5 Ethernet 0.1

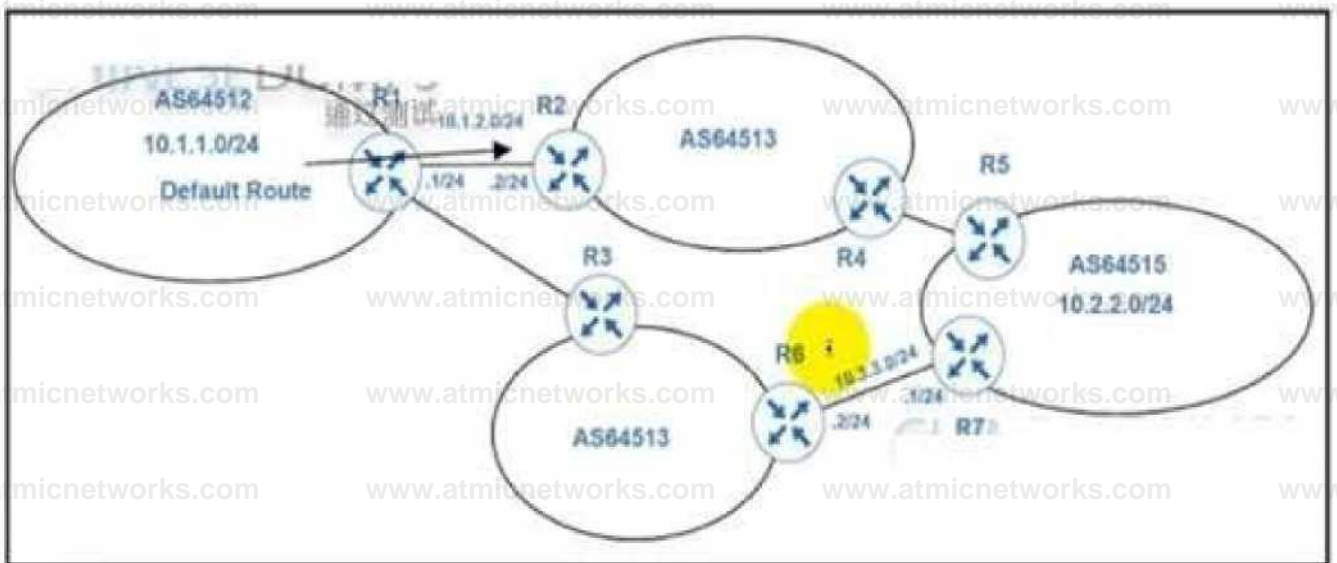
- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 344

Refer to the exhibit.



An engineer must configure PBR on R1 to reach to 10.2.2.0/24 via R3 AS64513 as the primary path and a backup route through default route via R2 AS64513. All BGP routes are in the routing table of R1. but a static default route overrides BGP routes. Which PBR configuration achieves the objective?

```
access-list 100 permit ip 10.1.1.0 0.0.0.255 10.2.2.0 0.0.0.255
```

```
|
```

```
route-map PBR permit 10
```

```
match ip address 100
```

```
set ip next-hop 10.3.3.1
```

```
access-list 100 permit ip 10.1.1.0 0.0.0.255 10.2.2.0 0.0.0.255
```

```
|
```

```
route-map PBR permit 10
```

```
match ip address 100
```

```
set ip next-hop recursive 10.3.3.1
```

```
access-list 100 permit ip 10.1.1.0 255.255.255.0 10.2.2.0 255.255.255.0
```

```
route-map PBR permit 10
```

```
match ip address 100
```

```
set ip next-hop recursive 10.3.3.1
```

```
access-list 100 permit ip 10.1.1.0 255.255.255.0 10.2.2.0 255.255.255.0
```

```
route-map PBR permit 10
```

```
match ip address 100
```

```
set ip next-hop 10.3.3.1
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Configuration Output:

```
aaa new-model
aaa group server tacacs+ admin
server name admin
!
ip tacacs source-interface GigabitEthernet1
aaa authentication login admin group tacacs+ local enable
aaa session-id common
!
tacacs server admin
address ip 10.11.15.6
key 7 01150F165E1C07032D
!
line vty 0 4
login authentication admin
```

Debug Output:

```
Oct 22 12:38:57.587: AAA/BIND(0000001A): Bind if
Oct 22 12:38:57.587: AAA/AUTHEN/LOGIN (0000001A): Pick method list 'admin'
Oct 22 12:38:57.587: AAA/AUTHEN/ENABLE(0000001A): Processing request action LOGIN
Oct 22 12:38:57.587: AAA/AUTHEN/ENABLE(0000001A): Done status GET_PASSWORD
Oct 22 12:39:02.327: AAA/AUTHEN/ENABLE(0000001A): Processing request action LOGIN
Oct 22 12:39:02.327: AAA/AUTHEN/ENABLE(0000001A): Done status FAIL - bad password
```

An administrator configured a Cisco router for TACACS authentication, but the router is using the local enable password instead. Which action resolves the issue?

- A. Configure the aaa authentication login admin group admin local enable command instead.
- B. Configure the aaa authentication login admin group tacacs* local enable none command instead.
- C. Configure the aaa authentication login admin group tacacs* local if-authenticated command instead.
- D. Configure the aaa authentication login default group admin local if-authenticated command instead.

Answer: C

Explanation:

Question: 346

Refer to the exhibit.

```
RwttllibM Ip tqp vpn^4 rd 1100 1M1 10 30 1 lt 023 fIOPttvDAgUte^MHK 1001 10 M VOW) WWH?WI PmII (S w»UM» Mrt
toHMFJ
```

```
A0VK%*«4 to upUto-oroup* ^t^SIJiU
```

```
1 2 5
```

```
4HO 164 W MU) «»H PK*«M own a RA-^li
```

```
172 16 254 226 iilnc 206451 trw 172 16 22 4 2)61172 16 224 236) Origin rap MK o fecjtottr too rM OHM W^MI
```

```
1'tented Conimufftiti At I IM 1001 WI UM nW MMwnS?
```

```
OKUI «VM 65003)65000
```

```
172 16254 226IINMC20646IhteI 10 131 12 3 71 (10 131 123 71) On^in IGF imVK 0 toO^vt 1M va*d f c r'j i c Atefr.ij
```

```
E.tented Commute) RT IIMIOOI L)UM '■pH Uteh B/ayi MUWM?
```

```
fMWt 6496 5 65003) 65005
```

```
''' A4TM A
```

```
1TJ 16 254 226 (iteM; 2M 0«* 372 16 2« 253 i R? 16 216 M3:
```

```
Qtw ISP mHK 0 IOCMpraf 100 WN cMOOMIKll &>1*nted Commuffty RI 11001001 rnpH Itais fl'MI taUbtVW
```

```
KU1164H5 6 500 31 6 5W
```

```
172 l« 354 JHim^lne J0045-VMI 172 16 J1( 252 dT? 1(216 2« Otl^I ISP m*ticQ tocl^rtf 100 viw rated wlmpl
```

```
Ei tented Commufftr RT 1100 1001 rrtHi UM *W MtaD*t'362
```

```
04H 65003) «MN
```

```
172 16 2W 226 mdrM 20645) IFM 10 77 255 57 10 77 255 P■ OwISP rteUKO MUM IK ¥ii*o MrttMrtontf E-tented Commui'itj
```

```
ST j 1« W
```

```
IM »M 41W *OUteV362 . 7--.11
```

```
«K5 66001) «SM
```

```
172 16 254 226 liMtoc 20645) trcfr 10 57 255 11 (10 57 255 11) QngnlQP tM*Wp toufefri 1K raid coited^MIMtel Owl
```

```
E.tented CommMy RT 1100 1001 001 UM Wwl MUDeU342
```

```
iW5i«5W3fm . .
```

```
172 14254 22* lmetoc 20M5J trwr f >2 16 224 2531177 16 224 253 .
```

```
O0g.fi IGF MK ft IKK(nf HWt 1*1 <tfpvd'4IHMI
```

```
frHiWtd Commiwirty HTHWI^dt
```

```
mpu l«MS iVOUt MUM1162
```

```
i«N3)4MU
```

```
172 14 2MJ2BMK2WS}RW 172 14 2 5* ZJ4H72W 2M JMt
```

```
Otani IOP mtix: 0 qcMMtl 100 nW cwrede>!tnul
```

```
ErtaAdK CaHMRuiy RT 11*0 IOOI
```

```
npH UMM ftWI notaMM3
```

```
45M4 lrt«tt^ ItBH iJUK-Ctw!!
```

```
172 l4 2»2Klffi«lcKN«4Wn> 172 14 22* 2MiT?2 «2 2*2MI OhgrtlGP
```

```
rt^xfl ^tuprltM V*-' :or'! r.hfiM , ,
```

```
E>Und*d Conmuntj RT 11001001
```

```
n^H Ubeis AMut MUBiiQTt
```

Refer to the exhibit. An engineer configured BGP and wants to select the path from 10.77.255.57 as the best path instead of current best path. Which action resolves the issue?

- A. Configure AS_PATH prepend for the desired best path
- B. Configure higher MED to select as the best path.
- C. Configure lower LOCAL_PREF to select as the best path.
- D. Configure AS_PATH prepend for the current best path

Answer: D

Explanation:

Question: 347

What is LDP label binding?

- A. neighboring router with label
- B. source prefix with label
- C. destination prefix with label
- D. two routers with label distribution session

Answer: C

Explanation:

For every 1GP IP prefix in its IP routing table, each LSR creates a local binding—that is, it binds a label to the IPv4 prefix. The LSR then distributes this binding to all its LDP neighbors. These received bindings become remote bindings. The neighbors then store these remote and local bindings in a special table, the label information base (LIB). Each LSR has only one local binding

Text Description automatically generated with medium confidence

Question: 348

Which table is used to map the packets in an MPLS LSP that exit from the same interface, via the same next hop, and have the same queuing policies?

- A. RIB
- B. FEC
- C. LDP
- D. CEF

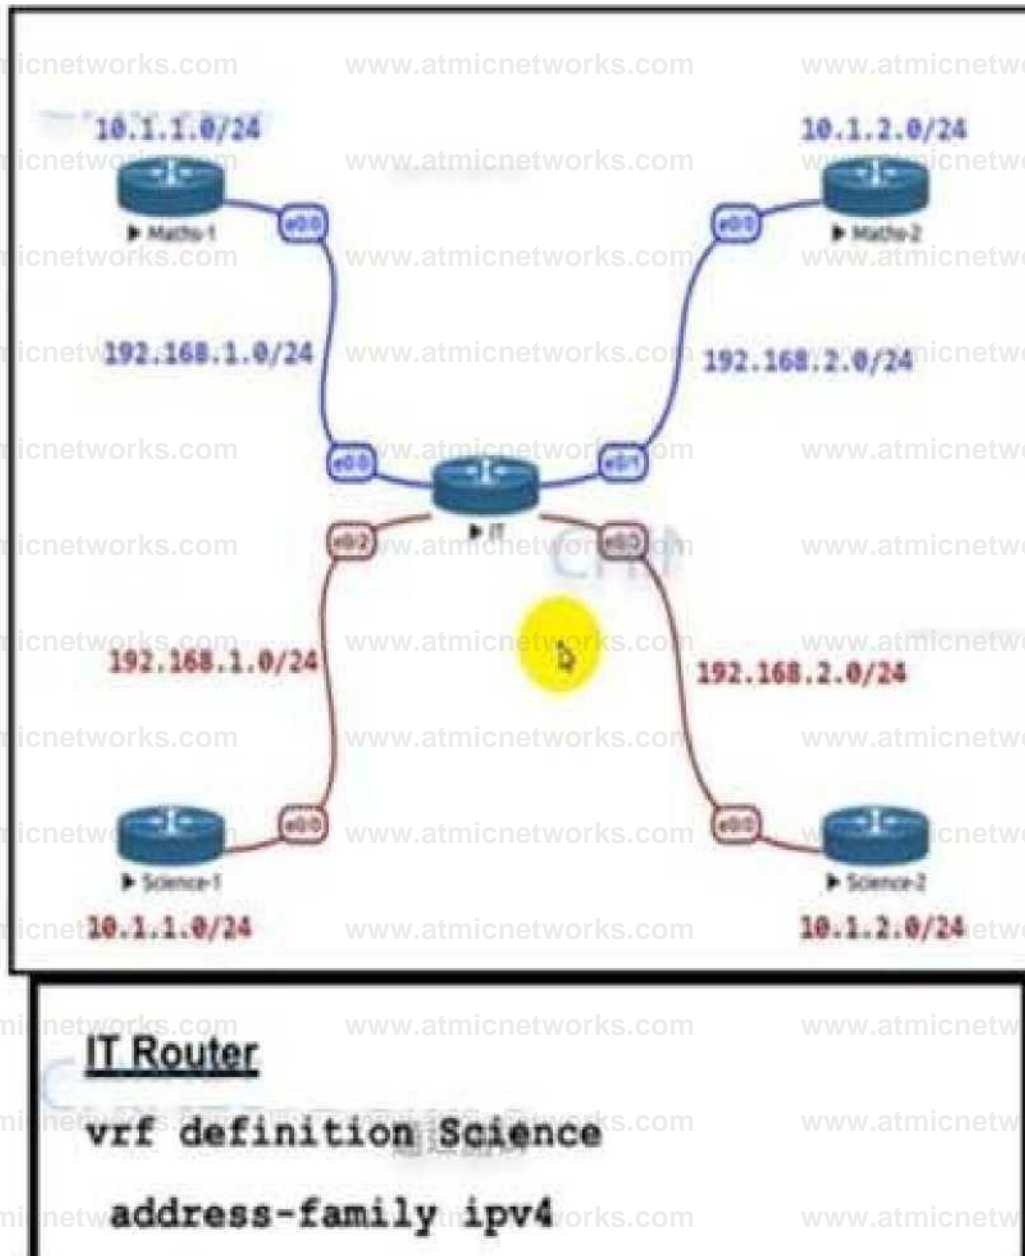
Answer:

B

Explanation:

Question: 349

Refer to the exhibit.



```
Interface E 0/2
```

```
Vrf forwarding Science
```

```
Ip address 192.168.1.1 255.255.255.0
```

```
No abut
```

```
Interface E 0/3
```

```
Vrf forwarding Science
```

```
Interface E 0/3
```

```
Vrf forwarding Science
```

```
Ip address 192.168.2.1 255.255.255.0
```

```
No shut
```

Refer to the exhibit. The IT router has been configured with the Science VRF and the interfaces have been assigned to the VRF. Which set of configurations advertises Science-1 and Science-2 routes using EIGRPAS 111?

```
mulef eigrp 111
```

```
uddrou-family ipv4 vrf Science Kutonornaus-system 1
```

```
network 192.163.1.0
```

```
network TS2.1ftH.Z0
```

```
router eigrp 111
```

```
address-family ipv4 vrf Science
```

```
network 192.133.1.0
```

```
network 192-1 MI2.0
```

```
routareigip 111
```

```
network 192.1 SB 1.0
```

```
network 192.163.2.0
```

```
router eigrp 1
```

```
nddiHB-family Ipv4 vrf Science sEjtoRomoumyiitwn 111 network 192403,1,0
```

```
network 192.133.2.0
```

A. Option A

B. Option B

C. Option C


D. Option D

Answer: D

Explanation:

Question: 350

Refer to the exhibit.



```

interface GigabitEthernet0/0
  description FTP SERVER no ip address ipv6 address 2001:DB8::F/33 ipv6
  enable ipv6 traffic-filter FTP-SERVER in
interface GigabitEthernet0/1 description FTP CLIENT no ip address >
  la/inc ipv6 address ioOUDBSi WOO::F/33 ipv6 enable ift^SUL^ ipv6 traffic-
  filter FTP-CLIENT in

ipv6 access-list FTP-CLIENT permit tcp host 2001:DB8:8000::1 host
2001:DB8::1 eq ftp pennit tcp host 2001:DB8:8000::1 host 200i:DB9::1 eq
ftp-data

```

```

ipvt access-list FTP-CLIENT
  permit tcp host 2001;$8:5000::1 host 200i:DB8::i eq ftp
  permit tcp host 2001:DB8:8000::i host 2001:DB8::l eq ftp-data
I
ipv6 access-list FTP-SERVER
  permit tcp host 2001:DB8::l host 2001:DBS:8000::1 eq ftp established
  permit tcp host 2001:DB0::1 host 2001:DB9:8000::1 eq ftp-data established

```

Refer to the exhibit. When an FTP client attempts to use passive FTP to connect to the FTP server, the file transfers fail

Which action resolves the issue?

- Configure active FTP traffic.
- Modify FTP-SERVER access list to remove established at the end.
- Modify traffic filter FTP-SERVER in to the outbound direction.
- Configure to permit TCP ports higher than 1023.

Answer: D

Explanation:

Question: 351

In a DMVPN network, the Spoke1 user observed that the voice traffic is coming to Spoke2 users via the hub router. Which command is required on both spoke routers to communicate directly to one another?

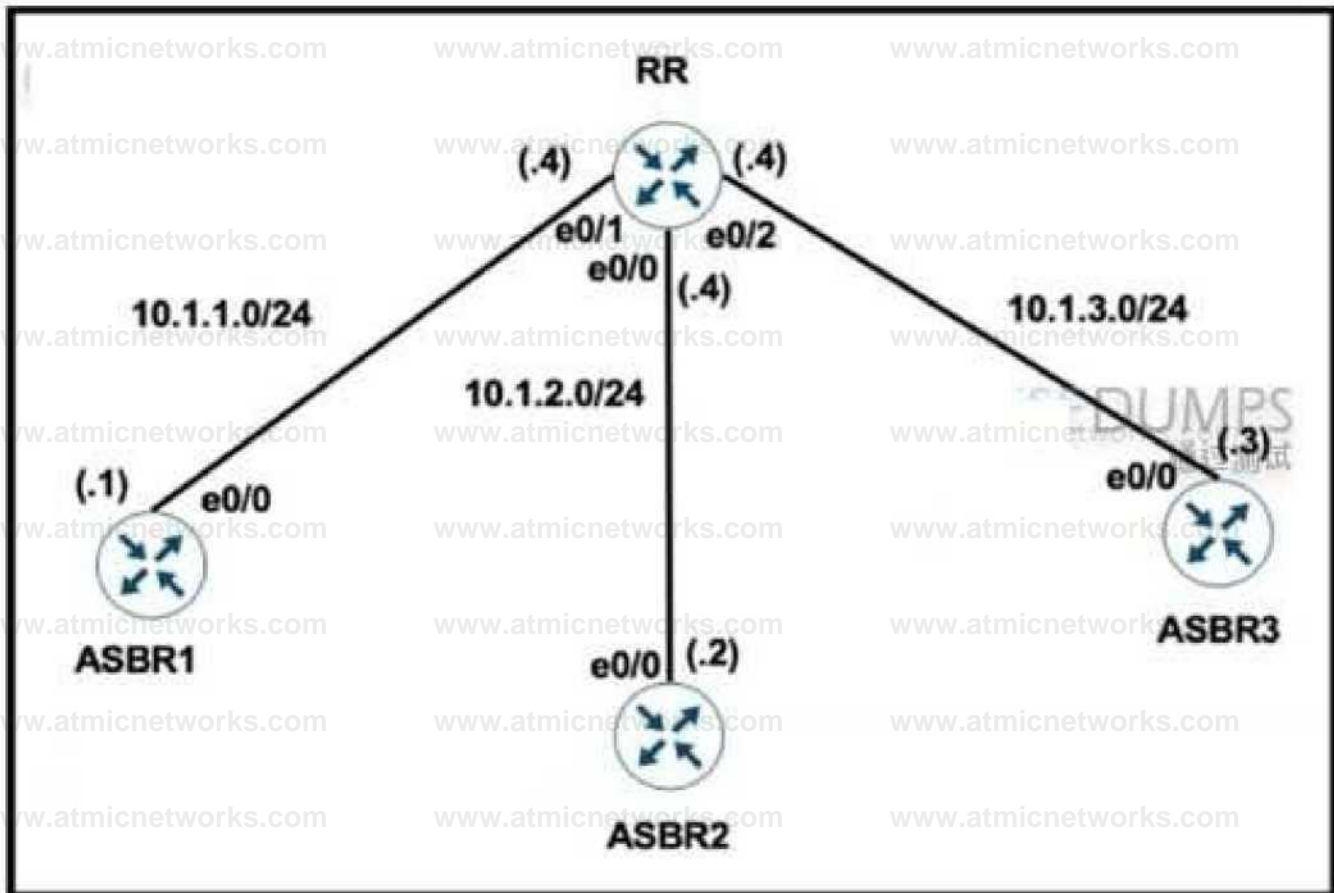
- A. ip nhrp map dynamic
- B. ip nhrp shortcut
- C. ip nhrp nhs multicast
- D. ip nhrp redirect

Answer: B

Explanation:

Question: 352

Refer to the exhibit.



RR Configuration:

```

router bgp 100
neighbor IBGP pser-graup
neighbor IBGP route-reflector-client
neighbor 10.1.1.1 remote** 100
neighbor 10.1.2.2 remains 100
neighbor 10.1.3.3 remete-a* 100

```

The network administrator configured the network to establish connectivity between all devices and notices that the ASBRs do not have routes for each other. Which set of configurations resolves this issue?

```

* Tauter bgp 100
neighbor 10.1.1.1 next-hap-self
neighbor 10,1.2.2 next-hop-self
neighbor 10.1.3.3 next-hop-seif

router bgp 100
neighbor IBGP updetewurcoLocptadtA

```

```
router bgp 100
 neighbor IEGP ne^i-hop-self

router bgp 100
 neighbor 10.1.1.1 peer-group IBGP
 neighbor 10.1.2.2 peer-group IBGP
 neighbor 10.1.3.3 peer-group IBGP
```

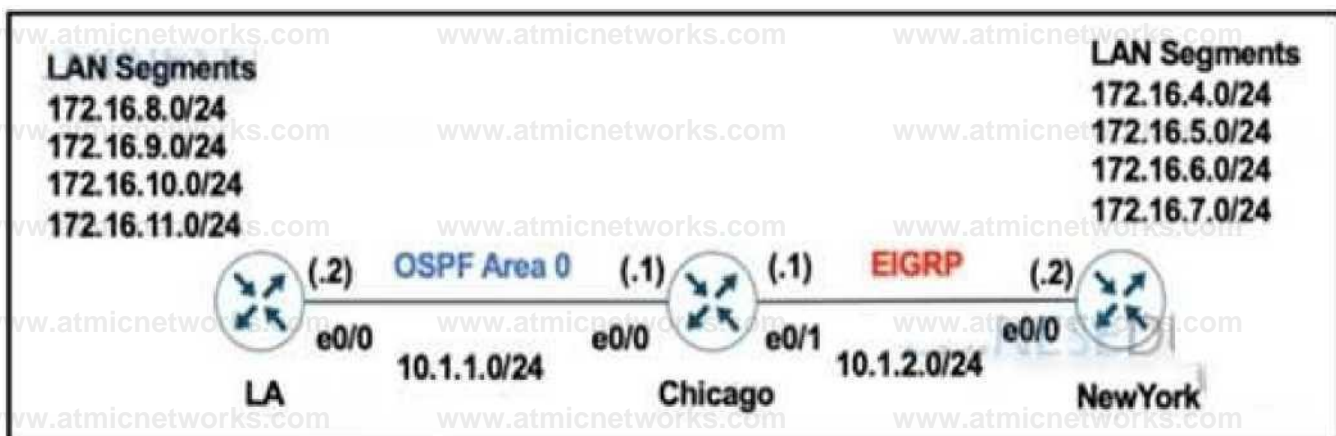
- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

Question: 353

Refer to the exhibit.



The network administrator configured the Chicago router to mutually redistribute the LA and New York routes with OSPF routes to be summarized as a single route in EIGRP using the longest summary mask:

```
router eigrp 100
 redistribute ospf 1 metric 10 10 10 10 10
router ospf 1
 redistribute eigrp 100 subnets
|
interface E 0/0
 ip summary-address eigrp 100 172.16.0.0 255.255.0.0
```

After the configuration, the New York router receives all the specific LA routes but the summary route. Which set of configurations resolves the issue on the Chicago router?

```
$ interface E 0/1
 ip summary-address eigrp 100 172.16.0.0 255.255.0.0

interface E 0/1
 ip summary address eigrp 100 172.16.8.0 255.255.252.0

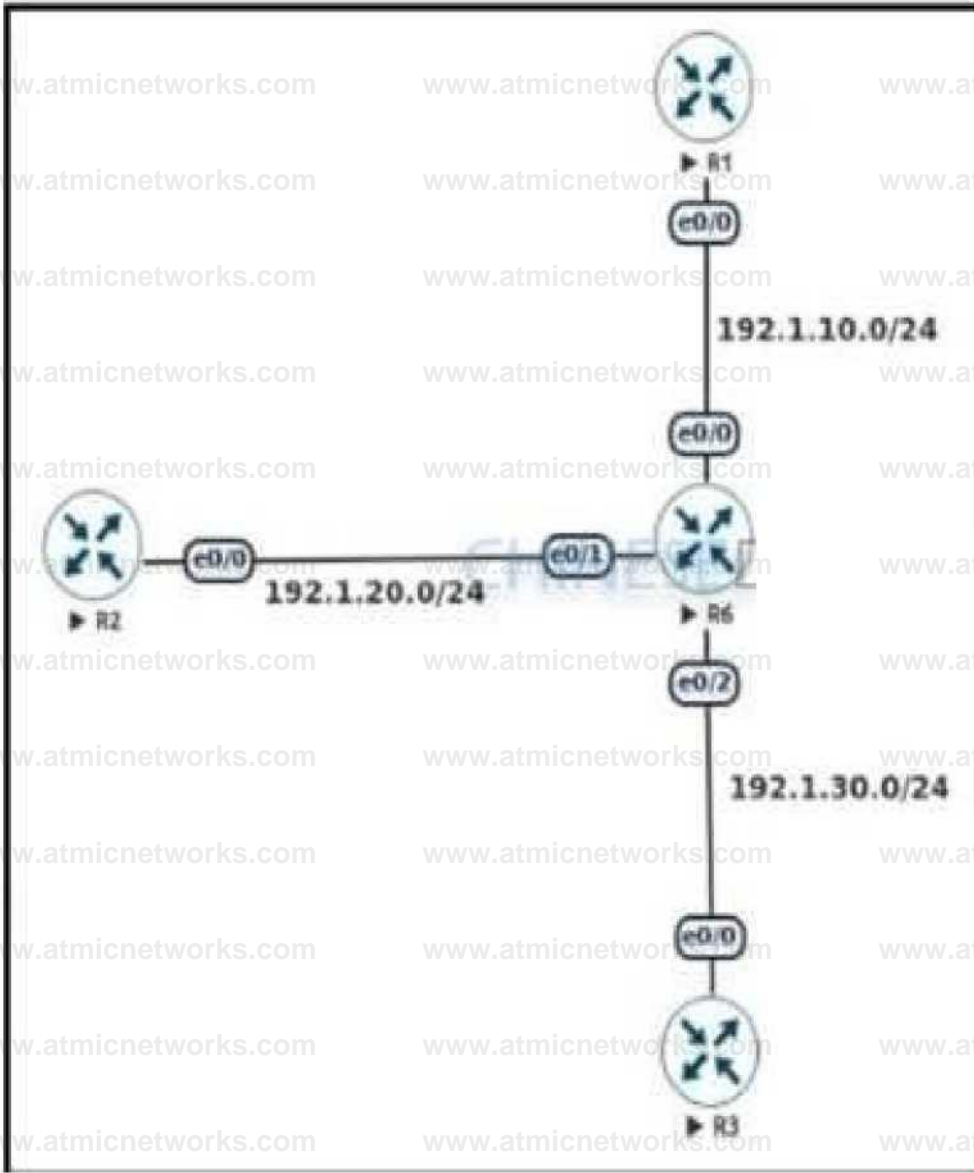
router eigrp 100
 summary-address 172.16.8.0 255.255.252.0

router eigrp 100
 summary-address 172.16.0.0 255.255.0.0
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:



An engineer must configure DMVPN Phase 3 hub-and-spoke topology to enable a spoke-to-spoke tunnel. Which NHRP configuration meets the requirement on R6?

IntiNfsce Tunnel 1

ip address 192.168.4.255 255.255.0 tunnel source a Oio

tunnel mode gre multipoint ip nhrp MtWCrk-ld 1

interface Tunnell

ip nhrp authentication Cisco123

ip nhrp map multicast dynamic

ip nhrp networked 1

ip nhrp holdtime 300

ip nhrp redirect

interface Tunnell

ip nhrp authentication Cisco123

ip nhrp map multicast dynamic

ip nhrp network-id 1

ip nhrp holdlime 300

ip nhrp shortcut

Interface Tunnel 1

Ip address 192.168 1.1 255.255 255.0 tunnel source a 04 tunnel mode gre

multipoint ip nhrp network-id 1

ip nhrp map 192.168.1.2 192.1 20.2

A. Option A

B. Option B

C. Option C

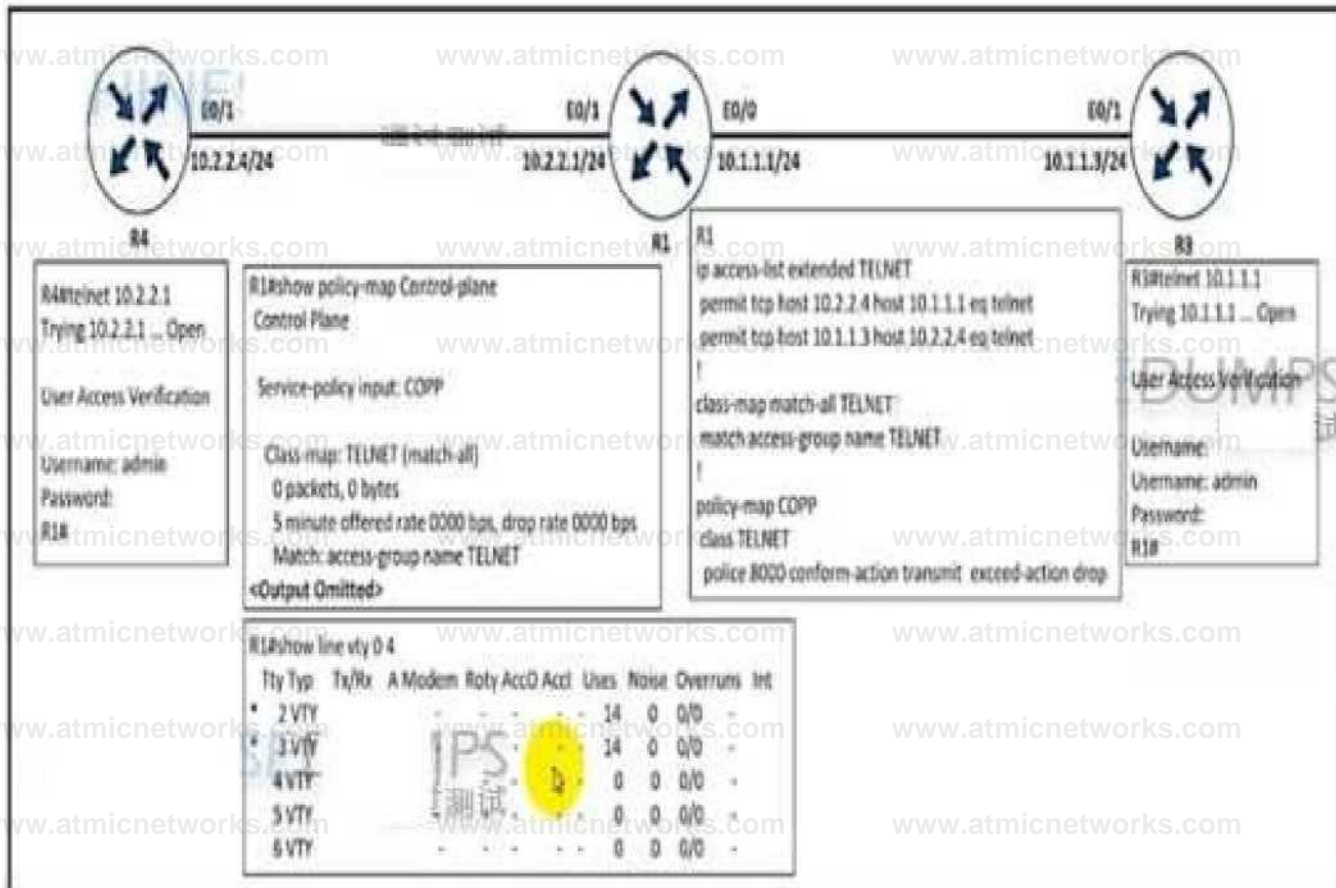
D. Option D

Answer: B

Explanation:

Question: 355

Refer to the exhibit.



An engineer implemented CoPP to limit Telnet traffic to protect the router CPU. It was noticed that the Telnet traffic did not pass through CoPP Which configuration resolves the issue?

policy-map COPP

class TELNET

police 8000 conform-action transmit exceed-action drop

policy-map COPP

class TELNET

ip access-list extended TELNET

permit tcp host 10.2.2.4 host 10.1.1.1 eq telnet

```
permit tcp host 10.1.1.5 host 10-1.1 3 eq telnet
```

```
ip access-list extended TELNET
```

```
permit tcp host 10 2,24 host 10.2.2.1 eq telnet
```

```
permit tcp host 101.1.3 host 10 1.1.1 eq telnet
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 356

Refer to the exhibit.



An engineer implemented CoPP but did not see OSPF traffic going through it. Which configuration resolves the issue?

- A. ip access-list extended OSPF permit ospf any any
- B. policy-map COPP class OSPF police 8000 conform-action transmit exceed-action transmit violate- action drop
- C. control-plane service-policy input COPP
- D. class-map match-all OSPF match access-group name OSPF

Answer: B

Explanation:

Question: 357

An engineer must override the normal routing behavior of a router for Telnet traffic that is destined to 10.10.10.10 from 10.10.1.0/24 via a next hop of 10.4.4.4, which is directly connected to the router that is connected to the 10.1.1.0/24 subnet. Which configuration reroutes traffic according to this requirement?

```
access-list 100 permit ip 104.0.0.0 0.0.0.255 host 104.0.0.23
```

```
route-map POLICY permit 10
```

```
match ip address 100
```

```
*1 ip 10.4.4.4
```

```
access-list 100 permit ip 104.1.0.0 0.0.0.255 host 104.1.0.23
```

```
route-map POLICY permit 10
```

```
match ip address 100
```

```
access-list 100 deny tcp 104.0.0.0 0.0.0.255 host 104.0.0.23
```

```
route-map POLICY permit 10
```

```
match ip address 100
```

```
deny ip 10.4.4.4
```

```
route-map POLICY permit 20
```

```
access-list 100 permit tcp 104.1.0.0 0.0.0.255 host 104.1.0.23
```

```
route-map POLICY permit 10
```

```
match ip address 100
```

A. Option A

B. Option B

C. Option C


D. Option D

Answer: B

Explanation:

Question: 358

Refer to the exhibit.



An engineer must redistribute networks 192.168.10.0/24 and 192.168.20.0/24 into OSPF from EIGRP. where the metric must be added when traversing through multiple hops to start an external route of 20 The engineer notices that the external metric is fixed and does not add at each hop. Which configuration resolves the issue?

```
R1 ex r ifig u:atc ■ ■ aaJist W (Mt ml t 192 J 63.10.0 0.0.0.255
R2 (configjiiieccmMlit 10 permit 192.163.20 0 0.0.0 2 5 5
```

```
R2 (config iftrotiio-map RD permit 10
R2[CONJigHX<ir^h^p inmate] I ip address 10
FtStconfig-route-iTjapiaset metric 20
FUtcoril^rouk-rnapiaset metric<ype type-l
```

```
R2(configiPiouTei oapf 10
R^millixoLiLi! ? ^redistribute eigrp 10 subnets route-map RD
```

```
R?i Config |iraccm'Il>t 10 permit 192.169.10.0 0.0.0.2 55
ftZIOOfifigPKHU-litt 10 permit 192.168 20 0 0.0.0 2 55 |
R2icofifig.isioute-mep RD pel mil 10
l^axil^cufe-mqp^mwtch ip Bdtrm 10
R2(CONl^*KMr1u4nap]AA4t metric 20
R2(config- route-rnap-i«et mblric-typo typo-1
```

R1(config)#route-map RD permit 10

R1(config-route-map)#set metric 20

R1(config)#route-map RD permit 10

R1(config-route-map)#set metric-type type-1

R1(config)#route-map RD permit 10

R1(config-route-map)#set metric 20

R1(config-route-map)#set metric-type type-1

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

R1 (config-router)#route-map RD permit 10

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 359

An administrator attempts to download the pack NBAR2 file using TFTP from the CPE router to another device

over the Gi0/0 interface. The CPE is configured as below:

```
hasmamo CPE
```

```
ip access-list extended WAN
```

```
remark ■> All UDP rules below for WAN ID: S420T92E35F99
```

```
permit udp any eq domain any
```

```
permit udp any any eq tftp
```

```
deny udp any any
```

```
interface GigabitEthe] nelO/O
```

```
ip access-group WAN in
```

```
tftp'&erver flash :pp-adw-csFIQDOv«1612.1^-37-53.0.0.pack
```

The transfer fails. Which action resolves the issue?

- A. Change the WAN ACL to permit the UDP port 69 to allow TFTP
- B. Make the permit udp any eq tftp any entry the last entry in the WAN ACL.
- C. Change the WAN ACL to permit the entire UDP destination port range
- D. Shorten the file name to the 8+3 naming convention.

Answer:

B

Explanation:

Question:
360

What is an MPLS LDP targeted session?

- A. session between neighbors that are connected no more than one hop away
- B. LDP session established between LSRs by exchanging TCP hello packets
- C. label distribution session between non-directly connected neighbors
- D. LDP session established by exchanging multicast hello packets

Answer:

C

Explanation:

Question:
361

Refer to the exhibit.

```
ip sla 1 iotp-ftcho 8 8.8.8 threshold 1000 timeout 3000 frequency 5
ip sla schedule 1 life forever start-time now
```

```
track J ip sla 1
```

```
ip route 0.0.0.0 0.0.0.0 203.0.113.1 name ISP1 track 1 ip route 0.0.0.0
0.0.0.0 198.51.100.1 name ISP2 rack J
```

An administrator configures a router to stop using a particular default route if the DNS server 8.8.8.8 is not reachable through that route. However, this configuration did not work as desired and the default route still works even if the DNS server 8.8.8.8 is unreachable. Which two configuration changes resolve the issue? (Choose two.)

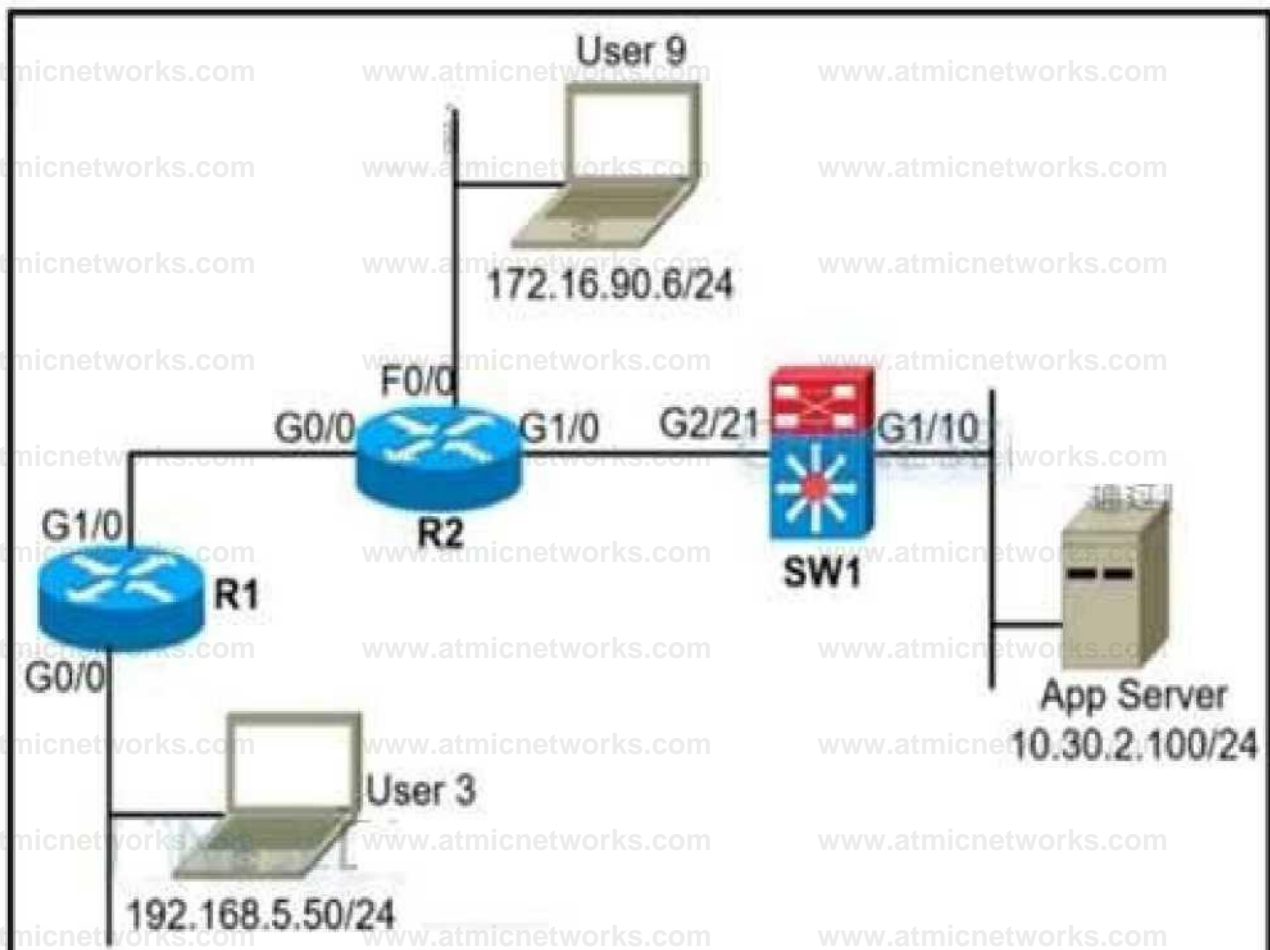
- A. Configure two static routes for the 8.8.8.8/32 destination to match the IP SLA probe for each ISP.
- B. Associate every IP SLA probe with the proper WAN address of the router.
- C. Reference the proper exit interfaces along with the next hops in both static default routes.
- D. Use a separate track object to reference the existing IP SLA 1 probe for every static route.
- E. Use a separate IP SLA probe and track object for every static route

Answer: A,E

Explanation:

Question: 362

Refer to the exhibit.



A network administrator must block ping from user 3 to the App Server only. An inbound standard access list is applied to R1 interface G0/0 to block ping. The network administrator was notified that user 3 cannot even ping user 9 anymore. Where must the access list be applied in the outgoing direction to resolve the issue?

- A. R2 interface G1/0
- B. R2 interface G0/0
- C. SW1 interface G1/10
- D. SW1 interface G2/21

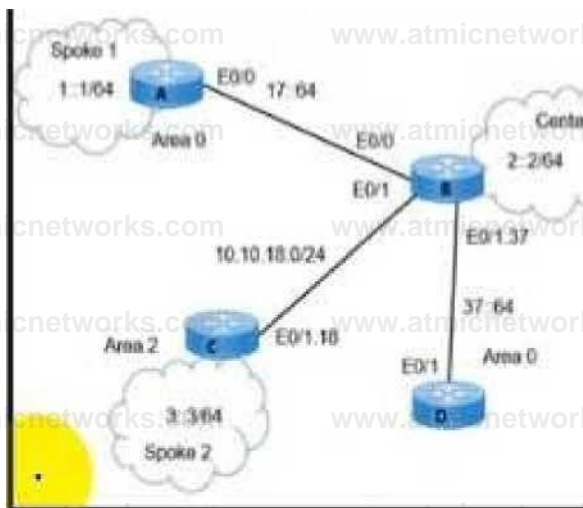
Answer:

D

Explanation:

Question: 363

Refer to the exhibit.



```
B(config-if)* do ah cun int *0/1 I b mt B(config-if)*
interface XthernetO/1 (oonfig-if)* ip address 78.1 1 0
255.255.255.0 E(config-if)* ipv6 enable
B(config-if)* ospfv3 1 ipv4 area X
```

```
C(config)I interface XthemetQ/1.10
C(config-subif)* encap dot1q 78 C(config-subif)* ip
add 78.1 1 7 255 255 255 0 C(config-subif)I ospfv3 1
ipv< area 0
```

```
D (conf ig-if)* do eh run int eO/1 I b lot
D(config-if)* interface XthometO/1
D(config-if)* no ip address
D (config-if)* ipv6 address 37::3/68
D(config-if)* ipv6 enable
D(config-if)*>> ipvd ospf 1 area 0
```

Refer to the exhibit. A network engineer receives a report that Spoke 1 users can perform bank transactions with the server located at the Center site, but Spoke 2 users cannot. Which action resolves the issue?

- A. Configure the Spoke 2 users IP on the router B OSPF domain
- B. Configure encapsulation dot1q 78 on the router C interface.
- C. Configure IPv6 on the routers B and C interfaces

D. Configure OSPFv2 on the routers B and C interfaces

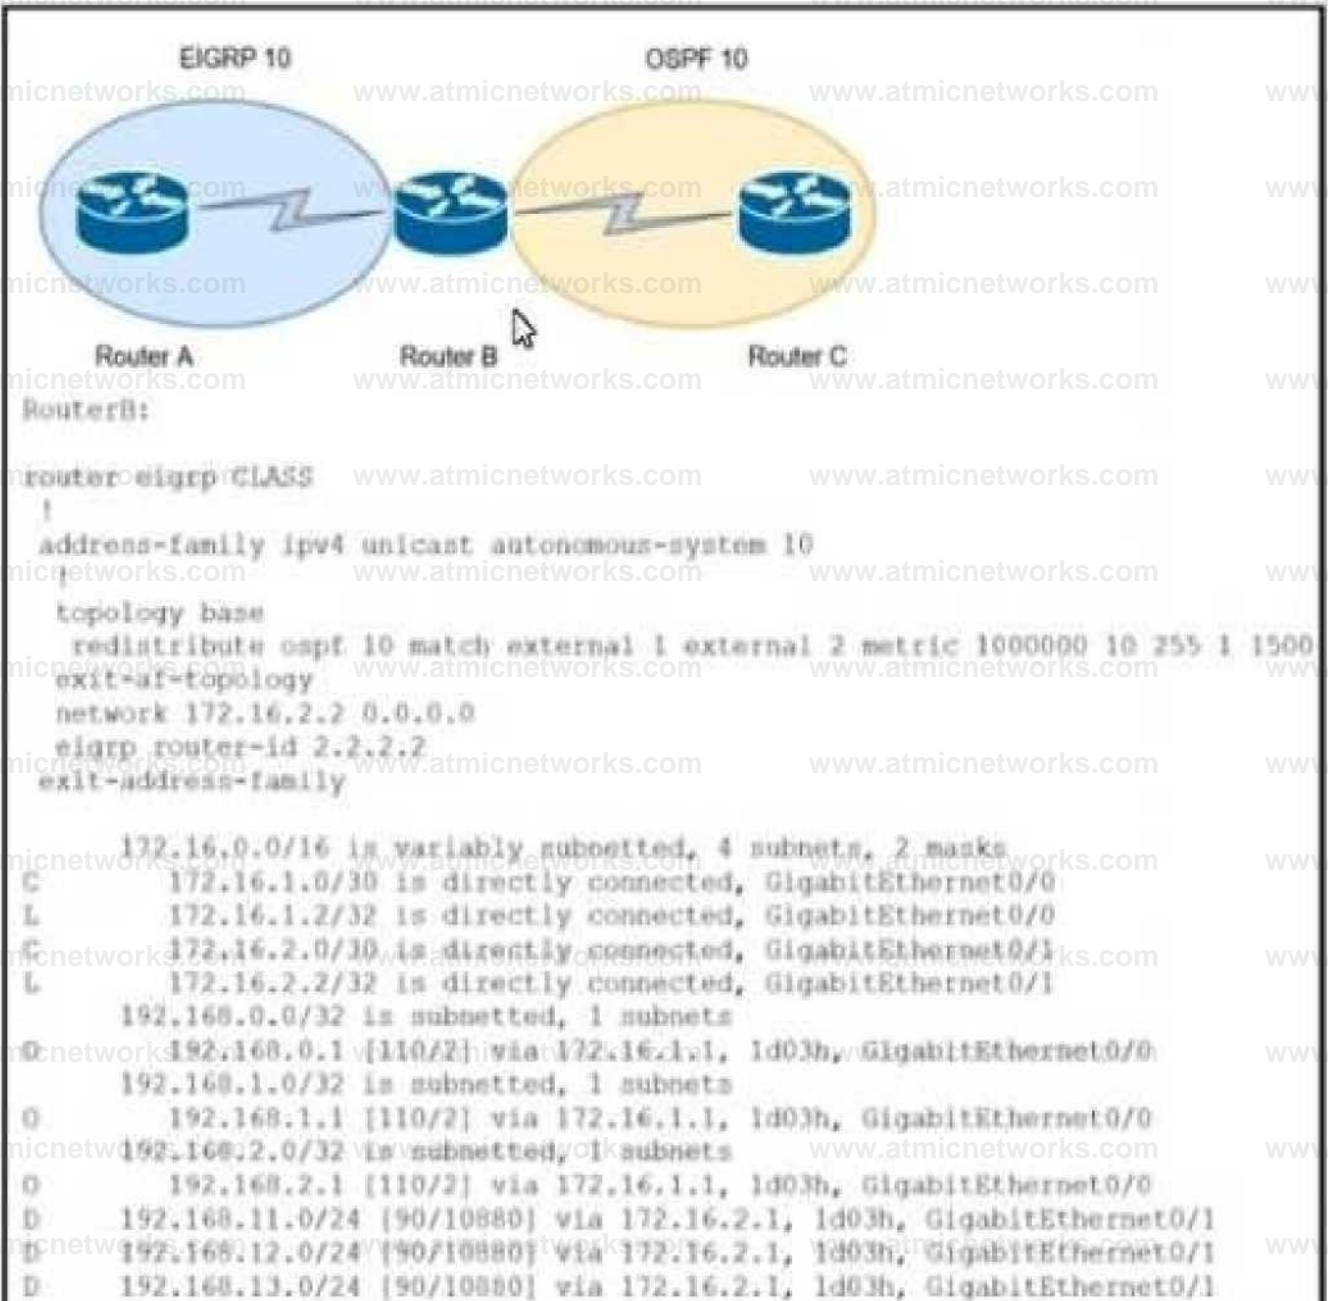
Answer:

C

Explanation:

Question: 364

Refer to the exhibit.



Refer to the exhibit. An engineer configured route exchange between two different companies for a migration project. EIGRP routes were learned in router C but no OSPF routes were learned in router

A. Which configuration allows router A to receive OSPF routes?

A. (config-router-af)#redistribute ospf 10 1000000 10 255 1 1500

B. (config-router-af-topology)#redistribute ospf 10 metric 1000000 10 255 1 1500

C. (config-router-af-topology)#redistribute connected

D. (config-router-af-topology)#no redistribute ospf 10 match external 1 external 2 metric 1000000 10 255 1 1500

Answer:
B

Explanation:

Question:
365

A network administrator cannot connect to a device via SSH. The line vty configuration is as follows:

```
line vty 0 4
  login local
  session-timeout 10
  transport preferred ssh
  transport input all
  output telnet ssh
  no exec
```

Which action resolves this issue?

- A. Increase the session timeout
- B. Change the stopbits to 10.
- C. Configure the transport input SSH
- D. initialize the SSH key

Answer:
D

Explanation:

Question: 366

DRAG DROP

Drag and drop the ICMPv6 neighbor discovery messages from the left onto the correct packet types ON the right.

Neighbor Solicitation	ICMPv6 Type 134
Neighbor Advertisement	ICMPv6 Type 137
Router Advertisement	ICMPv6 Type 135
Redirect Message	ICMPv6 Type 133
Router Solicitation	ICMPv6 Type 136

Answer:

Explanation:

Neighbor Solicitation

Router Advertisement

Neighbor Advertisement

Redirect Message

Redirect Message

Table Description automatically generated with medium confidence

generated with medium confidence

Question: 367

DRAG DROP

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

- FEC
- LSR
- LER
- LSR
- LDP

- routers in the core of the provider network known as P routers
- all traffic to be forwarded using the same path and same label
- routers that connect to the customer routers known as PE routers
- used for exchanging label mapping information between MPLS enabled routers
- path along which the traffic flows across an MPLS network

Answer:

Explanation:




Table Description automatically

generated

Question: 368

Refer to the exhibits.

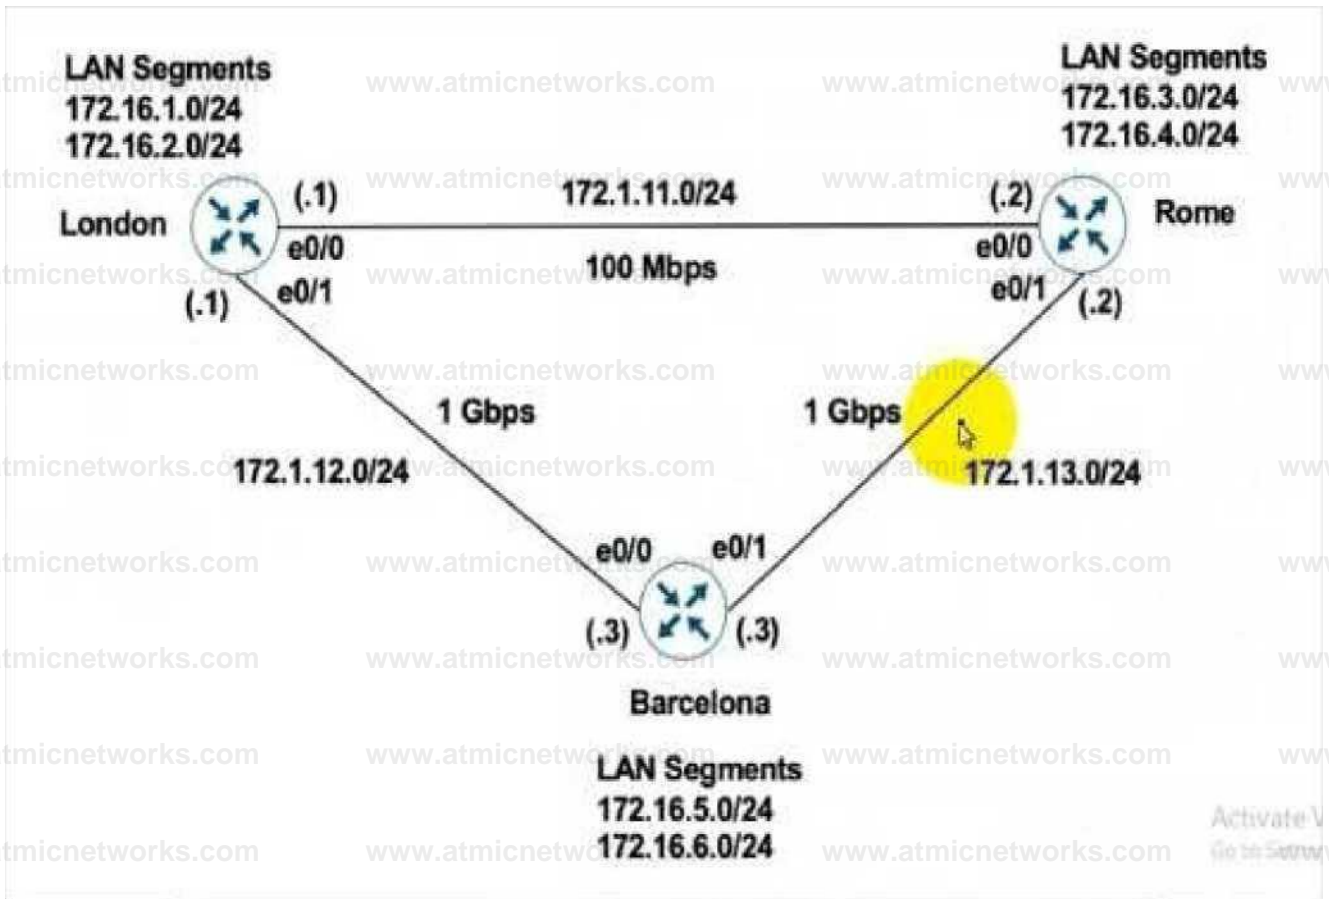
London - "show ip route" output

Gateway of last resort is not set

```
172.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
C   172.1.11.0/24 is directly connected, Ethernet0/0
L   172.1.11.1/32 is directly connected, Ethernet0/0
C   172.1.12.0/24 is directly connected, Ethernet0/1
L   172.1.12.1/32 is directly connected, Ethernet0/1
D   172.1.13.0/24 [90/76800] via 172.1.11.2, 00:00:50, Ethernet0/0
172.16.0.0/16 is variably subnetted, 8 subnets, 2 masks
C   172.16.1.0/24 is directly connected, Loopback0
L   172.16.1.1/32 is directly connected, Ethernet0/0
C   172.16.2.0/24 is directly connected, Loopback1
L   172.16.2.1/32 is directly connected, Loopback1
R   172.16.3.0/24 [120/1] via 172.1.11.2, 00:00:08, Ethernet0/0
R   172.16.4.0/24 [120/1] via 172.1.11.2, 00:00:08, Ethernet0/0
D   172.16.5.0/24 [90/156160] via 172.1.12.3, 00:00:50, Ethernet0/1
D   172.16.6.0/24 [90/156160] via 172.1.12.3, 00:00:50, Ethernet0/1
```

Rome - "show run | section router" output

```
router eigrp 111
 network 172.1.0.0
 network 172.16.0.0
 no auto-summary
```



London must reach Rome using a faster path via EIGRP if all the links are up but it failed to take this path. Which action resolves the issue?

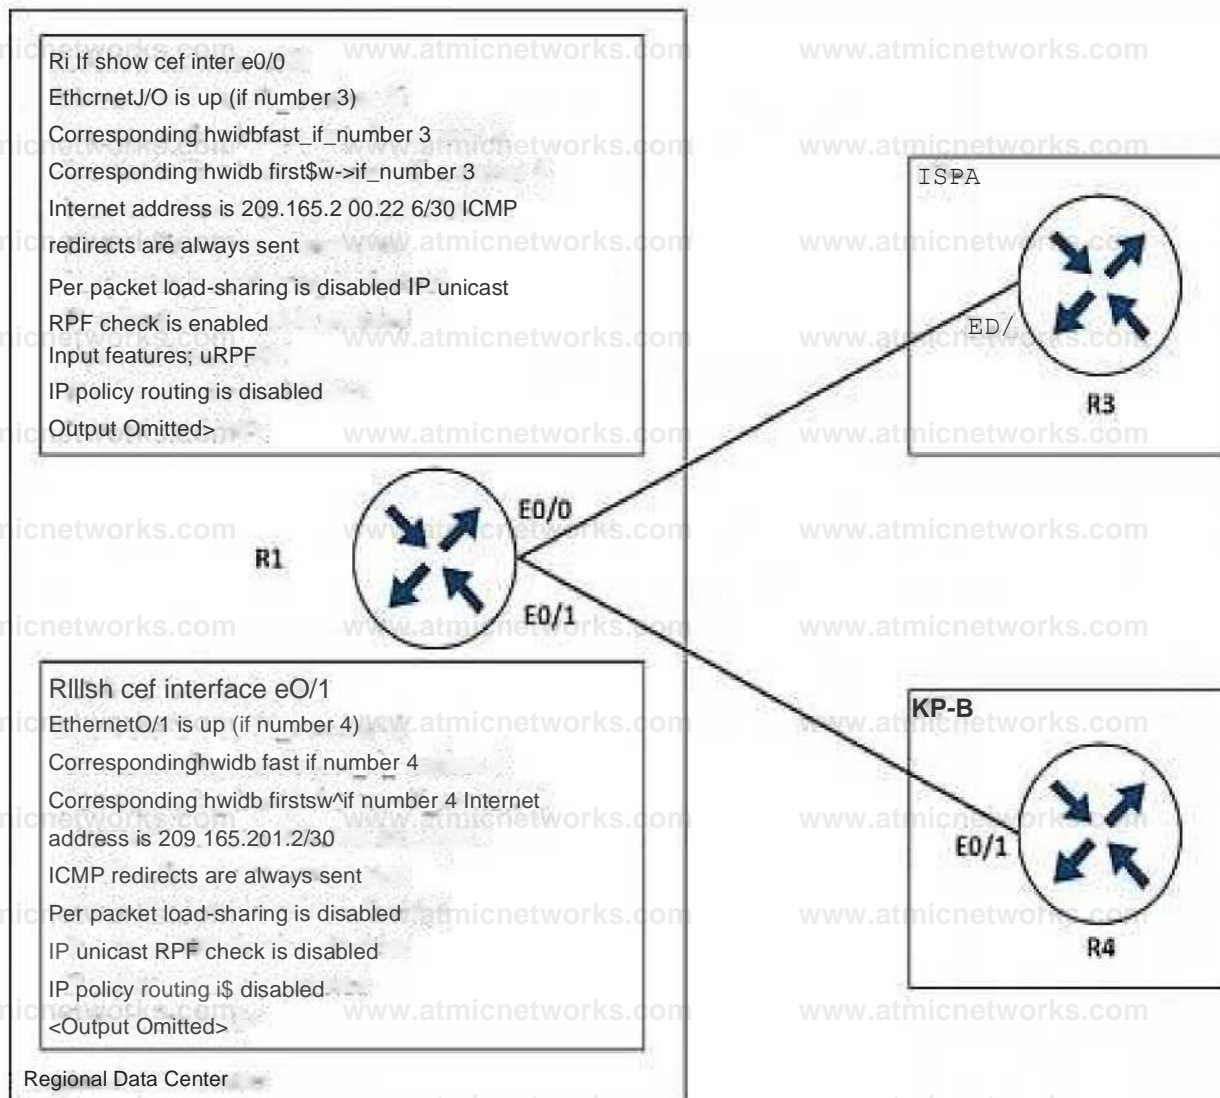
- A. Increase the bandwidth of the link between London and Barcelona
- B. Use the network statement on London to inject the 172.16.X.0/24 networks into EIGRP.
- C. Change the administrative distance of RIP to 150
- D. Use the network statement on Rome to inject the 172.16.X.0/24 networks into EIGRP

Answer: D

Explanation:

Question: 369

Refer to the exhibit.



Refer to the exhibit. The company implemented uRPF to address an antispoofing attack. A network engineer received a call from the IT security department that the regional data center is under an IP attack. Which configuration must be implemented on R1 to resolve this issue?

```
interface ethernet0/0
ip verify unicast reverse-path
```

```
interface ethernet0/1
ip verify unicast reverse-path
```

```
interface ethernet0/0
ip unicast RPF check receive-via any allow-default allow-self-ping
```

```
interface ethernet0/0
```

ip unicast RPF check reachable-via any allow-default allow-self-ping

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 370

What is a function of BFD?

- A. peer recovery after a Layer 3 protocol adjacency failure
- B. peer recovery after a Layer 2 adjacency failure

C. failure detection independent of routing protocols and media types

D. failure detection dependent on routing protocols and media types

Answer:

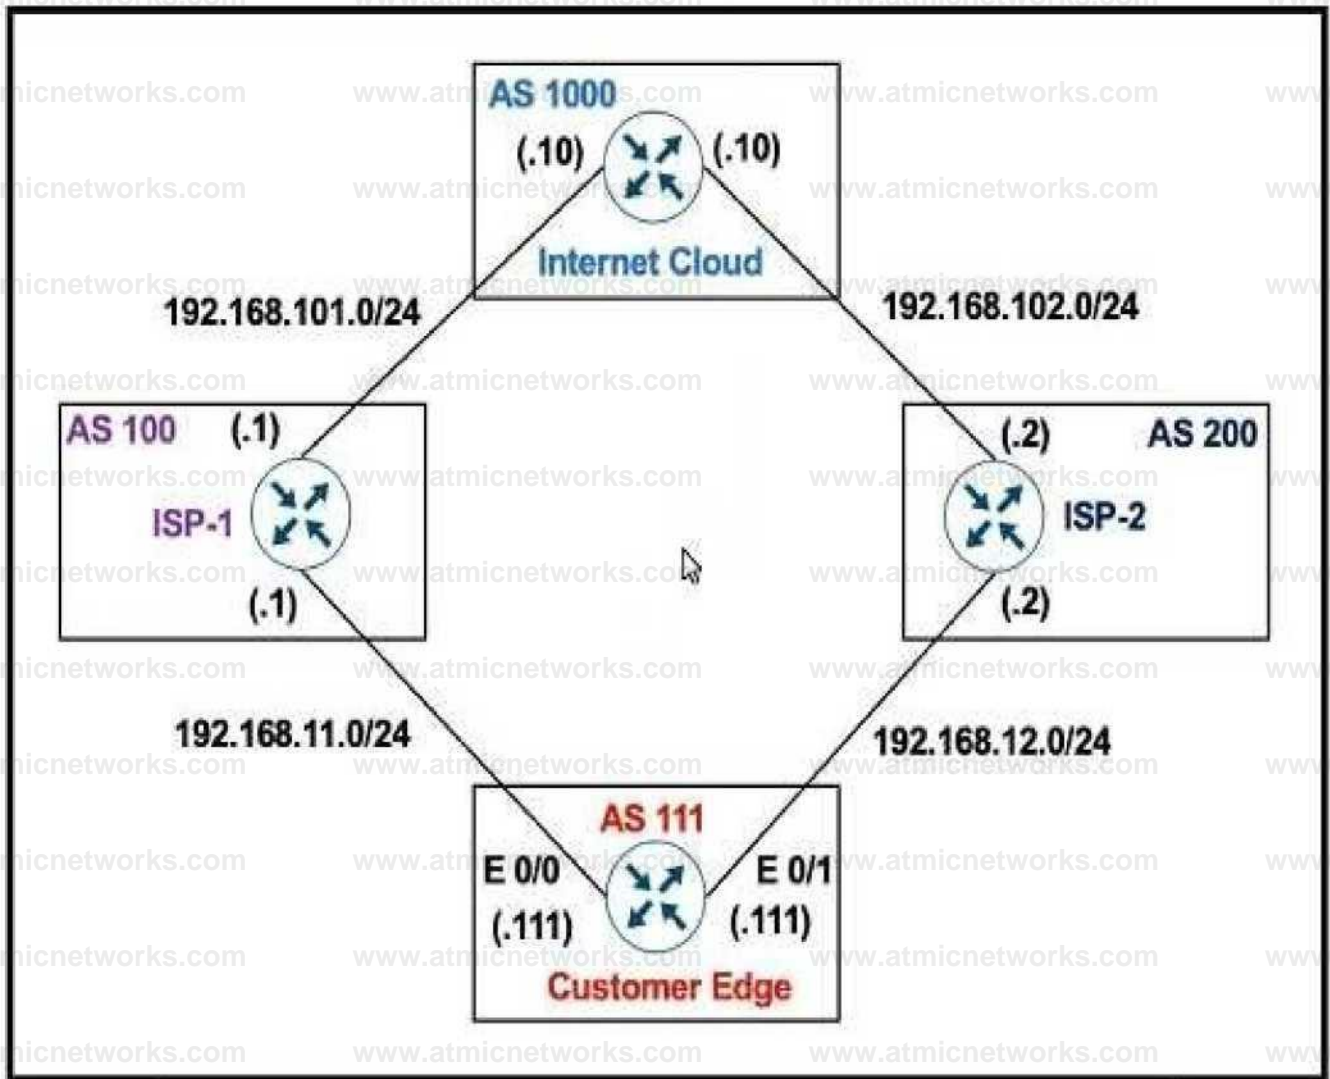
D

Explanation:

Question:

371

Refer to the exhibit.



ISP-1

```
ip as-path access-list 1 permit ^111
!
router bog 100
 neighbor 192.168.101.10 remote-as 1000
 neighbor 192.168.11.111 remote-as 111
 neighbor 192.168.11.111 filter-list 1 in
```

Refer to the exhibit. AS 111 must not be used as a transit AS, but ISP-1 is getting ISP-2 routes from AS 111.

Which configuration stops Customer AS from being used as a transit path on ISP-1?

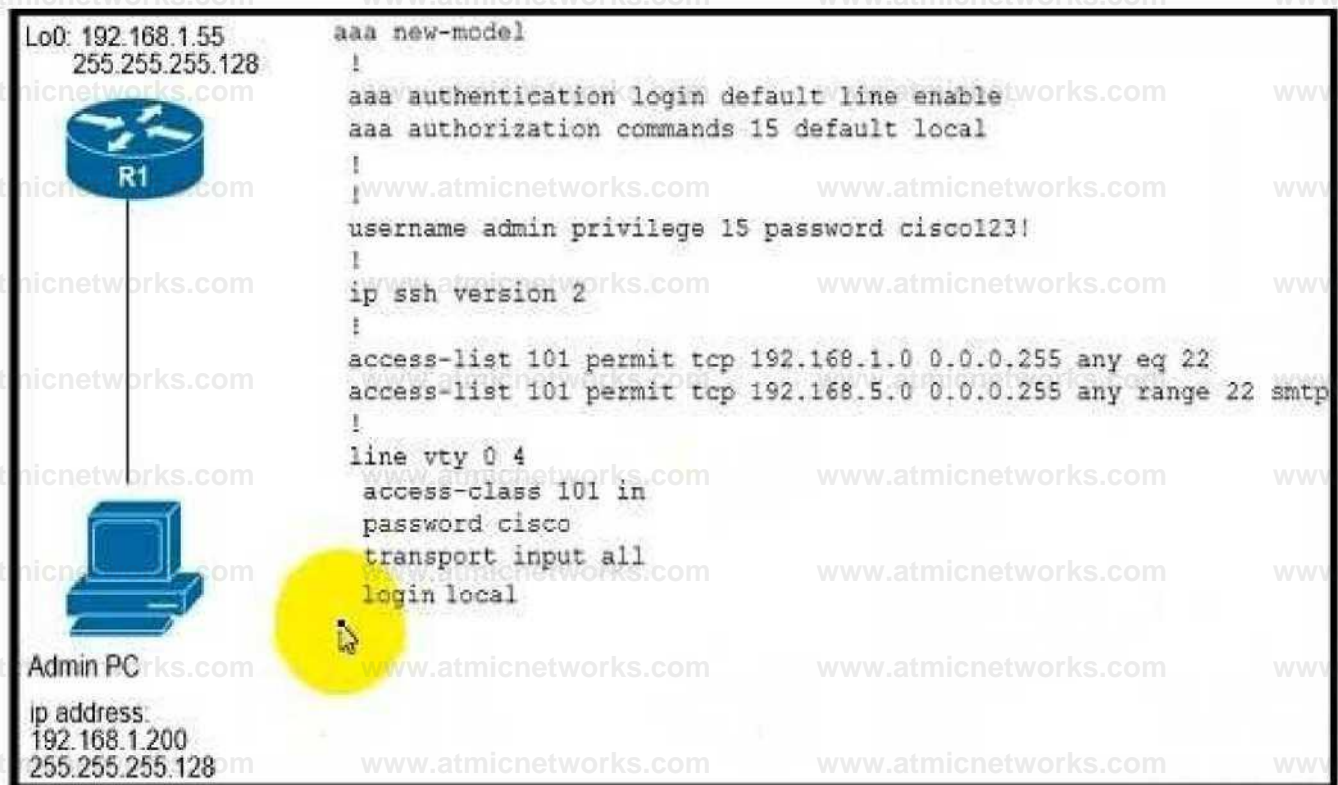
- A. ip as-path access-list 1 permit ^\$
- B. ip as-path access-list 1 permit_111_
- C. ip as-path access-list 1 permit."
- D. ip as-path access-list 1 permit ^111\$

Answer: A

Explanation:

Question: 372

Refer to the exhibit.



Refer to the exhibit. An engineer configured user login based on authentication database on the router, but no one can log into the router. Which configuration resolves the issue?

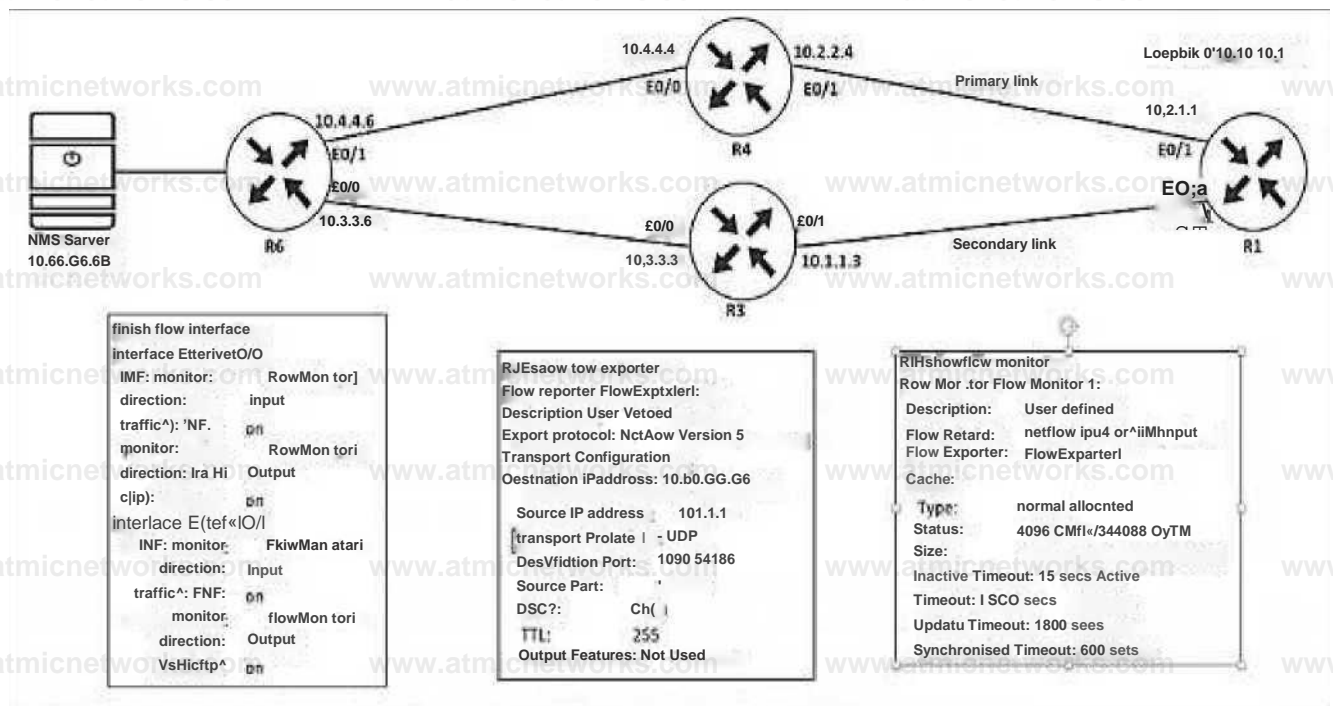
- A. aaa authentication login default enable
- B. aaa authorization network default local
- C. aaa authentication login default local
- D. aaa authorization exec default local

Answer: C

Explanation:

Question: 373

Refer to the exhibit.



Refer to the exhibit. An engineer configured NetFlow on R1, but the flows do not reach the NMS server from R1. Which configuration resolves this issue?

R1(config)#flow monitor FlowMonitor1

R1(config-flow-monitor)#destination 10.66.66.66

R1(config)#flow exporter FlowExporter1

R1(config-flow-exporter)#destination 10.66.66.66

R1(config)#interface Ethernet0/0

R1(config-if)#ip flow monitor FlowMonitor1 input

R1(config-if)#ip flow monitor FlowMonitor1 output

R1(config)#interface Ethernet0/1

R1(config-if)#ip flow monitor FlowMonitor1 input

R1(config-if)#ip flow monitor FlowMonitor1 output

A. Option A

B. Option B

C. Option C

D. Option D

Answer:
B

Explanation:

Question:
374

Refer to the exhibit.



```

R1#show route-map
route-map FROM->EIGRP, permit, sequence 10
  Match clauses:
    ip address (access-lists): 10
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
R1#show run | sec router
router eigrp 100
  network 10.96.69.0 0.0.0.3
  no auto-summary
  eigrp router-id 1.1.1.1
router ospf 100
  router-id 1.1.1.1
  log-adjacency-changes
  redistribute eigrp 100 subnets route-map FROM->EIGRP
  network 10.99.69.0 0.0.0.3 area 0
R1#show ip access-list
Standard IP access list 10
  10 permit 192.168.16.0, wildcard bits 0.0.3.255
  11 permit 192.168.0.0, wildcard bits 0.0.7.255
  20 deny any

```

Refer to the exhibit The engineer configured route redistribution in the network but soon received reports that R2 cannot access 192.168.7.0/24 and 192.168.15.0/24 subnets Which configuration resolves the issue?

```
R1(config)#ip access-list standard 10
```

```
R1 (config-std-nacl) 10 permit
```

```
R1 (config-std-nacl) #no 11 permit
```

```
R1 (config-std-nacl) 10 permit 192.160.0.0 0.0.3.255
```

```
R1 (config-std-nacl) 11 permit 192.16.0.0 0.0.3.255
```

```
R1 (config) ip access-list standard 10
```

```
R1 (config-std-nacl) 10 permit
```

```
R1 (config-std-nacl) #no 11 permit
```

```
R1 (config-std-nacl)#10 permit 192.168.0.0 0.0.7.255
R1 (conf ig-std-naclJ,# 1L permit 192. 168.8.0 0.0.3.255
R1 (configip access List standard 10
R1 (config-std-nacl}|no 10 permit
R1 (config-std-nacl}Ino 11 permit
R1 (config-std-nacl}|110 permit 192.168.0.0 0.0.3.255
R1 (config-std-nacl}|111 permit 192.166.8.0 0.0.7.255
R1 (config-std nacl)#nc 10 permit
Ri (config-std-nacl)#nc 11 permit
R1 (config-std-nacl)#10 permit 192.168.4.0 0.0.3.255
RHcimt iij-si.^ permit 192.168.12.0 0.0.a.255
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

Question: 375

An engineer received a ticket about a router that has reloaded. The monitoring system graphs show different traffic patterns between logical and physical interfaces when the router is rebooted. Which action resolves the issue?

- A. Configure the snmp ifindex persist command globally.
- B. Clear the logical interfaces with snmp ifindex clear command
- C. Configure the snmp ifindex persist command on the physical interfaces.
- D. Trigger a new snmpwalk from the monitoring system to synchronize interface OIDs

Answer: A

Explanation:

Question: 376

Refer to the exhibit.

```
R1#show policy map control plane
```

```
Control Plane
```

```
Service-policy input CoPP
```

```
Class-map SSH (mulch-alb
```

```
29 parcels. 2215 bytes
```

```
5 minute offered rate 0000 bps
```

```
Match accessed up 100
```

```
Class-map ANY (match all) 4$ packets, 3678 bytes 5 minute offered rate 0000 bps,
```

```
drop rate 0000 bps
```

```
Match access-group 199 drop
```

```
Class-map class default (match-any) 41 packets, 1687 bytes
```

```
5 minute offered rate 0000 bps drop rate 0000 bps Match any
```

```
R2#show access-list 100
```

```
Extended IP access list 100
```

```
10 deny tcp any any eq 22(14 matches)
```

```
20 permit rip host 192.168.12.1 any eq 22 (29 matches)
```

```
R2#show access-list 199
```

```
Extended IP access list 199
```

```
10 permit ip any any (51 matches)
```

Refer to the exhibit. Which action limits the access to R2 from 192.168.12.1?

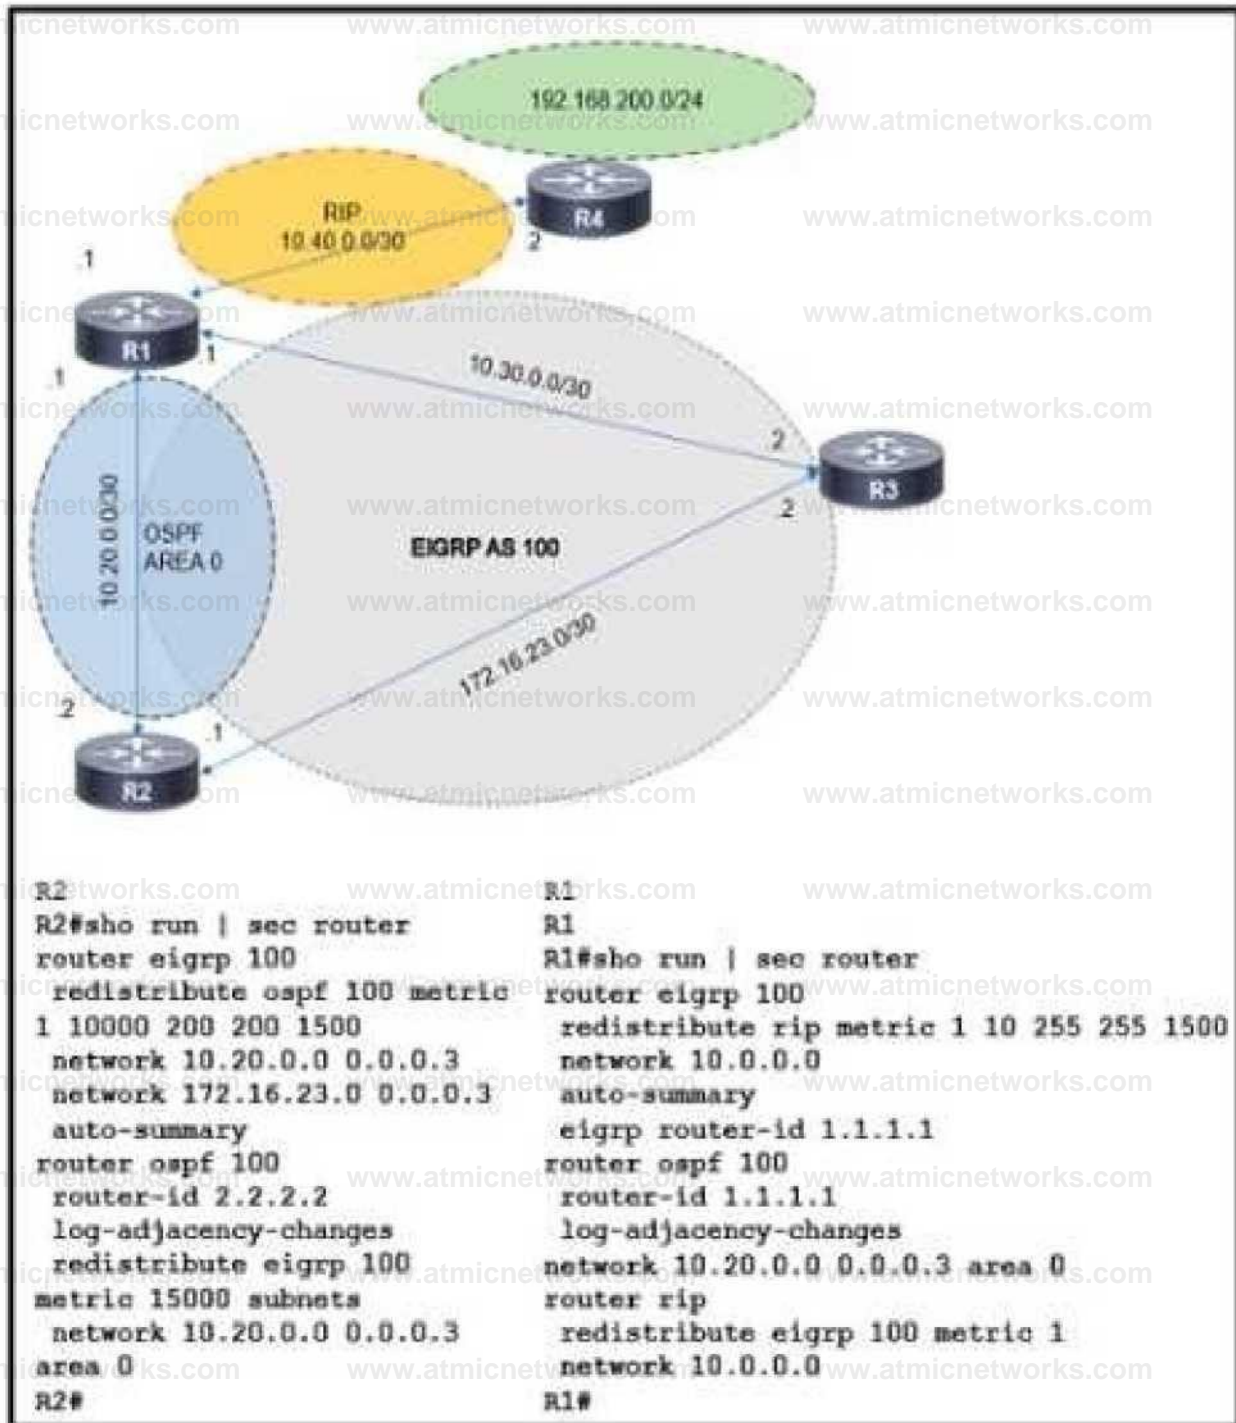
- A. Swap sequence 10 with sequence 20 in access-list 100.
- B. Modify sequence 20 to permit tcp host 192.168.12.1 eq 22 any to access-list 100.
- C. Swap sequence 20 with sequence 10 in access-list 100.
- D. Modify sequence 10 to deny tcp any eq 22 any to access-list 100.

Answer: C

Explanation:

Question: 377

Refer to the exhibit.



Refer to the exhibit The route to 192 168 200 0 is flapping between R1 and R2 Which set of

configuration changes resolves the flapping route?

```
R2(config)#router ospf 100
R2(config-router)#no redistribute eigrp 100
R2(config-router)#redistribute eigrp 100 metric 1 subnets
```

```
R1(config)#router rip
R1(config)#network 192.168.200.0 255.255.255.0 10.40.0.2
```

```
-. (config)#router eigrp 100
R2(config-router)#no redistribute rip 100
R2(config-router)#redistribute rip
```

```
R1(config)#router ospf 100
R1(config-router)#redistribute rip metric 1 metric-type 1
subnets
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

Question: 378

Refer to the exhibit.


```
R1 (config)# ip vrf CCNP
R1 (config-vrf)# rd 1:100
R1 (config-vrf)# exit
R1 (config)# interface Loopback0
R1 (config-if)# ip address 10.1.1.1 255.255.255.0
R1 (config-if)# ip vrf forwarding CCNP
R1 (config-if)# exit
R1 (config)# exit
R1# ping vrf CCNP 10.1.1.1
% Unrecognized host or address, or protocol not running.
```

Refer to the exhibit Which command must be configured to make VRF CCNP work?

```
Interface LMptockQ
ip address 10.11 1 255.255.255 0
vrf forwarding CCNP
```

```
interface Loopba^kO
ip address 10.1.1.1 255.255.255.0
```

```
jntsrffjcc LwpbackO
vrf forwarding CCNP
```

```
interface LoapbitkO
ip address 10.11 1 266.265 255 0
ip vrf forwarding CCNP
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 379

Refer to the exhibit.

```
ip sla 1
 icmp-echo 8.8.8.8
 threshold 1000
 timeout 2000
 frequency 5
ip sla schedule 1 life forever start-time now
track 1 ip sla 1
ip route 0.0.0.0 0.0.0.0 Ethernet0/0 203.0.113.1 name ISP1 track 1
ip route 0.0.0.0 0.0.0.0 Ethernet0/1 198.51.100.1 2 name ISP2
```

Refer to the exhibit. After recovering from a power failure, Ethernet0/1 stayed down while Ethernet0/0 returned to the up/up state. The default route through ISP1 was not reinstated in the routing table until Ethernet0/1 also came up. Which action resolves the issue?

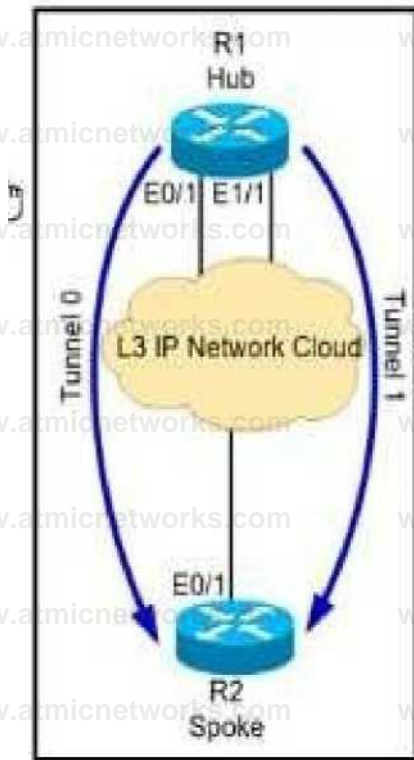
- A. Reference the track object 1 in both static default routes
- B. Remove the references to the interface names from both static default routes
- C. Configure the default route through ISP1 with a higher administrative distance than 2.
- D. Add a static route to the 8.8.8.8/32 destination through the next hop 203.0.113.1

Answer: D

Explanation:

Question: 380

Refer to the exhibit.



Refer to the exhibit. The hub and spoke are connected via two DMVPN tunnel interfaces. The NHRP is configured and the tunnels are detected on the hub and the spoke. Which configuration command adds an IPsec profile on both tunnel interfaces to encrypt traffic?


- A. tunnel protection ipsec profile DMVPN multipoint
- B. tunnel protection ipsec profile DMVPN tunnel1
- C. tunnel protection ipsec profile DMVPN shared
- D. tunnel protection ipsec profile DMVPN unique

Answer: C

Explanation:

Question: 381

Refer to the exhibit.



```
PE1# show run | sec router bgp
```

```
router bgp 65000
```

```
  bgp log-neighbor-changes
```

```
  neighbor 10.255.255.3 remote-as 65000
```

```
  neighbor 10.255.255.3 update-source Loopback0
```

```
|L/1/1 Uis
```

```
FE1* debug ip tcp transactions
```

```
Flip debug ip iemp
```

```
I* - .Jftlp. 8 J
```

```
♦Feb 22 14:04:12.374: TCP: sending SYN, seq 379810712, ack 0
```

```
`Feb 22 14:04:12.374: TCPO: Connection Co 10,255*255.3:17&, advertising
```

```
NSS Liu ft
```

```
`Feb 22 14:04:12.374: TCPO: stAte was CLOSED -> SYNSENT [21381
```

```
> 10,255.255,3(179)1
```

```
`Feb 22 14:04:12.375: ICMP: dst (10.255.255.1) administratively prohibited
```

```
UD re ashable rev fmn 10.0.12.2
```

```
`Feb 22 14:04:12.375: TCPO: ICMP destination unreachable
```

```
received
```

```
`Feb 22 14:04:12.375: Released port 21381 in Transport Port
```

```
Agent for TCP IP type 1 delay 140000
```

```
`Feb 22 14:04:12.375: TCPO: state was SYNSENT -> CLOSED [21381
```

```
> 10,255.255.3(179) ]
```

◆Feb 22 14:04:12.375: TCB 0xE35A92B8 destroyed

Refer to the exhibit. The administrator is troubleshooting a BGP peering between PE1 and PE3 that is unable to establish. Which action resolves the issue?

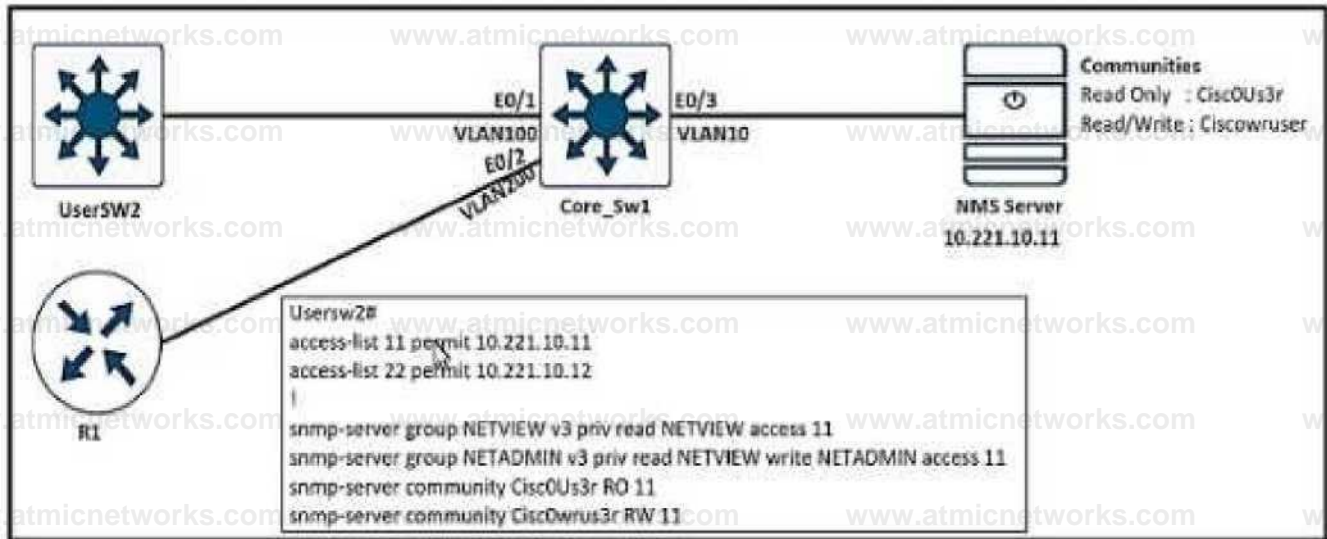
- A. P2 must have a route to PE3 to establish a BGP session to PE1
- B. Disable sending ICMP unreachables on P2 to allow PE1 to establish a session with PE3
- C. Ensure that the PE3 loopback address is used as a source for BGP peering to PE1
- D. Remove the traffic filtering rules on P2 blocking the BGP communication between PE1 and PE3

Answer: C

Explanation:

Question: 382

Refer to the exhibit.



Refer to the exhibit. An engineer configured SNMP Communities on UserSW2 switch, but the SNMP server cannot upload modified configurations to the switch. Which configuration resolves this issue?

- A. snmp-server community Ciscowruser RW 11
- B. snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 22
- C. snmp-server community CiscOU3r RW 11
- D. snmp-server group NETVIEW v2c priv read NETVIEW access 11

Answer: A

Explanation:

Question: 383

Refer to the exhibit.

RKsh run I section eigrp router eigrp 10 network 10*10*10.0 3.G.0.2S5 no auto-sumsry
 neighbor 10.10.10.2 FastEthernetO/O neighbor 10.10.10.3 FastEthexnetO/0

Rifshew Ip eigrp neighbors IP-EIGRP neighbors for process 10

Seq	Address	interface	Hold (sec)	Uptime	SRTT (ms)	RIO	Cnt
1	10.10.10.2	Fa0/0	10	00:01:01	42	232	0 £
0	10.10.10.3	Fa0/0	10	00:01:03	43	244	0 6

Refer to the exhibit The remote branch locations have a static neighbor relationship configured to R1 only R1 has successful neighbor relationships with the remote locations of R2 and R3, but the end users cannot communicate with each other. Which configuration resolves the issue'

R2

```
interface FastEthernetO^O.10
encapsulation dct1Q
jp address W.1C.W.2 25 5.2 55.25 5.0
```

R3

```
interface FastEthernet0/0.10
encapsulation dotfQ
ip address 10.10.10.3 250-255.255.0
```

R2

```
interface FostEthemet0.'0.10
encapsulation det IQ
i p ad drass 10.10.10.2 25 5.2 65.255.0
```

RS

```
interface Fa«EthernetO"O 10
encapsulation dot IQ ip address 10.10.10.3 25 5 255.25 5.0
```

TO

```
interface FastEthernetO/0.10
encapsulation dctIQIO
ip address 10.10.10.2 26 5 255.25 5.0
```

S

```
interface FastEthametO/O.W
encapsulation dodQ 10
ip address 10.10.10.3 255.2 66.25 5.0
```

R2and R3

interface Ethernet0
no ip address

RT

interface FastEthernet0/0
no ip split-horizon eigrp 10

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: E

Explanation:

Question: 384

Refer to the exhibit.

```
enable wa« 5 ^password*
```

```
UsflrpaniB risen prrvilngr- 15 socrol 5 poMMwi/> ijsRframM op&falctf password <p?)s>HrWJ lino vty 0 4
```

```
sdSslOn-llmetiul 240
```

```
password 7 ^password-
```

```
ire^pocet input tew
```

Refer to the exhibit. The authentication is not working as desired and the user drops into user-exec mode. Which configuration resolves the issue?

uannw-mnM


```
333 authentication login default local
330 authorization exec default local
```

```
line vty 0 4
login authentication default authorization exec default
```

```
aaa new-model
333 authentication login default local
333 authorization exec default local
```

```
!intvty04
login authentication default
authorization exec default
```

```
3M new-model
aaa authentication login default local
aaa authorization exec local
```

```
!intvty 0 4
login authentication local
automation exec default
```

```
aaa new-model
aaa authentication login default local
aaa authorization exec default local
```

```
line vty 0-1
login authentication default
authorization exec default
```

A. Option A

B. Option B

C. Option C


D. Option D

Answer: C

Explanation:

Question: 385

Refer to the exhibit.



```

R1(config)# ip verify drop-rate compute window 60
R1(config)# ip verify drop-rate compute interval 60
R1(config)# ip verify drop-rate notify hold-down 60
R1(config)# interface ethernet 0/0
R1(config-if)# ip verify unicast notification threshold 75000
R1(config-if)# snmp trap ip verify drop-rate
R1(config-if)# end

```

Refer to the exhibit. An engineer configured SNMP traps to record spoofed packets drop of more than 48000 a minute on the ethernet0/0 interlace. During an IP spoofing attack, the engineer noticed that no notifications have been received by the SNMP server. Which configuration resolves the issue on R1?

- A. ip verify unicast notification threshold 48000
- B. ip verify unicast notification threshold 8000
- C. ip verify unicast notification threshold 800
- D. ip verify unicast notification threshold 80

Answer: C

Explanation:

Question: 386

SIMULATION

Configure individual VRFs for each customer according to the topology to achieve these goals :

The screenshot shows a network configuration tool interface. On the left, there is a 'Topology Diagram' with two main sections: 'Customer RED' (top) and 'Customer GREEN' (bottom). The 'Customer RED' section shows two routers, R1 and R2, connected to each other and to a central 'BGP 3200 3200' device. The 'Customer GREEN' section shows two routers, R3 and R4, connected to each other and to the same central BGP device. On the right, a terminal window is open, showing a command prompt for router R1. The terminal output is mostly blacked out, but some text is visible: 'R1>', 'R1>', 'R1>', 'R1>', 'R1>', and 'R1>'. A large watermark 'CHINESEDUMPS 通过测试' is overlaid on the terminal window.

The screenshot shows a network configuration tool interface. On the left, there is a 'Tasks' section with a yellow circle around the word 'Tasks'. The text in this section reads: 'Configure individual VRFs for each customer according to the topology to achieve these goals.' followed by a list of three tasks: 1. VRF 'cu-red' has interfaces on routers R1 and R2. Both routers are preconfigured with IP addressing, VRFs, and BGP. Do not use the BGP network statement for advertisement. 2. VRF 'cu-green' has interfaces on routers R3 and R4. 3. BGP on router R1 populates VRF routes between router R1 and R2. 4. BGP on router R2 populates VRF routes between router R1 and R2. 5. LAN to LAN is reachable between SW1 and SW3 for VRF 'cu-red' and between SW2 and SW4 for VRF 'cu-green'. All switches are preconfigured. On the right, a terminal window is open, showing a command prompt for router R1. The terminal output is mostly blacked out, but some text is visible: 'R1>', 'R1>', 'R1>', 'R1>', 'R1>', and 'R1>'. A large watermark 'CHINESEDUMPS 通过测试' is overlaid on the terminal window.

R1

```
R1 R2 SW SW SW SW
R1>
R1> CHINESE DUMPS
R1> #S Sill lit
R1> R1>en
R1> R1sh run
Building configuration...
Current configuration : 1353 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password encryption
!
hostname R1
!
boot-start-marker
boot-end-marker
!
, I XI ED UM PS
no aaa new model
!
```

R1

R2

SW1

SW2

SW3

SW4

CHINESEDUMPS
通过测试



```
ip vrf cu-green  
rd 65000:200
```

```
ip vrf cu-red  
rd 65000:100
```

```
no ip domain lookup  
ip cef
```

```
no ipv6 cef
```

```
multilink bundle-name authenticated
```

CHINESEDUMPS
通过测试

CHINESEDUMPS
通过测试

R1

R2

SW1

SW2

SW3

SW4

CHINESEDUMPS

通过测试

```
interface Loopback0
 ip address 10.10.1.1 255.255.255.255
!
interface Ethernet0/0
 ip address 192.168.1.254 255.255.255.0
 duplex auto
!
interface Ethernet0/1
 ip address 192.168.20.254 255.255.255.0
 duplex auto
!
interface Ethernet0/2
 no ip address
 duplex auto
!
interface Ethernet0/2.100
 encapsulation dot1Q 100
 ip address 10.10.10.1 255.255.255.252
!
interface Ethernet0/2.200
 encapsulation dot1Q 200
 ip address 10.10.20.1 255.255.255.252
```

R1 R2 SW1 SW2 SW3 SW4

```
interface Ethernet0/2.200
 encapsulation dot1q 200
 ip address 10.10.201.1 255.255.255.252
 !
interface Ethernet0/3
 no ip address
 shutdown
 duplex auto
 !
router bgp 65000
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 !
 ip forward-protocol nd
 !
 !
 no ip http server
 no ip http secure-server
 !
 ipv6 ioam timestamp
 !
 control-plane
 !
 !
```

R2

R1 R2 SW1 SW2 SW3 SW4

```
R2>en
R2#Show run
Building configuration...
Current configuration : 1353 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R2
!
boot-start-marker
boot-end-marker
!
!
no aaa new-model
!
!
clock timezone PST -8 0
mmi polling-interval 60
no mmi auto-configure
```


R1

R2

SW1

SW2

SW3

SW4

CHINESEDUMPS

通过测试



```
ip vrf cu-green  
rd 65000:200
```

```
ip vrf cu-red  
rd 65000:100
```

```
no ip domain lookup
```

```
ip cef
```

```
no ipv6 cef
```

```
multilink bundle-name authenticated
```

CHINESEDUMPS

通过测试

R1 R2 SW1 SW2 SW3 SW4

```
CHINESEDUMPS  
interface Loopback0  
ip address 10.10.2.2 255.255.255.255  
!  
interface Ethernet0/0  
ip address 192.168.2.254 255.255.255.0  
duplex auto  
!  
interface Ethernet0/1  
ip address 192.168.22.254 255.255.255.0  
duplex auto  
!  
interface Ethernet0/2  
no ip address  
duplex auto  
!  
interface Ethernet0/2.100  
encapsulation dot1Q 100  
ip address 10.10.10.2 255.255.255.252  
!  
interface Ethernet0/2.200  
encapsulation dot1Q 200  
ip address 10.10.20.2 255.255.255.252
```

R1 R2 SW1 SW2 SW3 SW4

```
interface EthernetO/2.200
 encapsulation dot1Q 200
 ip address 001^5? 255.255.255.252
```

```
!
 iSMiit
 interface EtbernetO/3
```

```
no ip address shutdown duplex auto
```

```
router bgp 65000
```

```
bgp log-neighbor-changes no bgp default ipv4 unicast
```

```
ip forward-protocol nd
```

```
no ip http server no ip http secure-server
```

```
ipv6 loam timestamp
```

```
; CH1NESEDUMPS
```

```
control-plane
```

SW1

R1 R2 SW1 SW2 SW3 SW4

```
SW1>en
SW1#sh run DUMPS
Building configuration database..

Current configuration : 942 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20 22
!
i
version 15.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname SW1 ♦
boot start-marker
boot end-marker
! CHINESEDUMPS no aaa new-model J^J^
clock timezone PST -8 0
```

EDUMPS
iSMiS



```
R1 R2 SW1 SW2 SW3 SW4
no switchport
ip address 192.168.2.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3
!
ip forward-protocol nd
!
ip http server
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 192.168.2.254
ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
!
!
control-plane
!
```

SW2

RI R2 SW1 SW2 SW3 SW4

```
SW2>
SW2> SEDUMPS
SW2>en sama run
SW2lsho configuration
w

Current configuration : 944 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20
22
!
version 1*.* ..... rui
service timestamps debug datetime msec
service timestamps log datetime msec
no service password encryption
service compress-config
!
hostname
!
boot start marker
boot-end-maikoi
! CHINESE DUMPS
! aa»s
!
no aaa new model
```

R1 R2 SW1 SW2 SW3 SW4

```
spanning-tree mode pvsL
spanning -lree gteijd, j^ten-id

| ISMS
```

CHINESEDUMPS

asana

```
interface Ethernet0/0
```

```
interface Ethernet0/1
```

```
no switchport
```

```
ip address 192.168.22.1 255.255.255.0
```

```
interface Ethernet0/2
```

```
interface Ethernet0/3
```

```
interface Ethernet0/3
```



```

!
interface Ethernet0/1
  switchport mode access
  ip address 192.168.22.1 255.255.255.0
!
interface Ethernet0/2
!
interface Ethernet0/3
!
ip forward-protocol nd
!
ip http server
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 192.168.22.254
ip ssh server algorithm encryption aes128-ctr aes256-ctr aes192-ctr
ip ssh client algorithm encryption aes128-ctr aes256-ctr aes192-ctr
!
control-plane

```

www.atmicnetworks.com

www.atmicnetworks.com

SW3

& DUMPS SW3|show

run ^-^^ Building configuration.

Current configuration : 942 bytes

! Last configuration change at 04:43:09 PST Sat May 7 2022

version IS.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config

hostname SW3

boot start marker
boot-end marker

no aaa now model clock
timezone PST -8 0

JESE DUMPS
ISQ»IW

CHINESE DUMPS
WW

R1 R2 SW1 SW2 SW3 SW4

```
spanning-tree mode pvst
spanning-tree extend system-id
!
interface Ethernet0/0
  no switchport
  ip address 192.168.1.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3
```

R1 R2 SW1 SW2 SW3 SW4

```
no switchport
ip address 192.168.1.1 255.255.255.0
interface Ethernet0/1
interface Ethernet0/2
interface Ethernet0/3
ip forward-protocol nd
ip http server
ip http secure-server
ip route 0.0.0.0 0.0.0.0 192.168.1.254
ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
control-plane
```



R1 R2 SW1 SW2 SW3 SW4

SW4>en

SW4#show run

Building configuration...

Current configuration : 944 bytes

!

! Last configuration change at 04:43:09 PST Sat May 7 2022

!

version 15.2

service timestamps debug datetime msec

service timestamps log datetime msec

no service password-encryption

service compress-config

!

hostname SW4

!

boot-start-marker

boot-end-marker

!

!

!

no aaa new-model

clock timezone PST -8 0

!

R1

R2

SW1

SW2

SW3

SW4

```
* spanning-tree mode pvst
spanning-tree extend system-id
```

通过测试

3

```
!
interface Ethernet0/0
```

```
!
interface Ethernet0/1
```

```
no switchport
```

```
ip address 192.168.20.1 255.255.255.0
```

```
CHINESEDUMPS
interface Ethernet0/2
```

```
!
interface Ethernet0/3
```

CHINESEDUMPS
通过测试

R1 R2 \$W1 SW2 SW3 SW4

```
interface Ethernet0/1
no switchport
ip address 192.168.20.1 255.255.255.0

interface Ethernet0/2

interface Ethernet0/3

ip forward-protocol nd

ip http server
ip http secure-server

ip route 0.0.0.0 0.0.0.0 192.168.20.254

ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr

control-plane
```

Guidelines Topology Tasks

Configure individual VRFs for each customer according to the topology to achieve these goals

1. VRF 'cu-red' has interfaces on routers R1 and R2. Both routers are preconfigured with IP addressing, VRFs, and BGP. Do not use the BGP network statement for advertisement
2. VRF "cu-green has interfaces on routers R1 and R2.
3. BGP on router R1 populates VRF routes between router R1 and R2.
4. BGP on router R2 populates VRF routes between router R1 and R2

5. LAN to LAN is reachable between SW1 and SW3 for VRF "cu-red and between SW2 and SW4 for VRF "cu-green* AU switches are preconfigured.

Answer: See the solution below in Explanation.

Explanation:

Solution:

Use cu-red under interfaces facing SW1 & SW3:

On R1:

```
interface Ethernet0/0
```

```
ip vrf forwarding cu-red
```

```
ip address 192.168.1.254 255.255.255.0
```

Check reachability to SW1:

```
R1#ping vrf cu-red 192.168.1.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds: !!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

On R2:

```
interface Ethernet0/0
```

```
ip vrf forwarding cu-red
```

```
ip address 192.168.2.254 255.255.255.0
```

Check reachability to SW3:

```
R2#ping vrf cu-red 192.168.2.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds: !!!!!
```


Use vrf cu-green for SW2 & SW4:

On R1:

```
interface Ethernet0/1
```

```
ip vrf forwarding cu-green
```

```
ip address 192.168.20.254 255.255.255.0
```

Test reachability to SW2:

```
R1#ping vrf cu-green 192.168.20.1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

On R2:

```
interface Ethernet0/1
```

```
ip vrf forwarding cu-green
```

```
ip address 192.168.22.254 255.255.255.0
```

Test reachability to SW4:

```
R2#ping vrf cu-green 192.168.22.1
```

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

On R1:

```
interface Ethernet0/2.100
```

```
mpls ip
```

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

interface Ethernet0/2.200

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

mpls ip

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

!

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Configure BGP:

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

router bgp 65000

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

neighbor 10.10.10.2 remote-as 65000

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

neighbor 10.10.20.2 remote-as 65000

!

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

address-family vpnv4

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

neighbor 10.10.10.2 activate

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

neighbor 10.10.20.2 activate exit-address-family

!

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

address-family ipv4 vrf cu-green

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

redistribute connected

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

exit-address-family

!

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

address-family ipv4 vrf cu-red

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

redistribute connected

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

exit-address-family

!

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

R1(config)#ip vrf cu-red

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

R1(config-vrf)#route-target both 65000:100 !

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

R1(config)#ip vrf cu-green

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

```
R1(config-vrf)#route-target both 65000:200
```

```
On R2:
```

```
interface Ethernet0/2.100 mpls ip
```

```
!
```

```
interface Ethernet0/2.200
```

```
mpls ip
```

```
!
```

```
router bgp 65000
```

```
neighbor 10.10.10.1 remote-as 65000
```

```
neighbor 10.10.20.1 remote-as 65000
```

```
!
```

```
address-family vpnv4
```

```
neighbor 10.10.10.1 activate
```

```
neighbor 10.10.20.1 activate exit-address-family !
```

```
address-family ipv4 vrf cu-green
```

```
redistribute connected
```

```
exit-address-family
```

```
!
```

```
address-family ipv4 vrf cu-red
```

```
redistribute connected
```

```
exit-address-family
```

```
R2(config)#ip vrf cu-red
```

```
R2(config-vrf)#route-target both 65000:100 !
```

```
R2(config)#ip vrf cu-green
```

```
R2(config-vrf)#route-target both 65000:200
```

Verification:

From SW1 to SW3:

```
SW1#ping 192.168.1.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds: !!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

But can't Reach SW2 or SW4 in VRF cu-green:

```
SW1#ping 192.168.22.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds: U.U.U
```

```
Success rate is 0 percent (0/5)
```

```
SW1#ping 192.168.20.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
```

```
U.U.U
```

```
Success rate is 0 percent (0/5)
```

Same Test for SW2:

From SW2 to SW4:

```
SW2#ping 192.168.20.1
```

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

But can't Reach SW3 or SW1 in VRF cu-red:

SW2#ping 192.168.1.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:

U.U.U

Success rate is 0 percent (0/5)

SW2#ping 192.168.2.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:

U.U.U

Success rate is 0 percent (0/5)


Both R1 & R2 has separate tables for VRFs cu-red and cu-green.

Question: 387

SIMULATION

A network is configured with CoPP to protect the CORE router route processor for stability and DDoS protection. As a company policy, a class named class-default is preconfigured and must not be modified or deleted. Troubleshoot CoPP to resolve the

issues introduced during the maintenance window to ensure that:



Guidelines Topology Tasks

A network is configured with CoPP to protect the CORE router processor for stability and DDoS protection. As a company policy, a class named class-default is preconfigured and must not be modified or deleted. Troubleshoot CoPP to resolve the issues introduced during the maintenance window to ensure that:

1. Dynamic routing policies are under CoPP-CRITICAL and are allowed only from the 10.10.x.x range.
2. Telnet, SSH, and ping are under CoPP-IMPORTANT and are allowed strictly to/from 10.10.x.x.
3. All devices on the CORE are in the CORE network. Verify using EoPP on Interface 0.10.x.x range successfully. allow any other IP address do not NORMAL (Hint: Traceroute port range 33434-33464).

WAN

```
CHINESEDUMPS
interface Loopback0
ip address 10.10.2.2 255.255.255.255
interface Loopback1
ip address 172.16.2.2 255.255.255.255
```

WAN CORE MGMT

```
interface Loopback0
ip address 10.10.2.2 255.255.255.255
interface Loopback1
ip address 172.16.2.2 255.255.255.255
interface Ethernet0/0
ip address 10.10.12.1 255.255.255.0 duplex auto
interface Ethernet0/1 no ip address shutdown duplex auto
interface Ethernet0/2 no ip address shutdown duplex auto
interface Ethernet0/3 no ip address shutdown duplex auto
router eigrp 101
network 10.10.0.0
network 172.16.2.0
```



```
router-id 10.10.2.2
network 10.10.0.0 0.0.255.255
network 172.16.2.0 0.0.0.255
eigrp router-id 10.10.2.2
```

CORE

```
class-map match-all CoPP-CRITICAL
match access-group 122
class map match-all CoPP-IMPORTANT
match access-group 121

policy-map COPP
class CoPP-CRITICAL
police 1000000 50000 50000 conform action transmit exceed action
drop
class CoPP-IMPORTANT
police 300000 20000 20000 conform action transmit exceed-
action drop
class CoPP-NORMAL
police 64000 6400 64000 conform action transmit exceed ac
tion drop
class class-default
police 8000 1500 1500 conform action transmit exceed-
action drop
```

```
int^rra<U tk^QUM PS
ip address 10.10^13^6^.2 55.255.2 55
```

```
interlace Ethernet0/0
ip address 10.10.12.1 255.255.255.0
duplex auto
```

```
interface EthernetO/L ip
address 10.10.13.1 duplex If^puMPS
auto
```

```
duplex auto sS^^lijt
```

```
interlace Ethern&tQ/2 no ip address
shutdown duplex auto
```

```
J
interface EthornctO/3
no ip address shutdown
duplex auto
```

H1NESEDUMPS

```
$
router eigrp 101 network 10.10.0.0
0.0.255.255 eigrp router-id 10.10.1.1
```

```
ip forward-protocol nd
```

```
" 5 85 3Xg|g
```

```
!
CHINESEDUMPS
access-list 120 remark *** ACL for CoPP-Critical ***
access-list 121 remark *** ACL for CoPP-IMPORTANT
access-list 122 remark *** ACL for CoPP-NORMAL
!
control-plane
  service-policy input CoPP
!
```

CHINESEDUMPS
通过测试

MGMT

WAN CORE MGMT

```
interface Loopback0
  ip address 172.16.3.3 255.255.255.0

interface Ethernet0/0
  no ip address
  shutdown
  duplex auto

interface Ethernet0/1 ip address 10.10.13.3
  255.255.255.0 duplex auto

interface Ethernet0/2 - E >EDUMPS no ip addtes, -
  shutdown duplex auto

interface Ethernet0/3
  no Ip address shutdown duplex auto

router ^!ESSJUMPS
  network 10.10.0.0 <^17^,255
  network 112.16.3.0 0.0.0.255
```

WAN CORE MGMT

```
no ip address
shutdown DUMPS
duplex auto a^H

router eigrp 101
 network wlio.o*o 0.0,255*265
 network 172,16,1,0 0.0.0.255
 eigrp router-id 10*10.3*3

ip forward-protocol nd
!
no ip http server
no ip http sece re-server 1
ipv6 ioam timestamp
```

```
control plane
-CHINESEDUMPS
1 »$»«
```

Explanation:

CORE

policy-mao CoPP

class CoPP-CRITICAL

**Answer: See
the
solution below
in
Explanation.**

police 1000000 50000 50000 conform-action transmit exceed-action transmit

```
access-list 120 remark *** ACL for CoPP-Critical ***
access-list 120 permit ip 10.10.0.0 0.0.255.255 any
access-list 120 permit tcp any any
access-list 120 permit ip any 10.10.0.0 0.0.255.255
access-list 121 permit icmp 10.10.0.0 0.0.255.255 any
access-list 121 permit tcp 10.10.0.0 0.0.255.255 any eq 22
access-list 121 permit tcp 10.10.0.0 0.0.255.255 any eq telnet
access-list 122 remark *** ACL for CoPP-NORMAL
access-list 122 permit udp 10.10.0.0 0.0.255.255 any
access-list 122 permit udp any 10.10.0.0 0.0.255.255
access-list 122 permit udp any 10.10.0.0 0.0.255.255 range 33434 33464
access-list 122 permit udp 10.10.0.0 0.0.255.255 any range 33434 33464
!
control-plane
service-policy input
!
```

Text Description automatically generated with medium confidence

CORE# Copy run start

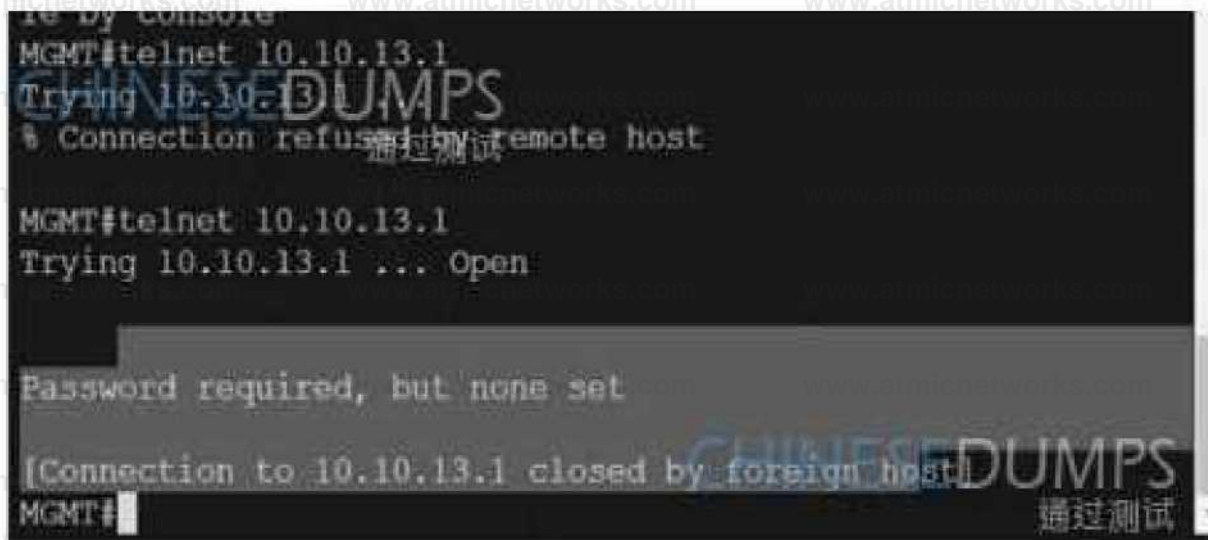
TESTING: -

CORE



Graphical user interface Description automatically generated with medium confidence

MGMT

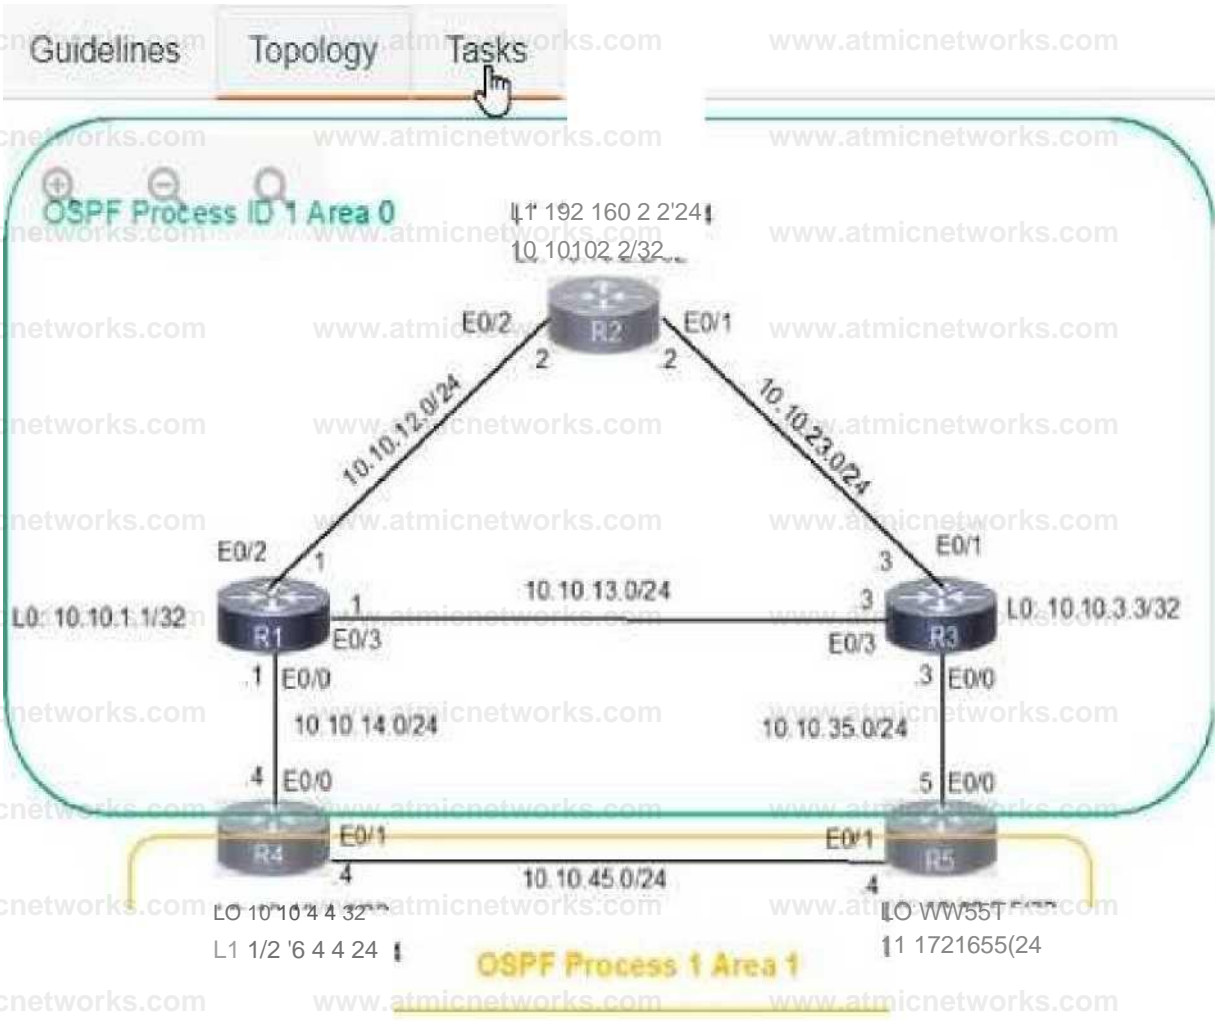


Graphical user interface, text Description automatically generated

Question: 388

SIMULATION

A network is configured with IP connectivity, and the routing protocol between devices started having problems right after the maintenance window to implement network changes. Troubleshoot and resolve to a fully functional network to ensure that:



Topology Diagram

Guidelines Topology Tasks

A network is configured with IP connectivity, and the routing protocol between devices started having problems right after the maintenance window to implement network changes. Troubleshoot and resolve to a fully functional network to ensure that:

1. Inter-area links have link authentication (not area authentication) using MD5 with the key 1 string CCNP.
2. R3 is a DR regardless of R2 status while R1 and R2 establish a DR/BDR relationship.
3. OSPF uses the default cost on all interfaces. Network reachability must follow OSPF default behavior for traffic within an area over intra-area VS inter-area links.
4. The OSPF external route generated on R4 adds link cost when traversing through the network to reach R2. A network command to advertise routes is not allowed.

R2 R4 R5

```
R2>en R2# R2J R2# r
R2| 1
R2f E2f E2#sh run Building configuration,.
Current configuration : 1279 bytes 1 version. IS.9
service timestamps debug datetime msec service timestamps log datetime msec no service password-
encryption i
hostname R2 i boot-sta rt-market boot-end-marker !!! no aaa new-model
```

j

```
clock timezone PST -S 0 mmi polling-interval 60 no mi auro-configure
```

Activate Windc
Go to Settings fer

R2 R4 R5

```
interface Loopback0
 ip address 10.10.2.2 255.255.255.255
 ip ospf 1 area 0
```

```
!
interface Loopback1
 ip address 192.168.2.2 255.255.255.0
 ip ospf 1 area 0
```

```
!
interface Ethernet0/0 no
 ip address shutdown
 duplex auto
```

```
!
interface Ethernet0/1
 ip address 10.10.23.2 255.255.255.0
 ip ospf 1 area 0
 duplex auto
```

```
!
interface Ethernet0/2
 ip address 10.10.12.2 255.255.255.0
 ip ospf 1 area 0
 duplex auto
```

```
!
interface Ethernet0/3 no ip address
 shutdown duplex auto
```

```
!
router ospf 1
 passive-interface default
 no passive-interface Ethernet0/1
 no passive-interface Ethernet0/2
```

Activate Windows

Go to Settings to activate!

R2 R4 R5

```
interface Ethernet0/3 no ip address shutdown duplex auto
```

```
!
router ospf "1
```

```
passive-interface default
no passive-interface Ethernet0/1
no passive-interface Ethernet0/2
```

```
ip forward-protocol nd I
```

```
no ip http server no ip http secure server
```

```
ipv6 ioam timestamp r r
```

```
control-plane
```

```
B I
```

```
i
```

```
line con 0
```

```
Activate
```

```
Goto Satti i
```

R4

```
R2 R4 R5
R4>
R4>
R4>
R4>
R4>en
R4#sh run
Building configuration...

Current configuration : 1479 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R4
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
!
!
!
clock timezone PST -8 0
mmi polling-interval 60
no mmi auto-configure
no mmi pvc
--More--
```

Activate V
Go to Setting

R2 R4 R5

```
key chain CCNP
  key 1
    key-string ccnp
    cryptographic-algorithm md5
!
!
!
!
!
!
ip address 172.16.4.4 255.255.255.0
!
interface Ethernet0/0
  ip address 10.10.14.4 255.255.255.0
  ip ospf authentication key-chain CCNP
  ip ospf 1 area 0
  duplex auto
!
interface Ethernet0/1
  ip address 172.16.45.4 255.255.255.0
  ip ospf 1 area 1
  duplex auto
!
interface Ethernet0/2
  no ip address
  shutdown
  duplex auto
!
interface Ethernet0/3
  no ip address
  shutdown
  duplex auto
```

Activate
Go to Sett

R2 Wd R5

```
router ospf 1
 redistribute connected subnets route-map to-ospf passive-inrerrace default
 no passive-interface Ethernet0/0
 no passive-interface EthemetO/1 E
```

```
ip forward-protocol nd
|
no ip http server
no ip http secure-server
I
|
ipv6 ioam timestamp 1
route-map to-ospf permit 10 match interface Loopbackl
```

```
I j control-plane !
```

```
i
```

```
i |
```

```
1
```

```
line con 0
 logging synchronous
1 me aux 0
```

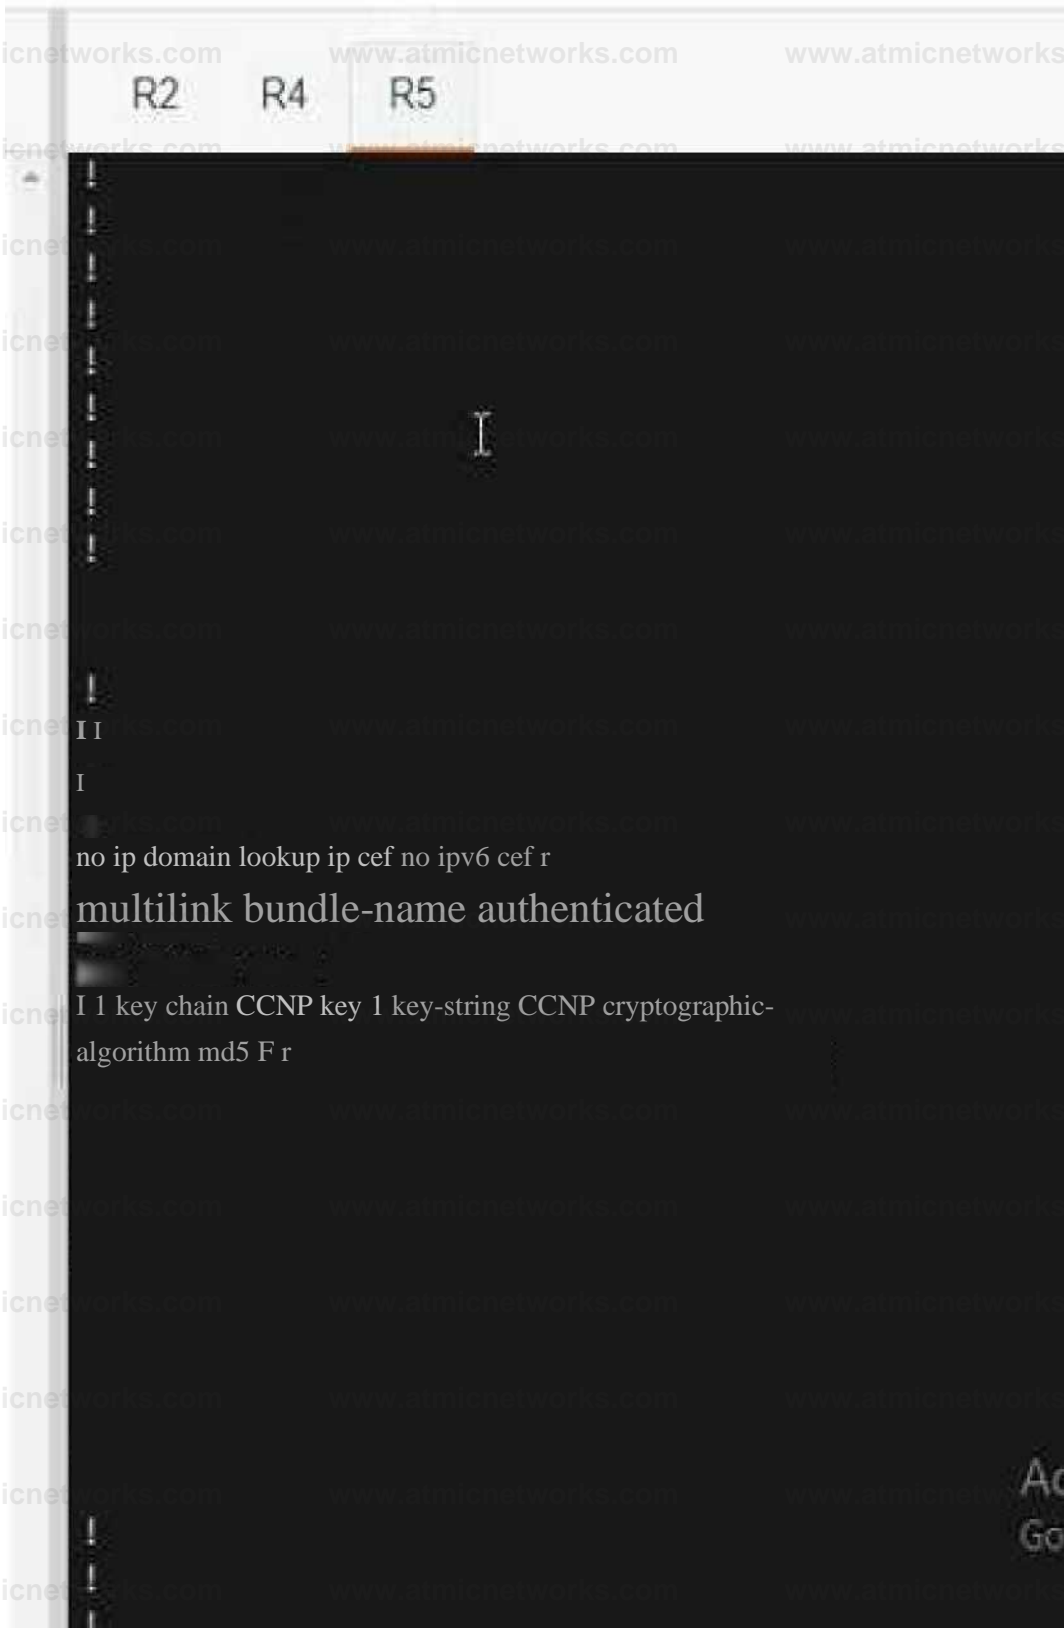
Activate W
Goto Settings

R5

```
R5>
R5>
R5>en R5i
R5#sh run
Building configuration...

Current configuration : 1496 bytes i
version 15*S
service rirftestrampG debug datetime msec service timestamps log datetime msec no service
password-encryption 1 hostname R5
i
boot-start-marker boot-end-marker I 1

no aaa new-model
{ [ r clock timezone PST -S 0 mmi polling-interval 60 Activate WI
no mmi auto-conf igure Gc to Settings
no mmi pvc
—More— |
```

R2 R4 R5

ii j

interface LoopbackJ

ip address 10.10.5.5 255,255,255.255

ip ospf 1 area 1 L

interface LoopbackI

ip address 172.16,5.5 255.255.255.0 i

interface Ethernet0/0

ip address 10.10.35.5 255.255.255.0

ip ospf authentication key-chain CCNP

ip ospf 1 area 0 duplex auto r

interface EthernetO/1

ip address 172.16.45.5 255.255.255.0

ip ospf 1 area 1

ip ospf cost 60 duplex auto 41 interface Ethernet0/2 no ip address shutdown

duplex auto

A

t

Gt

interface EthemetO/3

no ip address

R2 R4 R5

```
router ospf 1
 redistribute connected subnets route-map to-ospf passive-interface default
 no passive-interface EthernerO/O
 no passive-interface EthernetO/1

!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-
 server
!
ipv6 ioam timestamp
!
 route-map to ospf permit 10 match interface
      Loopback1
!
!
control-plane
!
!
!
!
!
line con 0
 logging synchronous
line aux 0
```

Activate Window
■ -o to Settings to act

**Answer: See
the
solution below
in
Explanation.**

Explanation:

R4

Int range et0/0 – 1

Ip ospf authentication message-digest

Ip ospf message-digest-key 1 md5 CCNP

Router ospf 1

Redistribute connected subnets route-map to-ospf metric-type 1

Copy run start

R5

Int range et0/0 – 1

Ip ospf authentication message-digest

Ip ospf message-digest-key 1 md5 CCNP

Interface eth 0/1

Ip ospf cost 10

Copy run start

VERIFICATION:-

```
R2#show ip ospf nei
R2#show ip ospf neighbor
```

Neighbor ID	ntface	Pri	State	Dead Time	Address
10.10.1.1	thernetO/2	1	FULL/BDR	00:00:38	10.10.12.1

```
10.10.3.3 themetO/I 1 FULL/BDR 00:00:3071 VdtO. 10!. 23 <3V S E
R2*| Go to Settings to activate Wine
```

Graphical user interface, text, application Description automatically generated

Question: 389

What are two characteristics of a VRF instance? (Choose two)

- A. It is defined by the VPN membership of a customer site attached to a P device.
- B. Each VRF has a different set of routing and CEF tables.
- C. All VRFs share customers routing and CEF tables.
- D. An interface must be associated to one VRF
- E. A customer site can be associated to different VRFs.

Answer: B,D

Explanation:

Question: 390

The network administrator configured CoPP so that all routing protocol traffic toward the router CPU is

limited to 1 mbps. All traffic that exceeds

this limit must be dropped. The router is running BGP and OSPF Management traffic for Telnet and SSH

must be limited to 500kbps.

access-list 100 permit tcp any any eq 179

access-list 100 permit tcp any any range 22 23

access-list 100 permit ospf any any

class-map CM-ROUTING

match access-group 100

class-map CM-MGMT

match access-group 100

policy-map PM-COPP

class CM-ROUTING

police 1000000 conform-action transmit

class CM-MGMT

police 500000 conform-action transmit

!

control-plane

service-policy output PM-COPP

No traffic is filtering through CoPP, which is resulting in high CPU utilization, which configuration resolves the issue ?

A. no access-list 100
access-list 100 permit tcp any any eq 179
access-list 100 permit ospf any any
access-list 101 Permit tcp any any range 22 23
class-map CM-MGMT
no match access-group 100
match access-group 101

B. control-plane
no service-policy output PM-COPP
service-policy input PM-COPP

C. No access-list 100
access-list 100 permit tcp any any eq 179
access-list 100 permit tcp any any range eq 22
access-list 100 permit tcp any any range eq 23
access-list 100 permit ospf any any

D. no access-list 100
access-list 100 permit tcp any any eq 179
access-list 100 permit ospf any any
access-list 101 Permit tcp any any range 22 23
class-map CM-MGMT
no match access-group 100
match access-group 101
control-plane
no service-policy output PM-COPP
service-policy input PM-COPP

Answer: D

Explanation:

Question: 391

An engineer is creating a policy that overrides normal routing behavior. If the route to a destination of 10.100.100.0/24 is withdrawn from the routing table, the policy must direct traffic to a next hop of 10.1.1.1. If the route is present in the routing table, then normal forwarding must occur. Which configuration meets the requirements?

- A. `access-list 100 permit ip any any!``route-map POLICY permit 10``match ip address 100``set ip next-hop recursive 10.1.1.1`
- B. `access-list 100 permit ip any 10.100.100.0 0.0.0.255!``Route-map POLICY permit 10``match ip address 100``set ip default next-hop 10.1.1.1`
- C. `access-list 100 permit ip any 10.100.100.0 0.0.0.255!``route-map POLICY permit 10``match ip address 100``set ip next-hop 10.1.1.1!``route map POLICY permit 20`
- D. `access-list 100 permit ip any 10.100.100.0 0.0.0.255!``route map POLICY permit 10``match ip address 100``Set ip next-hop recursive 10.1.1.1!``route-map POLICY permit 20`

Answer: D

Explanation:

Dallas_Router:

```
interface GigabitEthernet0/0/0.364
description Guest_Wifi_10.66.46.0/23
encapsulation dot1Q 364
ip address 10.66.46.1 255.255.254.0
ip helper-address 10.192.104.212
ip helper-address 10.191.103.140
ip access-group GUEST-ACCESS in
ip access-group GUEST-ACCESS-OUT out
no ip redirects
no ip unreachable
no ip proxy-arp
```

ip access-list extended GUEST-ACCESS

```
remark Internet Access Only
permit udp any any eq bootpc
permit udp any any eq bootps
deny ip any 10.0.0.0 0.255.255.255
deny ip any 172.16.0.0 0.15.255.255
deny ip any 192.168.0.0 0.0.255.255
deny ip any 224.0.0.0 31.255.255.255
deny ip any 169.254.0.0 0.0.255.255
deny ip any 127.0.0.0 0.255.255.255
deny ip any 192.0.2.0 0.0.0.255
deny ip any host 0.0.0.0
permit ip 10.66.42.0 0.0.0.255 any
permit ip 10.66.46.0 0.0.0.255 any
```

ip access-list extended GUEST-ACCESS-OUT

```
remark Used to block inbound traffic to Guest Networks
permit udp any any eq bootps
permit udp any any eq bootpc
permit udp any any eq domain
permit udp any any
permit icmp any any
permit tcp host 10.192.103.124 eq 15871 any
permit tcp any any established
deny ip any 10.0.0.0 0.255.255.255
deny ip any 172.16.0.0 0.15.255.255
deny ip any 192.168.0.0 0.0.255.255
deny ip any 224.0.0.0 31.255.255.255
deny ip any 169.254.0.0 0.0.255.255
deny ip any 127.0.0.0 0.255.255.255
deny ip any 192.0.2.0 0.0.0.255
deny ip any host 0.0.0.0
```

After a new regional office is set up, not all guests can access the internet via guest Wi

Fi. Clients are getting the correct IP address from guest Wi-Fi VLAN 364. which action resolves the issue ?

- A. Allow 10.66.46.0/23 in the outbound ACL
- B. Allow DNS traffic through the outbound ACL
- C. Allow DNS traffic through the inbound ACL
- D. Allow 10.66.46.0/23 in the inbound ACL

Answer: C

Explanation:

Question: 393

An engineer configures PBR on R5 and wants to create a policy that matches traffic destined toward

10.10.10.0/24 and forward 10.1.1.1. The traffic must also have its IP precedence set to 5. All other traffic should be forward toward 10.1.1.2 and have its IP precedence set to 0. Which configuration meets the requirements?

- A.

```
access-list 1 permit 10.10.10.0 0.0.0.255
access-list 2 permit any
route-map CCNP permit 10
match ip address 1
set ip next-hop 10.1.1.1
set ip precedence 5
route-map CCNP permit 20
match ip address 2
set ip next-hop 10.1.1.2
set ip precedence 0
route-map CCNP permit 30
```
- B.

```
access-list 100 permit ip any 10.10.10.0 0.0.0.255
route-map CCNP permit 10
match ip address 100
set ip next-hop 10.1.1.1
set ip precedence 0
route-map CCNP permit 20
set ip next-hop 10.1.1.2
set ip precedence 5
route-map CCNP permit 30
```
- C.

```
access-list 1 permit 10.10.10.0 0.0.0.255
route-map CCNP permit 10
match ip address 1
set ip next-hop 10.1.1.1
set ip precedence 5
route-map CCNP permit 20
set ip next-hop 10.1.1.2
set ip precedence 0
```

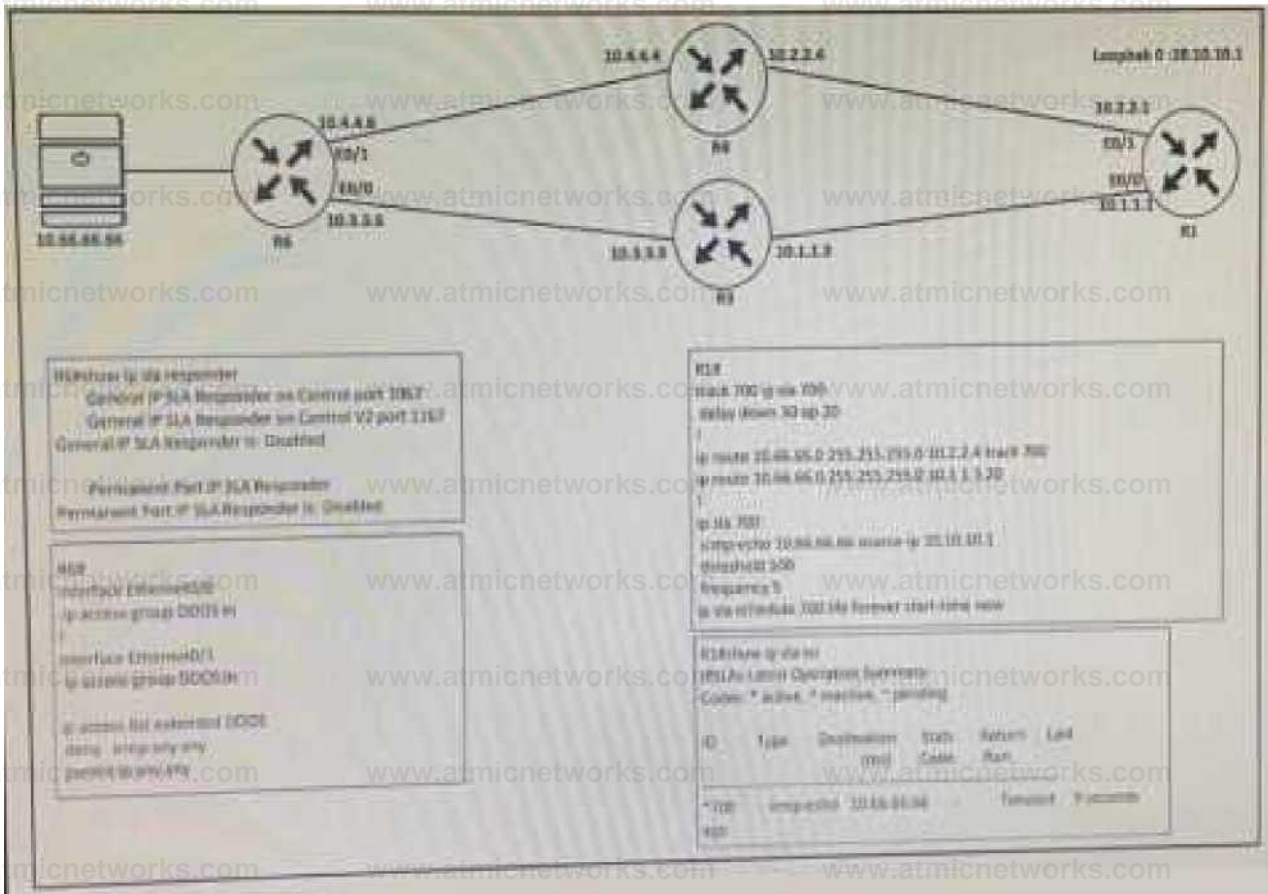
D. access-list 100 permit ip any 10.10.10.0 0.0.0.255 route-map CCNP permit 10 match ip address 100 set ip next-hop 10.1.1.1 set ip precedence 5 ! route-map CCNP permit 20 set ip next-hop 10.1.1.2 set ip precedence 0

Answer: D

Explanation:

Question: 394

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. tacass server enable-password Cisco@123
- C. tacacs-server enable-password Cisco@123
- D. enable-password Cisco@123

Answer: D

Explanation:

Question: 395

Which MPLS value is combined with the IP prefix to convert to a VPNv4 prefix?

- A. 16-byte Route Distinguisher
- B. 8-byte Route Target
- C. 16-byte Route Target
- D. 8-byte Route Distinguisher

Answer:

D

Explanation:

Question:

396

What is a function of the IPv6 DHCP Guard feature for DHCP messages?

- A. Only access lists are supported for matching traffic.
- B. All client messages are always switched regardless of the device role.
- C. It blocks only DHCP request messages.
- D. If the device is configured as a DHCP server, no message is switched.

Answer:

B

Explanation:

Question:

397

An engineer creates a default static route on a router with a hop of 10.1.1.1. On inspection, the engineer finds the router has two VRFs, Red and Blue. The next hop is valid for both VRFs and exists in each assigned VRF. Which configuration achieves connectivity?

A)

```
ip route vrf BLUE 0.0.0.0 0.0.0.0 10.1.1.1
```

B)

```
ip route vrf BLUE 0.0.0.0 0.0.0.0 10.1.1.1
```

C)

```
ip route vrf BLUE 0.0.0.0 0.0.0.0 10.1.1.1
```

D)

```
ip route vrf BLUE 0.0.0.0 0.0.0.0 10.1.1.1
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

Question: 398

Refer to the exhibit.

```
R1(config)#ipv6 prefix-list PRE-PEND-PREFiX permit 2001:db8:0:a:764
R1 (config)# route-map PRE-PEND permit 10
R1(config-route-map)#match ipv6 address prefix-list PRE PEND PREFIX
R1(config-route-map)#set as-path prepend 65412
R1 (config)#router bgp 65412
R1 (config-router )#address-family ipv6
R1 (config-router-af)#neighbor 2001:db8:0:2::2 route-map PRE-PEND out
```

R1 has a route map configured, which results in a loss of partial IPv6 prefixes for the BGP neighbor, resulting in service degradation. How can the full service be restored?

- A. The neighbor requires a soft reconfiguration, and this will clear the policy without resetting the BGP TCP connection.
- B. The prefix list requires all prefixes that R1 is advertising to be added to it, and this will allow additional prefixes to be advertised.
- C. The route map requires a deny 20 statement without set conditions, and this will allow additional prefixes to be advertised.
- D. The route map requires a permit 20 statement without set conditions, and this will allow additional prefixes to be advertised.

Answer: D

Explanation:

Question: 399

A customer is running an mGRE DMVPN tunnel over WAN infrastructure between hub and spoke sites. The existing configuration allows NHRP to add spoke routers automatically to the multicast NHRP mappings. The customer is migrating the network from IPv4 to the IPv6 addressing scheme for those spokes' routers that support IPv6 and can run DMVPN tunnel over the IPv6 network. Which configuration must be applied to support IPv4 and IPv6 DMVPN tunnel on spoke routers?

- A. Tunnel mode ipv6ip 6to4
- B. Tunnel mode ipv6ip isatap

C. Tunnel mode ipv6ip auto-tunnel

D. Tunnel mode ipv6ip 6rd

Answer: C

Explanation:

Question: 400

Refer to the exhibit.

```
R1(config)#ip access-list standard EIGRP-FILTER
R1(config-std-nacl)#permit 10.10.10.0 0.0.0.255
R1(config)#router eigrp 10
R1(config-router)#distribute-list route-map EIGRP in
!
R1(config)#route-map EIGRP permit 10
R1(config-route-map)#match ip address EIGRP-FILTER
!
R1#show ip route eigrp
D 10.10.10.0/24
```

An engineer must filter incoming EIGRP updates to allow only a set of specific prefixes. The distribute list is tested, and it filters out all routes except network 10.10.10.0/24. How should the engineer temporarily allow all prefixes to be learned by the routers again without adjusting the existing access list?

A. A permit 20 statement should be added before completing the ACL with the required prefixes, and then the permit 20 statement can be removed.

B. A permit any statement should be added before completing the ACL with the required prefixes and then the permit any statement can be removed.

C. A continue statement should be added within the permit 10 statement before completing the ACL with the required prefixes, and then the continue statement can be removed.

D. An extended access list must be used instead of a standard access list to accomplish the task

Answer: C

Explanation:

Question: 401

Refer to the exhibit.

A network engineer receives a fault ticket about traffic drops from BANK SITE to BANK

Users can reach BANK SITE Y from router RA as a source.

Routers RB and RD are acting as route reflectors.

Which configuration resolves the issue?

A. RC(config)#router bgp 65201RC(config-router)#neighbor 10.10.10.4 route-reflector-client

B. RF(config)#router bgp 65201RF(config-router)#neighbor 10.10.10.6 route-reflector-client

C. RC(config)#router bgp 65201RC(config-router)#neighbor 10.10.10.2 route-reflector-client

D. RB(config)#router bgp 65201RB(config-router)#neighbor 10.10.10.3 route-reflector-client

Answer: A

Explanation:

Question: 402

Refer to the exhibit.

```
CPE# show ntp associations
address      ref clock    st  when  poll reach  delay
offset disp
-10.1.255.40 .INIT.      16  -    64    0 0.000
0.000 15937.
* syn.peer, † selected, + candidate, - outlier, x failed,
- configured

CPE# debug ip icmp
*Feb 20 22:49:32.913: ICMP: dst (10.0.12.1) port unreachable rcv
from 10.1.255.40
*Feb 20 22:50:37.918: ICMP: dst (10.0.12.1) port unreachable rcv
from 10.1.255.40
*Feb 20 22:51:44.951: ICMP: dst (10.0.12.1) port unreachable rcv
from 10.1.255.40
```

An administrator is troubleshooting a time synchronization problem for the router time to another Cisco IOS XE-based device that has recently undergone hardening. Which action resolves the issue?

- A. Allow NTP in the ingress ACL on 10.1.225.40 by permitting UDP destined to port 123.
- B. Ensure that the CPE router has a valid route to 10.1.255.40 for NTP and rectify if not reachable.
- C. NTP service is disabled and must be enabled on 10.1.225.40.
- D. Allow NTP in the ingress ACL on 10.1.255.40 by permitting TCP destined to port 123.

Answer: C

Explanation:

Question: 403

Refer to the exhibit.

Ferine* ifiinr&r i^m		
tHfirtnc	IFMMfn DKV IMnd SMu	Prtrtid
Elhanwl IKI	c'n?0Rii TEQ tfrtuipiuii	u£
UKph«t1	iPsln it i Yt& mm4up	up
IK^IMU I	172 Ifl iM L ¥E & rnmrb jJ up	14
JHpfaKkJ	IES 1B.1S0.1 yes m«*I t*	UI
RIH sin* p T j'p uiNfli'tou		
E1GBP *^J Maiqbaii far JU^IJ		
H AtatM* iniMiHe* unit! jOfcte ^ATTnrOQ^G (w lrwICMHum		
o mu.1^ E^ M QMlifi? UHS «»&»		

K IP rima ip ff^ip M« 1,1 ^■
WRP4M T«*iatowTobtel«A5liyiCS^^
DK»a p ■ Pas^ A *^- U -Lpcxn. Q -^1?.* H*r^ r FW^WW." MS^

191 US miKI. 1 wciaK^i FTM
via area it ■J [4WWI2frai t mmiiM ■12
!H15 UX 1 HUCCCIWft. ^t>11 sass, soma 2

p ^ W f B8JJ0 0^4 I HJtEPiKflh- Q a
' I^K-JW'JBW.EIli^^

P2&1DITJWW ' suCWUDtB, FD sifl^&M >«' SOHN^S

p tR 16 1X! #H^1 wu*****! FD &1^a- »^ ^
^ cwi™w Ltnpta^I W

p 17'2 ia K.Wi 1 Sii«ei-»cni: FD K 12KW -^ * *
^rnnntfMI ILMBD^I

P igzi6&20
0^4 1 iwtMSKrtii rP* De 408660. wno 36

Routers R1 and R2 have established a network adjacency using EIGRP, and both routers are advertising subnets to its neighbor. After issuing the show ip EIGRP topology all-links command in R1, some prefixes are no showing R2 as a successor. Which action resolves the issue?



- A. Rectify the incorrect router ID in R2.

- B. Enable split-horizon.
- C. Configure the network statement on the neighbor.
- D. Resolve the incorrect metric on the link.

Answer: D

Explanation:

Question: 404

How does LDP operate in an MPLS network?

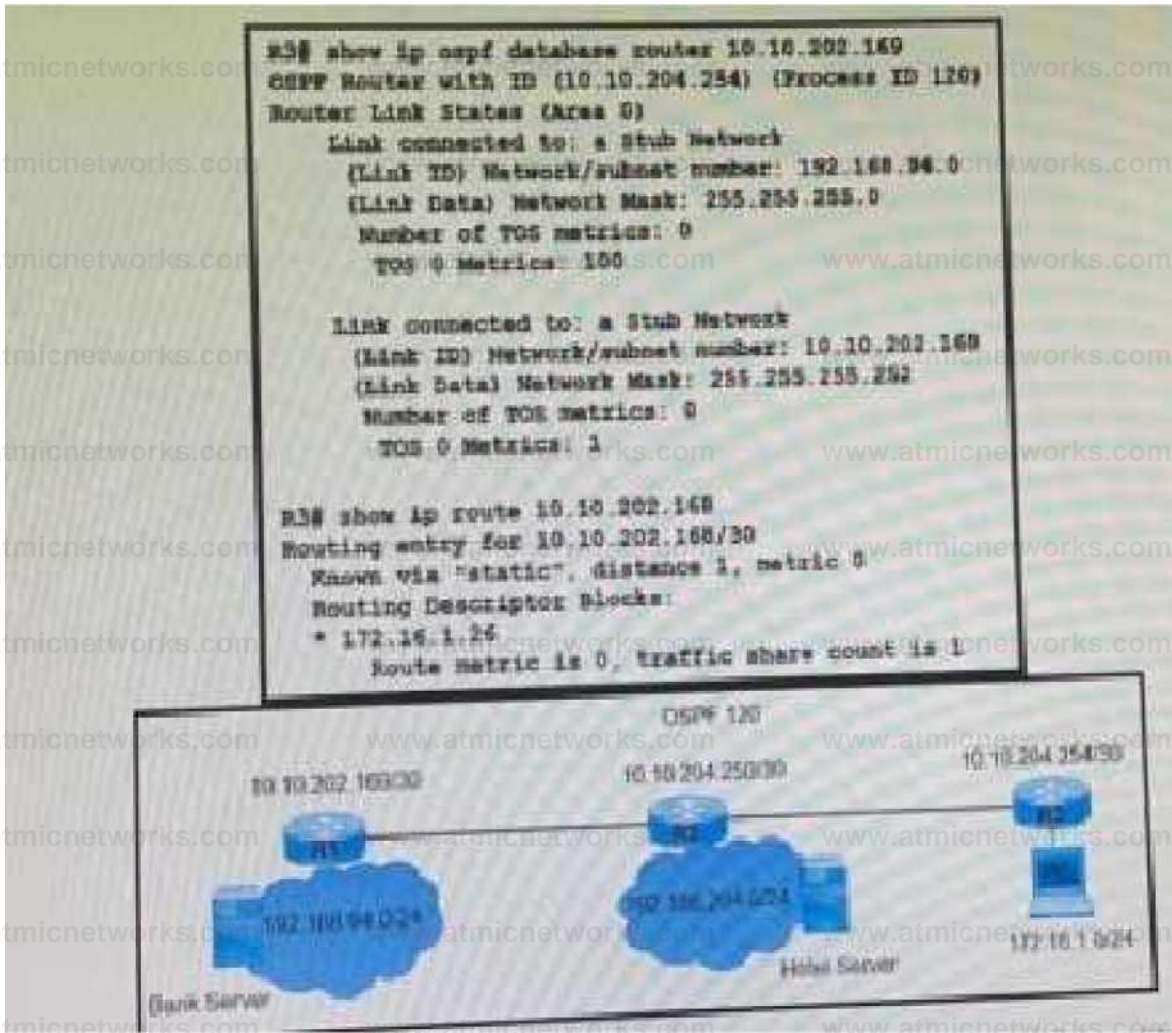
- A. When topology changes occur such as a router failure, LDP generates peer discovery messages that terminate the LDP session to propagate an LSP change.
- B. When an adjacent LSR receives LDP discovery messages, TCP two-way handshake ensures that the LDP session has unidirectional connectivity.
- C. Peer routers establish the LDP session, and the LDP neighbors maintain and terminate the session by exchanging messages
- D. LDP notification messages allow LERs to exchange label information to determine the next hops within a particular LSP.

Answer: D

Explanation:

Question: 405

Refer to the exhibit.



A network engineer finds that PC1 is accessing the hotel website to do the booking but fails to make payment. Which action resolves the issue?

- A. Allow stub network 10.10.202.168/30 on router R3 OSPF.
- B. Decrease the AD to 5 OSPF route 192.168.94.0 on R1.
- C. Increase the AD to 200 of static route 192.168.94.0 on R3.
- D. Configure a reverse route on R1 for PC1 172.16.1.0/24.

Answer: A

Explanation:

Question: 406

Refer to the exhibit.

```

R1#show ip bgp 10.10.10.4/32
BGP routing table entry for 10.10.10.4/32, version 31
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Not advertised to any peer
65201
 10.10.10.5 (metric 2) from 10.10.10.5 (10.10.10.5)
  Origin IGP, metric 0, localpref 100, valid, internal
65201
 10.10.10.2 (metric 2) from 10.10.10.2 (10.10.10.2)
  Origin IGP, metric 0, localpref 100, valid, internal, best

R2#show ip bgp 192.168.1.1/32
BGP routing table entry for 192.168.1.1/32, version 34
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Not advertised to any peer
65101
 10.10.10.4 (metric 156160) from 10.10.10.4 (10.10.10.4)
  Origin IGP, metric 0, localpref 100, weight 32768, valid, internal, best
65101
 10.10.10.3 (metric 2) from 10.10.10.3 (10.10.10.3)
  Origin IGP, metric 0, localpref 100, valid, internal

```

The diagram shows two Autonomous Systems (AS101 and AS201) connected via a central link. AS101 contains routers R1, R2, and R3. AS201 contains routers R4, R5, and R6. R1 is connected to R2, R2 to R3, and R3 to R4. R4 is connected to R5, R5 to R6, and R6 to R7. R7 is connected to a server icon. The central link connects R3 and R4. The diagram illustrates a network topology where traffic from the server to the AAA server (R1) is not using the primary path via R3-R2 link.

A customer reports that user traffic of bank XYZ to the AAA server is not using the primary path via the R3-R2 link. The network team observes:

No fiber is cut on links R2 and R3.

As101 and AS 201 routers established BGP peering.

Which configuration resolves the issue?

A)

```
^ / u - " . ! ^ r outturn J p BOP-Pat h pvrmill 10
R ^ COLLifl'inouifr™ irtmtwic lffD
FQ ion ^ dWn * ui ** MP ^ 1 ^
R ■ ' :-n' j itMiterJtf rwighbiW lfi 10 i0 ! rwuiMTup BWMML'1 CM*
```

B)

```
R6(config)#router bgp 65201
R6(config-router)#no neighbor 10.10.10.5 weight 32769
```

C)

```
R4(config)#router bgp 65201
R4(config-router)#no neighbor 10.10.10.5 weight 32769
```

D)

```
R1(config)#route-map BGP-Path permit 10
R1(config-route-map)# set local-preference 200
R1(config)#router bgp 65101
R1(config-router)# neighbor 10.10.10.2 route-map BGP-Path out
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 407

Refer to the exhibit.


IX* >hw m tv^nw^Fr e** IP EL^ lwdW hhh^/Vnn jajftH

lit JKWIJ 413 hsw 1 ft**

MWT<v*MV w*^{ic} Amiaw.ttl) A m ;n jOwi* i KCWMI, TO H *<KWWU

IMinrr/WitAll t-IrrrwKVU WW 10 At !□ |0,m J (WtaRW**n*il *4 MhwrUA w«» B. in<NK»d

A bank ATM site has difficulty connecting with the bank server. A network engineer troubleshoots the issue and finds that R4 has no active route to the bank ATM site. Which action resolves the issue?



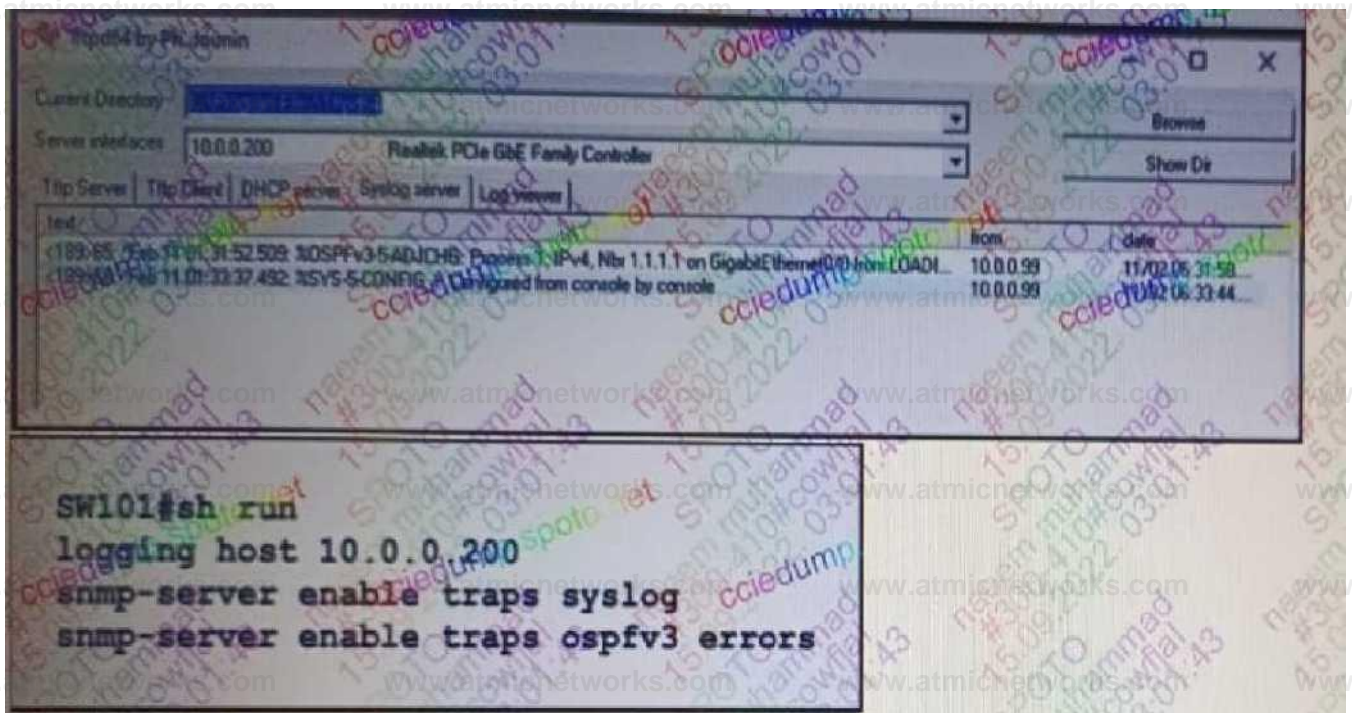
- A. Advertise 10.10.30.0/24 subnet in R1 EIGRP AS.
- B. EIGRP peering between R3 and R4 to be fixed.
- C. EIGRP peering between R1 and R2 to be fixed.
- D. Advertise 10.10.30.0/24 subnet in R3 EIGRP AS.

Answer: D

Explanation:

Question: 408

Refer to the exhibit.



An engineer configures SW101 to send OSPFv3 interfaces state change messages to the server. However, only some OSPFv3 errors are being recorded. which organization resolves the ..?

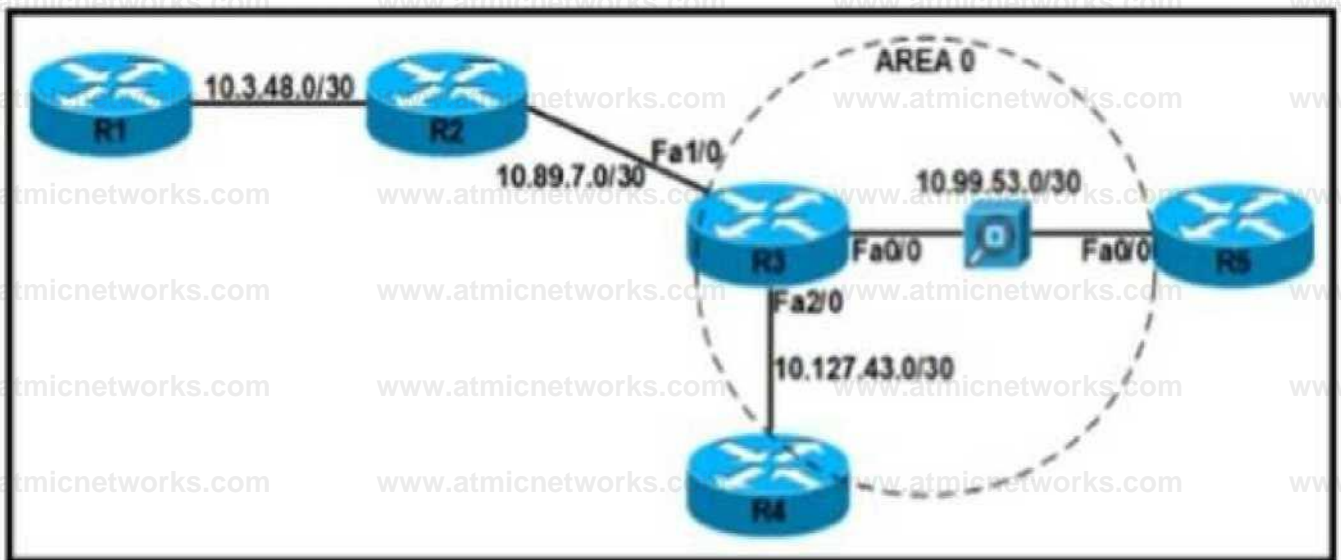
- A. snmp-server enable traps ospfv3 state-change if-state-change
- B. snmp-server-enable traps ospfv3 state-change restart-status-change
- C. snmp-server-enable traps ospfv3 state-change neighbor-state-change.
- D. snmp-server-enable traps ospfv3 state-change if-state-change neighbor-state-change

Answer: D

Explanation:

Question: 409

Refer to the exhibit.



The security department recently installed a monitoring device between routers R3 and R5, which a loss of network connectivity for users connected to R5. Troubleshooting revealed that the monitoring device cannot forward multicast packets. The team already updated R5 with the correct configuration. Which configuration must be implemented on R3 to resolve the problem by ensuring R3 as the DR for the R3-R5 segment?

A)

```

interface FastEthernet0/0
ip address 10.99.53.1 255.255.255.252
ip access-group 122 in
ip ospf network non-broadcast
ip ospf priority 0

router ospf 10
router-id 10.10.3.255
neighbor 10.89.53.2

access-list 122 permit 88 host 10.99.53.2 host 10.99.53.1
access-list 122 deny 88 any any
access-list 122 permit tcp any any
access-list 122 permit udp any any
access-list 122 permit icmp any any

```

B)

```
interface FastEthernet0/0
ip address 10.99.53.1 255.255.255.252
ip access-group 122 in
ip ospf network non-broadcast
ip ospf priority 0
!
router ospf 10
router-id 10.10.3.255
network 10.99.53.0 0.0.0.3 area 0
neighbor 10.99.53.2
!
access-list 122 permit 89 host 10.99.53.2 host 10.99.53.1
access-list 122 deny 89 any any
access-list 122 permit tcp any any
access-list 122 permit udp any any
access-list 122 permit icmp any any
```

C)

```
interface FastEthernet0/0
ip address 10.99.53.1 255.255.255.252
ip access-group 122 in
ip ospf network non-broadcast
ip ospf priority 100
!
router ospf 10
router-id 10.10.3.255
network 10.99.53.0 0.0.0.3 area 0
neighbor 10.99.53.2
!
access-list 122 permit 89 host 10.99.53.2 host 10.99.53.1
access-list 122 deny 89 any any
access-list 122 permit tcp any any
access-list 122 permit udp any any
access-list 122 permit icmp any any
```

D)

```
interface FastEthernet0/0
ip address 10.99.53.1 255.255.255.252
ip access-group 122 in
ip ospf network point-to-point
ip ospf priority 100
!
router ospf 10
router-id 10.10.3.255
network 10.99.53.0 0.0.0.3 area 0
neighbor 10.99.53.2
!
access-list 122 permit 60 host 10.99.53.2 host 10.99.53.1
access-list 122 deny 60 50 any
!
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 410

The network administrator must implement IPv6 in the network to allow only devices that not only have registered IP addresses but are also connecting from assigned locations. Which security feature must be implemented?

- A. IPv6 Snooping
- B. IPv6 Destination Guard

- C. IPv6 Prefix Guard
- D. IPv6 Router Advertisement Guard

Answer: A

Explanation:

Question: 411

What must be configured by the network engineer to circumvent AS_PATH prevention mechanism in IP/VPN Hub and Spoke deployment scenarios?

- A. Use allows in and as-override at all Pes.
- B. Use allows in and as-override at the PE-Hub.
- C. Use Allowas-in the PE_Hub
- D. Use as-override at the PE_Hub

Answer: D

Explanation:

Question: 412

A network engineer must configure a DMVPN network so that a spoke establishes a direct path to another spoke if the two must send traffic to each other. A spoke must send traffic directly to the hub if required. Which configuration meets this requirement?

**At the hub router interface tunneHO ip nhrp nhs multicast
dynamic ip nhrp nhs shortcut tunnel mode gre multipoint**

**On the spokes router interface tunneHO ip nhrp nhs multicast
dynamic ip nhrp nhs redirect tunnel mode gre multipoint**

**:• At the hub router interface tunneHO ip nhrp map multicast
dynamic ip nhrp redirect
tunnel mode gre multipoint**

**On me spokes router interface tunneHO ip nhrp map multicast
dynamic ip nhrp shortcut tunnel mode gre multipoint**

**At the hub router interface tunneHO ip nhrp nhs dynamic
multipoint ip nhrp nhs shortcut tunnel mode gre multicast**

**On the spokes router interface tunneHO ip nhrp nhs multicast
dynamic ip nhrp nhs redirect tunnel mode gre multicast**

ip vrf 1

ip vrf 2

1

int GigabitEthernetO/O no shut

1

int GigabitEthernet0/0.1 encapsulation dot1Q 1 ip vrf forwarding

1

ip address 10.1.2.55.2 55.2 55.2 55.0 i *

int GigabitEthernetO/O. 2 encapsulation dot 10 2 ip vrf

forwarding 2

ip address 10.2.2.1 255.255.255.0

A. Option

B. Option

C. Option

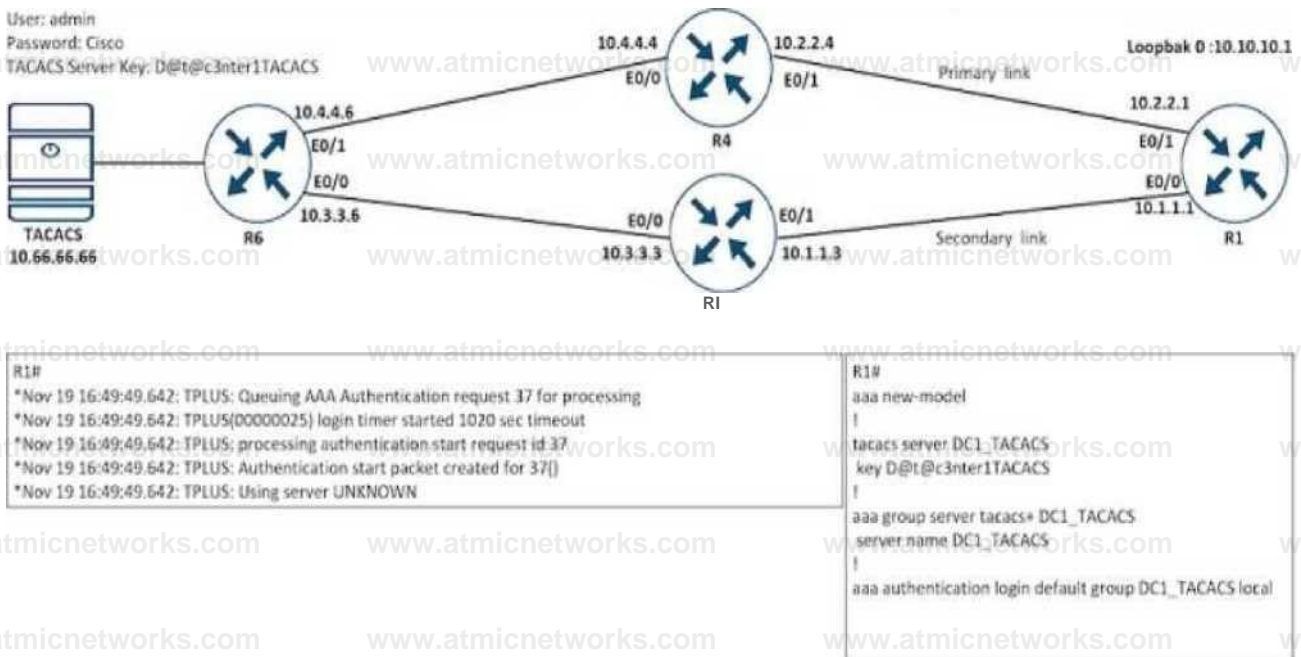
D. Option

Answer: B

Explanation:

Question: 413

Refer to the exhibit.



Refer to the exhibit R1 cannot authenticate via TACACS Which configuration resolves the issue?

a group server tacacs* DC_TACACS server name DC^TACACS

Ucaes server DC1JACACS address [pv410M66.66 key
D@I@c3nter1TACACS

aaa group server «oaes+ DC1_TACACS server name DC_TACACS

**tacacs server DC1_TACACS address Ipv4 10.60.66.66 key D @!@c
3^ er1 TACACS**

A. Option A

B. Option B

C. Option C

D. Option D

Answer:

B

Explanation:

Question:

414

Refer to the exhibit.

```
router ospfv3 1
router-id 10.1.1.1
address-family ipv4 unicast
passive-interface Loopback0
exit-address-family
address-family ipv6 unicast
passive-interface Loopback0
exit-address-family
interface Loopback0
ip address 10.1.1.1 255.255.255.255
ipv6 address 2001:DB8::1/64
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
interface GigabitEthernet2
ip address 10.10.10.1 255.255.255.0
ipv6 enable
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
```

An administrator must configure the router with OSPF for IPv4 and IPv6 networks under a single process. The OSPF adjacencies are not established and did not meet the requirement. Which action resolves the issue?

- A. Replace OSPF process 10 on the interface with OSPF process 1, and configure an additional router ID with IPv6 address.
- B. Replace OSPF process 10 on the interface with OSPF process 1, for the VpV6 address and remove process route ID with IPv6 address.
- C. Replace OSPF process 10 on the interface with OSPF process 1, and remove process 10 from the global configuration.

D. Replace OSPF process 10 on the interface with OSPF process 1 for the IPv4 address, and remove process 10 from the global configuration.

Answer: C

Explanation:

Question: 415

The summary route is not shown in the RouterB routing table after this below configuration on Router_A

```
interface eth^rnet o
description location ID:S4 289T9E09F39
Ip Address 192.16 911 25 5 255J5S.0 ip summary-address eigrp 1 172-16.80
0 255 255 240 0 .
```

Which Router_A configuration resolves the issue by advertising the summary route to Router B?

interface loopback 0

ip address 172.16.96 1 255.255.255.0

Interface Ethernet 0 ipaddress 192,168 3 1 255 255 255 0 ip summary-address eigrp 1 172.16 80.0 255.255 240.0

interface loopback 0

ip address 172,16,81,1 255 255 255 0

interface Ethernet 0

ipaddress 192 168,3 1 255 255255 0

ip summary-address eigrp 1 172 16 80 0 255 255 240,0

interface loopback 0

ip address 172.16791 2 5 5.255.2 55.0

interface Ethernet 0

ipaddress 142.16611 2 5 5.255.2 55 0

ip summary-address eigrp 1 172 16 80 (7^55.255 240.0

interface loopback 0

ip address 172.18.81.1 255 255 255 0

interface Ethernet 0

ip address 192768.3.1 255.255.255 0

ip summary-address eigrp 1 172.16,80,0 255,255-240.0

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 416

How do devices operate in MPLS L3VPN topology?

- A. P and associated PE routers with IGP populate the VRF table in different VPNs.
- B. CE routers connect to the provider network and perform LSP functionality
- C. P routers provide connectivity between PE devices with MPLS switching.
- D. P routers support PE to PE VPN tunnel without LSP functionality

Answer: C

Explanation:

Question: 417

Refer to the exhibit.

```
CPE(config)# lin c 0
CPE(config-line)# no exec
CPE(config-line)# end
CPE#
*Jan 31 23:07:22.655: %SYS-5-CONFIG_I: Configured from console
by console
CPE# wr
Building configuration...
[OK]
CPE# exit

CPE con0 is now available

Press RETURN to get started.

! Console stopped responding at this moment !
```

An administrator is attempting to disable the automatic logout after a period of inactivity. After logging out, the console stopped responding to all keyword inputs. Remote access through SSH still work resolves the issue?

- A. Configure the exec command on line con 0.
- B. Configure the absolute-timeout command on line con 0.
- C. Configure the default exec-timeout command on line con 0.
- D. Configure the no exec-timeout command on line con 0.

Answer: D

Explanation:

Question: 418

Refer to the exhibit.

```
RB#
*Sep 19 00:53:43.002: BGPNSF state: 10.10.10.3 went from nsf_not_active to
nsf_not_active
*Sep 19 00:53:43.006: BGP: 10.10.10.3 went from Established to Idle
*Sep 19 00:53:43.006: BGP: 10.10.10.3 Down User reset
*Sep 19 00:53:43.006: BGP: 10.10.10.3 closing
*Sep 19 00:53:43.106: BGP Router: unhandled major event code 128, minor 0

RD#show ip bgp neighbors 10.10.10.2
BGP neighbor is 10.10.10.2, remote AS 65101, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Active
  Last read 00:01:35, last write 00:01:35, hold time is 180, keepalive
  interval is 60 seconds
  Default minimum time between advertisement runs is 30 seconds
  Address tracking is enabled, the RIB does have a route to 10.10.10.2
  Connections established 11; dropped 11
  Last reset 00:01:36, due to Peer closed the session
  External BGP neighbor may be up to 3 hops away.
  Transport(tcp) path-mtu-discovery is enabled
  No active TCP connection
```

A NOC team receives a ticket that data traffic from RA to RF is not forwarded when the link between the RC-RE path goes down. All routers learn loopback IP through the IGP protocol. Which configuration resolves?

- A. RD(config)#router bgp B5201RD(config-router)# neighbor 10.10.10.2 update-source loopback 0
- B. RD(config-router)# neighbor bgp 65101RB(config-router)# neighbor 10.10.10.3 ebgp-multihop 3
- C. RB(config)# router bgp 65101RB(config)#neighbor 10.10.10.3 update-source loopback 0
- D. RD(config)# router bgp 65201RD(config-router)# neighbor 10.10.10.2 ebgp-multihop 3

Answer: B

Explanation:

Question: 419

Which feature is used by LDP in the forwarding path within the MPLS cloud?

A. IP forwarding

B. TTL

C. TDP

D. LSP

Answer: D

Explanation:

Question: 420

Refer to the exhibit. An engineer is trying to log in to R1 via R3 loopback address. Which action resolves the issue?

A. Add transport input SCP

B. Add transport input none

C. Remove the IPv6 traffic filter from R1, which is blocking the Telnet.

D. Remove the IPv6 traffic from R1, which is blocking the SSH

Answer: C

Explanation:

Question: 421

Refer to the exhibit.

```
ipv6 inspect udp idle-time 3600
ipv6 inspect name ipv6-firewall tcp
ipv6 inspect name ipv6-firewall udp
!
ipv6 access-list ipv6-internet
deny ipv6 any FEC0::/10
deny ipv6 any FF00::/8
permit ipv6 any FF02::/16
permit ipv6 any FF0E::/16
permit udp any any eq domain log
!
Interface gi0/1
ipv6 traffic-filter ipv6-internet in
ipv6 inspect ipv6-firewall in
ipv6 inspect ipv6-firewall out
```

A network administrator configured name resolution for IPv6 traffic to be allowed through an inbound access list. After the access list is applied to resolve the issue, name resolution still did not work. Which action does the network administrator take to resolve the name resolution problem?

A. Remove `ipv6 inspect ipv6-firewall in` from interface `gi0/1`

B. Add permit udp any eq domain any log in the access list.

C. inspect ipv6 inspect name ipv6-firewall udp 53 in global config.

D. Add permit any eq domain 53 any log in the access list.

Answer:

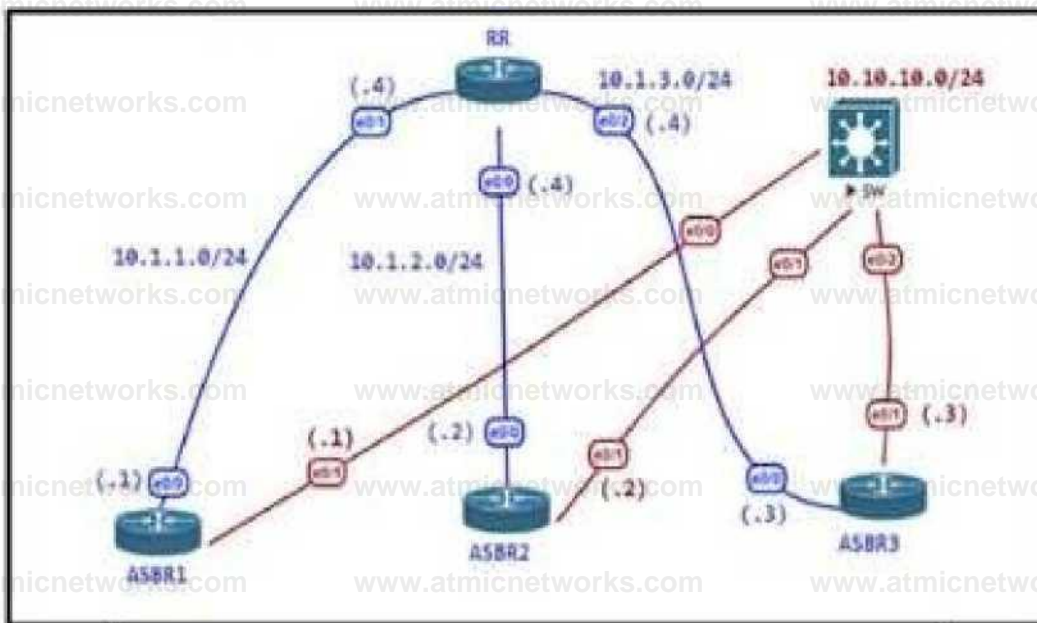
A

Explanation:

Question:

422

Exhibits:



RR

```

router bgp 100
  neighbor 10.1.1.1 remote-as 100
  neighbor 10.1.2.2 remote-as 100
  neighbor 10.1.3.3 remote-as 100

```

ASBR2

```

router bgp 100
  neighbor 10.1.1.4 remote-as 100

```

ASBR2

```
router bgp 100
  neighbor 10.1.1.4 remote-as 100
```

ASBR3

```
router bgp 100
  neighbor 10.1.2.4 remote-as 100
```

ASBR4

```
router bgp 100
  neighbor 10.1.3.4 remote-as 100
```

Refer to the exhibit The administrator configured the network devices for end-to-end reachability, but the ASBRs are not propagating routes to each other Which set of configurations resolves this issue?

```
router bgp 100
neighbor 10.1.1.1 route-reflector-client
neighbor 10.1 2.2 route-reflector-client
neighbor 10.1.3.3 route-reflector-cbent
```

```
router bgp 100
neighbor 10.1 1 1 update-source loopbackO
neighbor 10.1.2.2 update-source LoopbackO
neighbor 10.1.3.3 update-source LoopbackO
```

```
router bgp 100
neighbor 10.111 next-hop-self
neighbor 10.1.2.2 next-hop-self
neighbor 10 1 3.3 next-hop-self
```

```
router bgp 100
neighbor 10.1.1.1 ebgp-multihop
neighbor 10.1.2.2 ebgp-multihop
neighbor 10 13 3 ebgp-multihop
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

Question: 423

Refer to the exhibit.

```
interface GigabitEthernet2
```

```
no ip address
ip helper-address 192.168.255.3
no shutdown

interface GigabitEthernet2.10
encapsulation dot1q 210
ip address 192.168.210.1 255.255.255.0
ip ospf 1 area 0
no shutdown
```

With the partial configuration of a router-on-a-stick. Clients in VLAN 10 on Gi2 cannot obtain IP configuration from the central DHCP server is reachable by a successful ping from the router. Which action resolves the issue?

- A. Configure the ip ip dhcp pool f and network 192.168.210.0.255.255/0 commands.
- B. Configure the ip helper-address 192.168.255.3 command on the Gi2 10 subinterface.
- C. Configure a valid IP address on the Gi2 interface so that DHCP requests can be forwarded.
- D. Configure the ip dhcp excluded-address 192.168.255.3 command on the Gi1.10 subinterface.

Answer: B

Explanation:

Question: 424

Exhibit.

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
```

```
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 307200
```

```
Routing Descriptor Blocks:
```

```
10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
```

```
Composite metric is (307200/281600), Route is Internal
```

```
Vector metric:
```

```
Minimum bandwidth is 10000 Kbit
```

```
Total delay is 2000 microseconds
```

```
Reliability is 255/255
```

```
Load is 1/255
```

```
Minimum MTU is 1500
```

```
Hop count is 1
```

```
10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
```

```
Composite metric is (312320/286720), Route is Internal
```

```
Vector metric:
```

```
Minimum bandwidth is 10000 Kbit
```


```
Total delay is 2200 microseconds
```

```
Reliability is 255/255
```

```
Load is 1/255
```

```
Minimum MTU is 1500
```

```
Hop count is 1
```



Refer to the exhibit. A network is configured for EIGRP equal-cost load balancing, but the traffic destined to the servers is not load balanced. Link metrics from router R2 to R3 and R4 are the same. Which delay value must be configured to resolve the issue?

A. 208 on R3 E0/0

B. 120 on R4 E0/1

C. 120 on R3 E0/1

D. 2200 on R4 E0/1

Answer: C

Explanation:

Question: 425

A network administrator successfully established a DMVPN tunnel with one hub and two spokes using EIGRP. One of the requirements was to enable spoke-to-spoke tunnels through the hub router using EIGRP. Which configuration command must the engineer configure to meet the requirement?

A. no ip eigrp 1 mode multipoint

B. no ip eigrp 1 split-horizon

C. no ip eigrp 1 tunnel-redirect


D. no ip eigrp 1 mode mgre

Answer: B

Explanation:

Question: 426

Refer to the exhibit.



Spoke routers do not learn about each other's routes in the DMVPN Phase2 network. Which action resolves the issue?

- A. Remove default route from spoke routers to establish a spoke-to-spoke tunnel.
- B. Configure a static route in each spoke to establish a spoke-to-spoke tunnel.
- C. Rectify incorrect wildcard mask configured on the hub router network command.
- D. Disable EIGRP split horizon on the Tunnel0 interface of the hub router.

Answer: D

Explanation:

Question: 427

Refer to the exhibit.

```
RI#  
router ospf 1  
redistribute rip subnets network 131.108.1.0  
0.0.0.255 area 2 network 131.108.2.0 0.0.0.255  
area 2 distribute-list 1 out i  
access-list 1 permit 132.108.4.0 0.0.0.255
```

The R1 OSPF neighbor is not receiving type 5 external LSAs for 132.108.2.0/24 and 132.108.3.0/24 networks.
Which configuration command resolves the issue?

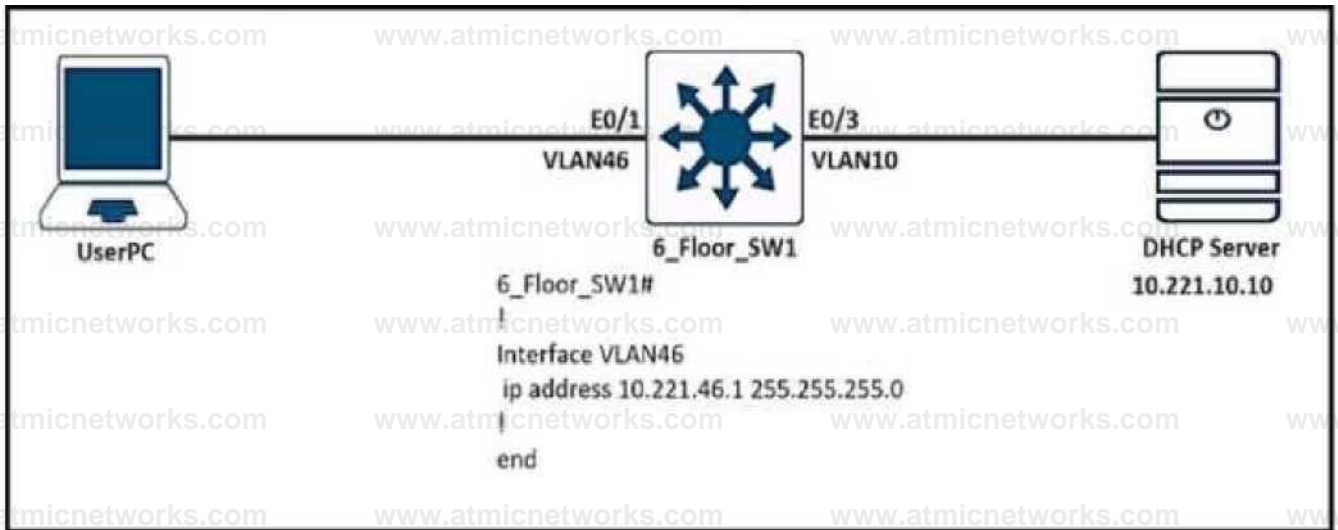
- A. access-list 1 permit 132.108.0.0 0.0.1.255
- B. access-list 1 permit 132.108.0.0 0.0.3.255
- C. access-list 1 permit 132.108.2.0 0.0.0.255
- D. access-list 1 permit 132.108.4.0 0.0.3.255

Answer: B

Explanation:

Question: 428

Refer to the exhibit.



Users in VLAN46 cannot get the IP from the DHCP server. Assume that all the parameters are configured properly in VLAN 10 and on the DHCP server. Which command on interface VLAN46 allows users to receive IP from the DHCP server?

- A. ip dhcp-address 10.221.10.10
- B. ip dhcp server 10.221.10.10
- C. ip helper-address 10.221.10.10
- D. ip dhcp relay information trust-all

Answer: C

Explanation:

Question: 429

Refer to the exhibit.

```
!
summary-address 10.1.0.0 255.255.0.0
!
```

The non area 0 routers in OSPF still receive more specific routes of 10.1.1.0.10.1.2.0.10.1.3.0 from area 0.
Which action resolves the issue?

- A. Configure route summarization on OSPF-enabled interfaces.
- B. Summarize by using the summary-address 10.1.0.0 255.255.252.0 command.
- C. Summarize by using the area range command on ABRs
- D. Configure the summary-address 10.1.0.0 255.255.252.0 command under OSPF process.

Answer: C

Explanation:

Question: 430

Refer to the exhibit.

```
flow exporter EXPORTER-1
 destination 172.16.10.2
 export-protocol netflow-v9
 transport udp 90
 exit
!
flow record v4_r1
 match ipv4 tos
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 collect counter bytes long
 collect counter packets long
!
flow monitor FLOW-MONITOR-1
 record v4_r1
 ip col
!
interface GigabitEthernet 0/0/0
 ip address 172.16.6.2 255.255.255.0
 ip flow monitor FLOW-MONITOR-1 input
!
```

An engineer configured NetFlow to capture traffic information through the router, but it iOS not working as expected. Which action captures the flow information from this router to the collector?

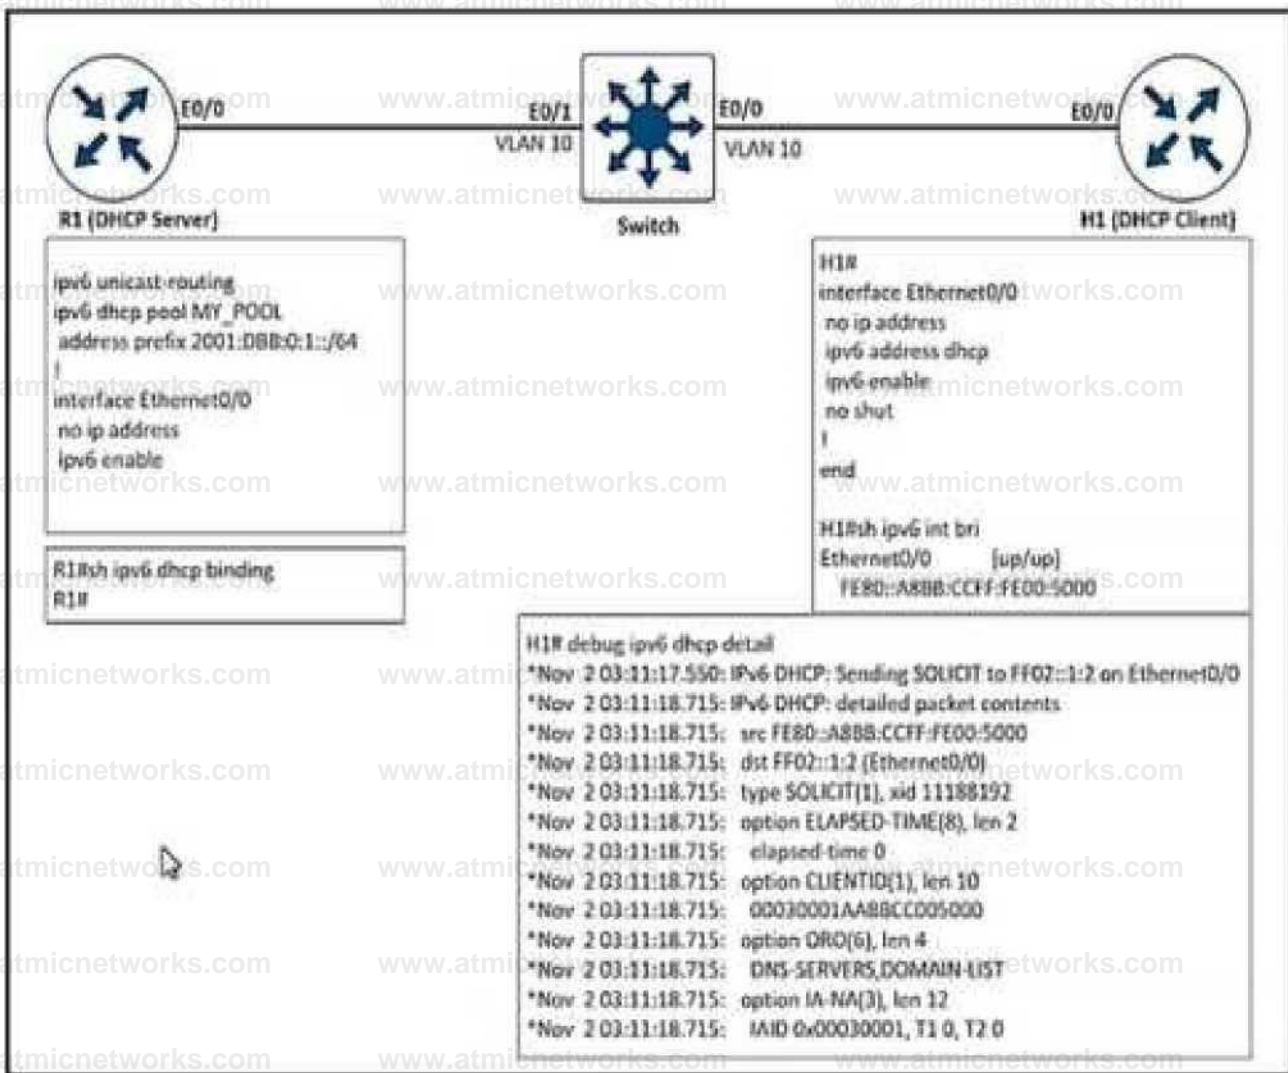
- A. Change the interface configuration FLOW-MONITOR-1 from input to output.
- B. Configure a flow exporter under flow FLOW-MONITOR-1.
- C. Configure more than one flow exporter destination addresses.
- D. Change the flow exporter transport protocol from UDP to TCP

Answer: B

Explanation:

Question: 431

Refer to the exhibit.



After the network administrator rebuilds the IPv6 DHCP server, clients are not getting the IPv6 address lease. Which action resolves the issue?

- A. Remove FE80 A8BB CCFF FE00 5000 assigned by the IPv6 DHCP server.
- B. Add IPv6 dhcp server MY_POOL under the interface ethernet 0/0 on H1.
- C. Add IPv6 dhcp server MY_POOL under the interface ethernet 0/0 on R1.
- D. Configure FF02::1:2 to discover all IPv6 DHCP clients.

Answer: C

Explanation:

Question: 432

Refer to the exhibit.

```
c^flttMMNppoicyl jirirv^kitai poAm MtBMpipHS career JJlfttOCOfl I
Cf>pK> ipwc lrwKlixmwt tram? np-OM rir.rr.15 TOWC pMatawt
|
crppio vMC praixv vppprat
»«lr.wi^tn.ie» Irani.' (
rarfiMn iirwo
UnMin tDU
9 UMM 10 0 01 MS MS MJ 0
10 MU MOO
^nw mnwUataiwMiif
ip rituji mip ra#K 311' rename
ijlllfj l^mjJJ
rpnhrp twlltxne 300
to ip ipu msnron v^p 1
ipip XVJM nmi INO
MW (000
tunnel tourte &gjMMnc*r»rt MW
tuiofr riuje M BWftOW
Curite* tty fDHW
■unM pratpdw) (p*ec pote Mpnproi i
ritftKfl F#KMOUM
10UMM lftVDtJM 2 55 755 0
|
■MrtKU failtirwttJIXI
9 UMM m IM 0 1 N57M2M0
i
Min <M I
MteU{ 1000 0004 MS
Mrt*9tlRMatlDI1.0J02K
1
```

A network administrator must configure DMVPN tunnels between the hub and spoke with dynamic spoke-to-spoke tunnel capabilities using EIGRP. Which tunnel interface command must the network administrator configure to establish an EIGRP peer?

- A. no ip next-hop-self eigrp 1
- B. ip next-hop-self eigrp 1
- C. no ip nhrp next-hop-self
- D. ip nhrp next-hop-self

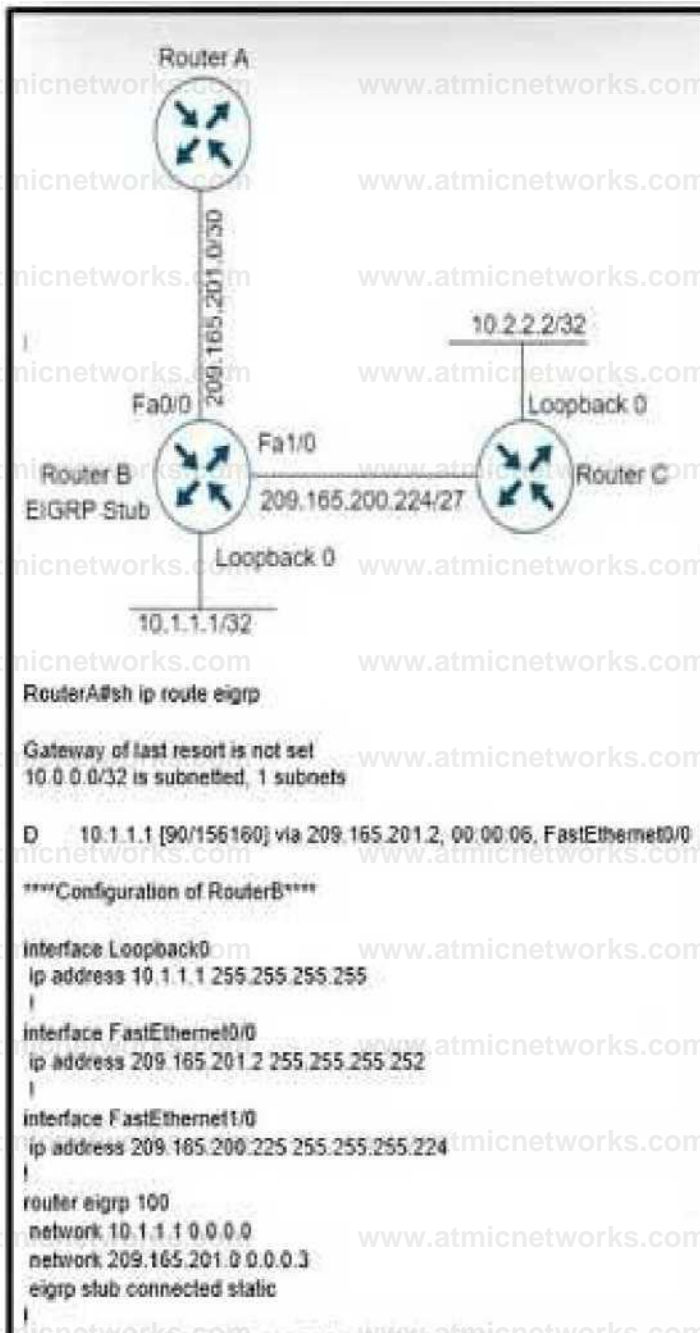
Answer:

C

Explanation:

Question: 433

Refer to the exhibit.



```

HttfeelMpMtf
to JMW W111 »$»1»5»$
^Ar f MfttWMKn
ip***TM M lts561 J2MJ55HM!a
HU:rKftFrEtMmtl*
ip MMM 2W tu 2M Wi TM W » 2?4 i
iMff r^p IM
M*M<1 'a.' I I I V WWW
MfiMt M lts 201 00001
e^rp alio uneeded sUte

toI** 10 2 2 2 255 2K2K2K2M IK200.221

```

Refer to the exhibit. Not all connected and static routes of router B are received by router A even though EIGRP neighborhood is established between the routers. Which configuration resolves the issue?

A)

```

router eigrp 100
network 209.165.200.224 0.0.07
redistribute static metric 1000 1 255 1 1500 eigrp stub
connected

```

B)

```

router eigrp 100
network 209.165.200.224 0.0 0.7

```

C)

```

router eigrp 100
network 209,165.200.224 0.0.0.31
redistribute static metric 1000 1 255 1 1500

```

D)

```

router eigrp 100
network 209,105-200-224 0,0.0,7 redistribute static
metric 1000 1 255 1 1500
eigrp stub static

```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 434

Which router attaches the VPN label to incoming packets from CE routing?

A. CE router

B. core router

C. P router

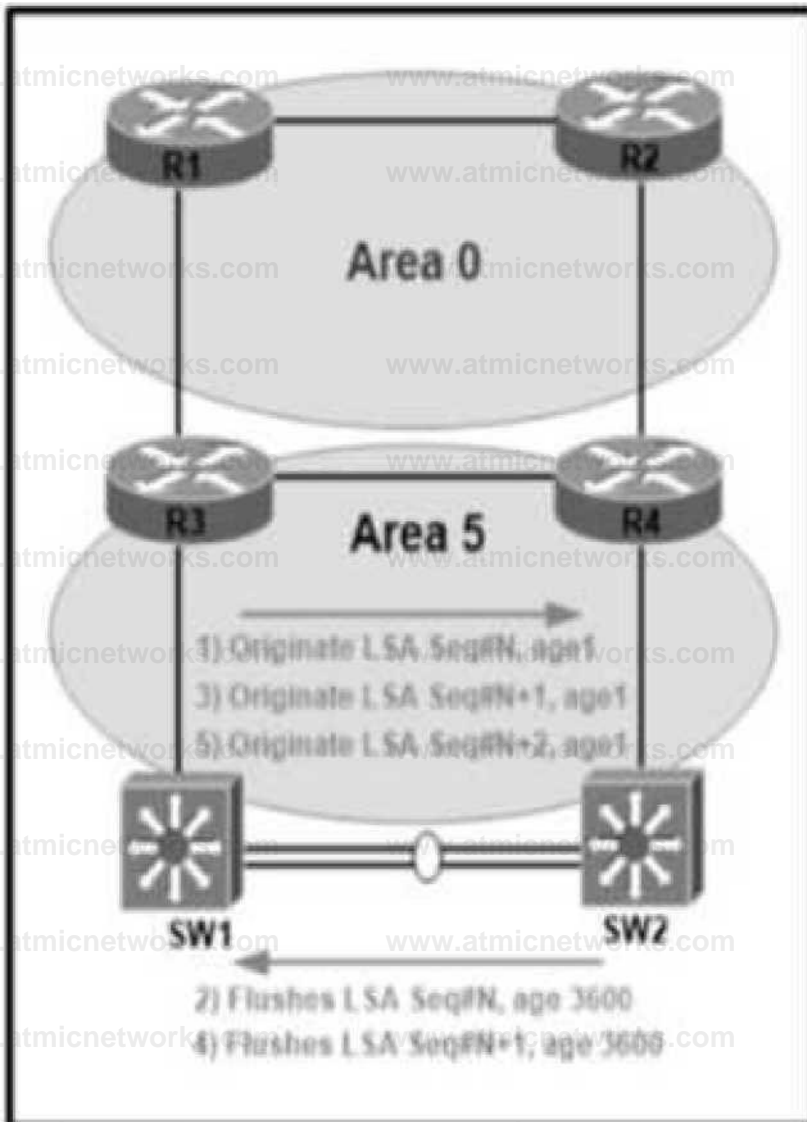
D. PE router

Answer: D

Explanation:

Question: 435

Refer to the exhibit.



An error message "an OSPF-4-FLOOD_WARNING" 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. Change area 5 to a normal area from a nonstub area
- B. Resolve different subnet mask issue on the link

C. Configure Layer 3 port channel on interfaces between switches

D. Resolve duplicate IP address issue in the network

Answer: D

Explanation:

Question: 436

Refer to the exhibit.

```
R2#sh 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 not set

C    192.168.10.0/24 is directly connected, Serial1/0
C    172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
C    172.16.160.0/19 is directly connected, Loopback1
C    172.16.128.0/19 is directly connected, Loopback0
C    172.16.224.0/19 is directly connected, Loopback3
C    172.16.192.0/19 is directly connected, Loopback2
D    172.16.0.0/16 is a summary, 00:01:27, Null0
```

An engineer must configure EIGRP between R1 and R2 with no summary route. Which configuration resolves the issue?

A)

R1 (config)#router eigrp 1
R1 (config-router)#no auto-summary

B)

R2 (config)#router eigrp 1
R2 (config-router)#no auto-summary

C)

R2 (config)#router eigrp 1
R2 (config-router)#auto-summary

D)

R1 (config)#router eigrp 1
R1 (config-router)#auto-summary

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 437

Refer to the exhibit.

R2# show ip ospf neighbor

82#

R2# debug ip ospf hello

*Feb 22 23:46:58.699: OSPF-1 HELLO Et1/1: Rev hello from 10.255.255.1
area 0 10.0.23.1

◆Feb 22 23:46:58.703: OSPF-1 HELLO Et1/1; Mismatched hello
parameters from 10.0.23.1

‘Feb 22 23:46:58.703: OSPF-1 HELLO Et1/1: Dead R 30 C 20, Hello R 10
C 10 Mask R 255.255.255.0 C 255.255.255.0

The connected routers do not show up as OSPF neighbors. Which action resolves the issue?

- A. Change the R1 dead timer to 20.
- B. Change the R2 dead timer to 20.
- C. Change the R2 hello timer to 20.
- D. Change the R1 hello timer to 20.

Answer: A

Explanation:

Question: 438

Refer to the exhibit.

```
ip prefix-list 1 permit 172.16.0.0/16
ip prefix-list 2 permit 192.168.2.0/24
!
route-map RED permit 10
match ip address prefix-list 1
set ip next hop 10.1.1.1
continue 20
exit
!
route-map RED permit 20
match ip address prefix-list 2
set ip next hop 10.2.2.2
end
```

The forwarding entries show that the next hop for prefixes from the 172.16.0.0/16 network is set to 10.2.2.2 instead of 10.1.1.1. Which action resolves the issue?

- A. Add set ip next hop 10.1.1.1 in route-map RED permit 20.
- B. Add the continue statement in route-map RED permit 10 instead of continue 20.
- C. Remove match ip address prefix-list 1 from route-map RED permit 10.
- D. Remove the continue 20 statement from route-map RED permit 10

Answer: D

Explanation:

Question: 439

Refer to the exhibit.

```
CPE# show ip route static
```

```
<output omitted>
```

```
S* 0.0.0.0/0 is directly connected, DialerO
```

```
S 198.51.100.0/24 [1/0] via 192.168.1.1
```

```
S 203.0.113.0/24 [1/0] via 192.168.2.1
```

```
CPE# show run | section router ospf
```

```
router ospf 1
```

```
redistribute static subnets
```

```
CPE# show ip ospf database | begin Type-5
```

Type-5 AS External Link States

```
Link ID ADV Router Age Seq# Checksum Tag 198.51.100.0
```

```
192.168.0.1 14 0x80000001 0x0007D0 0
```

```
203.0.113.0 192.168.0.1 14 0x80000001 0x009C5C 0
```

Refer to the exhibit. The default route is not advertised to the neighboring router. Which action resolves the issue?

A. Configure the redistribute static metric 200 subnets command under OSPF.

B. Configure OSPF on the Dialer0 interface.

C. Configure the network 0.0.0.0 255.255.255.255 area 0 command under OSPF.

D. Configure the default-information originate command under OSPF.

Answer:

D

Explanation:

Question:

440

Refer to the exhibit.



S1

Sloping 10.0.0.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echoes to 10.0.0.1, timeout is 2 seconds: i!!!H

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

Sifttelnet 10.0.0.1

Trying 10.0.0.1... Open

(Connection to 10.0.0.1 closed by foreign host)

R3

R3# hostname R3

enable password cisco

no aaa new-model

username admin password 0 cisco

interface Ethernet0/1

ip address 10.0.0.1 255.255.255.252 | line

con 0

logging synchronous

line aux 0

line vty 0 4

password cisco

login no

exec

transport input all

end

Refer to the exhibit. A network engineer cannot remote access R3 using Telnet from switch S1. Which action resolves the issue?

- A. Allow the inbound connection via the exec command on R3.
- B. Add the transport input telnet command on R3.
- C. Allow to use the ssh -l admin 10.0.0.1 command on the switch.
- D. Add the login admin command on the switch.

Answer: A

Explanation:

Question: 441

Refer to the exhibit.

```
R1#show ip rip database
10.0.0.0/8  auto-summary
10.1.1.0/24  directly connected, GigabitEthernet0/0
10.1.3.0/24
    [2] via 10.1.12.2, 00:00:03, GigabitEthernet1/0
10.1.12.0/24  directly connected, GigabitEthernet1/0
10.1.23.0/24
    [1] via 10.1.12.2, 00:00:03, GigabitEthernet1/0
```

Refer to the exhibit. A customer reports that networks in the 10.0.1.0/24 space do not appear in the RIP database.

What action resolves the issue?

- A. Remove summarization of 10.0.0.0/8.
- B. Permit 10.0.1.0/24 address in the ACL.
- C. Remove ACL on R1 blocking 10.0.1.0/24 network.
- D. Configure 10.0.1.0/24 network under RIP.

Answer: A

Explanation:

Question: 442

Refer to the exhibit.

100 JO 07J71\$ subnet led. 3 subnets

C 100.111.111.111 is directly connected.

D 100.2.2.2 (90/156160) via 10.1.1.2, 00.00.46. FastEthernet0/0

U 100.3.3.3 (190/150720) via 10.1.1.14, 00.00.44, FastEthernet0/0/8 is variably subnetted, 13 subnets, 4 masks

O 10.1.1.8/30 [90/30720] via 10.1.1.14, 00.00.44, FastEthernet0/0

C 10.1.12/30 is directly connected, FastEthernet1/0

C 10.1.1.1/30 is directly connected, FastEthernet0/0

D 10.1.4/30 (190/150720) via 10.1.1.14, 00.00.44, FastEthernet0/0

C 10.100.140/32 is directly connected, Loopback0

D EX 10.11.80/29 [170/33280] via 10.1.1.14, 00.00.44, FastEthernet1/0 (170/33280) via 10.1.12.0, 0.00.45, FastEthernet0/0

C 10.100.150/32 is directly connected, Loopback1

C 10.100.110/32 is directly connected, Loopback0

S 10.10.1.1 (M24 is a summary 00:00:48. Null)

C 10.100.1.30/32 is directly connected, Loopback30

C 10.100.120/32 is directly connected, Loopback30

C 10.200.1.0/24 is directly connected, FastEthernet0/0

DEX 10.24.24.24/24 [170/2174976] via 10.1.1.2, 00.00:46, FastEthernet0/0

Refer to the exhibit. R1 must advertise all loopback interface IP addresses to neighbors, but EIGRP neighbors receive a summary route. Which action resolves the issue?

- A. Redistribute connected routes into EIGRP
- B. EIGRP on loopback Interfaces.
- C. Disable auto summarization on R1.
- D. Remove the 10.100.1.0/24 static route.

Answer: D

Question: 443

Refer to the exhibit.



R1

```

service timestamps debug datetime msec
service timestamps log datetime msec
1 clock timezone EET 2 0
J end

```

R1#show clock

```

•23:50:13.297 EET Sat Nov 14 2020

```

R1#

```

•Nov 14 21:49:59.607: IP: s=10.1.1.1 (local), d=224.0.0.5 (EthernetO/O), len 80, local feature. Logical MN local(14), rtype 0, torus FALSE, sendself FALSE, mtu 0 fwdchk FALSE
•Nov 14 21:49:59.607: IP: s=10.1.1.1 (local), d=224.0.0.5 (EthernetO/O), len 80, sending broad/muhicast
•Nov 14 21:50:00.336: IP: s=10.23.4 (EthernetO/I), d=224.0.0.5, len 80, rcvd 0
•Nov 14 21:50:00.336: IP: s=10.2.2.4 (EthernetO/I), d=224.0.0.5, len 80, input feature, packet consumed. MCI Check(IOI).
rtyp 0, torus FALSE, sendsel FALSE, mtu 0, fwdchk FALSE

```

Refer to the exhibit. An engineer cannot determine the time of the problem on R1 due to a mismatch between the router local clock and logs. Which command synchronizes the time between new log entries and the local clock on R1?

- A. service timestamps debug datetime msec show-timezone
- B. service timestamps log datetime localdatetime msec
- C. service timestamps datebug datetime localtime msec
- D. service timestamps log datetime msec show-timezone

Answer: B

Question: 444

Refer to the exhibit.

The screenshot shows the Cisco DNA Assurance Center interface. The top navigation bar includes 'DESIGN', 'POLICY', 'PROVISION', and 'ASSURANCE'. The main heading is 'OSPF Adjacency Failed on Device "10.30.255.101" Interface TenGigabitEthernet1/0/23 with Neighbor 10.30.255.2'. Below this, the status is 'Open'. A green checkmark indicates a successful action: 'Check OSPF neighbors show ip ospf neighbor'. A table shows the neighbor details:

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.30.255.2	1	EXSTART/RCR	00:00:36	10.30.255.1	TenGigabitEthernet1/0/23

Below the table, there is a note: 'If the Neighbor is in "Init" state, Check if there is authentication configured using "show run | sec OSPF". Authentication type and keys should match on both routers.' A 'Run' button is visible at the bottom right.

Refer to the exhibit. An engineer is investigating an OSPF issue reported by the Cisco DNA Assurance Center. Which action resolves the issue?

- A. One of the neighbor links is down Bring the interface up by running shut and no shut
- B. One of the interfaces is using the wrong MTU Match interface MTU on both links
- C. An ACL entry blocking multicast on the interfaces Allow multicast through the interface ACL
- D. One of the interfaces is using the wrong authentication Match interface authentication on both links

Answer: B

Explanation:

Question: 445

What action is performed for untagged outgoing labels in an MPLS router?

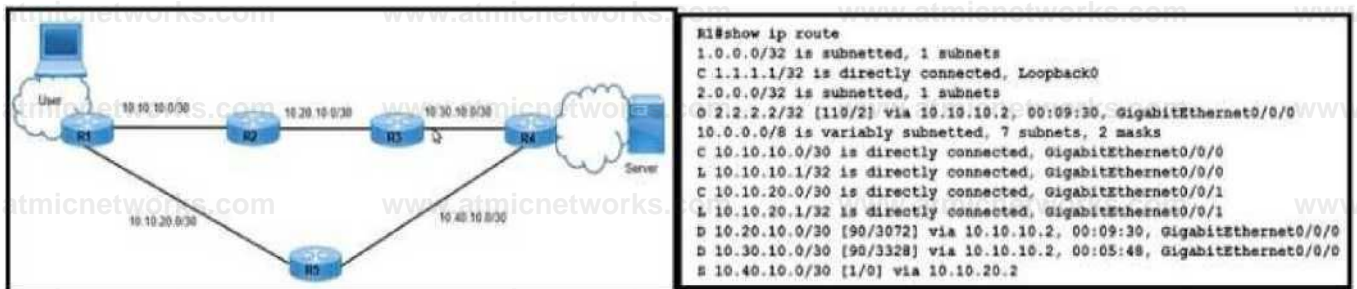
- A. Convert the incoming MPLS packet to an untagged packet and then do a FIB lookup
- B. Convert the incoming MPLS packet to an untagged packet and then do a RIB lookup.
- C. Convert the untagged packet to a labeled packet and forward it to the next router
- D. Convert the incoming MPLS packet to an IP packet and forward it to the next router.

Answer: C

Explanation:

Question: 446

Refer to the exhibit.



Routers R1, R2, R3, and R4 use EIGRP. However, traffic always prefers R1 to R5 backup links in nonfailure scenarios.

Which configuration resolves the issue?

- A)
R1(conf g)*no ip route 10.40.10.0 255.255.255.252 10.10.20.2
R1(config)#ip route 0.0.0.0 0.0.0.0 10.10.10.2

B) micnetworks.com

R1 (config) Mnt glgabl Ethernet 0/0/0

Rt(conf g-if ^bandwidth 10000000

C)

R1(config) sno ip route 10.40.10.0 255.255.255.252 10.10.20.2

RI(COntg>wip route 10.40.10.0 255.255.255.252 10.10.20.2 115

D)

R1 (config) si nt giga bi (Ethernet 0/0/0

R1(config-i1 ^bandwidth 10000

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 447

Refer to the exhibit.

```
R1#show ip route ospf
```

```
10.0.0.0/24 is subnetted, 7 subnets
```

```
O E1 10.4.9.0 [110/200] via 10.4.17.6, 00:06:43,  
FastEthernet0/0
```

```
O IA 10.4.27.0 [110/2] via 10.4.15.5, 00:06:44,  
FastEthernet0/1
```

```
O E1 10.4.49.0 [110/200] via 10.4.17.6, 00:06:43,  
FastEthernet0/0
```

```
O E1 10.4.59.0 [110/200] via 10.4.17.6, 00:06:43,  
FastEthernet0/0
```

Refer to the exhibit. An engineer configured two ASBRs, 10.4.17.6 and 10.4.15.5, in an OSPF network to redistribute identical routes from BGP. However, only prefixes from 10.4.17.6 are installed into the routing table on R1. Which action must the engineer take to achieve load sharing for the BGP- originated prefixes?

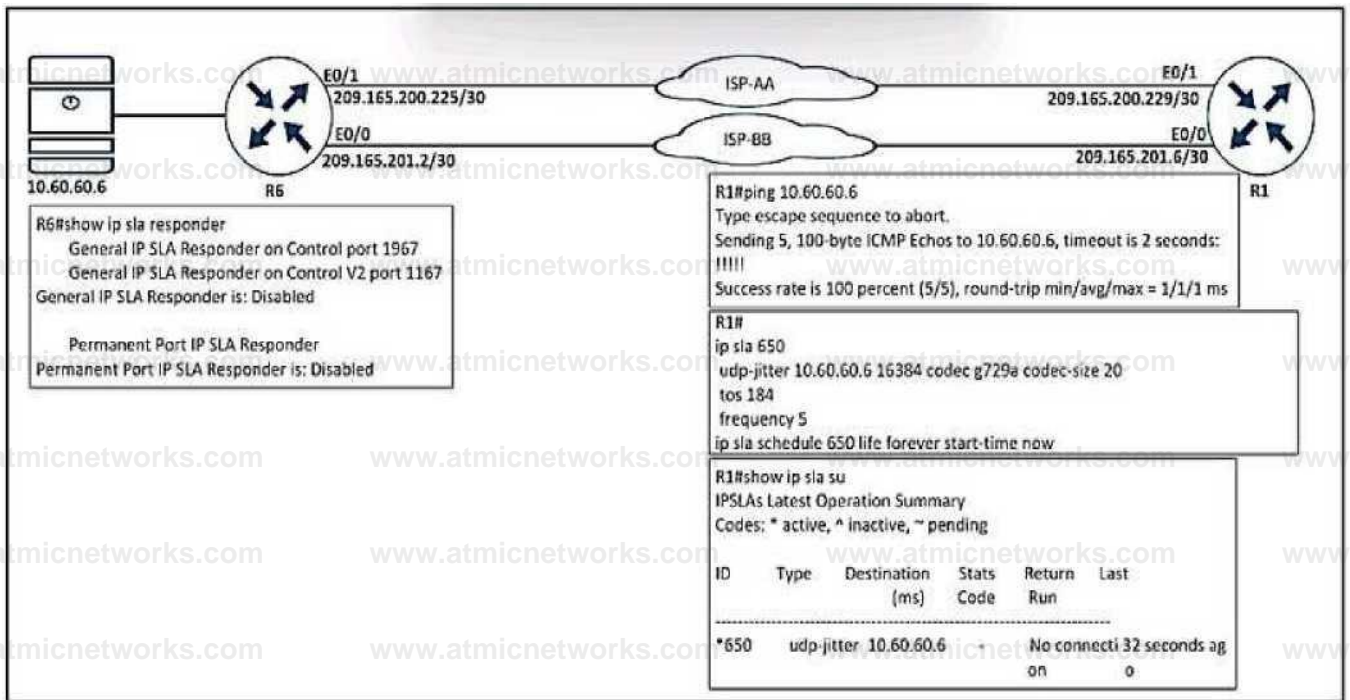
- A. The ASBRs are advertising the redistributed prefixes with the iBGP metric and must be modified to Type 1 on ASBR 10.4.17.6.
- B. The ASBRs are advertising the redistributed prefixes with a different admin distance and must be changed to 110 on ASBR 10.4.15.5.
- C. The admin distance of the prefixes must be adjusted to 20 on ASBR 10.4.15.5 to advertise prefixes to R1 identically from both ASBRs.
- D. The ASBRs are advertising the redistributed prefixes as Type 1 and must be modified to Type 2

Answer: D

Explanation:

Question: 448

Refer to the exhibit.



Refer to the exhibit. Which configuration resolves the IP SLA issue from R1 to the server?

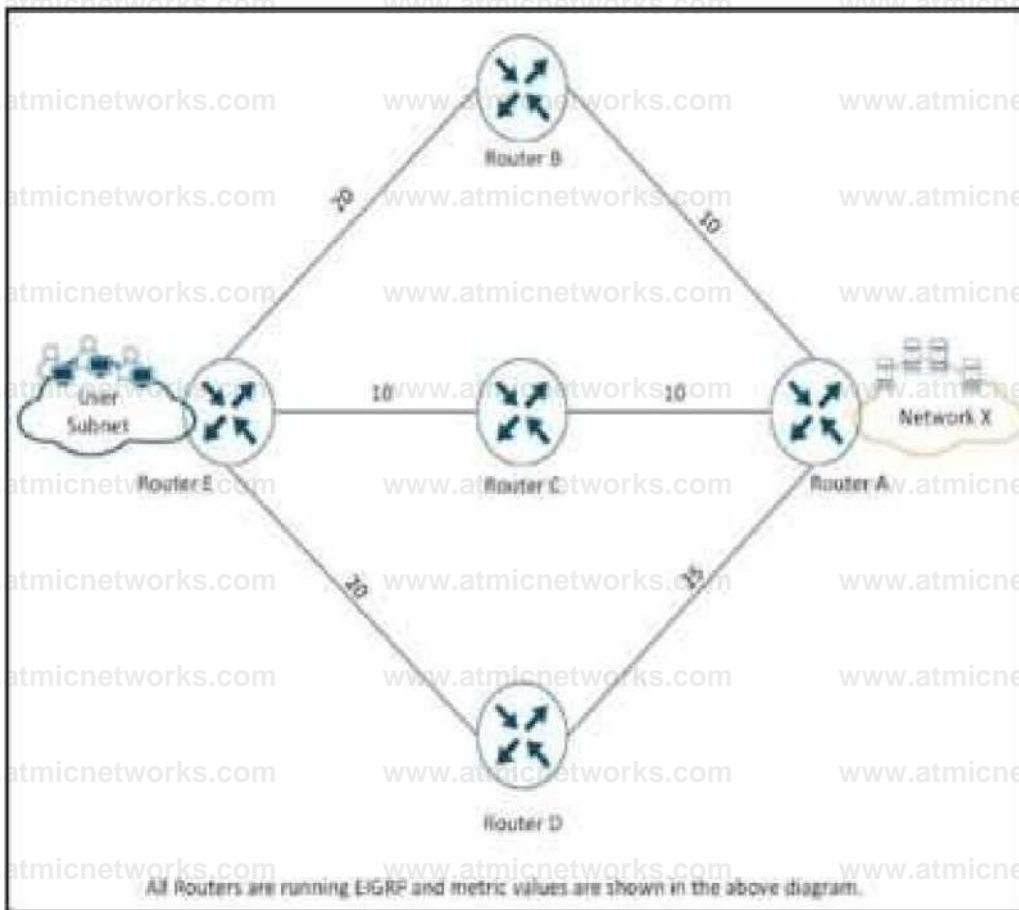
- A. R6(config)#ip sla responder
- B. R6(config)#ip sla responder udp-echo ipaddress 10.60.60.6 po 5000
- C. R6(config)#ip sla 650 R6(config-ip-sla)ff udp-jitter 10.60.60.6
- D. R6(config)#ip sla schedule 10 life forever start-time now

Answer: A

Explanation:

Question: 449

Refer to the exhibit.



Refer to the exhibit. The IT manager received reports from users about slow application through network x. which action resolves the issue?

- A. Use the variance 2 command to enable load balancing.
- B. Increase the bandwidth from the service provider.
- C. Move the servers into the users subnet.
- D. Upgrade the IOS on router E.

Answer: A

Explanation:

Question: 450

Refer to the exhibit.

U AJX11UL AIW1 V^ VL Wll*

UG WUHCXIV UV/ ^f

changed state to up

%LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/2, changed state to up

%LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/3, changed state to up

%OSPF-5-ADJCHG: Process 1, Nbr 10.1.1.2 on Ethernet0/0 from LOADING to FULL, Loading Done

%BGP-3-NOTIFICATION: received from neighbor 192.168.200.1 active 6/7 (Connection Collision Resolution) 0 bytes %BGP-5-NBR_RESET: Neighbor 192.168.200.1 active reset (BGP Notification received)

%BGP-5-ADJCHANGE: neighbor 192.168.200.1 active Down BGP Notification received

%BGP_SESSION-5-ADJCHANGE: neighbor 192.168.200.1 IPv4 Unicast topology base removed from session BGP Notification received

Refer to the exhibit. An engineer noticed that the router log messages do not have any information about when the event occurred. Which action should the engineer take when enabling service time stamps to improve the logging functionality at a granular level?

- A. Configure the debug uptime option
- B. Configure the msec option
- C. Configure the timezone option
- D. Configure the tog uptime option

Answer: D

Explanation:

Question: 451

Refer to the exhibit.

```
router ospfv3 1
router-id 10.1.1.1
address-family ipv4 unicast
passive-interface Loopback0
exit-address-family
address-family ipv6 unicast
passive-interface Loopback0
exit-address-family
interface Loopback0
ip address 10.1.1.1 255.255.255.255
ipv6 address 2001:DB8::1/64
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
interface GigabitEthernet2
ip address 10.10.10.1 255.255.255.0
ipv6 enable
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
```

An engineer noticed that the router log messages do not have any information about when the event occurred. Which action should the engineer take when enabling service time stamps to improve the logging functionality at a granular level?

- A. Replace OSPF process 10 on the interfaces with OSPF process 1 and configure an additional router IO with IPv6 address
- B. Replace OSPF process 10 on the interfaces with OSPF process 1. and remove process 10 from the global configuration
- C. Replace OSPF process 10 on the interfaces with OSPF process 1 for the IPv6 address and remove process 10 from the global configuration
- D. Replace OSPF process 10 on the interfaces with OSPF process 1 for the IPv4 address and remove process 10 from the global configuration

Answer: D

Explanation:

Question: 452

How is a preshared key "Test" for all the remote VPN routers configured In a DMVPN using GRE over IPsec set up?

- A. authentication pre-share Test address 0.0.0.0 0.0.0.0
- B. set pre-share Test address 0.0.0.0 0.0.0.0
- C. crypto Ipsec key Test address 0.0.0.0 0.0.0.0
- D. crypto isakmp key Test address 0.0.0.0 0.0.0.0

Answer: D

Explanation:

Question: 453

Refer to the exhibit.

H^shew ip route ospf

10,0,0,0/24 is 3 subnetLea_H 7 subnets

0 E2 10,4.5,0 1110/2001 via 10.4,1.7,6, 00:06:43, FastEthernet 0/0

[110/200] via 10.4.15.5, 00:06:43,

FastEthernet0/1

0 IA 10,4,27,0 [110/2] via 10,4,15,5, 00:06:44, FastEthernetO/1

0 E2 10,4.49.0 [110/200] via 10.4,17,6, 00:06:43, FaatEthemstO/O

An engineer configures two ASBRs 10 4 17.6 and 10 4 15 5 in an OSPF network to redistribute routes from EIGRP. However, both ASBRs show the EIGRP routes as equal costs even though the next-hop router 10 4 17 6 is closer to R1. How should the network traffic to the EIGRP prefixes be sent via 10 4.17.6?

- A. The administrative distance should be raised to 120 from the ASBR 10.4.15.5.
- B. The redistributed prefixes should be advertised as Type 1.
- C. The ASBR 10 4 17 6 should assign a tag to match and assign a lower metric on R1.
- D. The administrative distance should be raised to 120 from the ASBR 104.17.6.
- E. The administrative distance should be raised to 120 from the ASBR 104 15.5.
- F. The redistributed prefixes should be advertised as Type 1.
- G. The ASBR 10 4 17 6 should assign a tag to match and assign a lower metric on R1.
- H. The administrative distance should be raised to 120 from the ASBR 104 17 6.

Answer: B

Explanation:

Question: 454

An engineer configured routing between multiple OSPF domains and introduced a routing loop that caused network instability. Which action resolves the problem?

- A. Set a tag using the redistribute command toward a domain and deny inbound in the other domain by a matching tag.

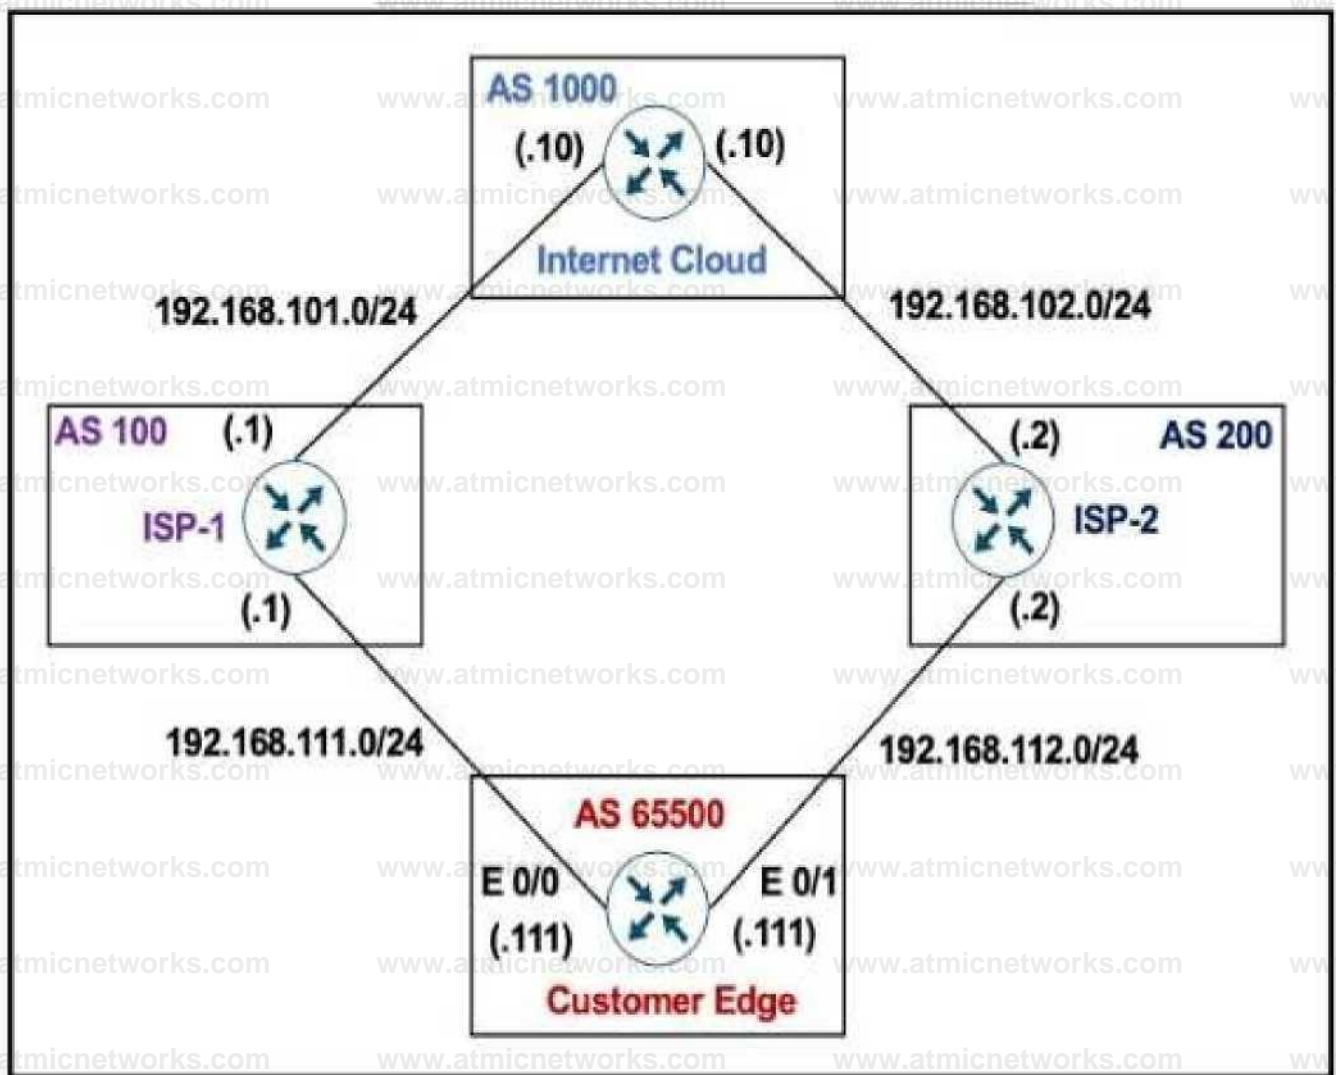
- B. Set a tag using the redistribute command toward a different domain and deny the matching tag when exiting from that domain
- C. Set a tag using the network command in a domain and use the route-map command to deny the matching tag when exiting toward a different domain
- D. Set a tag using the network command in a domain and use the route-map command to deny the matching tag when entering into a different domain

Answer: A

Explanation:

Question: 455

Refer to the exhibit.



The Customer Edge router (AS 65500) wants to use ASC100 as the preferred ISP for all external routes.

CuMbmtr Edge

```
route-map SETLP set local-preference 111
```

```
router bgp 65500
```

```
neighbor 192.168 111 i nnwte-H IK
```

```
neighbor 192.168.111.1 route-map SETLPout
```

```
neighbor 192.168 it! 2 remote-as 200
```

This configuration failed to send routes to AS 100 as the preferred path. Which set of configuration resolves the issue?

```
ro die-map SETLP set local preference 111
```

```
1
```

```
router bgp 56 too
```

```
neighbor 192.160.111.1 remote-as 1W
```

```
neighbor 192.168 111 1 route-map SETLPout
```

```
route-map SETLP
```

```
set local-preference 111
```

```
T
```

```
router bgp 56500
```

```
neighbor 192.168 1111 remote-as 1M
```

```
neighbor 192.168 111.1 route-map SETLP in
```

```
route-map SETPP
```

```
match as-path prepend 111 111
```

```
router bgp 66500
```

```
neighbor 192.168 111 1 remote-as 10E
```

```
neighbor 192.168 111.1 route-map SEIPP out
```

```
route-map SETPP set as-path prepend 100 10E
```

```
t
```

```
router bgp 65800
```

```
neighbor 192.168.111 1 remote-as 10E
```

```
neighbor 192.168 111 1 route-map SETPP in
```

A. Option A

B. Option B

C. Option C

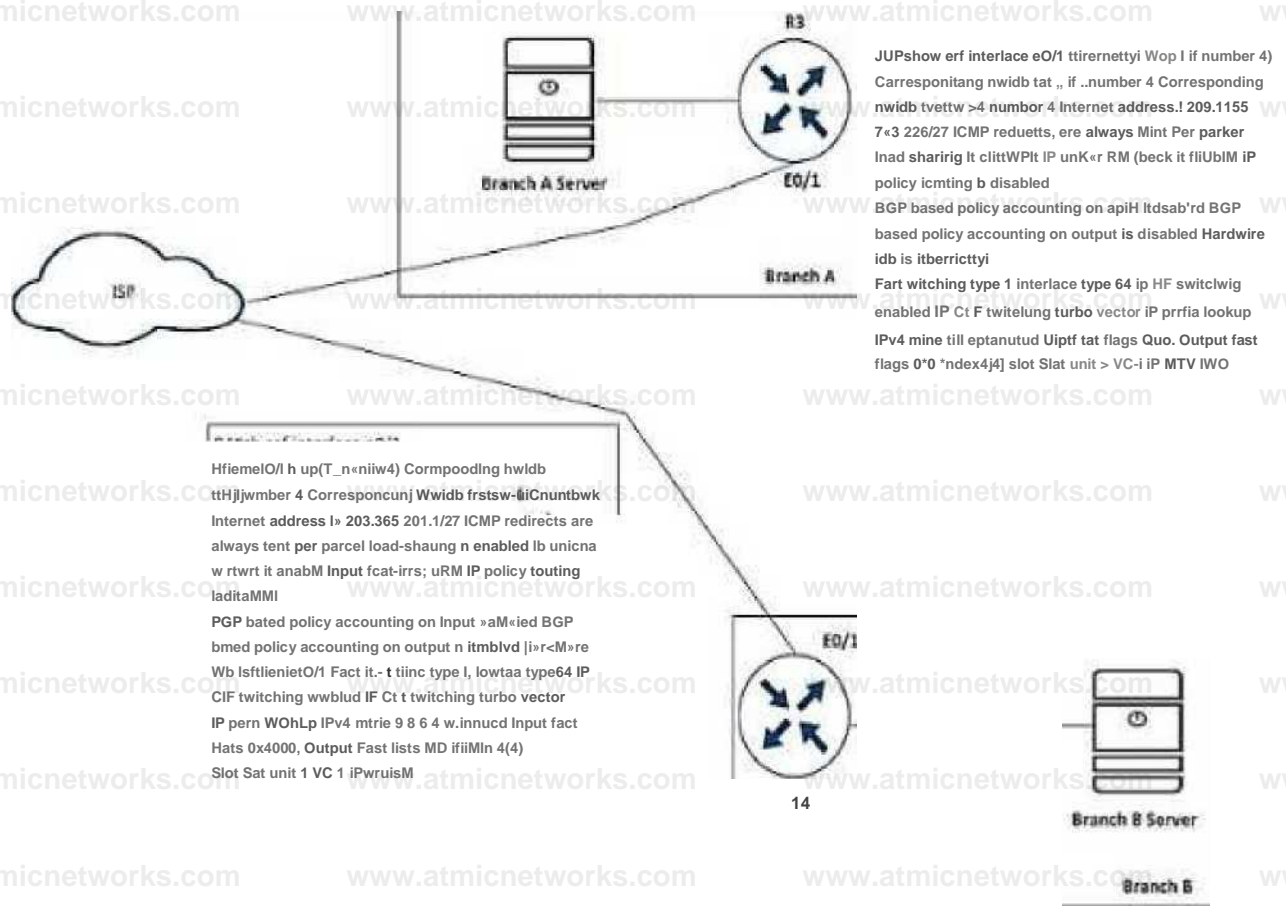
D. Option D

Answer: B

Explanation:

Question: 456

Refer to the exhibit.



Refer to the exhibit.

A shoe retail company implemented the uRPF solution for an antispoofing attack. A network engineer received the call that the branch A server is under an IP spoofing attack. Which configuration must be implemented to resolve the attack?

A)

R4

Interface ethernetO/1

ip unicast RPF check reachable-via any allow-default allow-self-ping

B)

R4

interface ethemetO/1

ip verify unicast source reachable-via any al low-default al low-self-pi ng

C)

R3

**interface ethernet0/1
ip verify unicast source reachable-via any all low-default allow-self-ping**

D)

R3

**interface ethernet0/1
ip unicast RPF check reachable-via any allow-dofault allow-self-ping**

A. Option A

B. Option B

C. Option C

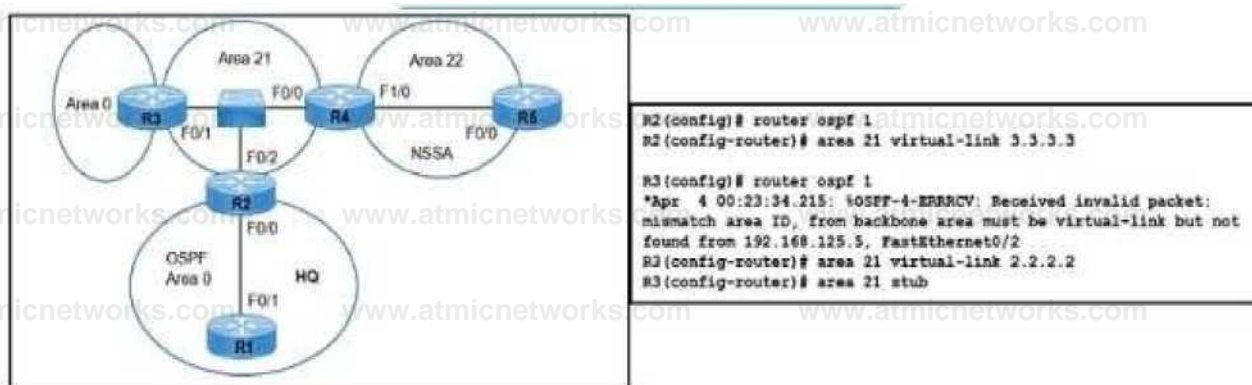
D. Option D

Answer: C

Explanation:

Question: 457

Refer to the exhibit.



Refer to the exhibit. A network engineer is troubleshooting a failed link between R2 and R3. No traffic loss is reported from router R5 to HQ. Which command fixes the separated backbone?

- A. R2(config-router)#no area 21 stub
- B. R2(config_router)#area 21 virtual-link 192.168.125.5
- C. R3(config-router)#area 21 virtual-link 192.168.125.5
- D. R3(config-router)#no area 21 stub

Answer: D

Explanation:

Question: 458

Refer to the exhibit.

```
R1# configure terminal
R1(config)# hostname CPE1
CPE1(config)# ip domain-name example.com
CPE1(config)# crypto key generate rsa
The name for the keys will be: CPE1.example.com
Choose the size of the key modulus in the range of 360 to 4096
for your
  General Purpose Keys. Choosing a key modulus greater than 512
may take
  a few minutes.

How many bits in the modulus [512]: 2048
% Generating 2048 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 2 seconds)

CPE1(config)# service password-encryption
CPE1(config)# username csadmin secret Secur3p4s$w0rd
CPE1(config)# line vty 0 4
CPE1(config-line)# transport input telnet ssh
CPE1(config-line)# login local
CPE1(config-line)# end
CPE1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
CPE1# ssh 10.0.0.1
% No user specified nor available for SSH client
```

CP£U copy running-config startup-config Destination filename [startup-config]? Building configuration...

[OK]

CPE1# ssh 10.0.0.1

% No user specified nor available for SSH client

Refer to the exhibit. An administrator must harden a router, but the administrator failed to test the SSH access successfully to the router. Which action resolves the issue?

A. Configure SSH on the remote device to log in using SSH

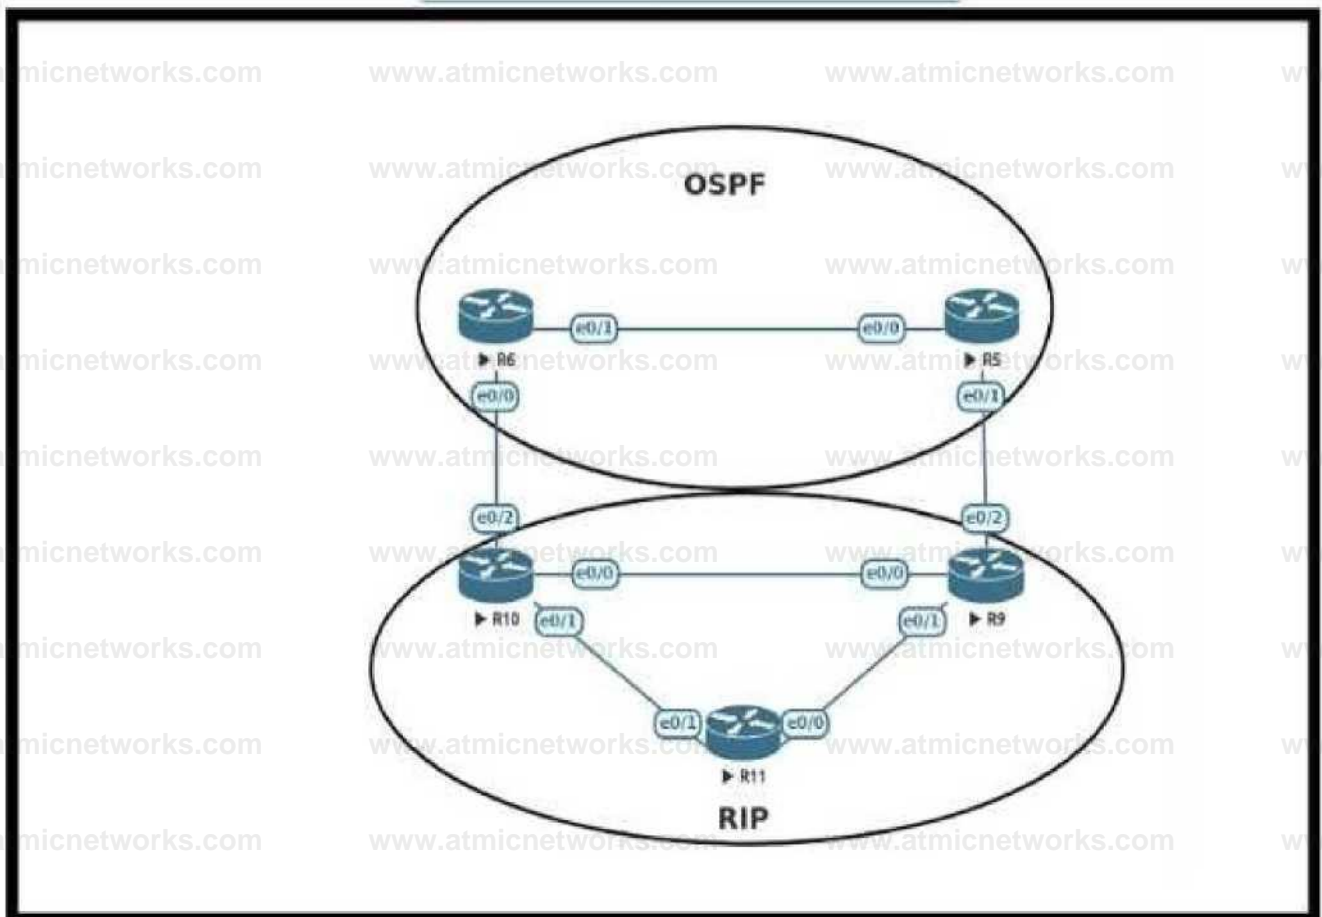
- B. SSH syntax must be ssh -l user ip to log in to the remote device
- C. Configure enable secret to log in to the device
- D. SSH must be allowed with the transport output ssh command

Answer: B

Explanation:

Question: 459

Refer to the exhibit.



An engineer must configure OSPF with R9 and R10 and configure redistribution between OSPF and RIP causing a routing loop Which configuration on R9 and R10 meets this objective?

- A)


```
router ospf 1
 redistribute rip subnets tag 20
 |
 route-map deny_tag20 deny 10
 match tag 20
 route-map deny_tag20 permit 20 t
 router ospf 1
 distribute-list route-map deny_tag20 in
```

B)

```
router ospf 1
 redistribute rip subnets tag 20
 1
 route-map deny_tag20 permit 10
 match tag 20
 route-map deny_tag20 permit 20
```

```
router ospf 1
 distribute-list route-map deny_tag20 in
```

C)

```
router ospf 1
 redistribute rip subnets tag 20

 route-map deny_tag20 deny 10
 match tag 20
 route-map deny_tag20 deny 20
```

```
router ospf 1
 distribute-list route-map deny_tag20 in
```

D)

```
router ospf 1
 redistribute rip subnets tag 20
 i
 route-map dony_tag20 deny 10
 match tag 20
 route-map deny_jag20 permit 20 i
 router rip 1
 distribute-list route-map deny_tag20 in
```

A. Option

B. Option

C. Option

D. Option

Answer: A

Explanation:

Question: 460

Refer to the exhibit.

A network administrator is troubleshooting OSPF adjacency issue by going through the console logs in the router, but due to an overwhelming log message stream it is impossible to capture the problem. Which two commands reduce console log messages to relevant OSPF neighbor problem details so that the issue can be resolved? (Choose two)

A. debug condition interface

B. debug condition ip

C. debug condition ospf neighbor

D. debug condition session-id ADJCHG


E. debug condition all

Answer: A,D

Explanation:

Question: 461

Refer to the exhibit.



Which policy configuration on R1 forwards any traffic that is sourced from the 192.168.130.0/24 network to R2?

A)

```
access-list 1 permit 192.168.130.0 0.0.0 255
```

```
interface Gi0/2  
ip policy route-map test  
route-map test permit 10
```

```
match ip address 1 set ip next-hop 172.20.20.1
```

B)

```
access-list 1 permit 192.168.130.0 0 0.0 255
```

```
interface Gi0/1  
ip policy route-map test
```

```
route-map test permit 10  
match ip address 1  
set ip next-hop 172.20.40.1
```

C)

```
access-list 1 permit 192.168.130.0 0.0.0.255
```

```
interlace Gi0/2  
ip policy route-map test
```

```
route-map test permit 10  
match ip address 1  
set ip next-hop 172.20.20.2
```

D)

```
access-list 1 permit 192.168.130.0 0.0.0.255 i
```

```
interface Gi0/1  
ip policy route map test
```

```
route map test permit 10
```

match ip address 1
set ip next-hop 172.20 40 2

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 462

Refer to the exhibit.

```

R4#
Interface FastEthernet1/0
ip address 10.1.1.14 255.255.255.252
ip access-group VENDOR in
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 EIGRPKEY
speed 100
full-duplex
!
interface loopback 100
ip address 10.199.100.1 255.255.255.255
!
router eigrp 100
network 10.1.1.8 0.0.0.3
network 10.1.1.12 0.0.0.3
no auto-summary
eigrp router-id 100.4.4.4
neighbor 10.1.1.13 FastEthernet1/0
redistribute connected
!
router bgp 65001
no synchronization
bgp log-neighbor-changes
network 100.4.4.4 mask 255.255.255.255
neighbor 10.1.1.13 remote-as 65001
no auto-summary
!
ip access-list extended VENDOR
permit tcp 192.168.32.0 0.0.7.255 host 10.199.100.1 eq 22 time-range VENDOR_ACCESS
!
time-range VENDOR_ACCESS
periodic weekend 22:00 to 23:00

```

Refer to the exhibit A network engineer received a call from the vendor for a failed attempt to remotely log in to their managed router loopback interface from 192.168.40.15 Which action must the network engineer take to resolve the issue?

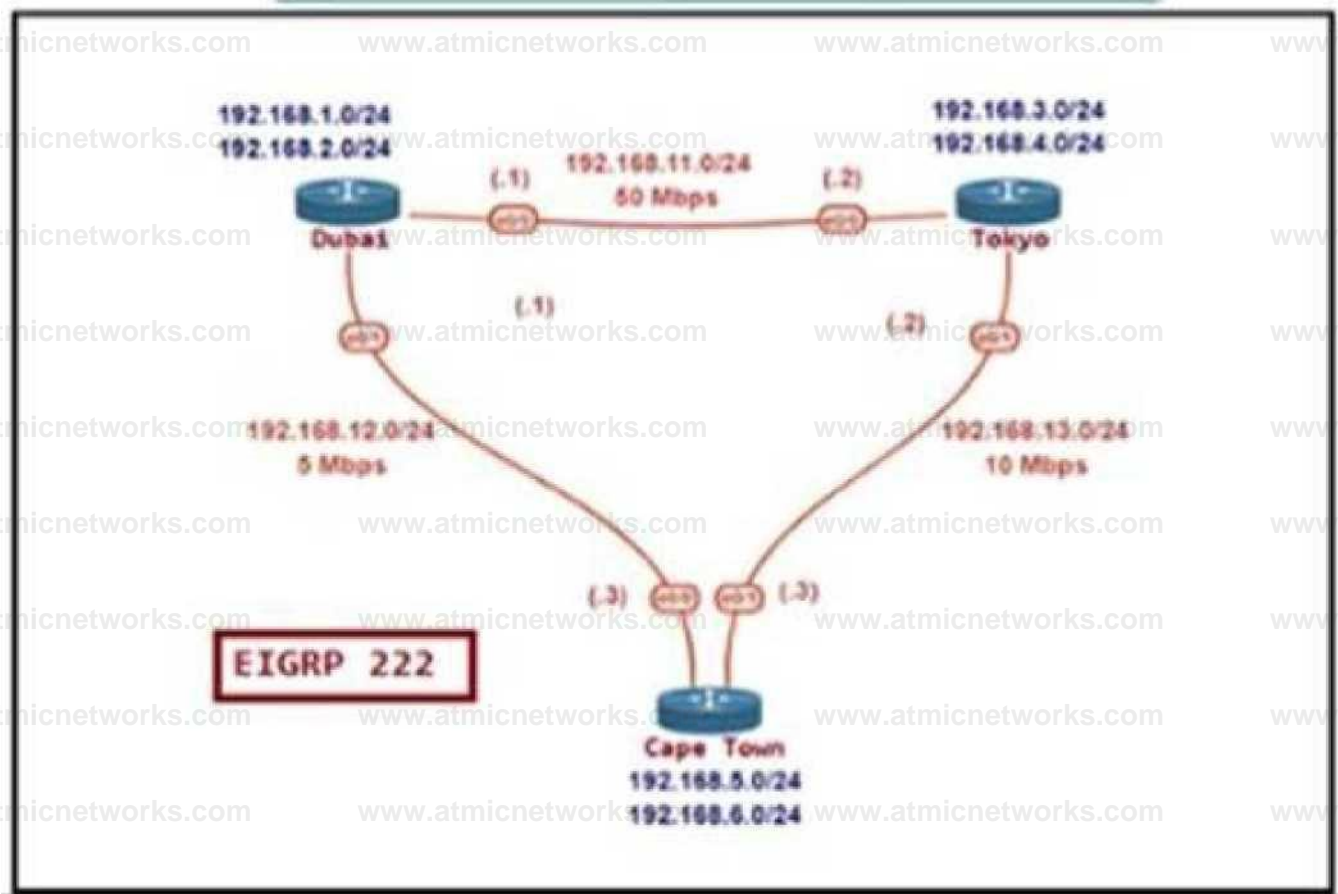
- A. The IP access list VENDOR must be applied to interface loopback 100
- B. The time-range configuration must be changed to use absolute instead of periodic
- C. The EIGRP configuration must be updated to include a network statement for loopback 100
- D. The source IP summarization must be updated to include the vendor source IP address

Answer:
C

Explanation:

Question:
463

Refer to the exhibit.



D 192.168.2.0/24 [90/4b3vwj] via i az. lvo. l z. l, w.ua. i r, cihernet0/0 D
192.168.3.0/24 [90/409600] via 192.168.13.2. 00:17:23, EthernetO/1 D
192.168.4.0/24 [90/409600] via 192.168.13.2, 00:17:23, EthernetO/1
192.168.5.0/24 is variably subnetted. 2 subnets, 2 masks
C 192.168.5.0/24 is directly connected, LoopbackO
L 192.168.5.1/32 is directly connected, LoopbackO
192.168.6.0/24 is variably subnetted, 2 subnets. 2 masks
C 192.168.6.0/24 is directly connected. Loopbackl
L 192.168.6.1/32 is directly connected. Loopbackl
D 192.168.11.0/24 [90/307200] via 192.168.13.2.00:17:40, EthernetO/1
[90/307200] via 192.168.12.1, 00:17:40, EthernetO/O 192.168.12.0/24 is
variably subnetted, 2 subnets, 2 masks
C 192.168.12.0/24 is directly connected, EthernetO/O
L 192.168.12.3/32 is directly connected, EthernetO/O
192.168.13.0/24 is variably subnetted. 2 subnets. 2 masks
C 192.168.13.0/24 is directly connected. EthernetO/1
192.168.13.3/32 is directly connected. EthernetO/1

The network administrator must configure Cape Town to reach Dubai via Tokyo based on the speeds provided by the service provider. It was noticed that Cape Town is reaching Dubai directly and failed to meet the requirement. Which configuration fixes the issue?

A)

Dubai

```
router eigrp 100  
variance 2
```

B)

Cape Town

```
router eigrp 100  
variance 2
```

C)

Cape Town

interface E 0/0 bandwidth 5000

interface E 0/1 bandwidth 10000

D)

Cape Town

interface E 0/0 bandwidth 5000

interface E 0/1 bandwidth 10000

Dubai

interface E 0/0 bandwidth 50000 interface E 0/1 bandwidth 5000

Tokyo

interface E 0/0 bandwidth 50000

interface E 0/1 bandwidth 10000

A. Option

B. Option

C. Option

D. Option

Answer: D

Explanation:

Question: 464

The network administrator configured the router for Control Plane Policing so that inbound SSH traffic is policed to 500 kbps. This policy must apply to traffic coming in from 10.10.10.0/24 and 192.168.10.0/24 networks.

```
access-list 100 permit ip 10.10.10.0 0.0.0.255 any
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 23
```

```
class-map CLASS-SSH
match access-group 100
```

```
policy-map PM-COPP
class CLASS-SSH
police 500 conform-action transmit
```

```
interface E0/0
service-policy input PM-COPP
```

```
interface E0/1
service-policy input PM-COPP
```

The Control Plane Policing is not applied to SSH traffic and SSH is open to use any bandwidth available. Which configuration resolves this issue?

```
no access-list 100
access-list 100 permit ip 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22
policy-map PM-COPP
class CLASS-SSH
no police 500 conform-action transmit
police 500 conform-action transmit exceed-action drop
```

```
interface E0/0
no service-policy input PM-COPP
```

```
interface E0/1
no service-policy input PM-COPP
```

```
control-plane
service-policy input PM-COPP
```

```
no access-list 100
access-list 105 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22
```

```
interface E0/0
no service-policy input PM-COPP
```

```
interface E0/1
```

no service-policy input PM-COPP

control-plane
service-policy input PM-COPP

no access-list 100
access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22

A)

no access-list 100

access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22

policy-map PM-COPP

class CLASSSSH

no police 50000 conform-action transmit

police 500000 conform-action transmit exceed drop

B)

interface E0/0

no service-policy input PM-COPP

interface E1

no service-policy input PM-COPP

control-plane

service-policy input PM-COPP

C)

no access-list 100

access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22

access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22

Interface E0/0

no service-policy input PM-COPP

Interface E0/1

no service-policy input PM-COPP

control-plane

service-policy input PM-COPP

D)

```
nc access-list 100  
access-list 100 permit tep 10.10.10.0 0.0.0.255 any eq 22  
access-list 100 permit top 192.168.10.0 0.0.0.255 any eq 22
```

A. Option

B. Option

C. Option

D. Option

Answer: C

Explanation:

Question: 465

Refer to the exhibit.

```
admin@linux:~$ telnet 198.51.100.64
Trying 198.51.100.64...
Connected to 198.51.100.64.
Escape character is '^]'.
User Access Verification

Password: admin
CPE> exit
Connection closed by foreign host.
admin@linux:~$ ssh 198.51.100.64
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Connection closed by 198.51.100.64 port 22
admin@linux:~$
```

Refer to the exhibit. An administrator can log in to the device using Telnet but the attempts to log in to the same device using SSH with the same credentials fail. Which action resolves this issue?

- A. Configure SSH service on the router
- B. Configure transport input all on the VTY lines to allow SSH
- C. Configure to use the Telnet user database for SSH as well
- D. Configure the VTY lines with login local

Answer:

A

Explanation:

Question:

466

Refer to the exhibit.

```
R2(config)* int tun0
```

```
Feb 23 00:42:06.179: 4LINEPROTO-S-UPMWN: Line protocol on  
Interface Tunnel0, changed state to down
```

```
R2(config-if)* ip address 192.168.12.2 255.255.255.0
```

```
R2(config-if)* tunnel source lo0
```

```
R2(config-if)* tunnel destination 10.255.255.1
```

```
Feb 23 00:42:15.845: 5LINEPROTO-S-UPDOWN: Line protocol on  
Interface Tunnel0, changed state to up
```

```
R2(config-if)* router eigrp 1
```

```
R2(config-router)* address-family ipv4 autonomous-system 1
```

```
R2(config-router-af)* net 192.168.12.2 0.0.0.0
```

```
Feb 23 00:43:05.730: 5DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor  
192.168.12.1 (Tunnel0) is up: new adjacency
```

```
Feb 23 00:43:05.993: 5VADJ-5-PARENT: Midchain parent maintenance for  
IP midchain out of Tunnel0 - looped chain attempting to stack
```

```
Feb 23 00:43:15.193: 5WIN-5-RECURDOWN: Tunnel0 temporarily  
disabled due to recursive routing
```

```
Feb 23 00:43:15.193: 4LINEPROTO-S-UPDOWN: Line protocol on  
Interface Tunnel0, changed state to down
```

An administrator is configuring a GRE tunnel to establish an EIGRP neighbor to a remote router. The other tunnel endpoint is already configured. After applying the configuration as shown, the tunnel started flapping. Which action resolves the issue?

- A. Stop sending a route matching the tunnel destination across the tunnel.
- B. Modify the network command to use the Tunnel0 Interface netmask.
- C. Advertise the Loopback0 interface from R2 across the tunnel.
- D. Readdress the IP network on the Tunnel0 on both routers using the /31 netmask.

Answer: A

Explanation:

Question: 467

A network administrator is troubleshooting a failed AAA login issue on a Cisco Catalyst c3560 switch. When the network administrator tries to log in with SSH using TACACS+ username and password credentials, the switch is no longer authenticating and is failing back to the local account. Which action resolves this issue?

- A. Configure ip tacacs source-interface GigabitEthernet 1/1
- B. Configure ip tacacs source-ip 192.168.100.55
- C. Configure ip tacacs-server source-ip 192.168.100.55
- D. Configure ip tacacs-server source-interface GigabitEthernet 1/1

Answer: A

Explanation:

Question: 468

Refer to the exhibit.

C-EI show snap mb if mb ifindex detail						
Interface	Index	Active	Persistent	Saved	Trap	Rate
Loopback1	9	yes	disabled	no	enabled	
GigabitEthernet1	1	yes	disabled	no	enabled	
GigabitEthernet1/3	3	yes	disabled	no	enabled	
Serial3/123	10	yes	disabled	no	disabled	
VTP-Sullo		yes	disabled	no	enabled	
Loopback0	7	yes	disabled	no	enabled	

Null0	€	y«	disabled	no	enabled
Lo:pback2	&	y«	disabled	no	enabled
Gi jab:tEthrrnet4	4	y^s	disabled	no	enabled
Gigabits theme t2	2	yes	disabled	no	enabled

Refer to the exhibit. After reloading the router an administrator discovered that the interface utilization graphs displayed inconsistencies with their previous history in the NMS. Which action prevents this issue from occurring after another router reload in the future?

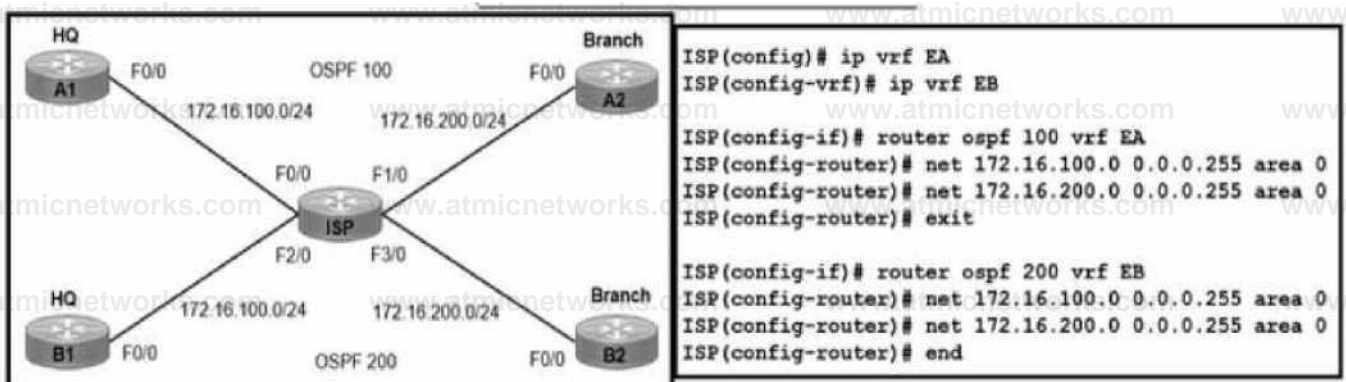
- A. Rediscover all the router interfaces through SNMP after the router is reloaded
- B. Save the router configuration to startup-config before reloading the router
- C. Configure SNMP to use static OIDs referring to individual router interfaces
- D. Configure SNMP interface index persistence on the router

Answer: D

Explanation:

Question: 469

Refer to the exhibit.



Refer to the exhibit. A network engineer is provisioning end-to-end traffic service for two different enterprise networks with these requirements

The OSPF process must differ between customers on HQ and Branch office routers, and adjacencies should come up instantly.

The enterprise networks are connected with overlapping networks between HQ and a branch office

Which configuration meets the requirements for a customer site?

A

A)

```
interface A
 ip address 172.16.100.2 255.255.255.0
 ip ospf network point-to-point
 ip ospf priority 1
 ip ospf area 0
 ip ospf cost 100
```

B)

```
interface A
 ip address 172.16.100.2 255.255.255.0
 ip ospf network point-to-point
 ip ospf priority 1
 ip ospf area 0
 ip ospf cost 100
```

C)

```
interface A
 ip address 172.16.100.2 255.255.255.0
 ip ospf network point-to-point
 ip ospf priority 1
 ip ospf area 0
 ip ospf cost 100
```

D)

```
interface A
 ip address 172.16.100.2 255.255.255.0
 ip ospf network point-to-point
 ip ospf priority 1
 ip ospf area 0
 ip ospf cost 100
```

ISP ' . ■ ^description TO-*EA2_Branch I\$P|wnflg4)t|p add 1K.1tm2
M&2S&2M.0 ISP|conng4»#fto shut

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 470

Refer to the exhibit.

```

R1#sh ipv6 route eigrp
IPv6 Routing Table - default - 1 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - IISP
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

```


```

R1#
R1#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(1)

```

H	Address	Interface	Hold	Uptime	SRTT	RTO	Q	Seq
		(sec)	(ms)					
0	Link-local address: FE80::C004:22FF:FE78:1	Fa1/0	11	00:04:22	1593	5000	0	15

```
R1#
```



```
R2#show run
interface Loopback10
no ip address
ipv6 address 2002::1/128
ipv6 eigrp 1
|
interface Loopback11
no ip address
ipv6 address 2002::2/128
ipv6 eigrp 1
|
interface Loopback12
no ip address
ipv6 address 2002::3/128
ipv6 eigrp 1
|
interface FastEthernet0/1
no ip address
duplex auto
speed auto
ipv6 address autoconfig
ipv6 eigrp 1
|
ipv6 router eigrp 1
stub summary
no shutdown
```

R1 cannot receive the R2 Interfaces with individual prefixes. What must be reconfigured to advertise R2 Interfaces to R1?

- A. EIGRP process on R2 by removing the stub command Keyword summary
- B. interface FastEthernet0/1 on R2 with an EIGRP summary for all three loopback prefixes

C. EIGRP process on R2 with the command stub summary receive-only

D. EIGRP process on R2 with the command stub summary connected

Answer:

D

Explanation:

Question: 471

What is a characteristic of IPv6 RA Guard?

A. RA messages are allowed from the host port to the switch

B. It is unable to protect tunneled traffic

C. It filters rogue RA broadcasts from connected hosts

D. It is supported on the egress direction of the switch

Answer:

C

Explanation:

Question:

472

Refer to the exhibit.

```
R1(config)#ip prefix-list EIGRP seq 10 permit 10.0.0.0/8
R1(config)#ip prefix-list EIGRP seq 20 deny 0.0.0.0/0 le 32
R1(config)#router eigrp 10
R1(config-router)#distribute-list prefix EIGRP in Ethernet0/0

R1#show ip route eigrp | include 10.
D EX 10.0.0.0/8 [170/2665332] via 192.168.10.1, 00:00:10,
Ethernet0/0
```

An engineer applies a prefix-list filter that filters most of the network 10 prefixes instead of allowing them. Which action resolves the issue?

- A. Modify the ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 9 command.
- B. Modify the command Modify the Ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 32 command.
- C. Modify the Ip prefix-list EIGRP seq 20 permit 0.0.0.0/0 le 32 command.
- D. Modify the ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9 command

Answer: C

Explanation:

Question: 473

Refer to the exhibit.

```
RtrA#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(1)/ID(10.1.6.1)
... snip ...
P 10.200.1.0/24, 1 successors, FD is 21026560
via 10.1.1.2 (21026560/20514560), Serial1/0
via 10.1.2.2 (46740736/20514560), Serial1/1
via 10.1.3.2 (46740736/46228736), Serial1/2
```

Which action makes 10.1.3.2 the feasible successor to reach 10.200.10/24 for location S42T447E33F95?

- A. Increase path bandwidth lower than 10112 and lower than 10122 between RtrA and the destination
- B. Increase path bandwidth higher than 10.122 and lower than 101.1.2 between RtrA and the destination.
- C. Increase path bandwidth higher than 10112 and lower than 10122 between RtrA and the destination
- D. Increase path bandwidth higher than 10.122 and higher than 10.1.1.2 between RtrA and the destination

Answer: A

Explanation:

Question: 474

Refer to the exhibit.

```
RI(config)#ip access-list standard EIGRP-FILTER
RI(config-std-nacl)#deny 10.10.10.0 0.0.0.0
RI(config-std-nacl)#permit 0.0.0.0 0.0.0.0
RI(config)#router eigrp 10
RI(config-router)#distribute-list route-map EIGRP in

R1(config)#route-map EIGRP permit 10
R1(config-route-map)#match ip address EIGRP-FILTER

R1#show ip route eigrp | include 10.10.10.
D 10.10.10.128/25
```

Refer to the exhibit. An engineer must filter EIGRP updates that are received to block all 10.10.10.0/24 prefixes. The engineer tests the distribute list and finds one associated prefix. Which action resolves the issue?

- A. There is a permit in the route map that allows this prefix. A deny 20 statement is required with a match condition

to match a new ACL that denies all prefixes

B. There is a permit in the ACL that allows this prefix into EIGRP. The ACL should be modified to deny 10.10.10.0/24.

C. There is a permit in the route map that allows this prefix. A deny 20 statement is required with no match condition to block the prefix.

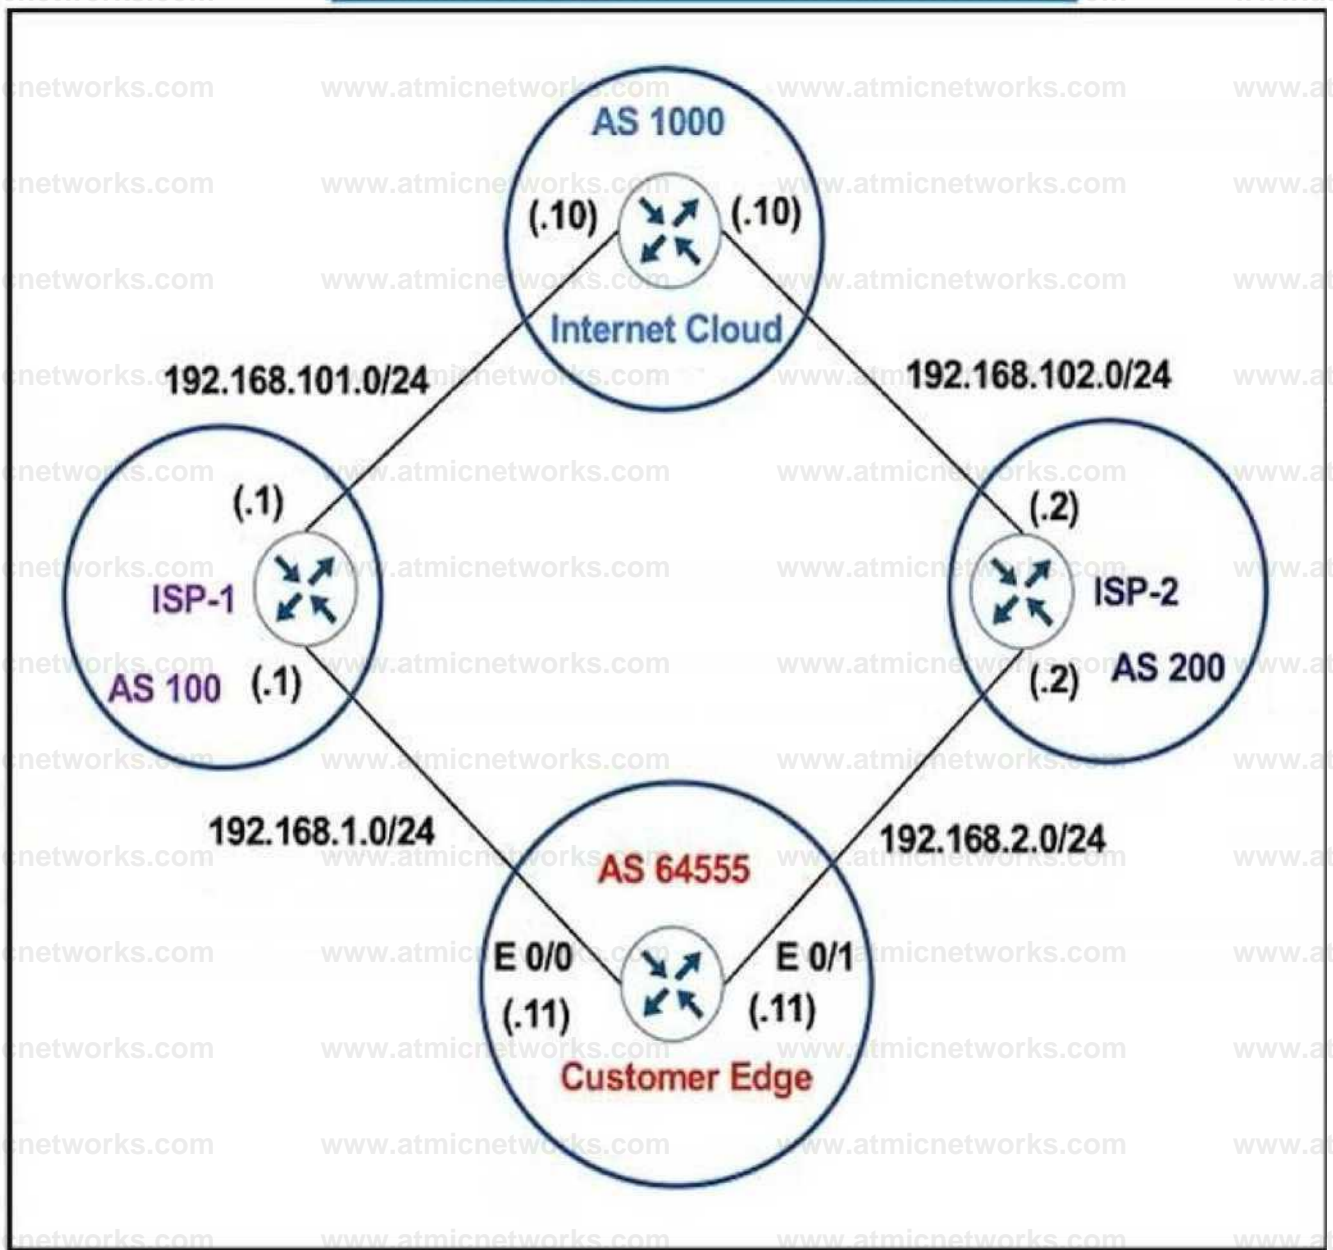
D. There is a permit in the ACL that allows this prefix into EIGRP. The ACL should be modified to deny 10.10.10.0/24.

Answer: B

Explanation:

Question: 475

Refer to the exhibit.



Refer to the exhibit. The Customer Edge router wants to use AS 100 as the preferred ISP for all external routes and ISP-2 as a backup.

Customer-Edge

```
route-map SETAS
set as-path prepend 111
```

```
router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS in
```

After this configuration, all the backup routes have disappeared from the BGP table on the Customer Edge router. Which set of configurations resolves the issue on the Customer Edge router?

A)

```
route-map SETAS set as-path prepend 111
```

```
router bgp 64555
```

```
neighbor 192.168.2.2 remote-as 100
```

```
neighbor 192.168.1.1 remote-as 200
```

```
neighbor 192.168.1.1 route-map SETAS in
```

B)

```
route-map SETAS
```

```
set as-path prepend 200
```

```
i
```

```
router bgp 64555
```

```
neighbor 192.168.1.1 remote-as 100
```

```
neighbor 192.168.2.2 remote-as 200
```

```
neighbor 192.168.2.2 route-map SETAS in
```

C)

```
route-map SETAS
```

```
set as-path prepend 200
```

```
i
```

```
router bgp 64555
```

```
neighbor 192.168.1.1 remote-as 100
```

```
neighbor 192.168.2.2 remote-as 200
```

```
neighbor 192.168.2.2 route-map SETAS out
```

D) **route-map SETAS**
set as-path prepend 111
1

router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS out

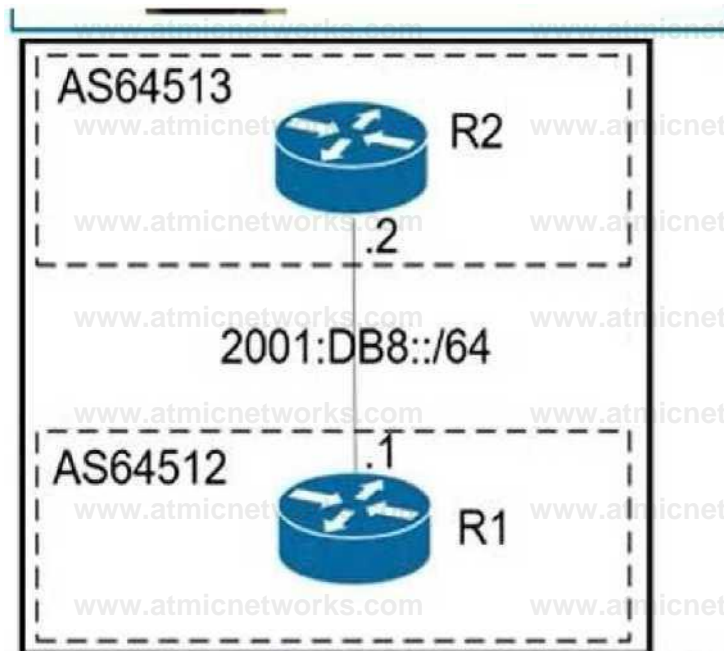
- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 476

Refer to the exhibit.



```
R1#show ipv6 access-list
```

```
IPv6 access list inbound-acl
```

```
permit tcp host 2001 :DB8::2 eq bgp host 2001 :DB8::1 (75 matches) sequence 20
```

```
permit tcp host 2001:DB8::2 host 2001:DB8::1 eq bgp (17 matches) sequence 30
```

```
deny ipv6 2001:DB8::/32 any (77 matches) sequence 40
```

```
permit ipv6 any (20 matches) sequence 1000
```

```
R1#ping ipv6 2001:DB8::2
```

```
Type escape sequence to abort.
```

```
Sending 5,100-byte ICMP Echos to 2001:DB8::2, timeout is 2 seconds:
```

```
Success rate is 0 percent (0/5)
```

```
R1#show ipv6 access-list
```

```
IPv6 access list inbound-acl
```

```
permit tep host 2001 :DB8::2 eq bgp host 2001:DB8::1 (77 matches)
sequence 20
```

```
permit tep host 2001:DB8::2 host 2001:DB8::1 eq bgp (19 matches)
sequence 30
```

```
deny ipv6 2001:DB8::/32 any (95 matches) sequence 40
```

```
permit ipv6 any (23 matches) sequence 1000
```

```
R1#
```

Refer to the exhibit. An engineer applied filter on R1. The interface flapped between R1 and R2 and cleaning the BGP session did not restore the BGP session and failed. Which action must the engineer take to restore the BGP session from R2 to R1?

- A. Apply the IPv6 traffic filter in the outbound direction on the interface
- B. ICMPv6 must be permitted by the IPv6 traffic filter
- C. Enable the BGP session, which went down when the session was cleared.
- D. Swap the source and destination IP addresses in the IPv6 traffic filter

Answer: B

Explanation:

Question: 477

What is the purpose of an OSPF sham-link?

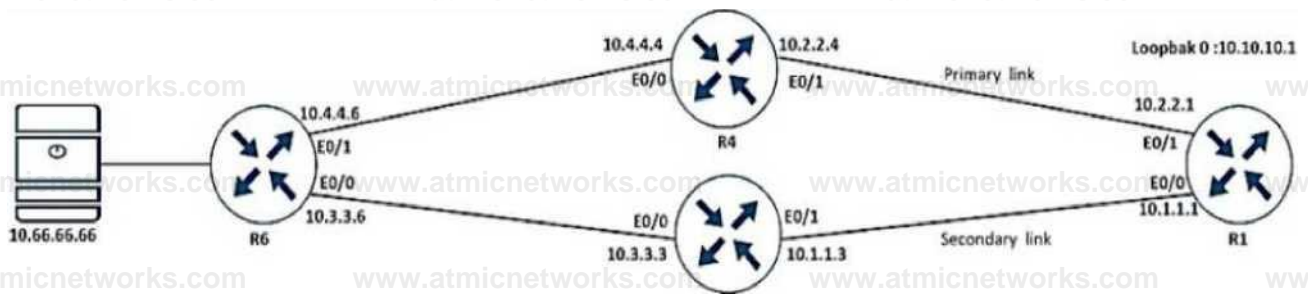
- A. to allow intra-area routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network
- B. to correct OSPF backdoor routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network
- C. to correct OSPF backdoor routing when OSPF is used as the PE-PE connection protocol in an MPLS VPN network
- D. to allow inter-area routing when OSPF is used as the PE-CE connection protocol in a MPLS VPN network

Answer: C

Explanation:

Question: 478

Refer to the exhibit.



```
R3* & R40
interface Ethernet0/1
 *P access-group DDOS in
 1
 ip access-list extended DDOS
 permit tcp any any
 deny udp any any range 1024 65535
 permit ip any any
```

```
R1 Ash flow interface
Interface Ethernet0/0 FNF monitor:
FlowMonitor1
 direction: input
 traffk(ip). on
FNF monitor: FlowMonitor1
 direction. Output traffk(ip) on
Interface Ethernet0/1 FNF monitor:
FlowMonitor1 direction: input traffk(ip)
 on
FNF: monitor: FlowMonitor1 direction'
Output traffk(ip): on
```

```
R1show flow exporter
Flow Exporter FlowExporter1: Descript^:
User defined
Export protocol: NetFlow Version 5
Transport Configuration:
Destination IPAddrsss: 1066.66.66
Source IP address' 10.1.11 Transport
Protocol UDP Destination Port: 1090
Source Port: 54186
DSCP 0x0
nu 255
Output Features: Not Used
```

```
R1show flow monitor
Flow Monitor FlowMonitor1:
Description User defined
Flow Record: netflow ipv4 original-input Flow
Exporter: FflowExporter1
Cache:
Type: normal
Status: allocated
Site 4096 entries / 344068 bytes
Inactive Timeout: 15 secs
Active Timeout 1800 secs Update Timeout:
1800 secs Synchronized Timeout: 600 secs
```

Refer to the exhibit An engineer configured NetFlow but cannot receive the flows from R1 Which two configurations resolve the issue? (Choose two)

A)

R1 (config inflow exporter FlowExporter1
R1 (config-flow-exporter ^destination 10.66.66.66

B)

R4(config)#ip access-list extended DDOS
R4(config-ext-nacl)s6 permit udp any host 10.66.66.66 eq 1090

C)

R3(config inflow exporter
FlowExporter!
R3(config-flow-exporter^destination
10.66.66.66

D)

R3(config)#ip access-list extended DDOS
R3(config-ext-nacl)*5 permit udp any host 10.66.66.66 eq 1090

E)

R4(config >=flo w exporter FlowExporter!
R4(config-flow-exporter)5destination 10.66.66.66

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

Answer: B,E

Explanation:

Question: 479

Refer to the exhibit.


```
CPE# copy flash:packages.conf ftp://192.0.2.40/  
Address or name of remote host [192.0.2.40]?  
Destination filename [packages.conf]?  
Writing packages.conf  
%Error opening ftp://192.0.2.40/packages.conf (Incorrect  
Login/Password)  
CPE#
```

Refer to the exhibit. An administrator must upload the packages.conf file to an FTP server. However, the FTP server rejected anonymous service and required users to authenticate. What are the two ways to resolve the issue?

(Choose two.)

- A. Use the ftp username and ip ftp password configuration commands to specify valid FTP server credentials.
- B. Use the copy flash:packages.conf scp: command instead and enter the FTP server credentials when prompted.
- C. Enter the FTP server credentials directly in the FTP URL using the ftp://username:password@192.0.2.40/ syntax.
- D. Create a user on the router matching the username and password on the FTP server and log in before attempting the copy.
- E. Use the copy flash-packages.conf ftp: command instead and enter the FTP server credentials when prompted.

Answer: A,C

Explanation:

Question: 480

A newly installed router starts establishing an LDP session from another MPLS router to which it is not directly connected.

Which LDP message type responds by target router to the initiating router using UDP protocol?

- A. notification message

- B. session message
- C. extended discovery message
- D. advertisement message

Answer: C

Explanation:

Question: 481

What is considered the primary advantage of running BFD?

- A. reduction in time needed to detect Layer 2 switched neighbor failures
- B. reduction in time needed to detect Layer 3 routing neighbor failures
- C. reduction in CPU needed to detect Layer 2 switch neighbor failures
- D. reduction in CPU needed to detect Layer 3 routing neighbor failures

Answer: B

Explanation:

Question: 482

An engineer configured VRF-Lite on a router for VRF blue and VRF red. OSPF must be enabled on each VRF to peer to a directly connected router in each VRF. Which configuration forms OSPF neighbors over the network 10.10.10.0/28 for VRF

router ospf 1 vrf blue network 10.10.10.0
0.0.0.15 area 0 router ospf 2 vrf red
network 192.168.0.0 0.0.0.3 area 0

router ospf 1 vrf blue network 10.10.10.0
0.0.0.240 area 0 router ospf 2 vrf red
network 192.168.0.0 0.0.0.252 area 0

router ospf 1 vrf blue network 10.10.10.0
0.0.0.252 area 0 router ospf 2 vrf red network
192.168.0.0 0.0.0.240 area 0

router ospf 1 vrf blue network 10.10.10.0 0.0.0.3
area 0 router ospf 2 vrf red
network 192.168.0.0 0.0.0,15 area 0

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 483

An engineer configured VRF-Lite on a router for VRF blue and VRF red. OSPF must be enabled on each VRF to peer to a directly connected router in each VRF. Which configuration forms OSPF neighbors over the network 10.10.10.0/28 for VRF blue and 192.168.0.0/30 for VRF red?

router ospf 1 vrf blue network 10.10.10.0
0.0.0.15 area 0 router ospf 2 vrf red network
192.168.0.0 0.0.0.3 area 0

router ospf 1 vrf blue network 10.10.10.0 0.0.0.240
area 0 router ospf 2 vrf red
network 192.168.0.0 0.0.0.252 area 0

router ospf 1 vrf blue network 10.10.10.0
0.0.0.252 area 0 router ospf 2 vrf red network
192.168.0.0 0.0.0.240 area 0

router ospf 1 vrf blue network 10.10.10.0 0.0.0.3
area 0 router ospf 2 vrf red
network 192.168.0.0 0.0.0,15 area 0

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 484

The network administrator is tasked to configure R1 to authenticate telnet connections based on Cisco ISE using RADIUS. ISE has been configured with an IP address of 192.168.1.5 and with a network device pointing towards R1 (192.168.1.1) with a shared secret password of Cisco123. If ISE is down, the administrator should be able to connect using the local database with a username and password combination of admin/cisco123.

The administrator has configured the following on R1:

```
aaa new-model
```

```
i
```

```
username admin password cisco123
```

```
.
```

```
radius server ISE1
```

```
address ipv4 192.168.1.5
```

```
key Cisco 23
```

```
i
```

```
aaa group server tacacs* RAD-SERV server name ISE1
```

```
i
```

```
aaa authentication login RAD-LOCAL group RAD-SERV
```

ISE has gone down. The Network Administrator is not able to Telnet to R1 when ISE went down.

Which two configuration changes will fix the issue? (Choose two.)

- line vty 0 4
login authentication RAD-LOCAL
- line vty 0 4
login authentication default
- line vty 0 4
login authentication RAD-SERV

aaa authentication login RAD-SERV group RAD-LOCAL local

aaa authentication login RAD-LOCAL group RAD-SERV local

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

Answer: C,E

Explanation:

Question: 485

What are the two reasons for RD and VPNv4 addresses in an MPLS Layer 3 VPN? (Choose two.)

A. RD is prepended to each prefix to make routes unique.

B. VPN RT communities are used to identify customer unique routes.

C. When the PE redistributes customer routes into MP-BGP, they must be unique.

D. They are on a CE device to use for static configuration.

E. They are used for a BGP session with the CE device.

Answer: A,C

Explanation:

Question: 486

Refer to the exhibit.

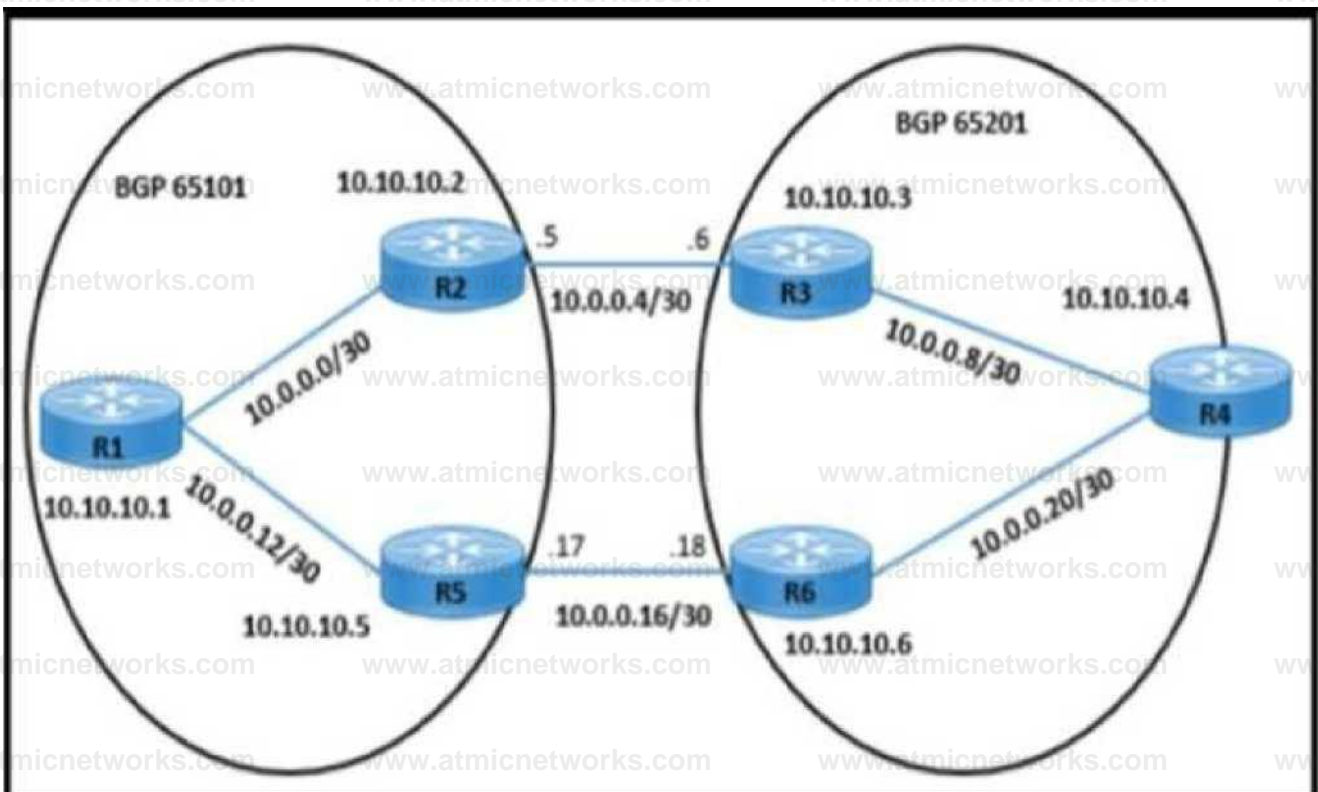
```

R3#
*Sep  5 07:29:34.031: %TCP-6-BADAUTH: No MD5 digest from 10.10.10.2(179) to
10.10.10.3(60942) (RST)
R2# show ip bgp neighbors 10.10.10.3
BGP neighbor is 10.10.10.3, remote AS 65201, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:02:19, last write 00:02:19, hold time is 180, keepalive interval is
60 seconds
Message statistics:
  InQ depth is 0
  OutQ depth is 0

      Sent      Rcvd
Opens:           2         2
Notifications:  0         0
Updates:         5         6
Keepalives:     10        9
Route Refresh:   0         0
Total:          17        17

Default minimum time between advertisement runs is 30 seconds
Address tracking is enabled, the RIB does have a route to 10.10.10.3
Connections established 2; dropped 2
Last reset 00:11:58, due to Peer closed the session
External BGP neighbor not directly connected.
Transport(tcp) path-mtu-discovery is enabled
No active TCP connection

```



The network operation team observes a traffic forwarding issue between R2 and R3:

Ping and traceroute of loopback IP address from R2 to R3 is successful.

iBGP peering in AS 65101 and AS 65201 is up.

Which configuration resolves the issue?

- A. Configure MD5 password authentication on R2.
- B. Advertise R2 and R3 loopback IPs in AS 65101 and AS 65201.
- C. Remove MD5 password authentication on R3.
- D. Set up eBGP multihop on R2 and R3 routers.

Answer: D

Explanation:


Question: 487

Refer to the exhibit.

```
RD*show ip bgp 192.168.1.1 Advertised to update-groups: 3 65101 10.10.10.2 (metric 2) from 10.10.10.2 (10.10.10.2) Origin IGP, metric 100, localpref 100, weight 65535, valid, external best 65101 10.0.0.17 (metric 2) from 10.10.10.6 (172.16.20.1) Origin IGP, metric 0, localpref 100, valid, internal
```

```
RBIshow ip bgp 192.168.1.1 BGP routing table entry for 192.168.1.1/32, version 10 Paths: (1 available, best *1, table Default-IP-Routing-Table) Advertised to update-groups: 2 Local 10.10.10.1 (metric 2) from 10.10.10.1 (192.168.1.1) Origin IGP, metric 0, localpref 100, valid, internal, best
```





Refer to the exhibit. A customer finds that traffic from the application server (192.168.1.1) to the HUB site passes through a congested path that causes random packet drops. The NOC team influences the BGP path with MED on RB, but RD still sees that traffic coming from RA is not taking an alternate route. Which configuration resolves the issue?

A)

```
RD(config)#router bgp 65201
RD(config-router)#no neighbor 10.10.10.2 weight 65535
```

B)


```
RB(config)#router bgp 65101
RB(config-router)#no neighbor 10.10.10.3 route-map HIGH-LP out
```

C)

```
RB(config)#router bgp 65101
RB(config-router)#neighbor 10.10.10.3 weight 50
```

D)

```
RC(config)#router bgp 65101
RC(config-router)#neighbor 10.10.10.6 route-map HIGH-LP out
```

Refer to the exhibit. R1 uses SP1 as the primary path. A network engineer must force all SSH traffic generated from R1

toward SP2. Which configuration accomplishes the task?

A)

```
ip access-list extended match_SSH permit  
tcp any any eq 22
```

```
route-map PBRSSH permit 10 match Ip  
address match_SSH set ip next-hop  
10.20.20.1
```

```
interface GigO/O ip policy route-map  
PBR_SSH
```

B)

**ip access-list extended match SSH permit tcp
any any eq 22**

r
route-map PER SSH permit 10 match
ip address match_SSH set ip next-hop
10.10.10.1

i
ip local policy route-map PBR SSH

c)
ip access-list extended match_SSH permit tcp
any any eq 22

t
route-map PBR_SSH permit 10 match ip
address match SSH set ip next-hop 10.20.20.1

ip local policy route-map PBR_SSH

D)
ip access-list extended match SSH permit tcp
any any eq 22

t

**route-map PBR SSH permit 10 match ip
address match_SSH set Ip next-hop
10.20,20.1**

interface Gig0/1

ip policy route-map PBR_SSH

A. Option

B. Option

C. Option

D. Option

Answer: C

Explanation:

Question: 489

Refer to the exhibit.



```
%DUAL-3-SIA: Route 10.10.1.1/32 stuck-in-active state in IP-EIGRP(0) 1: Cleaning up
%DUAL-3-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 192.168.1.1 (Serial0/0) is down:
stuck in active
```

Refer to the exhibit. An engineer notices a connectivity problem between routers R1 and R2. The frequency of this problem is high during peak business hours. Which action resolves the issue?


- A. Increase the MTU on the interfaces that connect R1 and R2.
- B. Increase the available bandwidth between R1 and R2.
- C. Decrease the EIGRP keepalive and hold down timers on R1 and R2.
- D. Set static EIGRP neighborship between R1 and R2.

Answer: B

Explanation:

Question: 490

Refer to the exhibit.



Refer to the exhibit. The R1 routing table has the prefixes for the FTP1 and FTP2 file servers. A network engineer must configure the R1 with these requirements:

Host1 must use the FTP1 fileserver.

Host2 must use the FTP2 fileserver.

Which configuration meets the requirement on R1?

A)

```
ip access-list extended FTP1
R1 permit ip host 10.10.10.11 host 10.10.100.11
ip access-list extended FTP2
R1 permit ip host 10.10.10.22 host 10.10.200.22
```

```
route-map PBR FTP
permit 10
```

```
match ip address FTP1
R1 set ip next-hop 10.10.40.1
route-map PBR FTP
permit 20
match ip address FTP2
R1 set ip next-hop 10.10.30.1
```

ip local policy route-map FBRJTF

B)

```
ip access-list extended FTP1 permit ip host 10.10.10.11 host
10.10.100.11 ip access-list extended FTP2 permit ip host 10.10.10.22
host 10.10.200.2 2 1 route-map PBR FTP permit 10 match ip address FTP1
R1 set ip next-hop 10.10.30.1
```

```
route-map PQR FTP permit 20 match ip address FTP2 R1 set ip next-hop
10.10.40.1
```

ip local policy route-map PBR FTP

C)

```
ip access-list extended FTP1_R1 permit ip host 10.10.10.1 boot
10.10.10.1 ip access-list extended FTP2_R1 permit ip host
10.10.10.22 host 10.10.200.22
```

```
route-map PBR FTP permit 10 match ip address FTP1_R1 set ip
next-hop 10.10.30.1
```

```
1
route-map PBR FTP permit 20 match ip address FTP2_R1 set ip next-
hop 10.10.40.1
```

```
1
Interface Gigabit Ethernet 0/0 ip policy route-map PBR FTP
```

D)

```
ip access-list extended FTP1_R1 permit ip host 10.10.10.11 any ip
access-list extended FTP2_R1 permit ip host 10.10.10.22 any route-
map PAR FTP permit 10 match ip address FTP1_R1 set ip next-hop
10.10.30.1
```

route-map PAR FTP permit 20 match ip address FTP2R1 set ip next-hop 10.10.40.1

interface GigabitEthernet W0 ip policy route-map PBR FTP

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 491

Which two technologies optimize MPLS infrastructure using bandwidth protection services when experiencing slow response? (Choose two.)

- A. IPLFA
- B. MPLS OAM
- C. VPLS
- D. SO-MPLS
- E. Fast-Route

Answer: A,E

Explanation:

Question: 492

Refer to the exhibit.

RIIshow time-range

time-range entry: timer (active)

periodic weekend 9:00 to 17:00

used in: IP ACL entry

used in: IP ACL entry

RKshow ip access-list interface gig0/0

Extended IP access list NO_Internet in

10 deny tcp any any eq www time-range timer (active) 20 deny tcp any any

eq 443 time-range timer (active) 30 permit ip any any



Refer to the exhibit. Users on a call center report that they cannot browse the internet on Saturdays during the afternoon. Which configuration resolves the issue?

A)

```
interface gig0/0  
ip access-group NO Internet out
```

B)

**ip access-list extended NOInternet 15 permit
tcp any any eq www**

C)

no lime-range timer

D)

time-range timer

**no periodic weekend 9:00 to 17:00 periodic
weekend 17:00 to 23:59**

A. Option A

B. Option B

C. Option C

D. Option D


Answer: D

Explanation:

Question: 493

Refer to the exhibit.

**Company XYZ
Network**



```
R101#show run | section ip sla 1
```

```
tcp-connect 2.2.2.2 3000 source-ip 1.1.1.1 threshold 1000
```

```
1AAA ip sla 2
```

```
icmp-jitter 2.2.2.2 source-ip 1.1.1.1 num-packets 100 interval 10
threshold 1000 timeout 1000 frequency 10
```

```
ip sla schedule 2 life forever start-time now R101#show ip sla summary
```

```
IPSLAs
Latest Operation Summary Codes: * active, A inactive, - pending
```

ID	Type	Destination	Stats (ma)	Return Code	Last Run
*1	tcp-connect	2.2.2.2	-	No connection	33 seconds ago
*2	icmp-jitter	2.2.2.2	RTT=4	OK	3 seconds ago

Refer to the exhibit While troubleshooting an issue on the network, an engineer notices that a TCP Connect operation failed on port 3000 between R101 and R201. Which command must be configured on R201 to respond to the R101 IP SLA configurations with a control connection on UDP port 1967?

A. ip sla responder udp-echo ipaddress 1.1.1.1 port 1967

B. ip sla responder tcp-connect ipaddress 1.1.1.1 port 3000

C. ip sla responder tcp-connect ipaddress 2.2.2.2 port 3001

D. ip sla responder

Answer: A

Explanation:

Question: 494

Refer to the exhibit.

```
ip access-list extended CoPP-ICMP
 permit icmp any any echo
!
ip access-list extended CoPP-BGP
 permit tcp any eq bgp any established
!
ip access-list extended CoPP-EIGRP
 permit eigrp any host 224.0.0.10
!
Class-map match-all CoPP-CLASS
 match access-group name CoPP-ICMP
 match access-group name CoPP-BGP
 match access-group name CoPP-EIGRP
!
```

Refer to the exhibit A CoPP policy is implemented to allow specific control traffic, but the traffic is not matching as expected and is getting unexpected behavior of control traffic. Which action resolves the issue?

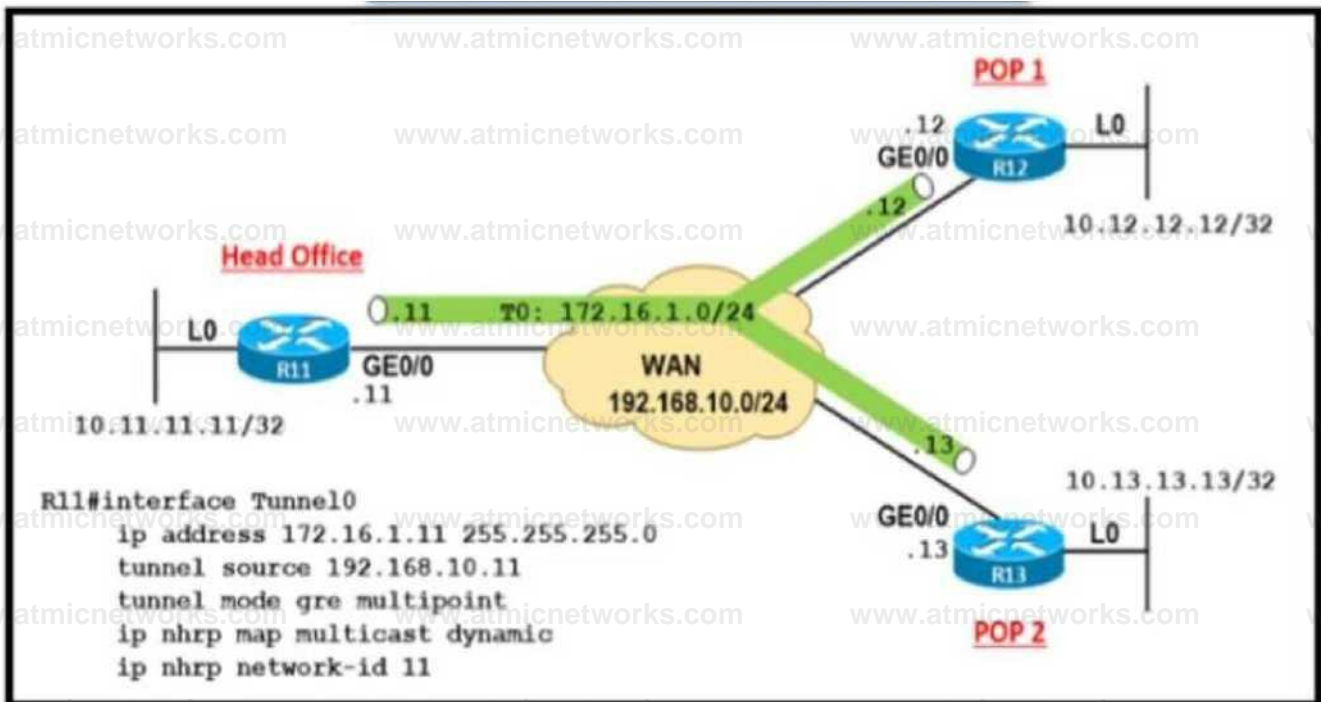
- A. Use match-any instruction in class-map
- B. Create a separate class map against each ACL
- C. Create a separate class map for ICMP traffic.
- D. Use default-class to match ICMP traffic

Answer: A

Explanation:

Question: 495

Refer to the exhibit.



Refer to the exhibit A company builds WAN infrastructure between the head office and POPs using DMVPN hub-and-spoke topology to provide end-to-end communication All POPs must maintain point-to-point connectivity with the head office Which configuration meets the requirement at routers R12 and R13?

ORI 2#

```
interface Tunnel0
```

```
ip nhrp map multicast 192.168.10.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 12
```

```
ip nhrp nhs 172.16.1.11
```

R13#

```
interface Tunnel0
```

```
ip nhrp map multicast 192.168.10.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 13
```

```
ip nhrp nhs 172.16.1.11
```

OR12#

```
interface Tunnel0
```

```
ip nhrp map multicast 172.16.1.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 12
```

```
ip nhrp nhs 192.168.10.11
```

R13#

```
interface Tunnel0
```

```
ip nhrp map multicast 172.16.1.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 13
```

```
ip nhrp nhs 192.168.10.11
```

Q Configure routers R12 and R13

as:

```
interface Tunnel0
```

```
ip nhrp map multicast 172.16.1.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 11
```

```
ip nhrp nhs 192.168.10.11
```

() Configure routers R12 and R13

as:

```
Interface Tunnel0
```

```
ip nhrp map multicast 192.168.10.11
```

```
ip nhrp map 172.16.1.11 192.168.10.11
```

```
ip nhrp network-id 11
```

```
ip nhrp nhs 172.16.1.11
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 496

Refer to the exhibit.

RF#traceroute 192.168.1.1


```
1 10.0.0.9 40 msec 28 msec 24 msec
2 * * *
3 * * *
```

RE#show ip prefix-list detail Prefix-list with the last deletion/insertion: Customer ip prefix-list Customer:

```
count: 2, range entries: 1, sequences: 5 - 10, reccount: 3
seq 5 deny 192.168.1.1/32 (hit count: 5, reccount: 1)
seq 10 permit 0.0.0.0/0 le 32 (hit count: 26, reccount: 1)
```

RC#show ip prefix-list detail Prefix-list with the last deletion/insertion: Customer ip prefix-list Customer:

```
count: 1, range entries: 1, sequences: 10 - 10, reccount: 4 seq 10 permit 0.0.0.0/0
le 32 (hit count: 7, reccount: 1)
```



Refer to the exhibit The enterprise users fail to authenticate with the TACACS server when a direct fiber link fails between RB and RD The NOC team observes

Users connected on AS65201 fail to authenticate with TACACS server 192 168 1 1

Users connected on AS65101 successfully authenticate with TACACS server 192 168 1 1

All AS65101 and AS65201 users are configured to authenticate with the TACACS server

Which configuration resolves the issue?

A)

```
RC(config)# ip prefix-list Customer seq 5 permit 192.168.30.1/32
```

B)

```
RC(config -router bgp 65101
```

```
RC(config-router)# neighbor 10.0.0.18 prefix-list Customer in
```

C)

```
RF(config)# ip prefix-list Customer seq 5 deny 192.168.1.1/32
```

D)

```
RF(config)# router bgp 65201
```

```
RF(config-router)# neighbor 10.0.0.17 prefix-list Customer out
```

A. Option A

B. Option B

C. Option C

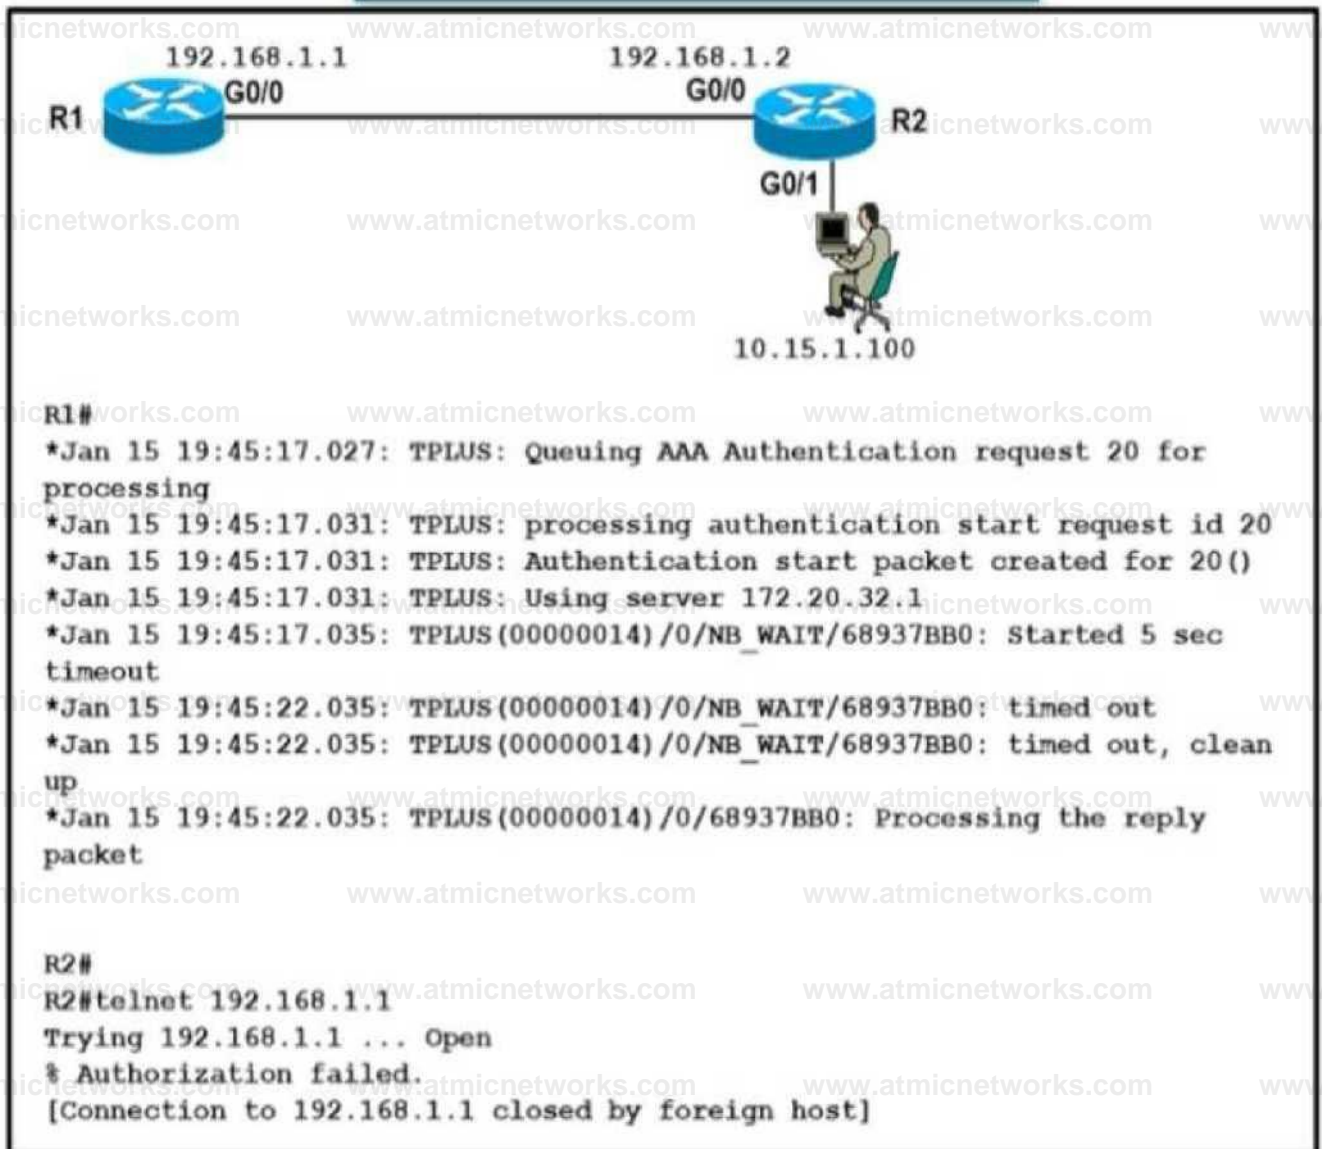
D. Option D

Answer: C

Explanation:

Question: 497

Refer to the exhibit.



Refer to the exhibit A network engineer is troubleshooting an AAA authentication issue for R1 from R2 When an engineer tries to open a telnet connection to R1 it opens the connection but shows a %Authorization failed error message on the terminal and closes the connection silently Which action resolves the issue?

- A. Resolve tacacs+ server host IP authentication miss configuration on the R1 router
- B. Resolve tacacs+ server reachability from the R1 router.
- C. Configure the tacacs+ server host IP on the R1 router

D. Configure authorization commands in the tacacs* server for the R1 router.

Answer: D

Explanation:

Question: 498

Refer to the exhibit.

R5#

```
•Sep 19 08:29:51.088: BGP: 10.10.10.2 open active, local address 10.0.0.14
```

```
*Sep 19 08:29:51.120: BGP: 10.10.10.2 read request no-op
```

```
•Sep 19 08:29:51.124: BGP: 10.10.10.2 open failed: Connection refused by  
remote host, open active delayed 12988ms (20000ms max, 60% jitter)
```

```
R2#show ip bgp neighbors 10.10.10.5
```

```
BGP neighbor is 10.10.10.5, remote AS 65101, internal link
```

```
BGP version 4, remote router ID 0.0.0.0
```

```
BGP state - Active
```

```
Last read 00:01:18, last write 00:01:18, hold time is 15, keepalive interval is 3  
seconds
```

```
Configured hold time is 15, keepalive interval is 3 seconds
```

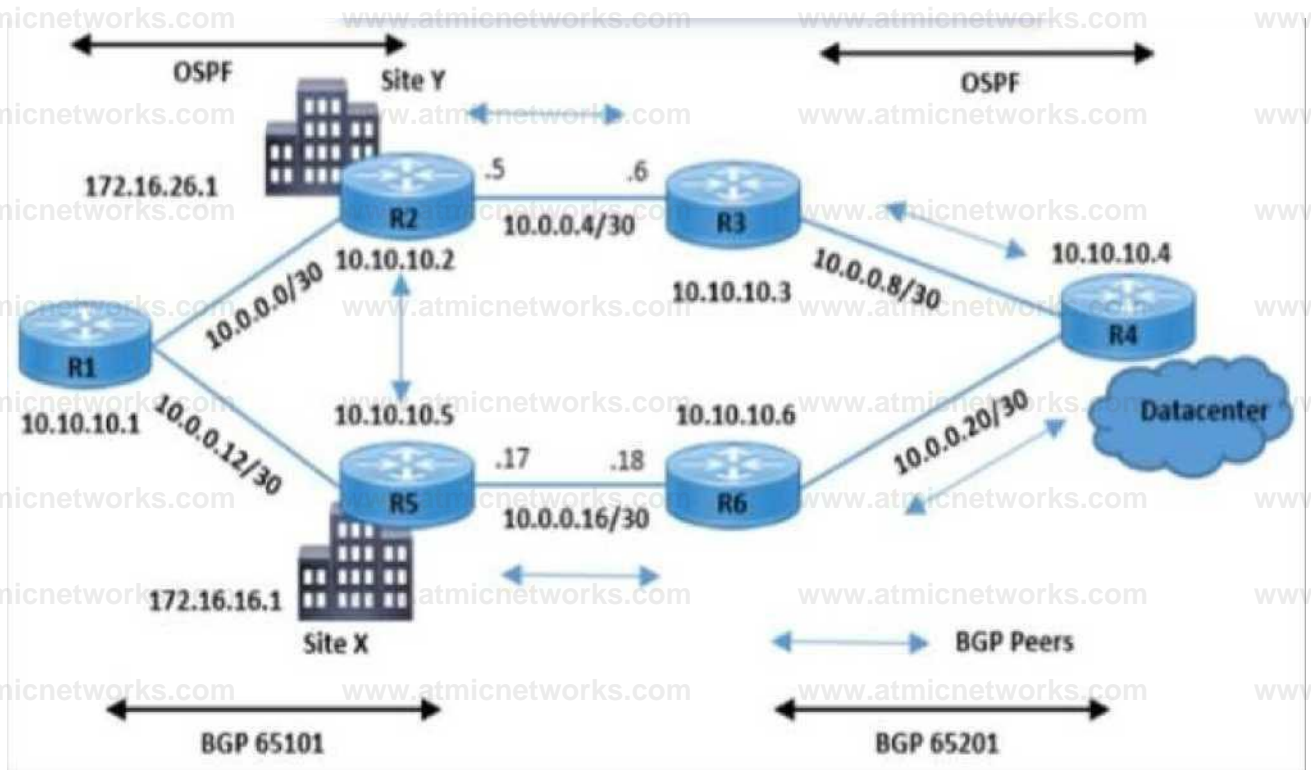
```
Minimum holdtime from neighbor is 0 seconds
```

```
Address tracking is enabled, the RIB does have a route to 10.10.10.5
```

```
Connections established 13; dropped 13
```

```
Last reset 00:01:18, due to User reset
```

```
Transport(tcp) path-mtu-discovery is enabled No active TCP connection
```



Refer to the exhibit A customer reported a failure and intermittent disconnection between two office buildings site X and site Y The network team finds that site X and site Y are exchanging email application traffic with the data center network Which configuration resolves the issue between site X and site Y?

A)

RC(config)s ip prefix-list Customer seq 5 permit 192.168.30.1/32

B)

RC(config -router bgp 65101

RCfconfig-routerp neighbor 10.0.0.18 prefix-list Customer in

C)

RF(configpno ip prefix-list Customer seq 5 deny 192.168.1.1/32

D)

RF(configrouter bgp 65201

RFconf# neighbor 10.0.0.17 prefix-list Customer out

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 499

Which router translates the customer routing information into VPNv4 routes to exchange VPNv4 routes with other devices through MP-BGP?

A. PE

B. CE

C. P

D. VPNv4 RR

Answer: A

Explanation:

Question: 500

Which router takes an active role between two LDP neighbors when initiating LDP session negotiation and LDP TCP connection establishment?

- A. with the higher IP address
- B. with the larger number of LDP TCP neighbors
- C. with the lowest IP address
- D. with one interface in the MPLS backbone

Answer: A

Explanation:

Question: 501

Which routing protocol is used by the PE router to advertise routes to a CE router without redistribution or static after removing the RD tag from the P router?

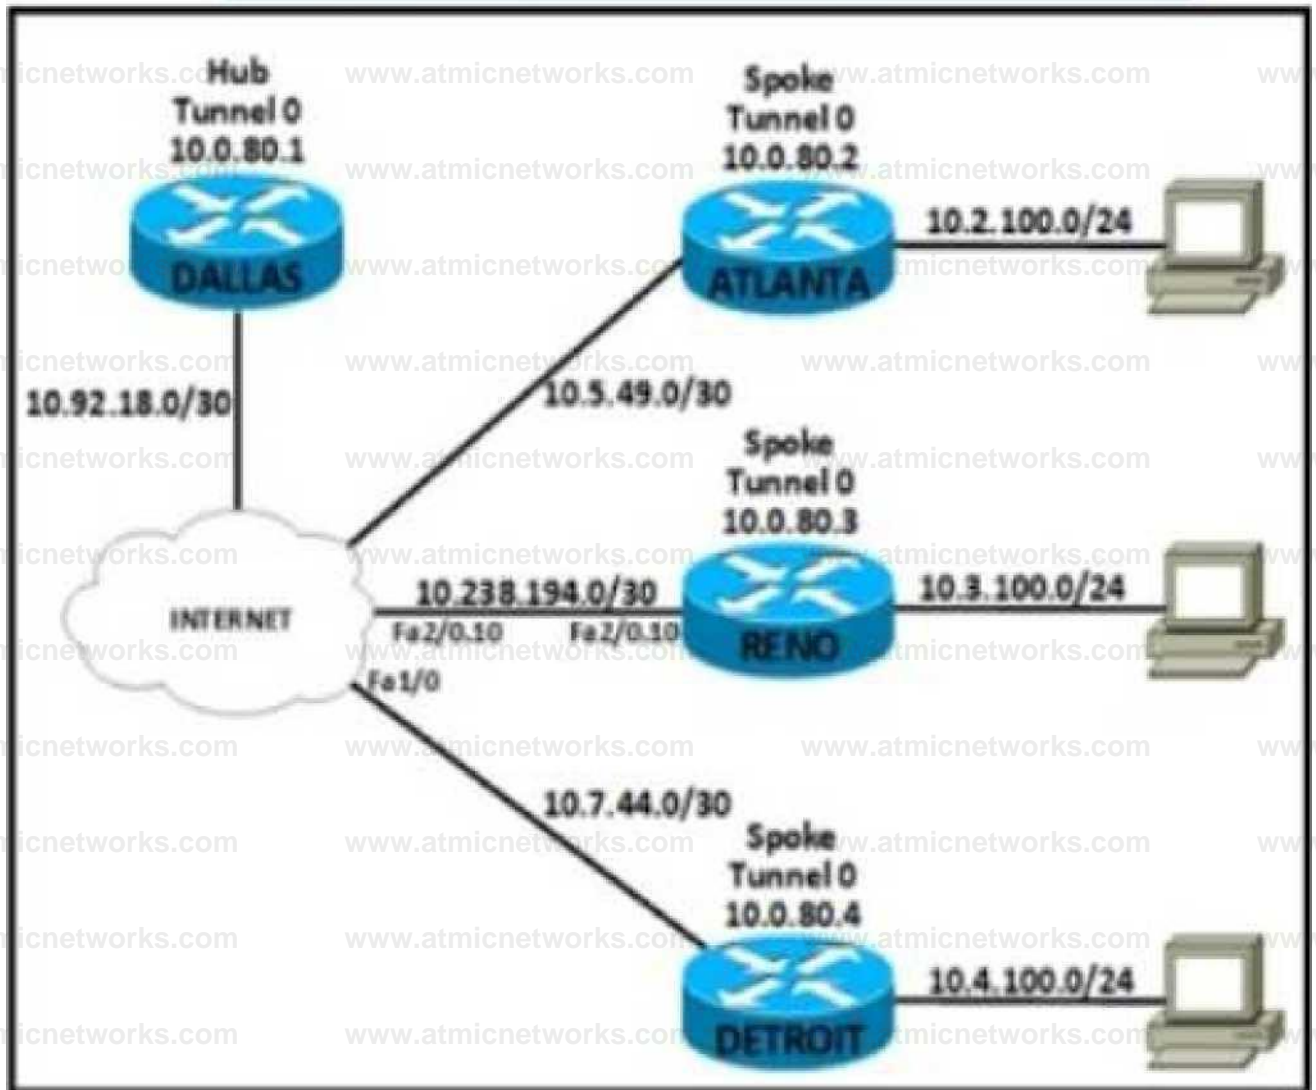
- A. IS-IS
- B. OSPF
- C. BGP IPv4
- D. MP-BGP

Answer: C

Explanation:

Question: 502

Refer to the exhibit.



Refer to the exhibit An engineer must connect the Reno and Detroit spokes using DMVPN phase 2 Hub tunnel configuration is

Dallas

interface Tunnel0

ip address 10.0.80 1 2 55 255 255 0

ip nhrp authentication cisco123 ip nhrp map multicast dynamic ip nhrp networked

5 tunnel source Sena 10 0 tunnel mode gre multipoint

Which configuration accomplishes the task?

Rgjo

interface Tunnel0

```
ip address 10 0 3 295 255 2 55 0 ip nhrp authentication  
cisco321 ip nhrp map multicast 10 92 18 2 ip nhrp map 10.0,80.1  
10.92.18.2 ip nhrp network-id 5 ip nhrp nhs 10.0.80.1 tunnel  
source 10.238 194 2 tunnel mode gre multipoint
```

Detroit

interface Tunnel0

```
ip address 10.0.8 0.4 255 255.255.0 ip nhrp authentication  
cisco321 ip nhrp map 10 0 80.1 10.92.18.2 ip nhrp map multicast  
10 92 18 2 ip nhrp network-id 5 ip nhrp nhs 10.0.80.1 tunnel  
source 10.7 44.2 tunnel mode gre multipoint
```

Reno

Interface Tunnel0

```
ip address 10 0.80.3 255 255.255.0
```

```
ip nhrp authentication JU On cisco 23
```

```
ip nhrp map multicast 10.92 18 2 ip nhrp map 10 92.18 2 10.0  
80 1
```

```
ip nhrp network-id 5
```

```
ip nhrp nhs 10.0.80.1
```

```
tunnel source 10 238.194.2
```

```
tunnel mode gre multipoint
```

Detroit

interface Tunnel0

```
ip address 10.0.80.4 2 55.2 55 2 55.0
```

```
ip nhrp authentication cisco 23 ip nhrp map 10 92.18 2 10.0 80  
1 ip nhrp map multicast 10.92 18 2
```

```
ip nhrp network-id 5
```

```
ip nhrp nhs 10.0.80,1 tunnel source 10 7.44.2 tunnel mode gre  
multipoint
```

Reno

interface Tunnel0

```
ip address 10 0.80.3 255 255 255 0 ip nhrp authentication  
cisco!23 ip nhrp map broadcast 10 9114.2 ip nhrp map  
10.0.80.1 10.92.18.2 ip nhrp networked 8 ip nhrp nhs  
10.0.80.1 tunnel source 10.238 194.2 tunnel mode gre  
multipoint
```

Detroit

interface Tunnel0

```
ipaddress 10 0 80 4 265 255 255 0 ip nhrp authentication  
cisco123 ip nhrp map 100.80.1 10.92.18 2 ip nhrp map  
broadcast 10.92 18 2 ip nhrp networked 5 ip nhrp nhs  
10.0.80.1 tunnel source 10.7.44.2 tunnel mode gre multipoint
```

Reno

```
interface Tunnel0
```

```
ip address 10.0.80.3 255.255.255.0
```

```
ip nhrp authentication cisco 123 ip nhrp map multicast 10.92.18.2  
ip nhrp map 10.0.80.1 10.92.18.2
```

```
ip nhrp networked
```

```
ip nhrp nh 10.0.80.1 tunnel source 10.238.194.2 tunnel mode  
pre multiport
```

Detroit

```
interface Tunnel0 ip address 10.0.30.2 255.255.255.0 ip nhrp  
authentication cisco 123 ip nhrp map 10.0.80.1 10.92.18.2 ip  
nhrp map multicast 10.92.18.2
```

```
ip nhrp network-ld 5 ip nhrp nh 10.0.80.1 tunnel source 10.7.44.2  
tunnel mode gre multipoint
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 503

Which control plane process allows the MPLS forwarding state to recover when a secondary RP takes over from a failed primary RP?

- A. MP-BGP uses control plane services for label prefix bindings in the MPLS forwarding table
- B. LSP uses NSF to recover from disruption in control plane service
- C. FEC uses a control plane service to distribute information between primary and secondary processors
- D. LDP uses SSO to recover from disruption in control plane service

Answer: C

Explanation:

Question: 504

What must a network architect consider for RTs when planning for a single customer full-mesh VPN in an MPLS Layer 3 network?

- A. RT must be globally unique within the same VPN
- B. RT must be globally identical within the same VPN
- C. RT values must be different from the RD values in the same VPN
- D. Each RT value must be identical to an RD value within the same VPN.

Answer: D

Explanation:

Question: 505

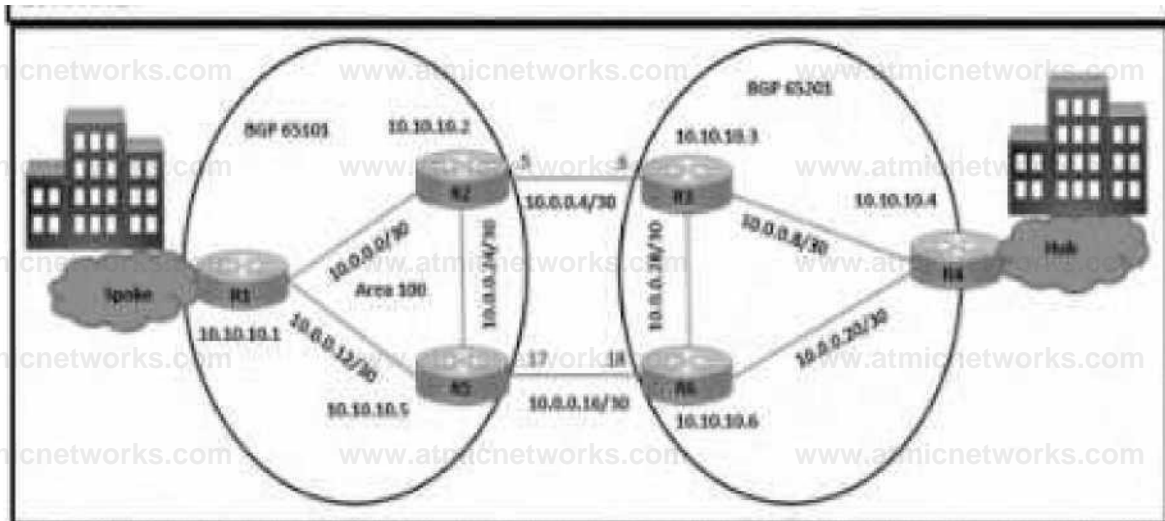
Refer to the exhibit.



```
'tep : 435.3150 091'HP 10 4 Q.17 <<■ t (tw Hua fa JUL1Y*
'tip 1 cs 3i si iH: KJ. lo.o.o.i? open attiva i<ai eitiui io 0.0 it
♦Sep 8 08'31:58 907; BGP ICO 017 reed request ne-op
+tep 9 05 31 50 911 KJ 10.0 0.17 went 3w Active :4 patent
'Sep 5 05:31SB 911BG? 10 0.0.11 lending OFW veriion 4, my as 65301, hoidfine
IM tecondi
'Sep S 05:3158.111: Kf 10,0,0,17 send nemge type 1 Length Unci header: 53
'Sep 1 aJ:11.4L.>?. KJ 10.0.0.11 rate glean
'Sep ■ 05:31 58 MI: HP 10.0,0,1? -qpet th* Matin
'Sep S 08:31:51.831: WWMSr abate: 10.0.0.17 went few. nafnot active fa
ml net active
```

Ki*

```
■tep 5 <3< ?1 033: MT 10.0 0 10 paeiiv* *p*n fa 10 5 0 17
'Sep 3 05'34:21 M3; BGJ 10.0.4.IB pairlve open Called - 13 C 0.17 La net ^iate-
tOUrca Inept*; ts I aiifeaa L14. IQ. 10 II
•Mp 1 05.34 22 DM: Hr 10.0.0.IB mate CMnectlaa attest failed. lw*1 addaaet 10.0.0.17
```



Refer to the exhibit. The traffic from spoke to hub is dropping. The operations team observes:

R2-R3 link is down due to the fiber cut.

R2 and R5 receive traffic from R1 in AS 65101.

R3 and R5 receive traffic from R4 in AS 65201.

Which configuration resolves the issue?

A)

```
R6(config)#router bgp 65101
R6(config-router)#no neighbor 10.0.0.17 update-source Loopback0
```

B)

```
R5(config)#router bgp 65101
R5(config-router)#no neighbor 10.0.0.18 update-source Loopback0
```

C)

```
R6(config)#router bgp 65201
R6(config-router)#neighbor 10.10.10.5 remote-as 65101
R6(config-router)#neighbor 10.10.10.5 update-source Loopback0
R6(config-router)#neighbor 10.10.10.5 ebgp-multihop 3
```

D)

```
R5(config)#router bgp 65101
R5(config-router)#neighbor 10.10.10.6 remote-as 65201
R5(config-router)#neighbor 10.10.10.6 update-source Loopback0
R5(config-router)#neighbor 10.10.10.6 ebgp-multihop 3
```

A. Option A

B. Option B

C. Option C

D. Option D


Answer: C

Explanation:

Question: 506

Refer to the exhibit.

```
R2#show ip eigrp topology 10.10.10.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.10.10.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD
  is 256005120
  Routing Descriptor Blocks:
  10.20.20.3 (FastEthernet0/1), from 10.20.20.3, Send flag is
  0x0
    Composite metric is (256005120/256002560), Route is
  External
  Vector metric:
    Minimum bandwidth is 10 Kbit
    Total delay is 200 microseconds
    Reliability is 10/255
    Load is 10/255
    Minimum MTU is 10
    Hop count is 1
  External data:
    Originating router is 10.10.10.1
    AS number of route is 1
    External protocol is OSPF, external metric is 0
    Administrator tag is 0 (0x00000000)
R1#sh run | s eigrp
router eigrp 1
router-id 10.1.1.1
network 10.2.2.0 0.0.0.255
no auto-summary
```



Refer to the exhibit. An engineer configured router R3 to redistribute the prefix 10.10.10.0/24 from OSPF into EIGRP. R1 has no connectivity to the prefix. Which action enables receipt of prefixes on R1?

- A. R3 is advertising the 10.20.20.0/24 prefix with a TTL of 1, R3 must set the TTL to 2 for this prefix.
- B. R1 does not have a neighbor relationship with R2. The EIGRP process should be cleared on R1.
- C. Duplicate router IDs on R1 and R3, R1 should modify its router ID.
- D. R1 is not receiving the next-hop IP address of R3. R2 must enable the network 10.20.20.0/24 within EIGRP.

Answer: B

Explanation:

Question: 507

Refer to the exhibit.

ConMiaii

```
flow exporter Flow to collector destination 192.160.100.17 vrf Mgmt-in tf transport udp 2601 export*protocol natflow-v5

flow monitor My-netflow
exporter Flow-to-collector record netflow ipv4 original -input

' end the management-interface it configured as follows:
interface GigabitEthernet0
description Management*Interface vrf forwarding Mgmt-intf
ip address 192.160.100.50 255.255.255.0 negotiation auto

rduer*sh flow exporter static j
Flow Exporter Flow-to-c...ector;
```

Packet (end statistics (last cleared w^d ago): Successfully sent: 0 (0 bytes)

Reason not given: 8696868 (11471678976 bytes)

Client send Statistics:

Client: Flow Monitor OeKB-netflow

Records added: 256781312
- failed to send: 256783312
Bytes added: 2783766194
- failed to send: 2783766384

router r

Refer to the exhibit. A network administrator configured NetFlow data, but the data is not visible at the NetFlow collector.

Which configuration allows the router to send the records?

- A. Configure the management interface in the global routing table to send the records.
- B. Configure a different interface to send the records.
- C. Configure the NetFlow collector to listen at export-protocol netflow-v5.

D. Rectify NetFlow collector reachability from the management interface.

Answer: B

Explanation:

Question: 508

DRAG DROP

An engineer must establish a connection between two CE routers for two customers with overlapping IP addresses

Customer_a is connected to interfaces Gig0/0, and Customer_b is connected to interfaces Gig0/1. Routers CE1 and CE2 are configured as follows:

```
ip vrf customer a
rd 1:1
route-target both 1:1
```

```
ip vrf customerb
rd 2:2
route-target both 2:2
```

Drag and drop the code snippets from the right onto the boxes in the configuration to establish the needed connection. Snippets may be used more than once.

```
CE1
interface Gig0/0
ip vrf forwarding [redacted]
ip address [redacted]
!
interface Gig0/1
ip vrf forwarding [redacted]
ip address [redacted]

CE2
interface Gig0/0
ip vrf forwarding [redacted]
ip address [redacted]
!
interface Gig0/1
ip vrf forwarding [redacted]
ip address [redacted]
```

- customer_a
- customer_b
- 192.168.1.1 255.255.255.0
- 192.168.1.2 255.255.255.0

Answer:

Explanation:

```
CE1
interface Gig0/0
 ip vrf forwarding customer_a
 ip address 192.168.1.1 255.255.255.0
!
interface Gig0/1
 ip vrf forwarding customer_b
 ip address 192.168.1.2 255.255.255.0

CE2
interface Gig0/0
 ip vrf forwarding customer_a
 ip address 192.168.1.1 255.255.255.0
!
interface Gig0/1
 ip vrf forwarding customer_b
 ip address 192.168.1.2 255.255.255.0
```

Question: 509

A network administrator opens a telnet connection to the router and gets the message:

R1#telnet 10.1.1.2

Trying 10.1.1.2 Open

(Connection to 10.1.1.2 closed by foreign host)

Router R2 is configured with enable secret and password commands. Which action resolves the issue?


- A. Configure the logging synchronous command on line vty.
- B. Configure the exec command on line vty.
- C. Configure the login local command on line vty
- D. Configure the enable password command on line vty.

Answer: C

Explanation:

Question: 510

Refer to the exhibit.



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 by sending traffic to ISP-A only with failover to ISP-B?

- A. Configure OSPF between R1, ISP-A, and ISP-B for dynamic failover if any ISP link to R1 fails
- B. Configure two static routes on R1, one pointing to ISP-A and another pointing 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

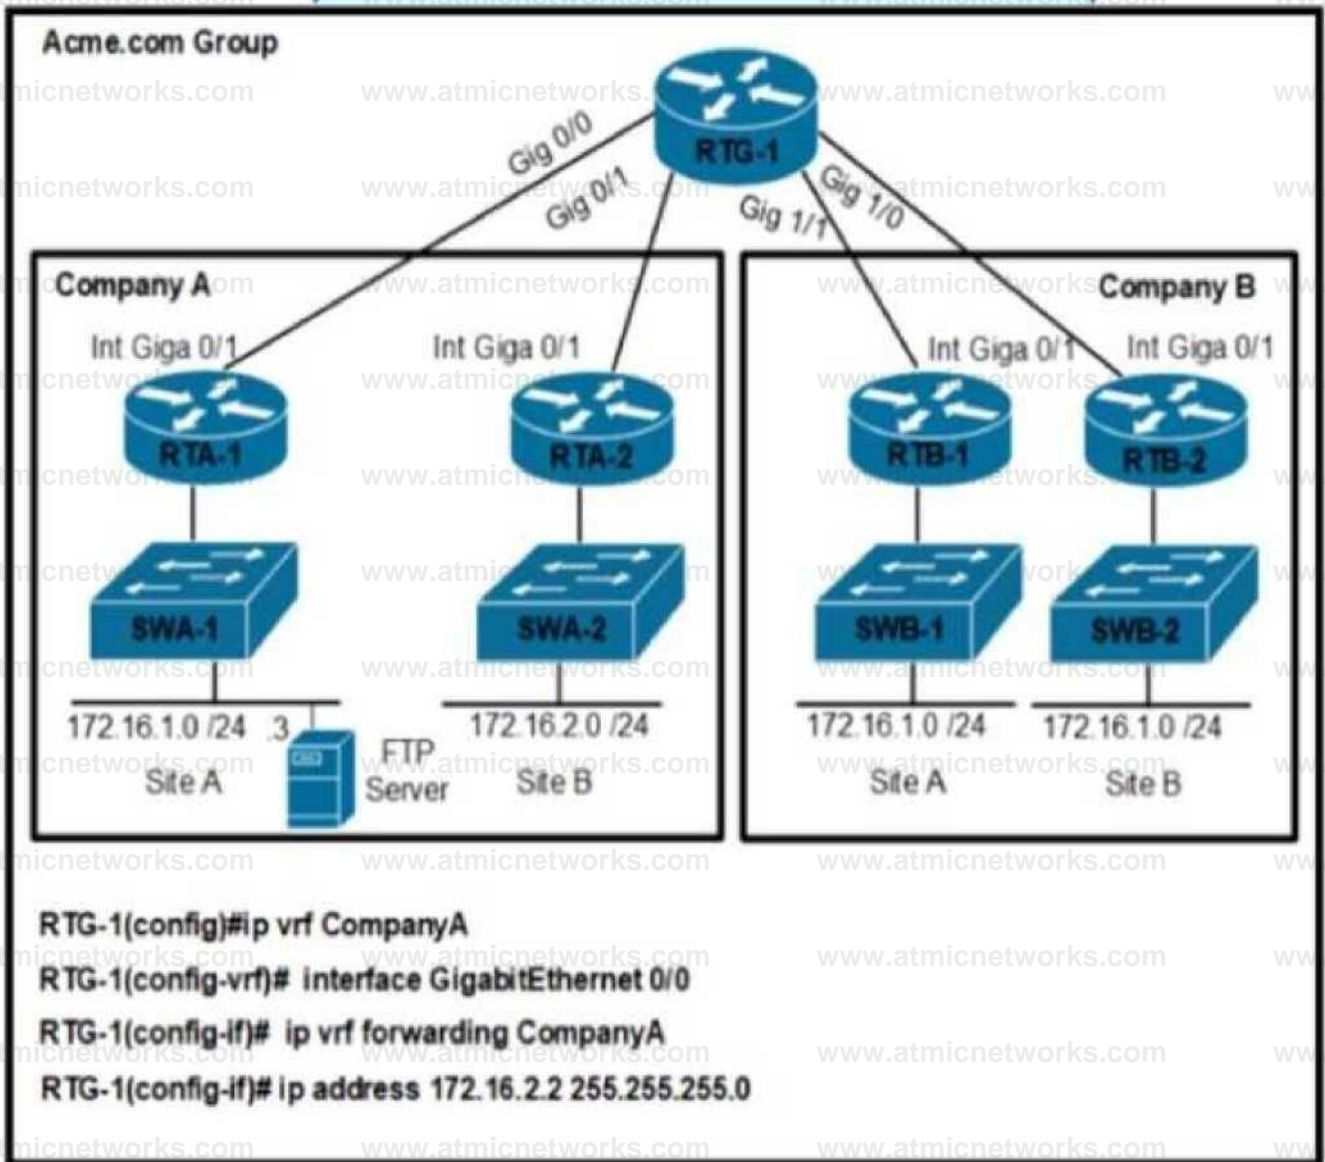
D. Configure two static routes on R1, one pointing to ISP-B with more specific routes and another pointing to ISP-A with summary routes

Answer: D

Explanation:

Question: 511

Refer to the exhibit.



Refer to the exhibit. An engineer must configure a per VRF for TACACS+ for company

A. Which configuration on RTG-1 accomplishes the task?

```

aaa new-model
aaa group server tacacs+ Tacacscluster server-
private 172.16.11 port 49 key routing ip tacacs
source-interface GigabitEthernet 0/0 Ip vrf forwarding
CompanyA

```

Oaaa new-model

```

aaa group server tacacs* Tacacscluster server-private

```

172.16.1.3 port 49 key routing ip tacacs source-
interface GigabitEthernet 0/1 ip vrf forwarding
CompanyA

aaa new-model

aaa group server tacacs* Tacacscluster server-private
172.16.11 port 49 key routing ip tacacs source-
interface GigabitEthernet 0/1 ip vrf CompanyA

0 aaa new-model

aaa group server tacacs* Tacacscluster server-
private 172.16.13 port 49 key routing ip tacacs
source-interface GigabitEthernet 0/0 ip vrf
CompanyA

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

Explanation:

Question: 512

A company is redesigning WAN infrastructure so that all branch sites must communicate via the head office and the head office can directly communicate with each site independently. The network engineer must configure the head office router by considering zero-touch technology when adding new sites in the same WAN infrastructure. Which configuration must be applied to the head office router to meet this requirement?

Interface Tunnel0 tunnel mode ip ip nhrp map multicast dynamic
interface Tunnel0 tunnel mode dvmrp ip nhrp redirect
Interface Tunnel0 tunnel mode Ip ip nhrp redirect
interface Tunnel0 tunnel mode gre multipoint Ip nhrp map
multicast dynamic

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

Question: 513

Which protocol must be secured with MD-5 authentication across the MPLS cloud to prevent hackers from introducing bogus routers?

- A. MP-BGP
- B. LSP
- C. RSVP
- D. LDP

Answer: A

Explanation:

Question: 514

Which technique removes the outermost label of an MPLS-tagged packet before the packet is forwarded to an adjacent LER?

A. label swap

B. explicit-null

C. label imposition

D. PHP


Answer:

D

Explanation:

Question: 515

Refer to the exhibit.



```

R7#sh Ip ospf Interface fa0/0
FastEthernet 0/0 la up, line protocol la up Internet Address 10.19.6.2/28, Area
0, Attached via Network Statenant
Process ID 1, Router ID 10.7.7.255, Network Typo POINT TO POINT, Cost: 1
Topology MTTD Cost Disabled Shutdown Topology Nam
0 1 no no Base

```

Refer to the exhibit. Router R4 is configured correctly with default OSPF values. A network engineer configured R7 for OSPF. R7 must not be elected as a DR for the segment between R4-R7. The adjacency between R4 and R7 failed to form.

Which configuration resolves the issue?

```

R7(config) ["interface fa0/0
R7(config-if)#ip ospf priority 255
R7(config-if)#ip ospf hello-Interval 10
R7(config-if)#ip ospf dead-Interval 30
R7(config-if)#ip ospf network broadcast

```



```
R7(config)#interface fa0/0
R7(config-if)#ip ospf priority 0
R7(config-if)#ip ospf hello-interval 10
R7(config-if)#ip ospf dead-interval 30
R7(config-if)#ip ospf network non-broadcast
```

```
R7(config[*]interf ace fa0/0
R7{config-if)#ip ospf priority 0
R7{config-if)#ip ospf hello-Interval 10
R7(config-if)#ip ospf dead-interval 40
R7(config-if)#ip ospf network broadcast
```

```
R7(config [-interlace fa0/0
R7{config-if)#ip ospf priority 255
R7| config-if )Mp ospf hello-Interval 10
R7(config-if)#ip ospf dead-interval 40
R7(config-if)#ip ospf network non-broadcast
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 516

Refer to the exhibit.

```
R1#show bgp ipv6 unicast 2001:db8::1/128
BGP routing table entry for 2001:db8::1/128, version 3
Paths: (1 available, best #1, table Global-IPv6-Table)
Not advertised to any peer
Local
  2001:db8:33:33::33 (metric 128) from 2001:db8:11:11::11 (1.1.1.1)
  Origin IGP, metric 0, localpref 100, valid, internal, best
  Originator: 3.3.3.3, Cluster list: 1.1.1.1
```

Refer to the exhibit. An engineer examines the BGP update for the IPv6 prefix 2001:db8::1/128, which should have been summarized into a /64 prefix. Which sequence of actions achieves the summarization?

- A. R1 is a route reflector client of a RR with a router ID of 1.1.1.1, and the originator of the prefix has a router ID of 3.3.3.3. Both routers belong to different ASs. The prefix is not advertised to any peer and must be advertised using the network statement on R3.
- B. R1 is a route reflector with a router ID of 3.3.3.3, and the originator of the prefix is a route reflector client, which has a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 1.1.1.1 for the prefix.
- C. R1 is a route reflector with a router ID of 1.1.1.1, and the originator of the prefix is a route reflector client, which has a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 1.1.1.1 for the prefix.
- D. R1 is a route reflector client of a RR with a router ID of 1.1.1.1, and the originator of the prefix has

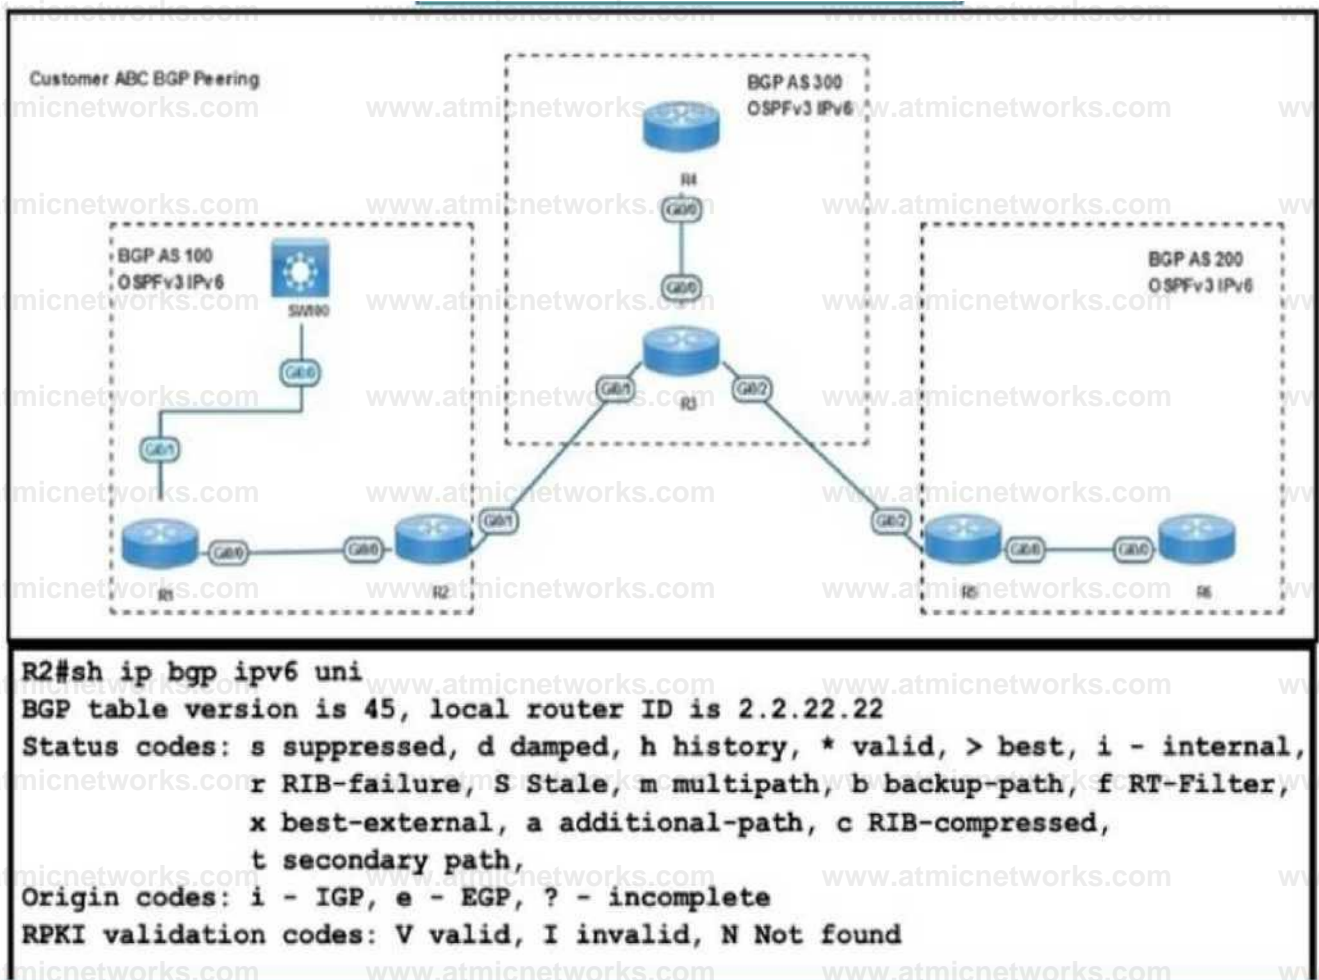
a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 3.3.3.3 for the prefix.

Answer: D

Explanation:

Question: 517

Refer to the exhibit.



```

t secondary
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

  Network          Next Hop           Metric LocPrf Weight Path
 *> 2001::5/128    2001::5            0         0 300 200 i
 *> 2001::4/128    2001::4            0         0 300 i
 *> 2002::2/128    ::                0        32768 i

R2#sh run | section bgp
router bgp 100
  address-family ipv6
  neighbor 2001::4 route-map Filter in

  ip as-path access-list 1 permit _300_[0-9]

  route-map Filter permit 10
  match as-path 1

```

Refer to the exhibit R2 has been receiving routes from R4 that originated outside AS300 A network engineer configured an AS-Path ACL to avoid adding these routes to the R2 BGP table but the routes are still present in the R2 routing table

Which action resolves the issue?

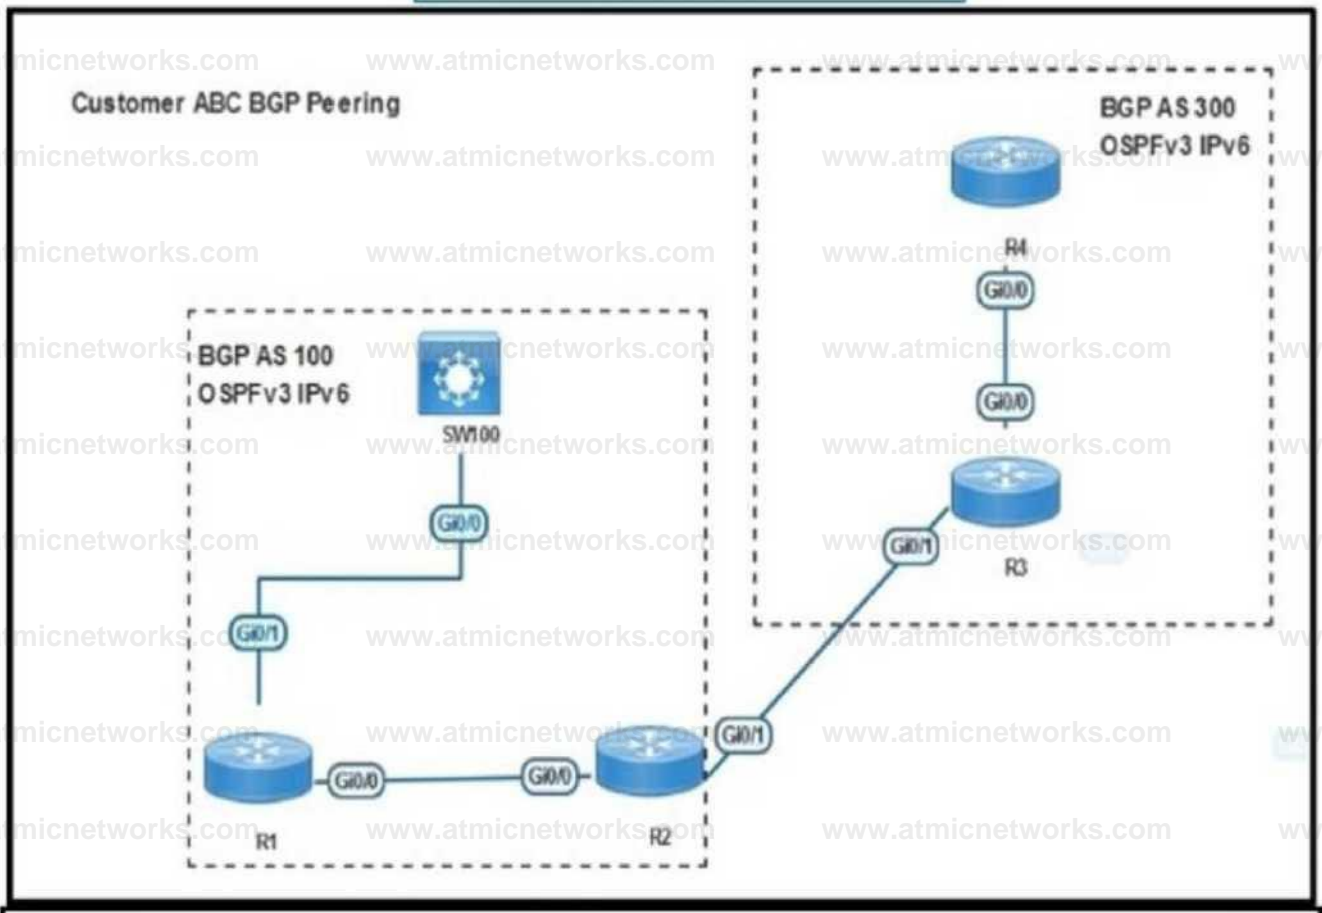
- A. Replace as-path access-list 1 with their as-path access-list 1 permit A300\$ command
- B. Replace as-path access-list 1 with their as-path access-list 1 permit ".300." command
- C. Replace as-path access-list 1 with their as-path access-list 1 permit A300_ command.
- D. Replace as-path access-list 1 with their as-path access-list 1 permit A300." command

Answer: B

Explanation:

Question: 518

Refer to the exhibit.



```
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:AAAI
          4      100        6      5      100 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?

(Ri

```
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
```

O_{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
```

RI
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

R1
router bgp 100
address-family Ipv6 redistribute connected

R2
router bgp 100 address-family ipv6
redistribute connected

A. Option A

B. Option B

C. Option C

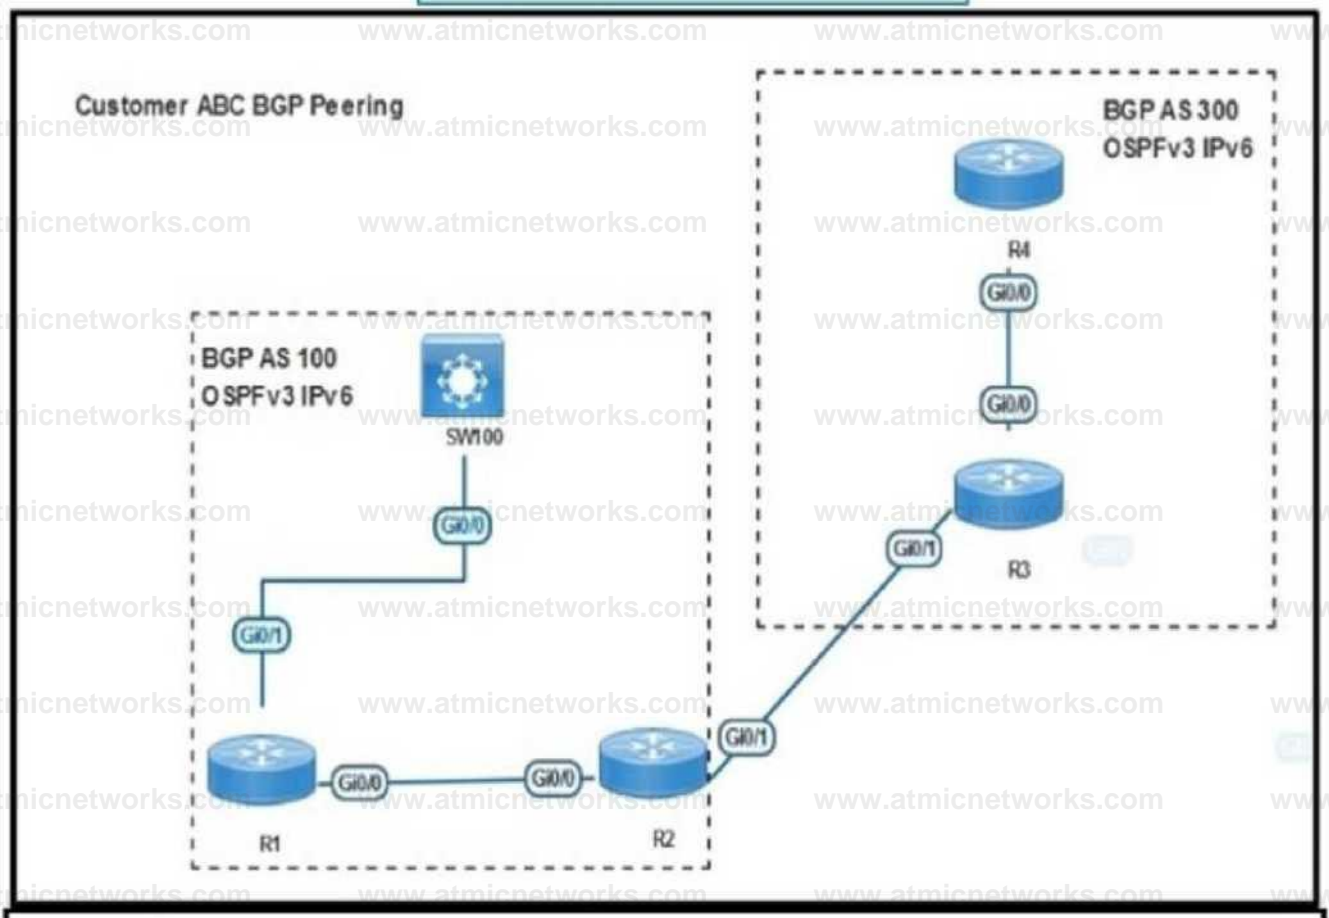
D. Option C

Answer:
A

Explanation:

Question: 519

Refer to the exhibit.



```
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:AAAI
          4      100      6      5      100 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  WeightPath
* 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
0 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?

(Ri

```
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
```

O_{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
```

RI
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

R1
router bgp 100
address-family Ipv6 redistribute connected

R2
router bgp 100 address-family ipv6
redistribute connected

A. Option A

B. Option B

C. Option C

D. Option C

Answer: A

Explanation:

Question: 520

Refer to the exhibit.

- Sep 3 231821 264EIGRP Neighbor (10 1 2 192) not yet found
- Sep 3 23 1918 675Going down Peer 10121 total=2 stub 0, iidb stub-0 ndall-0
- Sep 3 23 19 18 675 EIGRP Handle deallocation failure [1]
- Sep 3 23 1918 675EIGRP Neighbor 10 1 2 1 went down on Tunnell
- Sep 3 231922 943EIGRP New peer 10 1 2 1
- Sep 3 23 19 22 943 %DUAL 5 NBRCHANGE EIGRP IPv4 3111 Neighbor 101 21 (Tunnell) is up new adjacency

Refer to the exhibit. Which configuration command establishes an EIGRP neighbor adjacency between the hub and spoke?

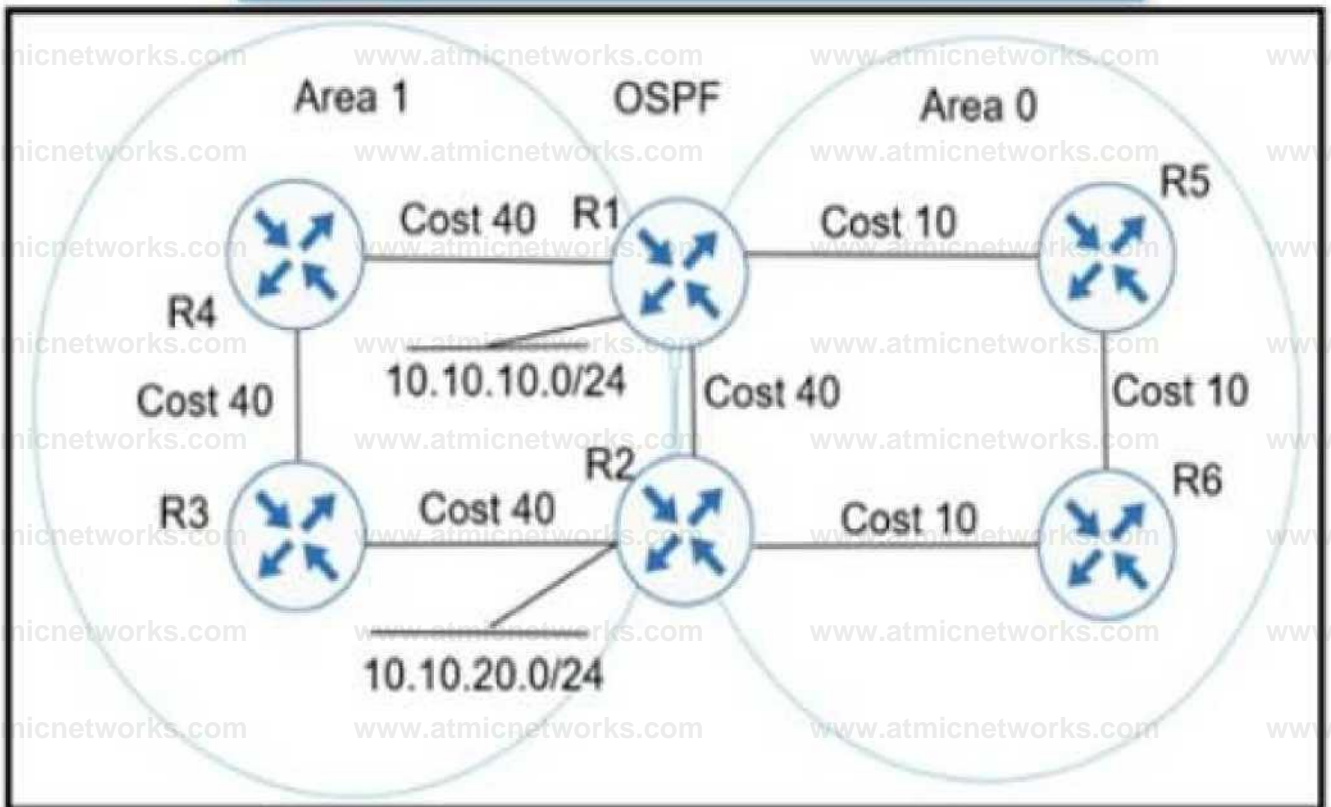
- A. connected 10.1.2.192 command on spoke router
- B. network 10.1.2.192 command on spoke router
- C. eigrp-peer 10.1.2.192 command on the hub router
- D. neighbor 10.1.2.192 command on hub router

Answer: D

Explanation:

Question: 521

Refer to the exhibit.



Refer to the exhibit Which action ensures that 10 10 10 0/24 reaches 10 10 20 0/24 through the direct link between R1 and R2?

- A. Configure R1 and R2 LAN links as nonpassive.
- B. Configure R1 and R2 links under area 1
- C. Configure OSPF link cost to 1 between R1 and R2
- D. Configure OSPF path cost to 3 between R1 and R2

Answer: B

Explanation:

Question:
522

Refer to the exhibit.

```
March 10 19:28:53.254 GMT: %SNMP-3-AUTHFAIL: Authentication
failure for SNMP request from host 10.1.1.1

snmp-server community public RO 15
snmp-server community private RW 16
!
logging snmp-authfail
!
access-list 15 permit 10.1.1.1
access-list 16 permit 10.1.1.2
```

Refer to the exhibit Which action resolves the issue?


- A. Configure host IP address in access-list 16
- B. Configure SNMPv3 on the router
- C. Configure SNMP authentication on the router
- D. Configure a valid SNMP community string

Answer:
D

Explanation:

Question: 523

Refer to the exhibit.



Refer to the exhibit After an engineer configured a new Cisco router as a DHCP server, users reported iwo primary issues:

Devices in the HR subnet have intermittent connectivity problems.

Workstations in the LEGAL subnet cannot obtain IP addresses.

Which configurations must the engineer apply to ROUTER_1 to restore connectivity for the affected devices?

```
interface GigabitEthernetO/O.5
 encapsulation dot1Q 5
 ip address 192.168.5.10 255.255.255 0
 ip helper-address 192.168.39.100
```

```
interface GigHbitEthernetO/OBO
 encapsulation dotlQ 80
 ip address 192.168.80.10 255 255,255.128
 ip helper-address 192 168.39.100
```


ip dhcpexcluded-address 192.168.5.1 192.168.5.10
ip dhcp excluded-address 192.168.80.1 192.168.80.10

ip dhcp pool LEGAL

network 192.168.80.0 255.255.255.128

default-router 192.168.80.10

ip dhcp pool HR

network 192.168.5.0 255.255.255.0

default-router 192.168.5.10

interface GigabitEthernet0/0.5

encapsulation dot1q 5

ip address 192.168.5.10 255.255.255.0

ip helper-address 192.168.39.100

interface GigabitEthernet0/0.80

encapsulation dot1q 80

ip address 192.168.80.10 255.255.255.128

ip helper-address 192.168.39.100

ip dhcpexcluded-address 192.168.80.1 192.168.80.10

ip dhcp pool LEGAL

network 192.168.80.0 255.255.255.128

default-router 192.168.80.10

ip dhcp pool HR

interface GigabitEthernet0/0.5

encapsulation dot1q 5

ip address 192.168.5.10 255.255.255.0

ip helper-address 192.168.39.100

interface GigabitEthernet0/0.80

encapsulation dot1q 80

ip address 192.168.80.10 255.255.255.128

ip helper-address 192.168.39.100

1

ip dhcp excluded-address 192.168.5.1 192.168.5.10

ip dhcp excluded-address 192.168.80.1 192.168.80.10

ip dhcp pool LEGAL

network 192.166.80.0 2 55 255.255.128

default-router 192.168.80.10

ip dhcp pool HR network 192.168.5.0 25 5.2 55.2 55.0 default-router

192.168.5.10

interface Gigabit EtherneW/0.5

encapsulation dot1Q 5

ip address 192.168.5.10 255.255 2 55 0

Ip helper-address 192.168 3 9 100

interface GigabitEthernet0/0.80

encapsulation dot1Q 80

ip address 192.168.80.10 255 255 25 5J28

Ip helper-address 192 168 39 100

ip dhcp excluded-address 192.168,5,1 192.168.5.5

ip dhcp excludcd-address 192.168.80 1 192.168.80,110

ip dhcp pool LEGAL

network 192.168.80.0 255.255.255.128

default-router 192,168.80.10

ip dhcp pool HR

A. Option A

B. Option B

C. Option C

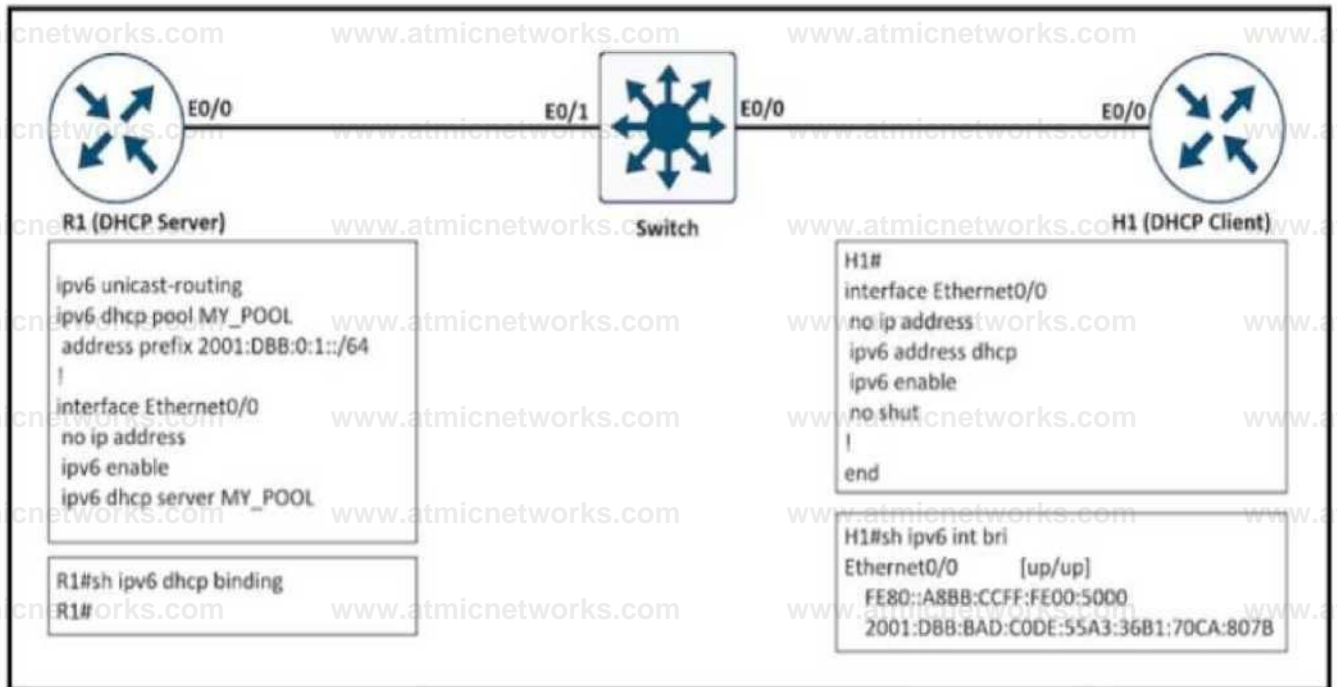
D. Option D

Answer: A

Explanation:

Question: 524

Refer to the exhibit.



Refer to the exhibit. The client server but the show command does not show the IPv6 DHCP bindings on the server. Which action resolves the issue?

- A. Extend the DHCP lease time because R1 removed the IPv6 address earlier after the lease expired.
- B. Configure H1 as the DHCP client that manually assigns the IPv6 address on interlace e0/0.
- C. Use the 2001:DBB:BAD:CODE::/64 prefix for the DHCP pool on R1.
- D. Configure authorized DHCP servers to avoid IPv6 addresses from a rogue DHCP server.

Answer: C

Explanation:

Question: 525

What is a MPLS PHP label operation?

- A. Downstream node signals to remove the label.
- B. It improves P router performance by not performing multiple label lookup.
- C. It uses implicit-NUL for traffic congestion from source to destination forwarding
- D. PE removes the outer label before sending to the P router.

Answer: A

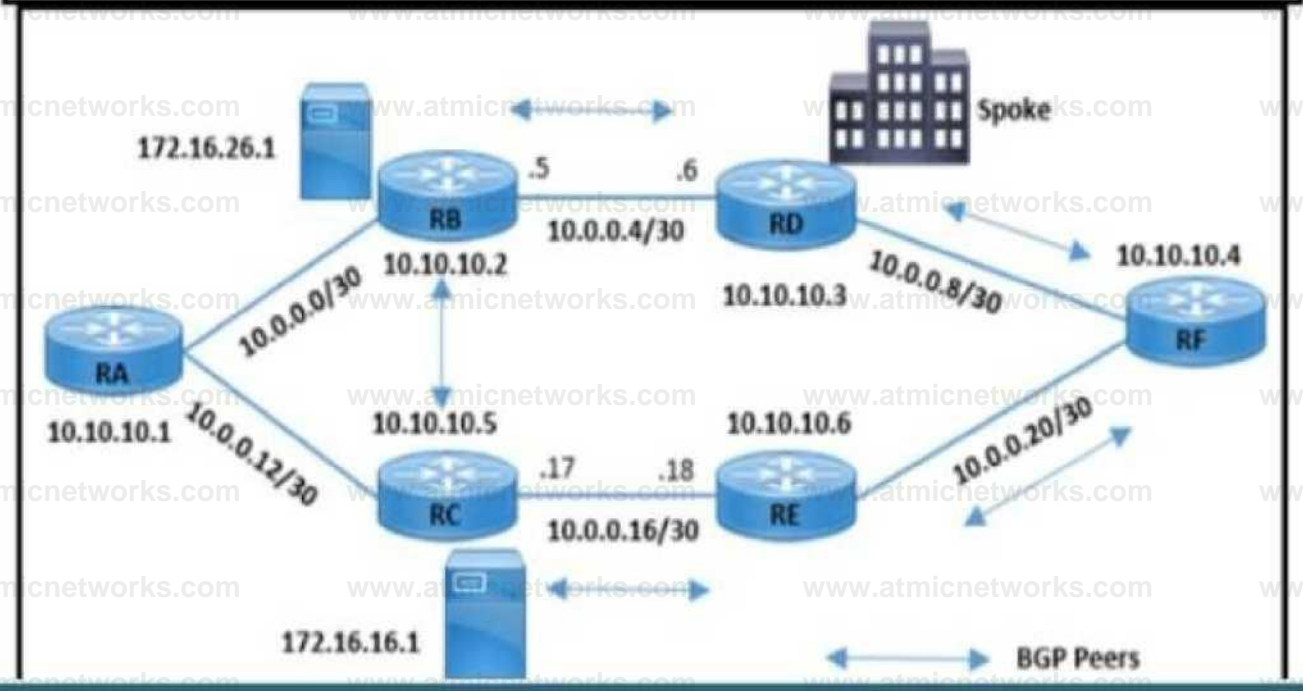
Explanation:


Question: 526

Refer to the exhibit.

```
RB#show ip bgp 172.16.16.1
BGP routing table entry for 172.16.16.1/32, version 11
Paths: (1 available, no best path)
Not advertised to any peer
Local
 10.10.10.5 (metric 3) from 10.10.10.5 (172.16.16.1)
   Origin IGP, metric 0, localpref 100, valid, internal, not synchronized

RD#traceroute 172.16.16.1
Tracing the route to 172.16.16.1
 1 10.0.0.10 [MPLS: Label 29 Exp 0] 64 msec 56 msec 60 msec
 2 10.0.0.21 60 msec 56 msec 72 msec
 3 * * *
```





Refer to the exhibit A customer reported an issue with a fiber link failure between RC and RE Users connected through the spoke location face disconnection and packet drops with the primary email server (172.16.16.1) but have no issues with the backup email server (172.16.26.1). All the router loopback IPs are advertised through the OSPF protocol. Which configuration resolves the issue?

```
RB(config)#router bgp 65101
RBfconfig-router^no synchronization
```

```
RC( config ^router bgp 65101
RC(config-router ^neighbor 10.10.10.2 next-hop-self
```

```
RBI config brouter bgp 65101
RB(config-router^nelghbor 10.10.10.5 next-hop-self
```

```
RC( config )#router bgp 65101
RC(conhg-router)^no synchronization
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer:

B

Explanation:

Question: 527

Refer to the exhibit.

```
interface Tunnel0
 ip address 172.23.5.10 255.255.255.0
 no ip redirects
 ip mtu 1420
 ip nhrp authentication C@trts81
 ip nhrp map multicast 192.168.200.1
 ip nhrp map 172.23.5.1 192.168.200.1
 ip nhrp network-id 10
 ip nhrp holdtime 300
 ip nhrp shortcut
 ip ospf network broadcast
 ip ospf priority 0
 tunnel source 192.168.100.146
 tunnel mode gre multipoint
 tunnel key 100
```

A network engineer is adding a new spoke router into an existing DMVPN Phase 3 tunnel with a hub router to provide secure communication between sites. Which additional configuration must the engineer apply to enable the tunnel to come up?

- A. ip nhrp registration no-unique
- B. ip nhrp server-only non-caching
- C. ip nhrp responder tunnel
- D. ip nhrpns 172.23.5.1

Answer: D

Explanation:

Question: 528

Refer to the exhibit.

```
R1#sh track brief
Track Type      Instance      Parameter      State Last Change
1 ip sla        10            reachability   Down 00:03:52

R1#show ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 10

Owner:

Tag:

Operation timeout (milliseconds): 5000

Type of operation to perform: icmp-echo
Target address/Source interface: 10.10.10.10/GigabitEthernet0/0
↻

Schedule:

  Operation frequency (seconds): 60 (not considered if randomly scheduled)
  Next Scheduled Start Time: Pending trigger

  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

Distribution Statistics:
```

```
Operation timeout (milliseconds): 5000
Type of operation to perform: icmp-echo
Target address/Source interface: 10.10.10.10/GigabitEthernet0/0
↔
Schedule:
  Operation frequency (seconds): 60 (not considered if randomly scheduled)
  Next Scheduled Start Time: Pending trigger
  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
  Distribution Statistics:
```

Refer to the exhibit A network engineer notices that the configured track option is down Which configuration resolves the issue*?

- A. ip sla schedule 10 start-time now
- B. ip sla schedule 10 start-time pending life forever
- C. ip sla schedule 10 no timeout
- D. ip sla schedule 10 no threshold

Answer: A

Explanation:

Question: 529

Refer to the exhibit.

```
RI (config) interface GigabitEthernet 0/0  
RI(config-if)#ip address 10.10.10.10 255.255.255.252  
RI(config-if)#ospfv3 1 ipv4 area 0
```

```
R2(config)#interface GigabitEthernet 0/0  
R2(config-if)#ip address 10.10.10.11 255.255.255.252  
R2(config-if)#ospfv3 10 ipv4 area 0  
R2(config-if)#ospfv3 network broadcast
```

Refer to the exhibit. An engineer is troubleshooting an OSPF adjacency issue between directly connected routers R1 and R2. Which configuration resolves the issue?

A)

```
R1 (config) interface GigabitEthernet 0/0  
RI(config-if)#ospfv3 network broadcast
```

B)

```
R2(config) ^interface GigabitEthernet 0/0  
R2(config-if) ~ip address 10.10.10.9 255.255.255.252
```

C)

```
R1 (config) ^interface GigabitEthernet 0/0  
R1(config-if) =ospfv3 10 ipv4 area 0
```

D)

**R2(config ^interface GigabitEthernet 0/0
R2(config-if)=no ospfv3 network broadcast**

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 530


Refer to the exhibit.

```

R3#show ip sla statistics
IPSLAs Latest Operation Statistics
IPSLA operation id: 10
Type of operation: icmp-echo
    Latest RTT: 24 milliseconds
Latest operation start time: *21:26:43.211 UTC Sat Sep 18 2021
Latest operation return code: OK
Number of successes: 75
Number of failures: 0
Operation time to live: Forever

IPSLA operation id: 20
Type of operation: icmp-echo
    Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: *21:26:47.499 UTC Sat Sep 18 2021
Latest operation return code: No connection
Number of successes: 128
Number of failures: 459
Operation time to live: Forever

```



Refer to me exhibit Traffic from R3 to the central site does not use alternate paths when R3 cannot reach 10 10 10 2 Traffic on R3 destined to R4 takes an alternate route via 10 10 10.6 when 10 10 10 4 is not accessible from R3 Which configuration switches traffic destined to 10 10 10 2 from R3 on the alternate path''

A. R3(config)#ip route 192.168.10.1 255.255.255.255 10.10.10.2 track 20

B. R2(config)#ip route 10.10.10.3 255.255.255.255 10.0.0.6

C. R3(config)#track(20 ip sla 20 reachability

D. R6(config)#ip route 10.10.10.3 255.255.255.255 10.0.0.30

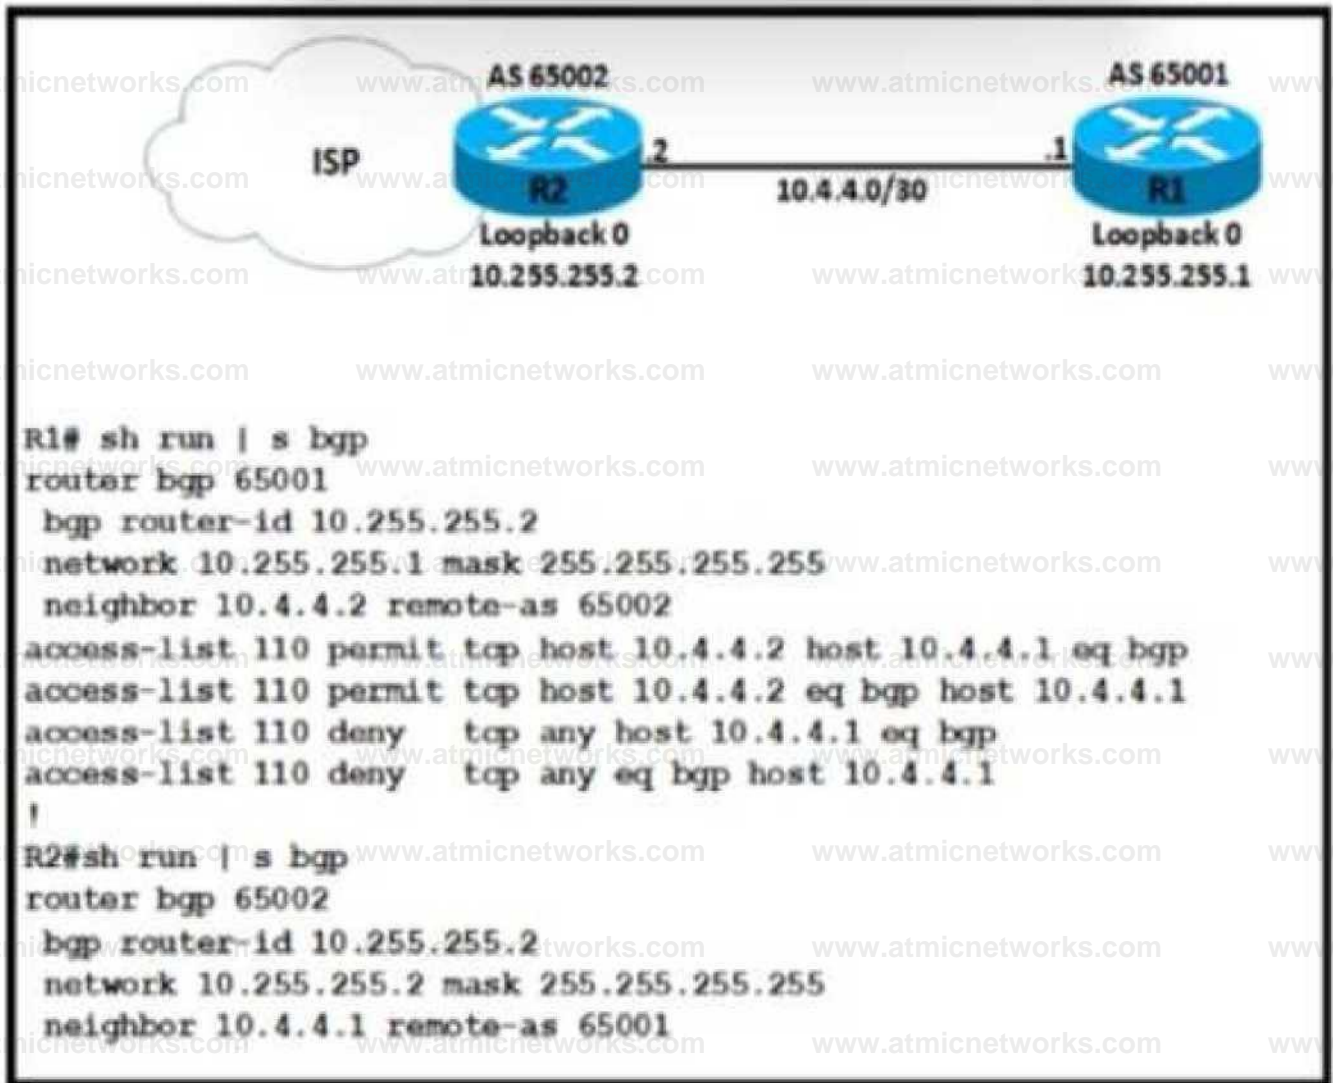
Answer:

A

Explanation:

Question: 531

Refer to the exhibit.



Refer to the exhibit A network engineer notices that R1 and R2 cannot establish an eBGP peering. The following messages appear in the log:

```

*Dec 21 12 08 59 991 BGP br lopo global 10 4 4 2 IPv4 Unicast base (0x6A883998 1) NSF delete stale NSF not active
*Dec 21 12 08 59 995 BGP br lopo global 10 4 4 2 IPv4 Unicast base (0x44397103 1) NSF no state paths state is NSF not active
*Dec 21 12 08 59 995 BGP ix lopo global 10 4 4 2 IPv4 Unicast base (0x6A883098 1) Resetting ALL counters
*Dec 21 12 09 09 819 BG 3 NOTIFICATION sent to neighbor 10 4 4 2 passive 23 (BGP xlen/peer wrong) 4 bytes OAFF02
*Dec 21 12 09 00 823 BGP 4 MSGDUMP unsupported of mal formatted message received from 10 4 4 2
*Dec 21 12 09 12 443 BGP SESSION 5 ADJCHANGE neighbor 10 4 4 2 IPv4 Unicast topology base removed from session BGP Notification received
*Dec 21 12 09 00 191 BGP tx global 10 4 4 2 Open active delayed 12288ms (35000ms max 60% jM)
  
```

Which configuration must the engineer apply to R1 to restore the eBGP peering?

A)
 router bgp 65001
 bgp router-id 10.255.255.2

neighbor 10.4.4.2 remote-as 65002

access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179

access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1

access-list 110 deny tcp any host 10.4.4.1 eq 179

access-list 110 deny tcp any eq 179 host 10.4.4.1

B)

router bgp 65001

bgp router-id 10.255.255.2

neighbor 10.4.4.2 remote-as 65002

access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179

access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1

access-list 110 deny udp any host 10.4.4.1 eq 179 access-list 110

deny udp any eq 179 host 10.4.4.1

C)

router bgp 65001

bgp router-id 10.255.255.1

neighbor 10.4.4.2 remote-as 65002

access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179

access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1

access-list 110 deny tcp any host 10.4.4.1 eq 179

access-list 110 deny tcp any eq 179 host 10.4.4.1

D)

router bgp 65001

bgp router-id 10.255.255.1

neighbor 10.4.4.2 remote-as 65002

access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179

access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1

access-list 110 deny udp any host 10.4.4.1 eq 179

access-list 110 deny udp any eq 179 host 10.4.4.1

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

Question: 532

How is the LDP router ID used in an MPLS network?


- A. The MPLS LDP router ID must match the IGP router ID.
- B. If not configured, the operational physical interface is chosen as the router ID even if a loopback is configured.
- C. The loopback with the highest IP address is selected as the router ID.
- D. The force keyword changes the router ID to the specified address without causing any impact.


```

R4#show ip flow export
Flow export v9 is enabled for main cache
Export source and destination details :
VRF ID : Default
Source(1)      10.0.0.10 (GigabitEthernet2/0)
Destination(1) 192.168.10.1 (656)
Version 9 flow records
254 flows exported in 41 udp datagrams
0 flows failed due to lack of export packet
0 export packets were sent up to process level
41 export packets were dropped due to no fib
0 export packets were dropped due to adjacency issues
0 export packets were dropped due to fragmentation failures
0 export packets were dropped due to encapsulation fixup failures

R4#show ip flow interface
GigabitEthernet2/0
ip flow ingress

```



Refer to the exhibit An enterprise operations team must monitor all application server traffic in the data center The team finds that traffic coming from the hub site from R3 and R6 rs monitored successfully but traffic destined to the application server is not monitored Which action resolves the issue?

A)

R4(config)#ip flow ingress

B)

R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# no shutdown

C)

R4 (confirm GigabitEthernet 2/0)
R4(config)# interface GigabitEthernet 2/0
R4(config-if)# ip address 10.1.1.1 255.255.255.0
R4(config-if)# no shutdown

D)

R5(config)# interface GigabitEthernet 0/0
R5(config-if)# ip address 10.1.1.1 255.255.255.0
R5(config-if)# no shutdown

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 534

Refer to the exhibit.

Router A line con 0 exec-

```
timeout 60 0 logout-warning 15
logging synchronous login
transport output all stopbits
1
```

Refer to the exhibit After a misconfiguration by a junior engineer, the console access to router A is not working Which configuration allows access to router A?

A)

```
RouterA(config)#aaa new-model
RouterA(config)#aaa authentication login my-auth-list tacacs*
```

B)

```
RouterA(config)#line console 0
RouterA(config)#password cisco
RouterA(config)#end
```

C)

```
RouterA(config)#line console 0
RouterA(config)#password cisco
RouterA(config)#login local
RouterA(config)#end
```

D)

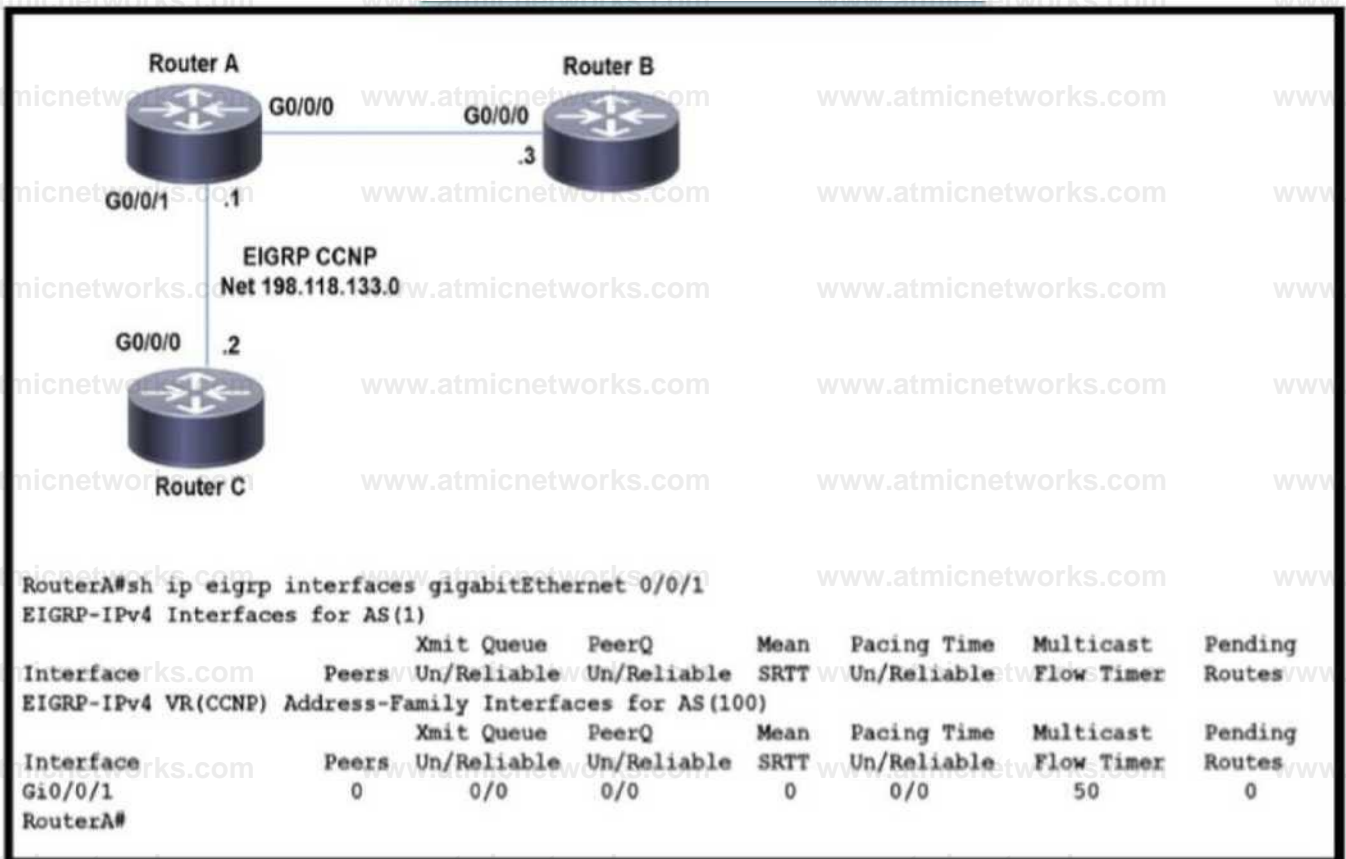
- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 535

Refer to the exhibit.



Refer to the exhibit EIGRP adjacency between router A and router C is not working as expected Which two configurations resolve the issue? (Choose two)

A)

Router C
 router eigrp CCNP
 address-family ipv4 unicast autonomous-system 100
 topology base
 exit-af-topology
 network 198.18.133.0
 exit-address-family

B)

Router C

router eigrp CCNP

address-family ipv4 unicast autonomous-system 100

af-interface GigabitEthernet0/0/0 hold-time 90 exit-af-
interface topology base exit-af-topology exit-address-family

C)

Router A

router eigrp CCNP

address-family ipv4 unicast autonomous-system 100

af-interface GigabitEthernet0/0/1

hello-interval 15 topology base exit-af-topology network
192.18.133.0 exit-address-family

D)

Router A

router eigrp CCNP

address-family ipv4 unicast autonomous-system 100

topology base

exit-af-topology network 198.18.133.0

exit-address-family

E)

Router A router eigrp CCNP address-family ipv4 unicast autonomous-system 10 af-interface

GigabitEthernet0/0/1 hello-interval 15 hold-time 90 exit-af-interface topology base exit-af-topology

network 198.18.133.0 exit-address-family

A. Option A

B. Option B

C. Option C

D. Option D

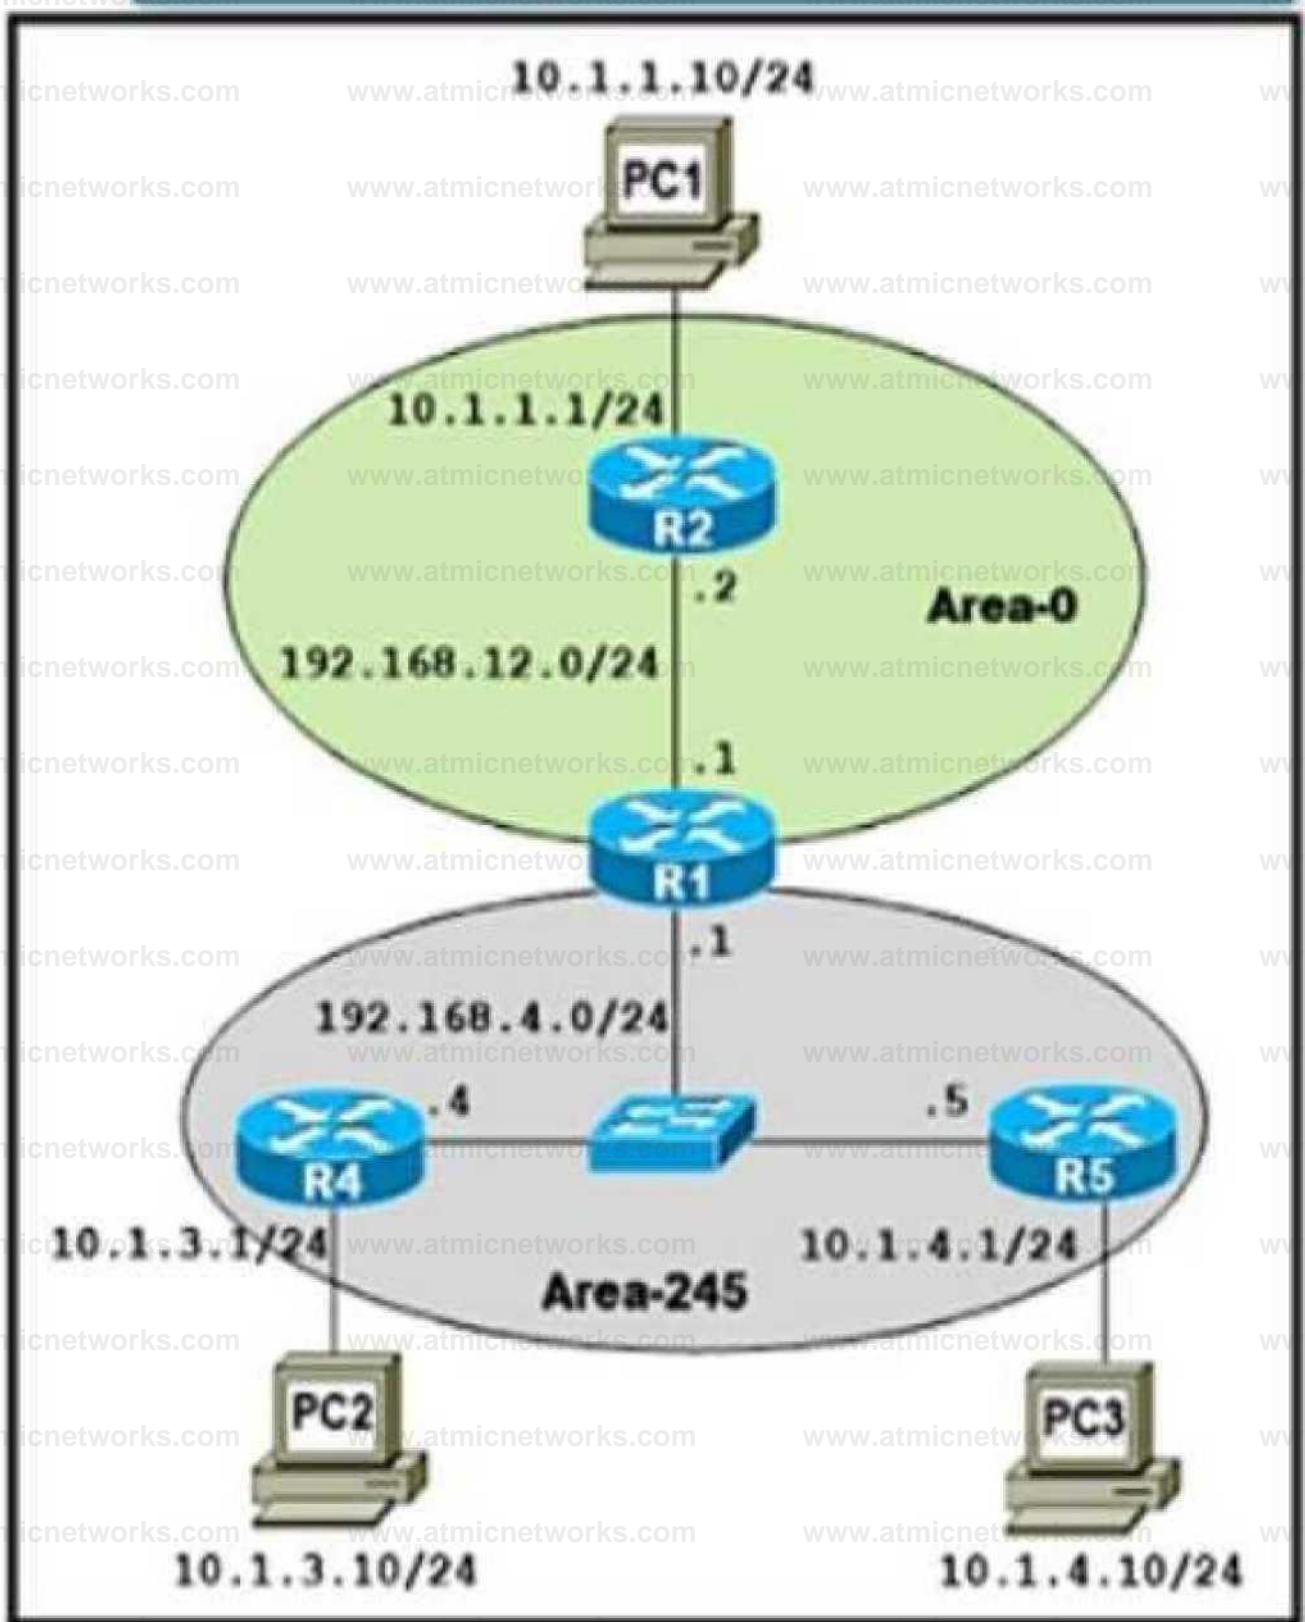
E. Option E

Answer: B,C

Explanation:

Question: 536

Refer to the exhibit.



Refer to the exhibit A network administrator is troubleshooting to reduce the routing table of R4 and R5 to learn only the default route to communicate from Inter-Area and Intra-Area networks Which configuration resolves the issue?

A)

R-1" default area 245

R-4ddefault area 245 default-cost

R-5^default area 245 default-cost

R- '=area 245 stub no-summary

B)

R-1=area 245 stub no-summary

R tsarea 245 stub

R-5=area 245 stub

C)

R-l=default area 245 default-cost

R- --default area 245

R-S^default area 245

D)

R-1#area 245 stub

R-4#area 245 stub no-summary

R-5#area 245 stub no-summary

A. Option A

B. Option B

C. Option C

D. Option D

Answer:

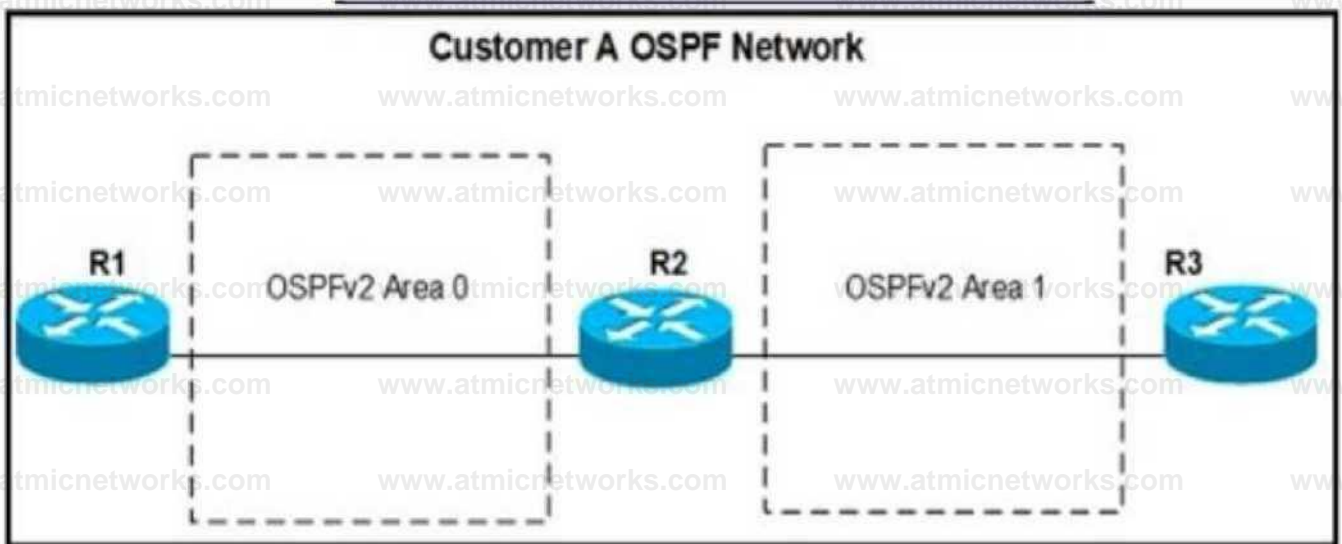
D

Explanation:

Question:

537

Refer to the exhibit.



Refer to the exhibit. An engineer must ensure that R3 sees only type 1 and 2 LSAs in area 1. Which command must the engineer apply on R2?

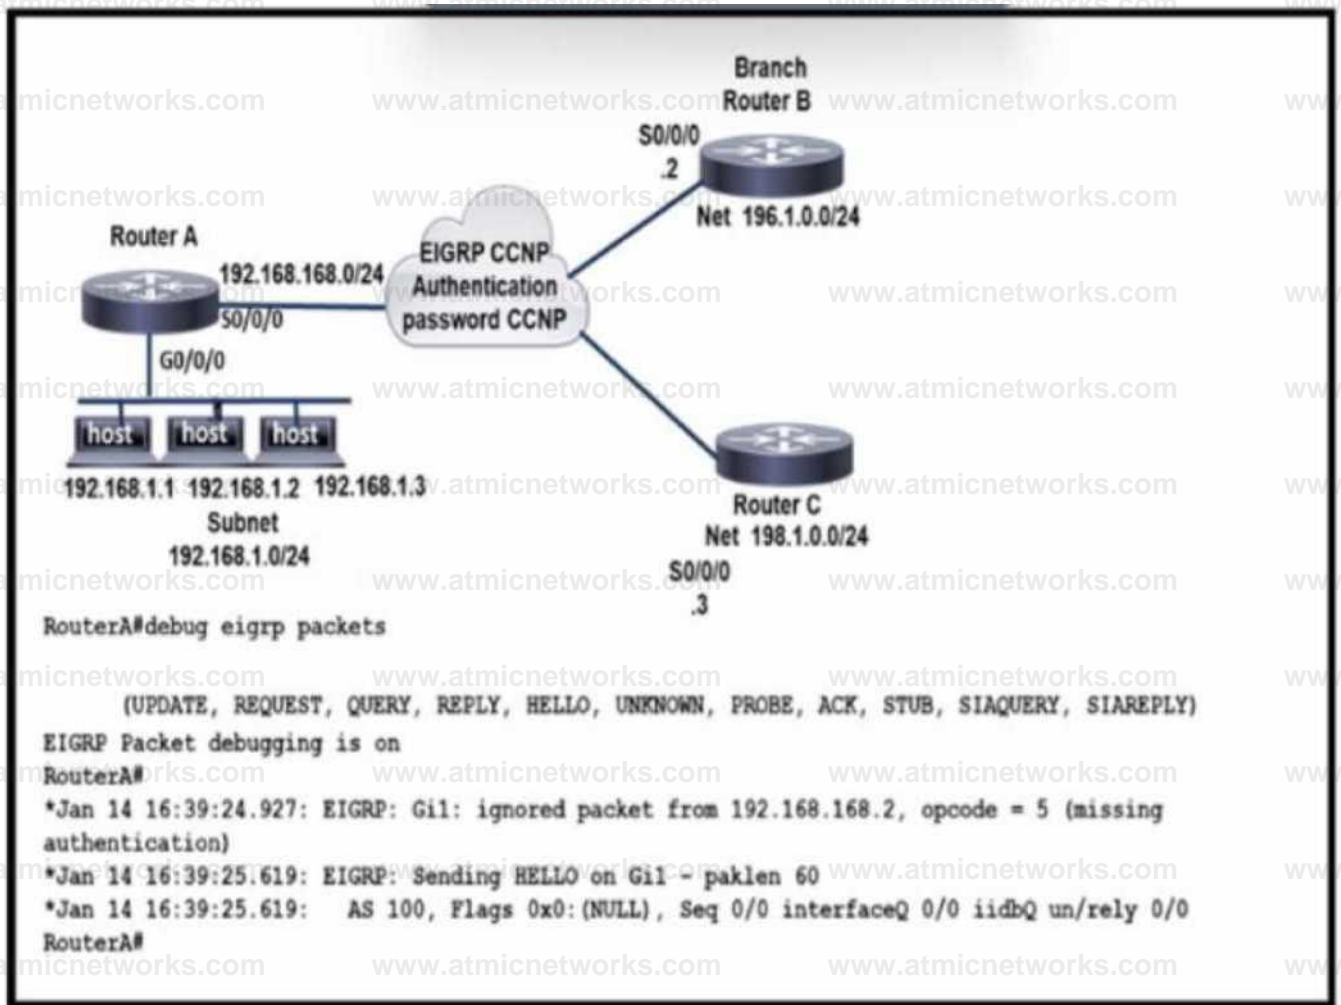
- A. Area 1 stub nssa
- B. Area 1 nssa no-summary
- C. Area 1 stub no-summary
- D. Area 1 stub

Answer: C

Explanation:

Question: 538

Refer to the exhibit.



Refer to the exhibit. The services at branch B are down. An engineer notices mal router A and router B are not exchanging any routes Which configuration resolves the issue on router B?

- A)
- router eigrp 100 network 192.168.168.0
 - key chain CCNP key 1
 - key-string EIGRP
 - Interface serlalo/O/O
 - Ip address 192.168.168.2 255.255.255.0 ip
 - authentication mode eigrp 100 md5 ip authentication

key-chain eigrp 100 EIGRP negotiation auto

B)

router eigrp 100 network 192.168.168.0

key chain EIGRP key 1

key-string CCNP

Interface seria 10/0/0

Ip address 192.168.168.2 255.255.255.0 ip

authentication mode eigrp 100 md5

negotiation auto

C)

router eigrp 100 network 192.168.168.0

key chain EIGRP key 1

key-string CCNP

Interface serl0/0/0

Ip address 192.168.168.2 255.255.255.0

ip authentication mode eigrp 100 md5 ip

authentication key-chain eigrp 100 EIGRP negotiation

auto

D)

```
router eigrp 100 network 192.168.168.0
```

```
key chain EIGRP key 1
```

```
key-string CCNP
```

```
interface serial0/0/0
```

```
ip address 192.168.168.2 255.255.255.0
```

```
ip authentication key-chain eigrp 100 EIGRP  
negotiation auto
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 539

1.

Which two label distribution methods are used by routers in MPLS? (Choose two)

- A. targeted hello message
- B. LDP discovery hello message
- C. LDP session protection message
- D. downstream unsolicited
- E. downstream on demand

Answer: D,E

Explanation:

Question: 540

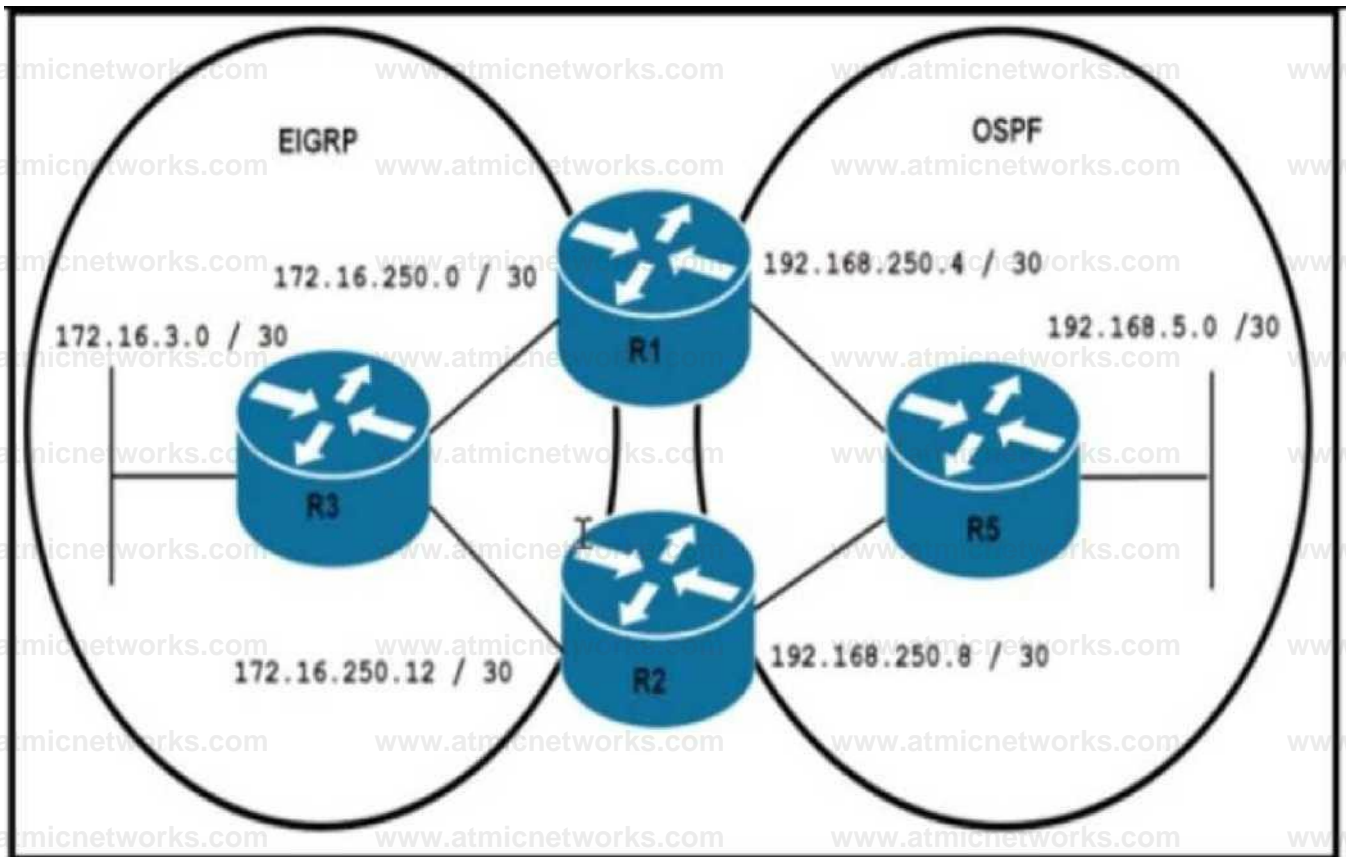
Which two protocols are used by a P router to transfer VPN traffic between PE routers in an MPLS network? (Choose two.)

- A. BGP
- B. OSPF
- C. MP-BGP
- D. LDP
- E. RSVP

Answer: C,D

Explanation:

<pre> R1#show running-config begin router router eigrp router eigrp 100 network 172.16.250.0 0.0.0.255 redistribute ospf 1 metric 11111 redistribute eigrp 100 subnets network network 192.168.250.0 0.0.0.255 area 0 </pre>	<pre> R5#traceroute 172.16.3.1 Type escape sequence to abort. Tracing the route to 172.16.3.1 VRF info: (vrf in name/id, vrf out name/id) 0 192.168.250.9 66 msec 1 192.168.250.6 6 msec 2 192.168.250.9 8 msec 3 172.16.250.2 33 msec 4 172.16.250.14 88 msec 5 172.16.250.2 11 msec R5# </pre>
<pre> R2#show runn I begin router eigrp router eigrp 100 network 172.16.250.0 0.0.0.255 redistribute ospf 1 metric 11111 t router ospf 1 redistribute eigrp 100 subnets network 192.168.250.0 0.0.0.255 area 0 t ip forward-protocol nd </pre>	



Refer to the exhibit. An engineer is troubleshooting a routing loop on the network to reach the

172.16.3.0/16 from the OSPF domain. Which configuration on router R1 resolves the Issue?

A)

```
router ospf 1
 redistribute eigrp 100 subnets route-map LOOPFILT
route-map LOOPFILT deny 10
 match ip address 15
route-map LOOPFILT permit 20
access-list 15 permit 172.16.0.0 0.0.255.255
```

B)

```
router eigrp 100
 redistribute ospf 1 metric 11111 route-map LOOPFILT i
route-map LOOPFILT deny 10
 match ip address 15
route-map LOOPFILT permit 20 i
access-list 15 permit 172.16.0.0 0.0.255.255
```

C)

router ospf 1

redistribute eigrp 100 route-map LOOPFILT

route-map LOOPFILT deny 10

match ip address 15 I

access-list 15 permit 172.16.0.0 0.0.255.255

D)

router eigrp 100

redistribute ospf 1 metric 11111 route-map LOOPFILT

route-map LOOPFILT deny 10 match ip address 15

access-list 15 permit 172.16.0.0 0.0.255.255

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 542

Refer to the exhibit.

```
R1# show running-config | begin router eigrp
router eigrp 100
network 172.16.250.0 0.0.0.3
redistribute ospf 10 metric 11111
```

```
R2# show running-config | begin router eigrp
router eigrp 103
network 172.16.2.0 0.0.0.3
network 172.16.2.16 0.0.0.15
network 172.16.2.32 0.0.0.15
```

```
router ospf 10
redistribute eigrp 100 metric 100 subnets route-map CCNP
network 172.16.1.0 0.0.0.3
area 0
```

```
redistribute static metric 100 111: route-map CCNP
```

```
ip forward-protocol nd
```

```
ip forward-protocol nd
```

```
no ip http server
no ip http secure-server
```

```
no ip http server
no ip http secure-server
ip route 172.16.2.48 255.255.255.240 172.16.2.1
```

```
route-map CCNP deny 10
match route-type local
```

```
route-map "CM?" permit 10
match ip address 10
set tag 200
```

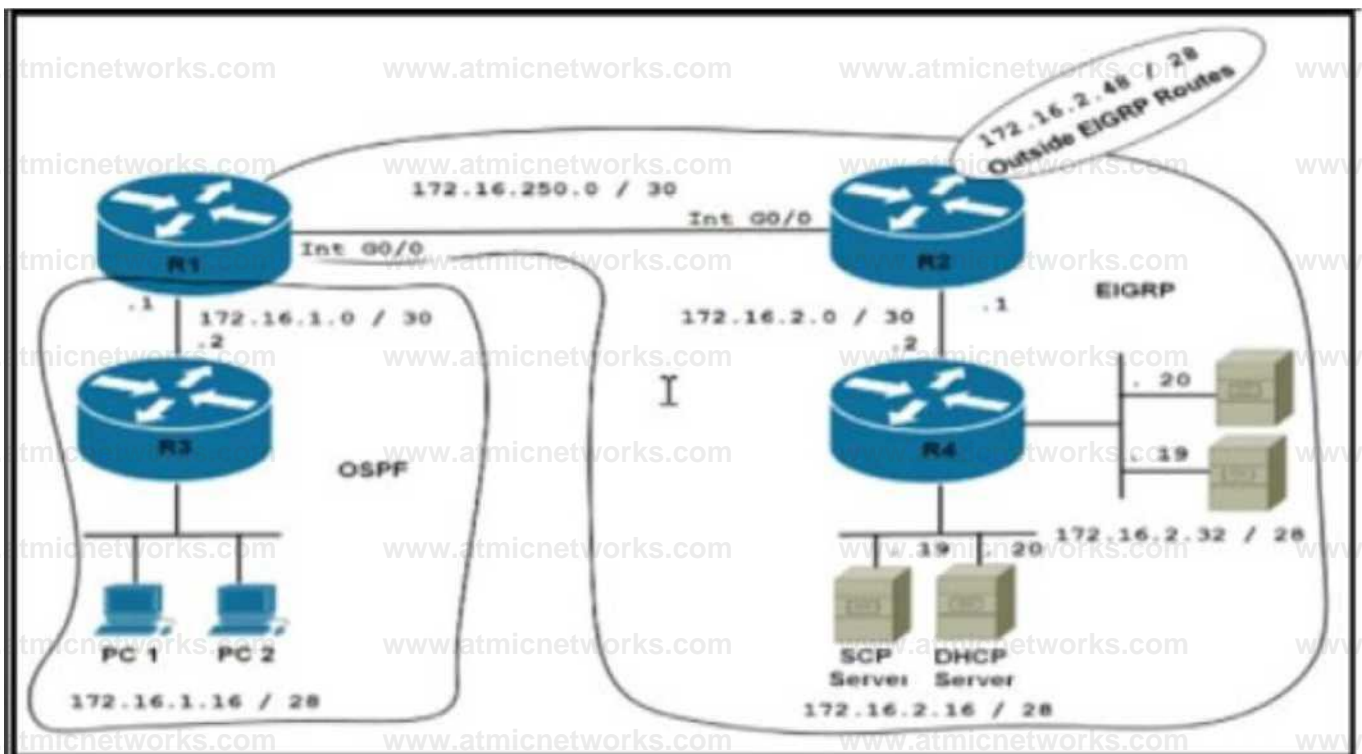
```
access-list 10 permit 172.16.2.32
```

```
access-list 10 permit 112.16.2.48 0.0.0.15
```

R3# show ip route

Gateway of last resort is 172.16.1.1 over FastEthernet0/0

```
172.16.0.0/16 is variably subnetted, 1 subnets, 1 neighbor
172.16.1.0/30 is directly connected, GigabitEthernet0/1
L 172.16.1.16/28 is directly connected, Loopback1
L 172.16.1.17/32 is directly connected, Loopback1
C 172.16.1.32/26 is directly connected, Loopback2
L 172.16.3.33/32 is directly connected, Loopback2
S 172.16.1.48/28 (1/0) via 172.16.1.18
Sat
```



Refer to the exhibit. Which configuration resolves the route filtering issue on R1 to redistribute all the routes except 172.16.2.48/28?

A)

```
R1 (config ^route-map CCNP deny 10
```

```
RI (config-route-map )#no match route-type local R1(config-route-map)#match route-type external type-
```

```
1
```

```
RI(config ^route-map CCNP permit 20
```

B)

```
RI (config ^route-map CCNP deny 10
```

```
Ri(conf|g-route-mapHfno match route-type local R1(conf|g-route-map)# match route-type level-2
```

```
R1(contig »route-map CCNP permit 20
```

C)

```
RI (config ^route-map CCNP deny 10
```

```
R1 (conflg-route-map}*no match route-type local R1(conflg-route-mappmatch route-type external R1
```

```
(config isroute-map CCNP permit 20
```

D)

```
R1 (config grouts-map CCNP deny 10
```

```
R1 (config-route-map )#no match route-type local
```

```
R1 (conflg-route-map)smateh route-type external type-2
```

```
R1 (config ^route-map CCNP permit 20
```

A. Option A

B. Option B

C. Option C


D. Option D

Answer: D

Explanation:

Question: 543

Refer to the exhibit.



Refer to the exhibit R2 can reach Loopback222, but R1 SW1 and PC1 cannot communicate with 172.16.222.254 R1 and R2 configurations are shown here

```
R1#show run | sec router eigrp router eigrp VR1
```

```
address-family ipv4 unicast autonomous-system 1
```

```
topology base ex it-attopology network 172,16.1.1 0.0 0.0  
network 192.168.100.0 network 192.168.200.0 network 192  
168.255.91 00.0.0 exit-address-family
```

```
R2(config)#do show run | sec router eigrp router eigrp 1 network  
172 16 1.2 0 0.0.0 network 172 16 2 22.0 0 0 0.255 network  
192.168 222.254 0.0.0.0
```

Which EIGRP configuration command resolves the issue?

A. R2(config-router) # redistribute static

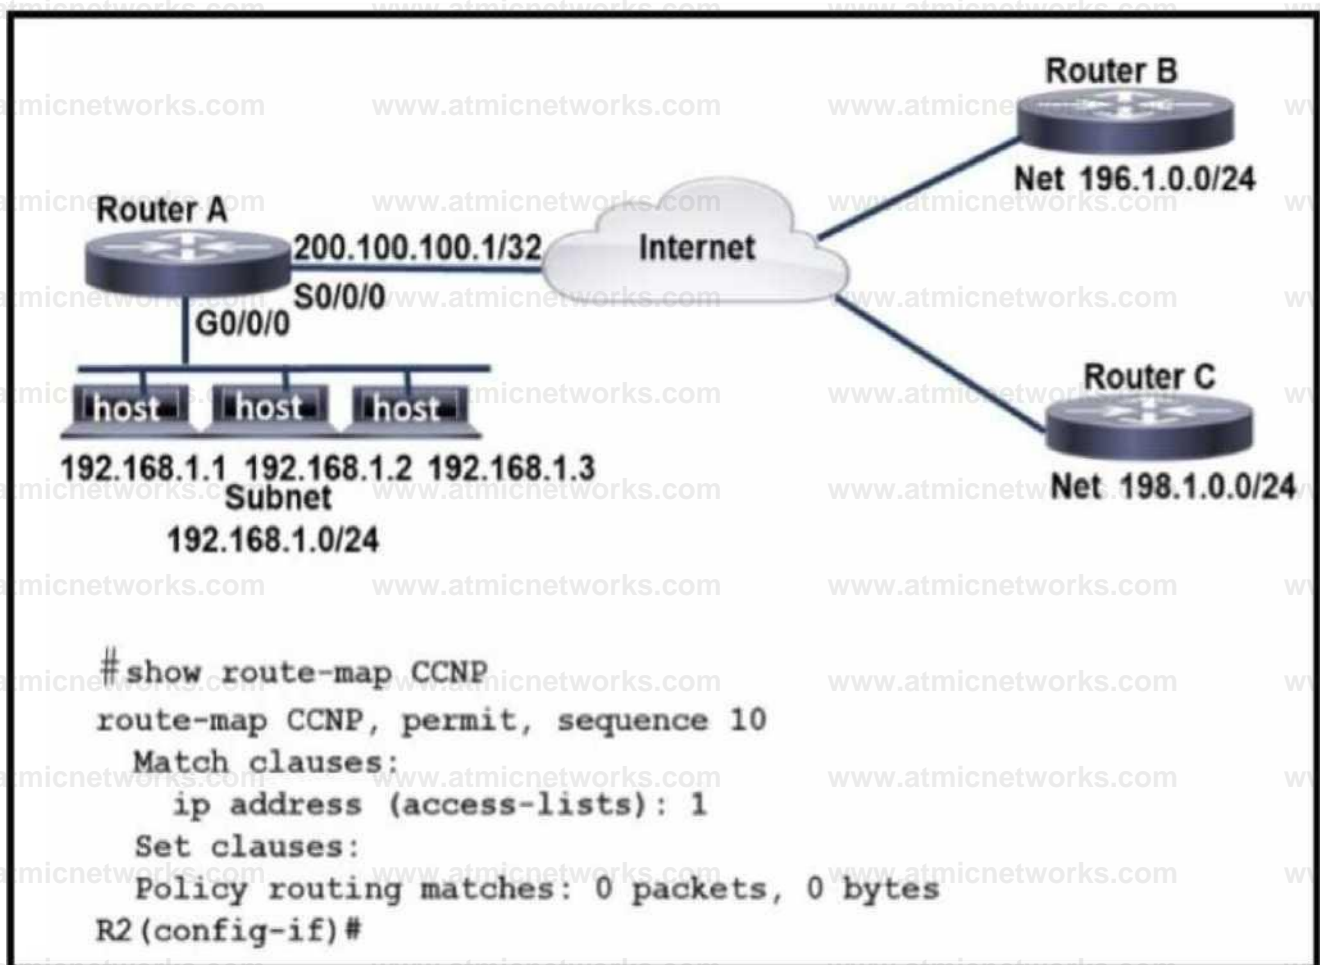
- B. R1(config-router)# network 172.16.222.254 0.0.0.0
- C. R1 (config-router)# network 172.16.222.264 255.255.255.255
- D. R1(config-router)# redistribute static

Answer: A

Explanation:

Question: 544

Refer to the exhibit.



Refer to the exhibit. An engineer configures router A to mark all inside to outside traffic from

network 192.168.1.0, except from host 192.168.1.1. with critical IP precedence. The policy did not work as expected

Which configuration resolves the issue?

A)

```
Router A(config)#access-list 1 deny host 192.168.1.1
```

```
Router A(config)#route-map CCNP permit 10
```

```
Router A(config)#match ip address 1
```

```
Router A(config)#set ip precedence critical
```

```
Router A(config)#route-map CCNP permit 20
```

```
Router A(config)#interface g0/0/0
```

```
Router A(config-if)#ip address 192.168.1.4 255.255.255.0
```

```
Router A(config-if)#ip policy route-map CCNP
```

B)

```
Router A(config)#access-list 1 deny host 192.168.1.1
```

```
Router A(config)#access-list 1 permit any any
```

```
Router A(config)#route-map CCNP deny 10
```

```
Router A(config)#match ip address 1
```

```
Router A(config)#set ip precedence critical
```

```
Router A(config)#route-map CCNP permit 20
```

```
Router A(config)#interface g0/0/0
```

```
Router A(config-if)#ip address 192.168.1.4 255.255.255.0
```

```
Router A(config-if)#ip policy route-map CCNP
```

C)

```
Router A(config)#access-list 1 deny host 192.168.1.1
```

```
Router A(config)#access-list 1 permit any any
```

```
Router A(config)#route-map CCNP permit 10
```

```
Router A(config)#match ip address 1
```

```
Router A(config)#set ip precedence critical
```

```
Router A(config)#route-map CCNP permit 20
```

```
RouterA(config)# ip precedence critical
RouterA(config)# interface g0/0/0
RouterA(config-if)# ip address 192.168.1.4 255.255.255.0
RouterA(config-if)# ip policy route-map CCNP
```

D)

```
RouterA(config)# access-list 1 deny host 192.168.1.1
RouterA(config)# access-list 1 permit any any
RouterA(config)# route-map CCNP permit 10
RouterA(config)# route-map CCNP match ip address 1
RouterA(config)# route-map CCNP set ip precedence critical
RouterA(config)# interface g0/0/0
RouterA(config-if)# ip address 192.168.1.4 255.255.255.0
RouterA(config-if)# ip policy route-map CCNP
```

A. Option

B. Option

C. Option

D. Option

Answer: A

Explanation:

Question: 545

configuration on the hub router meets this requirement?

- A. interface Tunnel0 tunnel mode gre multipoint
- B. interface Tunnel0 tunnel mode dmvrp
- C. interface Tunnel0 tunnel mode ipsec ipv4
- D. interface Tunnel0 tunnel mode ip

Answer: A

Explanation:

Question: 546

Refer to the exhibit.

```
Router# show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes, 0
overruns, xml disabled, filtering disabled)
No Active Message Discriminator.
No Inactive Message Discriminator.
Console logging: level debugging, 8 messages logged, xml disabled,
filtering disabled
Monitor logging: level debugging, 0 messages logged, xml disabled,
filtering disabled
Buffer logging: level debugging, 8 messages logged, xml disabled,
filtering disabled
Exception Logging: size (8192 bytes)
Count and timestamp logging messages: disabled
Persistent logging: disabled
```

Refer to the exhibit. A network engineer lost remote access to the router due to a network problem. The engineer used the console to access the router and noticed continuous logs on the console terminal. Which configuration limits the number of log messages on the console to critical and higher severity level messages?

- A. term no monitor
- B. logging console 2
- C. no logging console
- D. logging console 5

Answer: D

Explanation:

Question: 547

Refer to the exhibit.

```
RouterA#show snmp community
Community name: ILMI
Community Index: ILMI
Community SecurityName: ILMI
storage-type: read-only active

Community name: ccnp
Community Index: ccnp Community SecurityName: ccnp
storage-type: nonvolatile active access-list: 4

RouterA#show ip access-lists
Standard IP access list 4
10 permit 172.16.1.1
20 permit 172.16.2.2
30 permit 172.16.3.3
Extended IP access list BRANCHES
10 permit ip 172.16.4.4 any (95 matches)
20 deny ip any any (95 matches)
```

Refer to the exhibit The SNMP server with IP address 172.16.4.4 cannot access host router A Which configuration command on router A resolves the issue?

- A. snmp-server community ccnp
- B. access-list 4 permit 172.16.4.0 0.0.0.3
- C. access-list 4 permit host 172.16.4.4


D. snmp-server host 172.16.4.4 ccnp

Answer: D

Explanation:

Question: 548

Refer to the exhibit.



```
UIUrteM TunnelO
Ip address 171.11 11 111 111 111.0
Ip nhrp authentication ttlIKO
Ip nhrp MP iMlticaat dynamic
Ip nhrp network-id 10 ip ospf network
broadcast Ip ospf priority 211 tunrial sour co
1*2 100 100 11 tunnal Baria ijra Multipoint
tunrial toy 100
```

```
interface TunnelO
Ip address 172 11 1,2 211 211 21! 0 Ip nhrp
authentication till ICO Ip nhrp nap nulleast
1*2.100 100 11 Ip nhrp nap 172.1* !I 1*2
100.100.11 Ip nhrp network-id 10 ip ospf
network broadcast Ip ospf priority 0 tunnel
source 1*2 IM 100 12 tunnel destination
IM.lt.IM.li tunnel key 100
```

```
interface TunnelO ip address 17).11.1.3 255 2
5!.IM 0 Ip nhrp authentication ttlItCO ip nhrp
aap miltleast 1*2.1*0.100.11 ip nhrp nap
1721*51 1*2.1*1.100 11 Ip nhrp network-lid 10
Ip ospf network broadcast Ip ospf priority 0
tunnel source 1*21**10013 tunnel destination
1*2.10.100.11 tunnel key 100
```

Refer to the exhibit. An organization is installing a new L3 MPLS link to establish DM VPN Phase 2 tunnels between the hub and two spoke routers Which additional configuration should the engineer implement on each device to achieve optimal routing between the spokes?

A)

**interface TunnelO no tunnel destination
192.168.100 11 tunnel mode mpls**

traffic-eng

B)

interface Tunnel0

ip ospf priority 1

ip ospf network non-broadcast

C)

interface Tunnel0 no tunnel destination

192.168.100.11 tunnel mode gre

multipoint

D)

interface Tunnel0 ip ospf priority 253 ip

ospf network point-to-multipoint

A. Option

B. Option

C. Option

D. Option

Answer:

C

Explanation:

Question:

549

Refer to the exhibit.

```
R4#show ip rout*
Gateway of last resort is not set


172.16.0.0/16 is variably subnetted, 8 subnets, 3 masks
C   172.16.2.0/30 is directly connected, GigabitEthernet0/1
L   172.16.2.2/32 is directly connected, GigabitEthernet0/1
C   172.16.2.16/28 is directly connected, Loopback1
L   172.16.2.17/32 is directly connected, Loopback1
C   172.16.2.32/28 is directly connected, Loopback2
L   172.16.2.33/32 is directly connected, Loopback2
S   172.16.2.48/28 is directly connected, Loopback2
D   172.16.250.0/30 [90/3072] via 172.16.2.1, Id08h, GigabitEthernet0/1


R3#sho ip route
Gateway of last resort is not set

R   172.16.0.0/16 (120/10) via 192.168.1.1, 00:00:03, GigabitEthernet0/1
C   192.168.1.0/24 is variably subnitted, 2 subnets, 2 masks ia directly
L   192.166.1.0/30 connected, GigabitEthernet0/1 is directly connected,
C   192.168.1.2/32 GigabitEthernet0/1 variably subnetted, 2 subnets, 2 masks
L   192.166.2.0/24 is directly connected, Loopback!
C   192.168.2.0/24
L   192.168.2.33/32 xs directly connected, Loopback!
C   192.168.3.0/24 io variably subnetted, 2 subnets, 2 masks
L   192.168.3.0/24 is directly connected, Loopback!
L   192.168.3.17/32 io directly connected, Loopback!

R1#sho running-config | begin router eigrp router eigrp 100
router eigrp 100
network 172.16.250.0 0.0.3 redistribute rip
router rip redistribute eigrp 100 metric 10
network 192.168.1.0
ip forward-protocol nd
route-map REDIST permit 10
match ip address 15
route-map CCNP deny 10
match route-type external
route-map CCNP permit 20
access-list 15 permit 192.168.0.0 0 0.0.255

R3#traceroute 172.16.2.33
Type escape sequence to abort.
Tracing the route to 172.16.2.33
VRF info: (vrf in name/id, vrf out name/id) 1 192.168.1.1 27
msec 31 msec 16 msec 2 * * *
 3 * * *
 4 * * *
R3#
```





Refer to the exhibit Users from the 192.168.2.0/24 network cannot connect to the 172.16.2.32/28 network Which configuration resolves the issue?

A)

```
R4(config)#ip route 0.0.0.0 0.0.0.0 172.16.2.1
```

B)

```
R1 (config)#route-map REDIST permit 10
```

```
R1 (config-route-map)#match ip address 15
```

```
R1 (config-route-map)# set metric 1000 10 255 1 1500
```

```
R1 (config-route-map)#exit
```

```
R1 (config)# access-list 15 permit 192.168.2.0 0.0.255.255
```

C)

```
R1 (config-router)#router eigrp 100
```

```
R1 (config-router)#redistribute rip
```

```
R1 (config-router)#default-metric 10000 100 255 100 1500
```

D)

R1 (config ^router eigrp 100

R1(config-router)#network 192.168.0.0

A. Option

B. Option

C. Option

D. Option

Answer: A

Explanation:

Question: 550

Refer to the exhibit.

XOUtil tip 1 variance 2

Rltshow ip eigrp topology 172.16.100.5 255.255.255.255

XP-EIGRP (AS 1): Topology entry for 172.16.100.5/32

State x* Passive, Query origin flag is 1, 1 Successor(s), FD xs 409600 Routing

Descriptor Blocks:

10.4.1.5 (Ethernetl/O), from 10.4.1.5, Send flag is 0x0 Composite metric xs (409600/128256), Route is Internal Vector metric:

Minimum bandwidth xs 10000 Kbit

Total delay xs 6000 microseconds

Reliability is 255/255

Load IS 1/255

Minimum MTU is 1500

Hop count is 1

10.4.1.6 (Serial2/0), from 10.3.1.6, Send flag is 0x0 Composite metric is (435200/409600), Route is Internal Vector metric:

Minimum bandwidth is 10000 Kbit

Total delay is 7000 microseconds

Reliability is 255/255

```
Load is 1/255
Minimum MTU is 1500

Hop count is 1

10.3.1.6 (Serial2/0), from 10.3.1.6, Send flag is 0x0
Composite metric is (435200/409600), Route is Internal
Vector metric:

Minimum bandwidth is 10000 Kbit

Total delay is 7000 microseconds
Reliability is 255/255

Load is 1/255

Minimum MTU is 1500

Hop count is 2
```

Refer to the exhibit. A network engineer troubleshooting a packet drop problem for the host 172.16.100.5 notices that only one link is used and installed on the routing table, which saturates the bandwidth. Which action must the engineer take to resolve the high bandwidth utilization problem and share the traffic toward this host between the two available links?

- A. Set the eigrp variance equal to 4 to install a second route with a metric not larger than 4 times of the best metric.
- B. Change the EIGRP delay metric to meet the feasibility condition.
- C. Set the eigrp variance equal to 3 to install a second route with a metric not larger than 3 times of the best metric.
- D. Disable the eigrp split horizon loop protection mechanism.

Answer: B

Explanation:

Question: 551

Refer to the exhibit.

RKshow ip bgp 10.0.0.0/8

BGP routing table entry for 10.0.0.0/8, version 0

Paths: (1 available, no best path)

Not advertised to any peer

Refresh Epoch 1

100

192.168.10.20 (inaccessible) from 192.168.20.20 (192.168.20.20) Origin

incomplete, metric 0, localpref 100, valid, internal rx pathid: 0, tx

pathid: 0

Refer to the exhibit. An engineer is troubleshooting a prefix advertisement issue from R3, which is not directly connected to R1. Which configuration resolves the issue?

A)

R1 (config ^router bgp 64512

R1(config-router)#neighbor 192.168.10.20 next-hop-self

B)

R1 (config ^router bgp 64512

R1(config*router)#neighbor 192.168.20.20 next-hop-self

C)

R2(config)#router bgp 64512

R2(config-router^neighbor 192.168.20.10 next-hop-self

D)

```
R2(config)#router bgp 64512
R2(config-router)#neighbor 192.168.10.20 next-hop-self
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

Question: 552

Refer to the exhibit.



Device time has drifted from Cisco DNA Center > Issue Instance

Excessive time lag between Cisco DNA Center and device "SW1.ap.com"

[Open](#)

Description

The time on Cisco DNA Center and Device "SW1.ap.com" has drifted too far apart. Cisco DNA Center cannot process the device data accurately if the time difference is more than 3 minutes.

[Go to SW1.ap.com](#)

Last Occurred: Jan 12, 2022 2:42 AM

Refer to the exhibit. Which action resolves the issue?

- A. Establish connectivity between the NTP server and the switch.
- B. Configure the local time on Cisco DNA Center
- C. Configure the local time on the SW1 device
- D. Establish connectivity between the NTP server and Cisco DNA Center.

Answer:

C

Explanation:

Question:

553

Refer to the exhibit.

```
!
ip sla 1
  icmp-echo 192.168.2.1 source-interface GigabitEthernet0/0/1
  timeout 1000
  threshold 1000
  frequency 30
ip sla schedule 1 life forever start-time now
!
track 1 ip sla 1 reachability
```


Refer to the exhibit An engineer observes that every time the ICMP packet is lost at a polling interval, track 1 goes down, which causes unnecessary disruption and instability in the network. The engineer does not want the traffic to be rerouted if the loss of ICMP packets is negligible. If the packet loss is persistent for a longer duration, the track must go down and the traffic must be rerouted. Which action resolves the issue?

- A. Change the IP SLA schedule to run only at certain intervals.
- B. Increase the threshold value from 1000 to 1500.
- C. Increase the timeout value from 1000 to 1500
- D. Define a delay timer under track 1.

Answer: D

Explanation:

Question: 554

Refer to the exhibit.

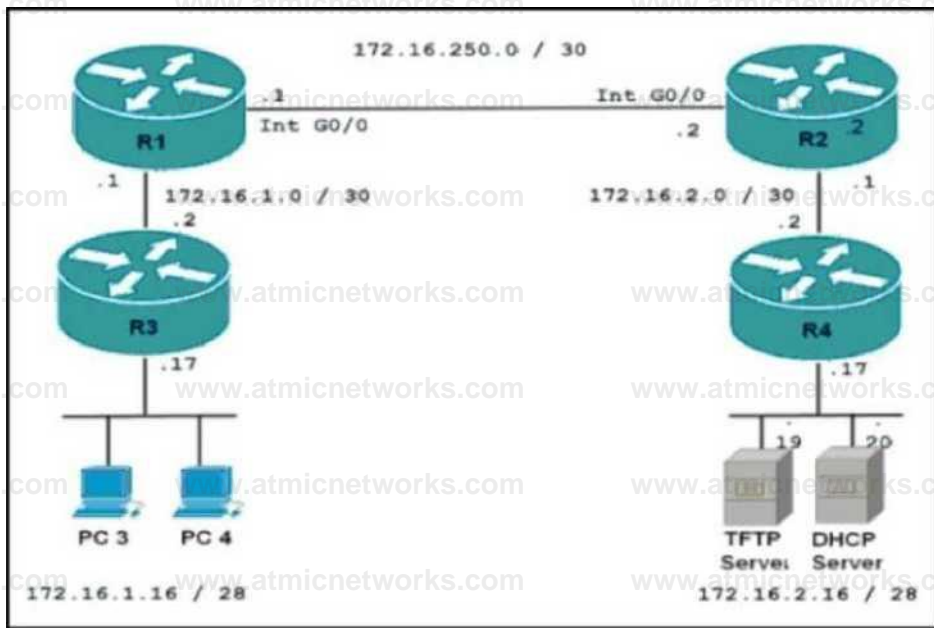
R3*copy tftp flash:

Address or name of remote host [172.16.2.19]?

Source filename [c2600-i-mz.121.T.bin]? c2600-i-mz.l21-l-t.bin

Destination filename [c2600-i-mz.121-1.T.bin]?

Loading c2600-i-mz.121-1.T.bin from 172.16.2.19(via GigabitEthernet0/0): • %Error copying tftp://172.16.2.19/c2600-i-mz.121-1.T.bin (Not enough space on device) R3#



Refer to the exhibit. The engineer is getting an error when trying to transfer a new IOS file to the router. Which action resolves the issue?

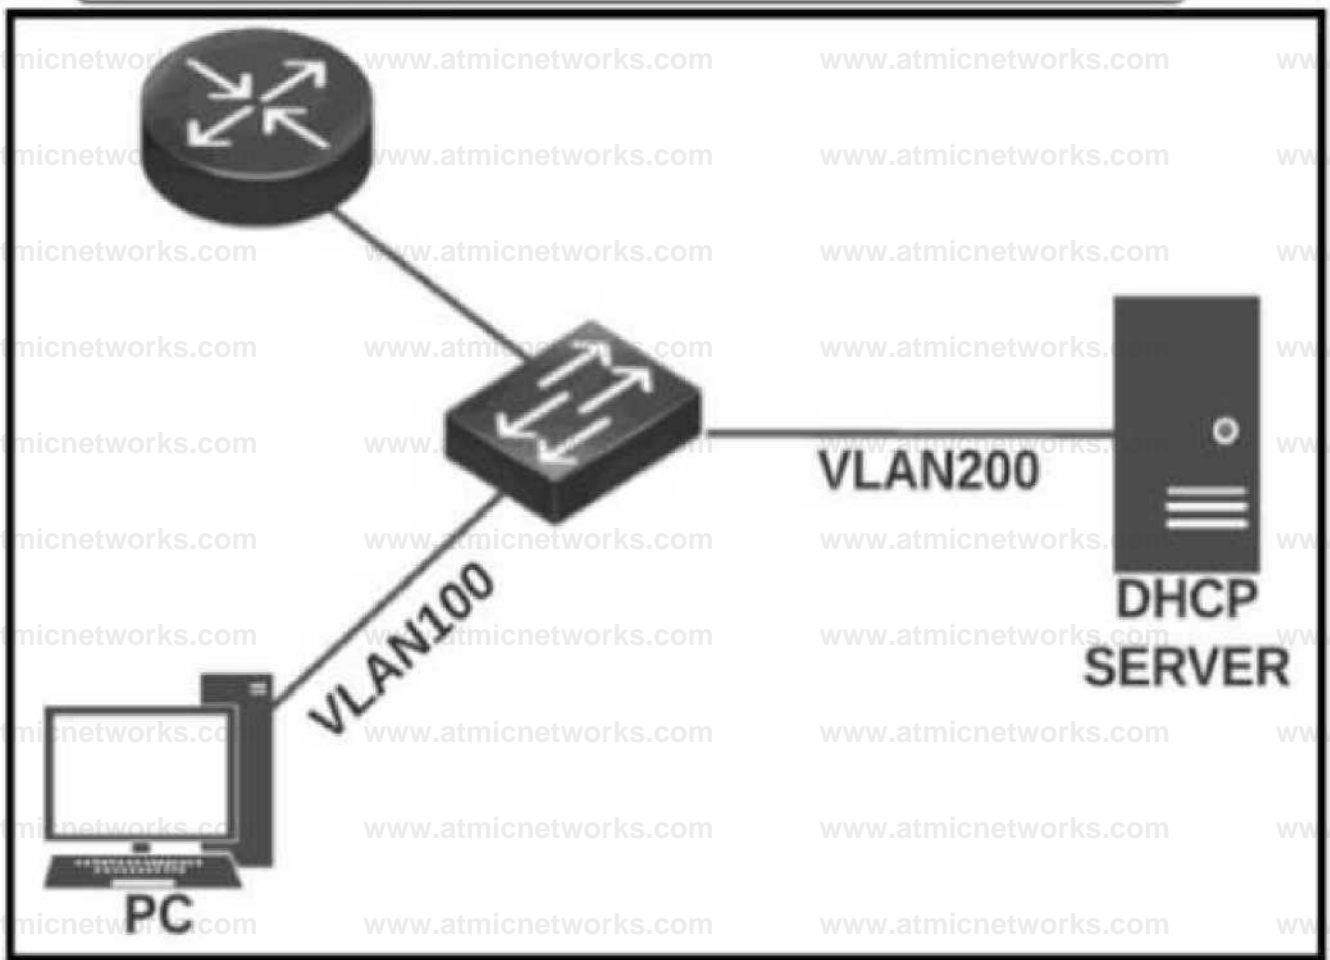
- A. Delete some files on the router flash memory.
- B. Delete some files on the router NVRAM.
- C. Remove any access-list filtering the TFTP file transfer.
- D. Split the file into parts to transfer them one by one.

Answer: A

Explanation:

Question: 555

Refer to the exhibit.



Refer to the exhibit. APC is configured to obtain an IP address automatically, but it receives an IP address only from the 169.254.0.0 subnet. The DHCP server logs contained no DHCPDISCOVER message from the MAC address of the PC.

Which action resolves the issue?

- A. Configure an ip helper-address on the router to forward DHCP messages to the server.
- B. Configure DHCP Snooping on the switch to forward DHCP messages to the server.
- C. Configure a DHCP reservation on the server for the PC.
- D. Configure a static IP address on the PC and exclude it from the DHCP pool.

Answer: A

Explanation:

Question: 556

Refer to the exhibit.

```
GigabitEthernet2 is up, line protocol is up
Internet Address 172.16.1.42/30, Interface ID 8, Area 1
Attached via Network Statement
Process ID 1, Router ID 172.16.100.7, Network Type BROADCAST, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
0 1 no no Base
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 172.16.100.7, Interface address 172.16.1.42
Backup Designated router (ID) 172.16.100.5, Interface address 172.16.1.41
Timer intervals configured, Hello 10, Dead 40, wait 40, Retransmit 5
oob-resync timeout 40
Hello due in 00:00:01
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 172.16.100.5 (Backup Designated Router)
Suppress hello for 0 neighbor(s)
Cryptographic authentication enabled
Sending SA: Key 1, Algorithm HMAC-SHA-256 - key chain ospf
Rollover in progress, 1 neighbor(s) using the old key(s):
key id 1 algorithm MD5
CSR103#
CSR103#
CSR103#sh ip ospf nei
Neighbor ID Pri State Dead Time Address Interface
172.16.100.3 1 FULL/DR 00:00:30 172.16.1.25 GigabitEthernet3
172.16.100.5 1 FULL/BDR 00:00:16 172.16.1.41 GigabitEthernet2
CSR103#
CSR103#
*Jan 11 16:49:35.311: %SYS-6-LOGOUT: User admin has exited tty session 1(10.228.200.250)
*Jan 11 16:49:45.396: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.100.5 on GigabitEthernet2 from
FULL to DOWN, Neighbor Down: Dead timer expired
```

Refer to the exhibit. Which configuration resolves the issue?

A)

```
router ospf 1
area 1 authentication message-digest
int GigabitEthernet 2
ip ospf message-digest-key 1 md5 cisco
```

B)

```
int GigabitEthernet 2
ip ospf message-digest-key 1 md5 cisco
ip ospf authentication message-digest
```

C)

```
int GigabitEthernet 2
ip ospf key 1 cisco
ip ospf authentication
```

D)

```
key chain ospf
key 1
key-string 7 02050D480809
cryptographic-algorithm hmac-sha-1
interface GigabitEthernet2 ip ospf
authentication key-chain ospf
```

A. Option A

B. Option B

C. Option C

D. Option D

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Explanation:

www.atmicnetworks.com

Question: 557

www.atmicnetworks.com

Refer to the exhibit.

www.atmicnetworks.com

Network > Device* MO

www.atmicnetworks.com

Router CSR301 .ap.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Jih 10. 2022 11 05 AM Syilitfi fttwutctt

Device Health: 10 Mrwu«kM<»_0 nm
Urvx* HMOTI ■ CJHjUtfceeon w M%

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Answer:

D

Ev*

tor'll

```

atomic-aggregate, best
  Extended Community: RT:1:4099
  rx pathid: 0, tx pathid: 0x0
  Updated on Jul 28 2022 15:17:49 UTC

router#
router#sh ip bgp 10.140.217.0/24
% Network not in table
router#

router#sh ip bgp 10.140.217.0/24
BGP routing table entry for 10.140.217.0/24, version 685
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    5          11
  Refresh Epoch 1
  65396, (aggregated by 65396 10.140.210.2), imported path from
  1:4099:10.140.217.0/24 (Guest_VN)

    10.140.212.5 from 10.140.212.5 (10.140.210.2)
      Origin IGP, metric 0, localpref 100, valid, external,
atomic-aggregate, best
  Extended Community: RT:1:4099
  rx pathid: 0, tx pathid: 0x0
  Updated on Jul 31 2022 18:32:12 UTC

```

Refer to the exhibit. In Cuco DNA Center, a network engineer identifies that BGP-learned networks are repeatedly withdrawn from peers. Which configuration must the engineer apply to resolve the Issue?

A)

```

router bgp 100
  bgp gracefulrestart

```

B)

```

router bgp 100
  bgp dampening

```

C)

route-map Dampening permit 10 set dampening 15
750 2000 60

router bgp 100
neighbor 10.140.212.5 route-map Dampening in

D)

route-map Dampening permit 10 set dampening 15
750 2000 60 router bgp 100 neighbor 10.140.212.5
route-map Dampening out

A. Option A

B. Option B

C. Option C

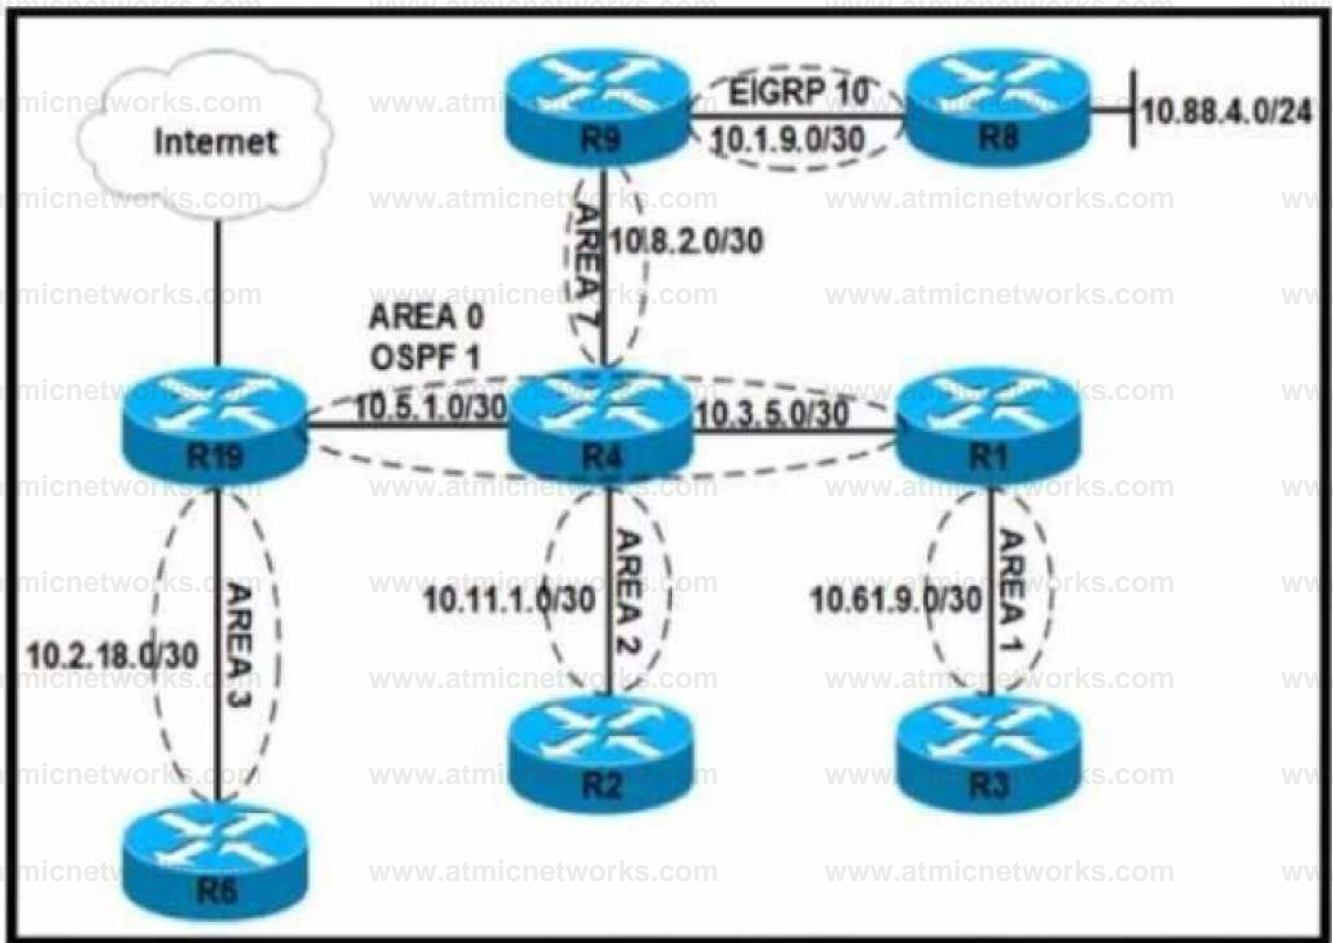
D. Option D

Answer: D

Explanation:

Question: 558

Refer to the exhibit.



Refer to the exhibit After an engineer modified the configuration for area 7 to permit type 1 2 and 7 LSAs only users connected to router R9 reported that they could no longer access the internet. Which configuration restores internet access to users on R9 and permits only LSA type 1,2, and 7?

A) R4»
 router ospf 1
 area 0 nssa default-information-originate
 network 10.5 1 0 0 0.0 3 area 0
 network 10.8.2 0 0.0.0.3 area 7

R9»
 router ospf 1
 area 7 MM
 redistribute eigrp 10 subnets
 network 10.8.2 0 0.0.0 3 area 7

B)

R4#

```
router ospf 1
  area 7 nssa no-summary network 10.5.1.0 0.0.0.3 area
0
  network 10.8.2.0 0.0.0.3 area 7
```

R9#

```
router ospf 1
  area 7 MM
  redistribute eigrp 10 subnets ^.^.^.^A a A A A A A 4 * _
```

c)

R#

```
router ospf 1
  area 7 nssa
  network 10.5.1.0 0.0.0.3 area 0
  network 10.8.2.0 0.0.0.3 area 7
```

R9#

```
router ospf 1
  area 7 nssa
  redistribute eigrp 10 subnets
  network 10.8.2.0 0.0.0.3 area 7
```

d)

R4s

```
router ospf 1
  ana 0 area 7 stub no-summary network 10.5.1.0 0.0.0.3
area 0 network 10.8.2.0 0.0.0.3 area 7
```

R9s

```
router ospf 1
  area 7 stub
  redistribute eigrp 10 subnets network 10.8.2.0 0.0.0.3
```

area 7

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Option B configures area 7 as a not-so-stubby area (NSSA) with the no-summary option on the area border router (ABR) R4. This allows the injection of external routes from EIGRP into the OSPF domain as type 7 LSAs, while preventing the propagation of inter-area summary LSAs into area 7. The nosummary option also generates a default summary route for area 7, which can be used by R9 to reach the internet².

Option A configures area 7 as a stub area, which does not allow any external routes or type 7 LSAs in the area. This prevents R9 from learning the EIGRP routes and accessing the internet³.

Option C configures area 7 as a NSSA without the no-summary option on R4. This allows the injection of type 7 LSAs into the area, but also allows the propagation of inter-area summary LSAs into the

area. However, this option does not generate a default summary route for area 7, which means R9 has no route to reach the internet4.

Option D configures area 7 as a NSSA with the default-information originate option on R4. This allows the injection of type 7 LSAs into the area, but also allows the propagation of inter-area summary LSAs into the area. The default-information originate option generates a type 7 default route for area 7, which can be used by R9 to reach the internet. However, this option is redundant and less efficient than Option B, because it injects both a type 3 and a type 7 LSA for the default route into the area5.

Question: 559

Refer to the exhibit.

```
R3#show ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	no route	
0.0.0.0/8	drop	
0.0.0.0/32	receive	
127.0.0.0/8	drop	
172.16.1.0/30	172.16.3.254	GigabitEthernet0/2
	172.16.4.254	GigabitEthernet0/3
172.16.3.252/30	attached	GigabitEthernet0/2
172.16.3.252/32	receive	GigabitEthernet0/2
172.16.3.253/32	receive	GigabitEthernet0/2
172.16.3.254/32	attached	GigabitEthernet0/2
172.16.3.255/32	receive	GigabitEthernet0/2
172.16.4.252/30	attached	GigabitEthernet0/3
172.16.4.252/32	receive	GigabitEthernet0/3
172.16.4.253/32	receive	GigabitEthernet0/3
172.16.4.254/32	attached	GigabitEthernet0/3
172.16.4.255/32	receive	GigabitEthernet0/3
172.16.222.254/32	172.16.4.254	GigabitEthernet0/3
192.168.100.0/24	172.16.3.254	GigabitEthernet0/2
192.168.200.0/24	172.16.3.254	GigabitEthernet0/2
192.168.222.0/24	172.16.4.254	GigabitEthernet0/3
224.0.0.0/4	drop	
224.0.0.0/24	receive	
Prefix	Next Hop	Interface
240.0.0.0/4	drop	
255.255.255.255/32	receive	

An engineer recently implemented uRPF by configuring the ip verify unicast source reachable-via rx command

on interface gi0/3 The engineer noticed right after implementing F that an inbound packet on the gi0-3 interface with a source address of 172.16.3.251 was dropped. Which action resolves the issue?

- A. Configure uRPF loose mode to forward the packet.
- B. Permit the 172.16.3.251 in the inbound ACL on interface gi0/3.
- C. Permit the 172.16.3.251 in the inbound ACL on interface gi0/3 to allow 172.16.3.251.
- D. Configure uRPF strict mode to forward the packet

Answer: A

Explanation:

Option A configures uRPF loose mode to forward the packet. This mode allows the router to check if there is a route in the routing table that matches the source IP address of the incoming packet, regardless of the interface that is used to reach the source. This mode is suitable for networks that have asymmetric routing, where the incoming and outgoing interfaces for a packet may differ².

Option B permits the 172.16.3.251 in the inbound ACL on interface gi0/3. This option does not resolve the issue, because it only allows the packet to pass the ACL check, but not the uRPF check. The packet will still be dropped by uRPF if there is no route to reach the source IP address via the same interface³.

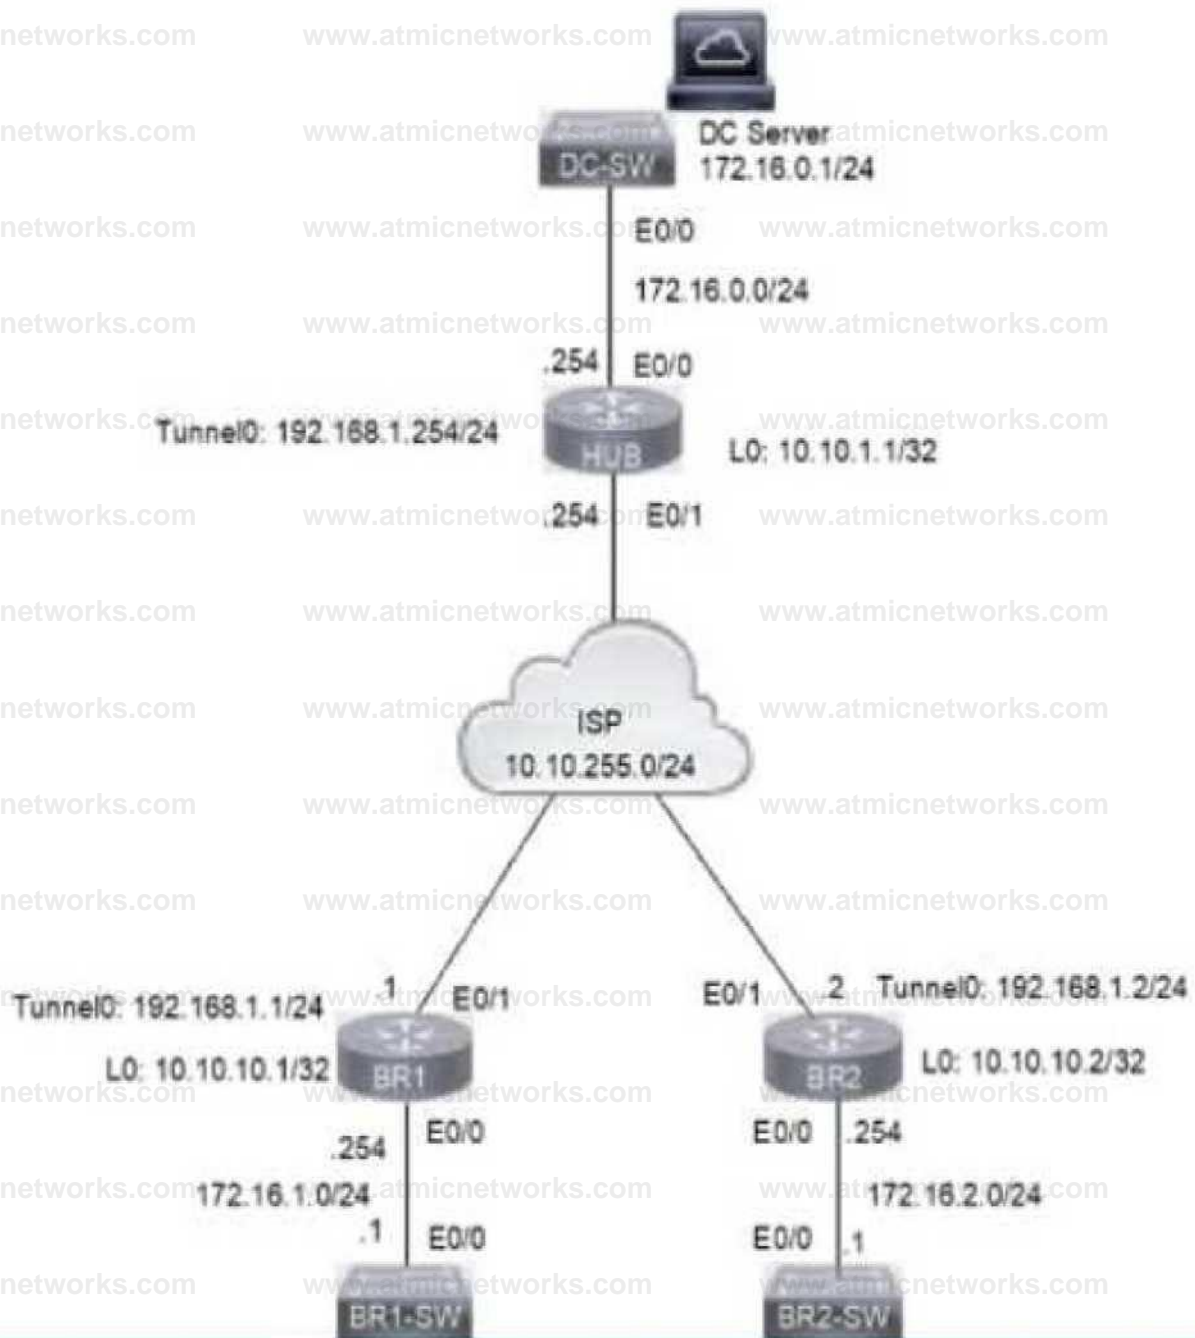
Option C permits the 172.16.3.251 in the inbound ACL on interface gi0/3 to allow 172.16.3.251. This option is redundant and incorrect, because it repeats the same IP address twice in the ACL statement. It also does not resolve the issue for the same reason as Option B³.

Option D configures uRPF strict mode to forward the packet. This option does not resolve the issue, because it is the same mode that was already configured on the interface. Strict mode requires that the router has a route to reach the source IP address via the same interface where the packet was received. If this condition is not met, the packet will be dropped⁴.

Question: 560

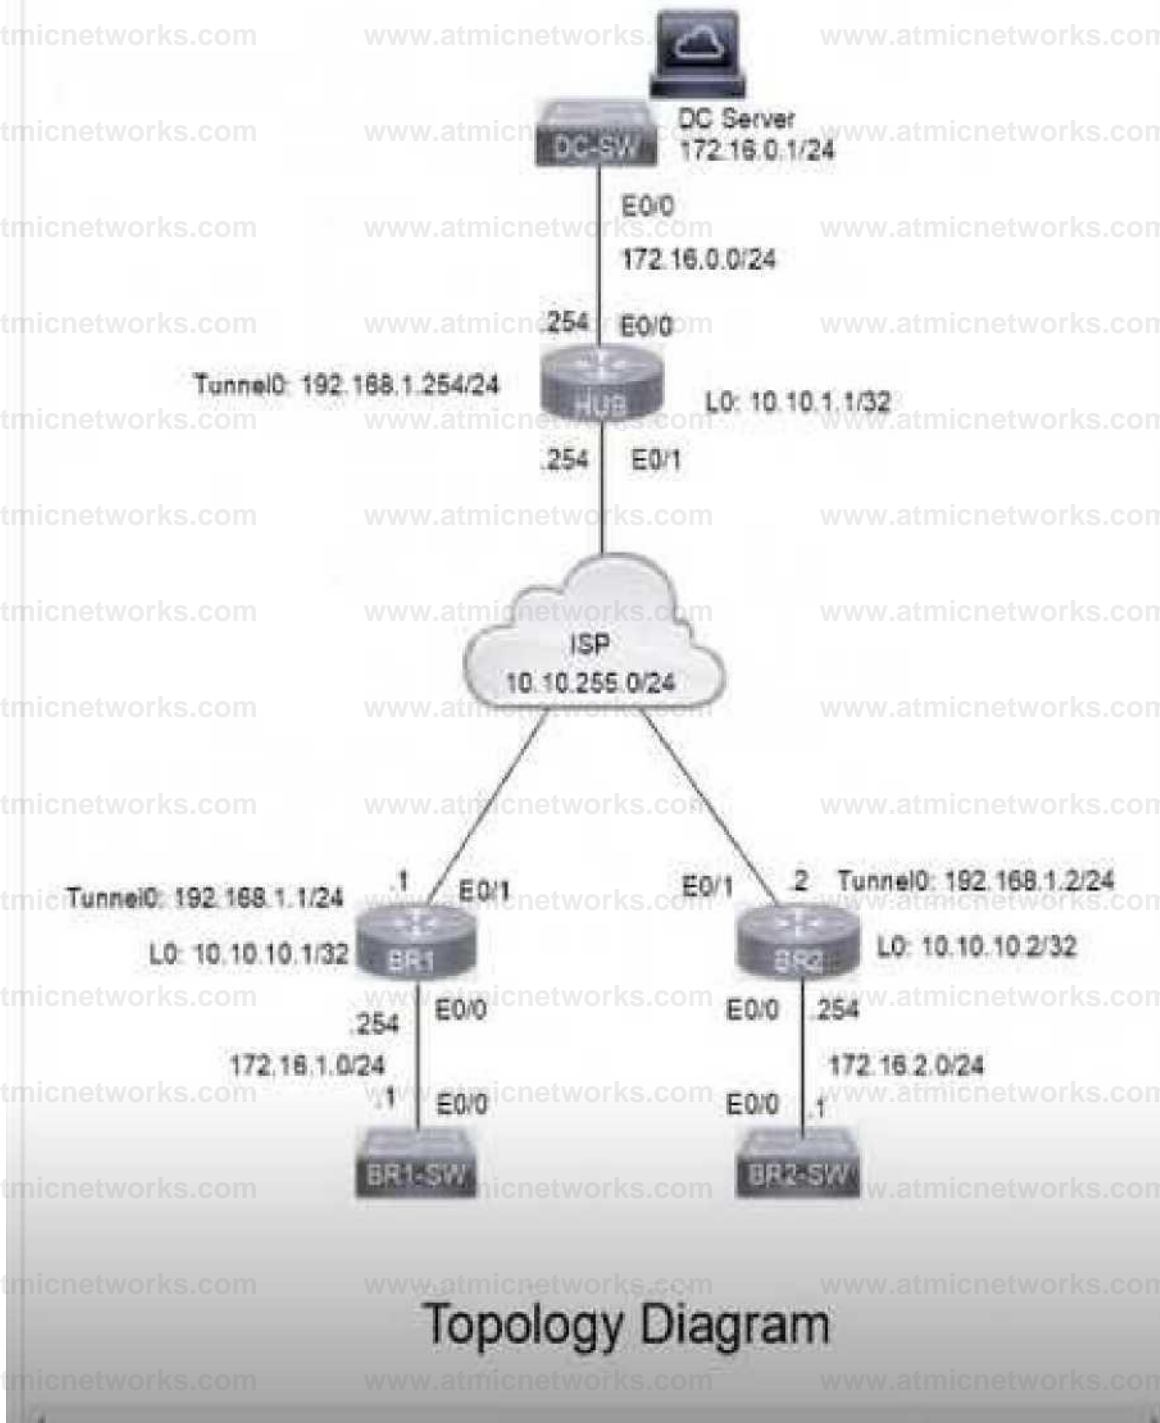
SIMULATION

A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS 100. The NHRP password is ccnp123, and the network ID and tunnel key is EIGRP ASN Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical interface as the tunnel source to achieve these goals:



A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB. IP connectivity, crypto policies, profiles, and EIGRP AS 100. The NHRP password is **ccnp123**, and the network ID and tunnel key is **EIGRP ASN**. Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical interface as the tunnel source to achieve these goals:

1. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR1.
2. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR2.
3. Ensure that packet fragmentation is done before encryption to account for GRE and IPsec header and allow a maximum TCP segment size of 1360 on an IP MTU of 1400 on the tunnel interfaces of both branch routers.
4. Apply an IPsec profile to the tunnel. Verify that direct spoke-to-spoke tunnel is functional between branch routers BR1



A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS 100. The NHRP password is ccnp123. and the network ID and tunnel key is EIGRP ASN. Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical Interface as the tunnel source to achieve these goals:

1. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR1.
2. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR2.
3. Ensure that packet fragmentation is done before encryption to account for GRE and IPsec header and allow a maximum TCP segment size of 1360 on an IP MTU of 1400 on the tunnel interfaces of both branch routers.
4. Apply an IPsec profile to the tunnel. Verify that direct spoke-to-spoke tunnel is functional between branch routers BR1 and BR2 by using traceroute to Ethernet 0/0 IP address to get a full score.

[Submit feedback about this item](#)

**Answer: See the
answer solution in
Explanation.**

Explanation:

ON BR1

```
Current Configuration :
```

```
interface Tunnel0
```

```

ip address 192.168.1.1 255.255.255.0
no ip redirects
ip situ 1400
ip nhrp nucbenticirian ccnp123
ip nhrp »np 192.168.1.254 10.10.255.254
ip nhrp MF BuiltiGflirt ■".-0,755,25-1
ip oh=p network id 100
ip nhrp holdtisiw 5
ip nhrp nhs 192.168.1.254
ip nhrp ahartent
ip ten nd] U Bl-sins 1360
delay 1000
tunnel mu roe 10.10.255.1
tunnel destination 10.10.255.254
tunnel key 100
end

```

```

BRI (config) *
ER1(config)#

```

ON BR2

DC-SW HUB BRI BR1-SW JS2 9R2-SW

```
UpDa Tine -> Up or Down Tune for a Tunnel
```

```
Interface: Tunnel0, IPv4 NHftP details Typt:Spokc, KURP Peers:
```

```
I tDt Petr KBMft Mdr Fe-r TtMhel Arid $tn^ Uc^M Tm AtErh
```

```
1 LD.lJh2SS.254 192.163.1-254 STOP 00111120 5
```

```
ERZ(config I(do show run int tu 0 BmiJing coiiflgiifstlon...
```

```
Current configuration ; 404 bytes
```

```
inter^CC Tunnel0
```

```
Ip iddiMi 192.162.1.2 255.25S.2SS.-3 no ip redirects
```

```
>9 ME'J 1400
```

```
io nhrp authcntiction conpLSS
```

```
ip nhrp itap 192.16?,1.254 10.10.255.254
```

```
ip nhrp cap taultiaat 10.10.355.251
```

```
ip nhrp netuort-id 100
```

```
ip nhrp hbldcinc 5
```

```
ip nhrp nhs 192. Ltt. 1.25--
```

```
ip nhrp shortcut
```

```
.c zap adjust-msi ISSA delay 1000 tunnel source 10.10.10.2 tunnel
```

```
dstmuriion 10-10-355,25^ tunnel key UM end
```

Question: 561

How are CE advertised routes segmented from other CE routers on an MPLS PE router?

- A. with a combination of VRF-Lite and MP-BGP
- B. by pushing MPLS labels advertised by LDP on customer routes
- C. by enabling multiple instances of BGP. one for each CE router
- D. by assigning CE-facing interfaces to different VRFs

Answer: D

Explanation:

In an MPLS PE router, CE advertised routes are segmented from other CE routers by assigning CE-facing interfaces to different Virtual Routing and Forwarding (VRF) instances¹². A VRF is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time. By associating a network interface with a VRF, the network layer (Layer 3) has a different view of the network topology. This allows the segmentation of routing paths for traffic from different customers or different types of traffic. In the context of MPLS, VRFs are used to create separate routing instances for each customer on a PE router³.

Reference:

Implementing Cisco Enterprise Advanced Routing and Services (ENARSI) training videos

CCNP Enterprise Advanced Routing ENARSI 300-410 Official Cert Guide

Implementing Cisco Enterprise Advanced Routing and Services source documents or study guide

MPLS routes to advertise from CE to PE - Cisco Community

Configuring Route Exchange Between PE and CE Devices - CloudEngine 12800 and 12800E V200R005C10 Configuration Guide - VPN - Huawei

Customer edge router - Wikipedia

Question: 562

Refer to the exhibit.

Network > Device 360

●		DEVICE_AVAILABILITY:REACHABLE	Event	8:52:52.443 PM
●	Notice	DUAL:NBRCHANGE	Syslog	8:46:37.210 PM
●	Notice	DUAL:NBRCHANGE	Syslog	8:46:37.207 PM

● DUAL_NBRCHANGE Jan 11, 2022 8:46:37 PM

Detailed Information

Severity	Notice
Mnemonic	NBRCHANGE
Facility	DUAL
Message Text	682: *Jan 11 15:41:03.036: EIGRP-IPv4 88: Neighbor 172.16.33.2 (GigabitEthernet2/10) is down: authentication mode changed
Message Type	Syslog

R1 test its directly connected EIGRP peer 172.16.33.2 (SW1). Which configuration resolves the issue1?

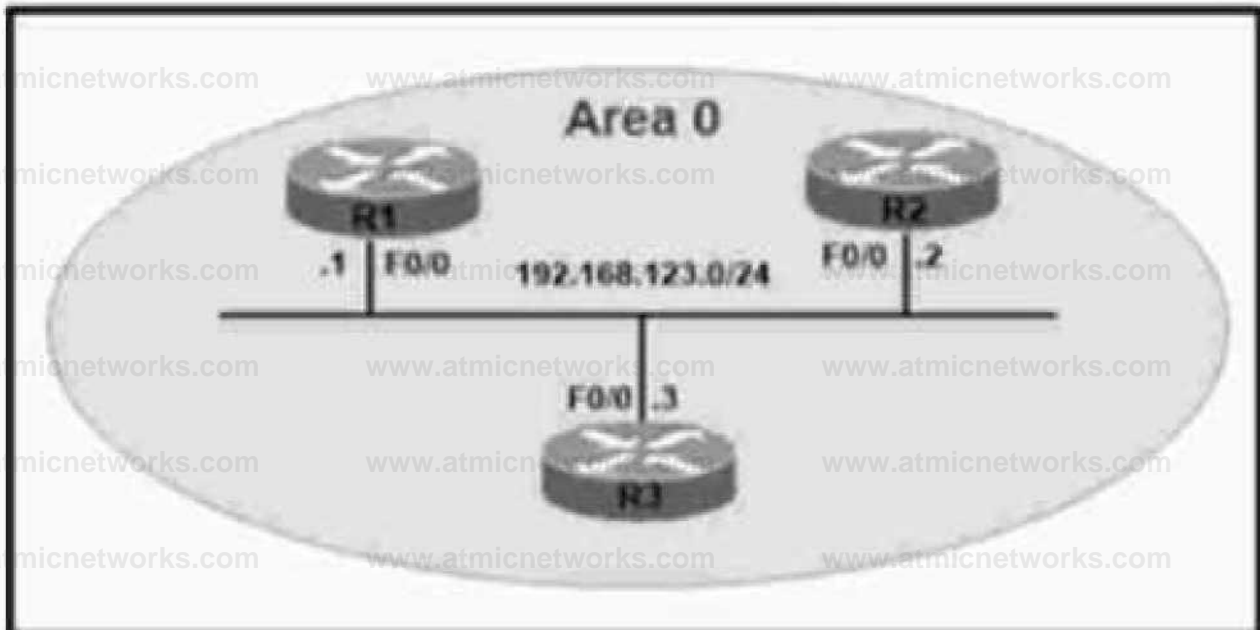
- A. key chain EIGRP key 1 key-string Cisco!interface Gigabit Ethernet 2 IP authentication mode elgrp 88 md5 IP authentication key-chain elgrp 88 EIGRP
- B. key chain EIGRP key1 key-string Cisco!interface Gigabit Ethernet 2.10 IP authentication mode eigrp 88 md5 IP authentication key-chain eigrp 88 Cisco
- C. key chain EIGRP key 1 key-string Cisco linterface Gigabit Ethernet 2.10 IP authentication mode elgrp 88 md5 IP authentication key-chain eigrp 88 EIGRP
- D. key chain EIGRP key1 key-string Cisco linterface Gigabit Ethernet 2 IP authentication mode eigrp 88 md5 IP authentication key-chain elgrp 88 Cisco

Answer: B

Explanation:

Question: 563

Refer to the exhibit.



Router R1 Output:

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.123.2	1	FULL/BDR	00:00:32	192.168.123.2	FastEthernet0/0
192.168.123.3	1	FULL/DR	00:00:31	192.168.123.3	FastEthernet0/0

Router R2 Output:

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.123.1	1	FULL/DROTHER	00:00:31	192.168.123.1	FastEthernet0/0
192.168.123.3	1	FULL/DR	00:00:32	192.168.123.3	FastEthernet0/0

Router R3 Output:

R3#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.123.1	1	FULL/DROTHER	00:00:31	192.168.123.1	FastEthernet0/0
192.168.123.3	1	FULL/DR	00:00:32	192.168.123.3	FastEthernet0/0

Router R3 Output:

R3#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.123.1	1	FULL/DROTHER	00:00:36	192.168.123.1	FastEthernet0/0
192.168.123.2	1	FULL/BDR	00:00:39	192.168.123.2	FastEthernet0/0

Refer to the exhibit. An administrator wanted to make R1 always elected as DR, R2 as BDR, and R3 as DROTHER but could not achieve the desired results. Which two configurations resolve the issue? (Choose two.)


- A. On the R1F0/0 interface, configure OSPF priority to255.
- B. On the R2F0/0 interface, configure OSPF priority to201.
- C. On the R1F0/0 interface, configure OSPF priority to202.
- D. On the R3F0/0 interface, configure OSPF priority to201.
- E. On the R2F0/0 interface, configure OSPF priority to200.

Answer: A,D

Explanation:

Question: 564

Refer to the exhibit.



Refer to the exhibit. A network engineer configured routers R1 and R2 with MP-BGP. The engineer noticed that the routers cannot exchange any IPv6 routes, however, the IPv4 neighbor relationship is working fine. Which configuration must the engineer apply to router R2 to exchange IPv6 routes?

- A. `ipv6cef | Interface Loopback10 | ipv6 address 2001: DB8:128::2/128 | Interface GigabitEthernet1/0 | ipv6 address 2001:DB8:1::2/64 | router bgp 65002 | no bgp default ipv4-unicast | neighbor 2001: DB8:1::1 remote-as 65001 | address-family ipv6 | network 2001:DB8:128::2/128 | neighbor`

2001:DB8:1::1 activate

B. Ipv6 unicast-routing ipv6 cef ! interface Loopback100 ipv6 address 2001: DB8:12C:2 129 ! interface GigabitEthernet1/0 ipv6 address 2001: DB8:1::2 64 description AS65001 ID B466:A83D:3D7::1 ! router bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 ! address-family ipv4 neighbor 2001:DB8:1::1 activate

C. Ipv6 unicast-routing ipv6 cef ! interface Loopback100 ipv6 address 2001:DB8:128::2/128 ! interface GigabitEthernet1/0 ipv6 address 2001:DB8:1::2/64 ! router bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 ! address-family ipv6 network 2001:DB8:128::2/128

D. Ipv6 unicast-routing ipv6 cef ! interface Loopback100 ipv6 address 2001: DB8:128::2/128 ! interface GigabitEthernet1/0 ipv6 address 2001:DB8:1::2/64 ! router bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 ! address-family ipv6 network 2001:DB8:128::2/128 neighbor 2001:DB8:1::1 activate

Answer: D

Explanation:

Question: 565

What is the role of LDP in MPLS networks?

- A. It enables label binding that exchanges route descriptors
- B. It creates MPLS packet forwarding along with the IGP routes.
- C. It disables label binding information to exchange with peer LSRs.
- D. It enables label binding information to exchange with peer LSRs

Answer: D

Explanation:

The Label Distribution Protocol (LDP) plays a crucial role in MPLS (Multiprotocol Label Switching) networks. It enables label switch routers (LSRs) to exchange label binding information, which is essential for supporting hop-by-hop forwarding in an MPLS network¹.

LDP establishes LSPs (Label Switched Paths) that follow the existing IP routing table². This process is particularly well-suited for establishing a full mesh of LSPs between all of the routers on the network².

When a packet arrives at a router in an MPLS network, the router looks at the incoming label, looks up the label in a table, and then forwards the packet to the next hop¹. This method of label distribution is also called hop-by-hop forwarding¹.

Reference:

MPLS LDP (Label Distribution Protocol) - NetworkLessons.com

MPLS Label Distribution Protocol (LDP) - Cisco

What is Label Distribution Protocol (LDP)? - Metaswitch

Label Distribution Protocol - Wikipedia

What is MPLS LDP in networking? – CCNA-Classes

Question: 566

Refer to the exhibit.

```
ip sla 10
  icmp-echo 10.1.1.10
  timeout 2000
  threshold 2000
  frequency 40
ip sla schedule 10 life forever start-time now
!
track 1 ip sla 10 reachability
```

An engineer configured IP SLA to monitor a next hop on a router for reachability. When the next hop is unreachable, the router is executing tracking and failing over another route, but packet loss is experienced because the reachability is flapping. Which action resolves the issue?

A. Append delay up 0 down 0 to the track command

- B. Increase the timeout of the sla probe to 6000
- C. Append delay up 50 down 60 to the track command
- D. Increase the frequency of the sla probe to 60.

Answer: C

Explanation:

IP SLA (Internet Protocol Service Level Agreement) is a feature that allows you to measure network performance such as latency, jitter, packet loss, and so on. In this case, it's being used to monitor the reachability of a next hop on a router123.

When the next hop is unreachable, the router is executing tracking and failing over to another route. However, packet loss is experienced because the reachability is flapping. This could be due to the router switching back and forth between the primary and backup routes too quickly.

To resolve this issue, you can introduce a delay in the tracking process. This can be done by appending a delay to the track command. Option C suggests appending a delay of 50 seconds for the up state and 60 seconds for the down state. This means that the router will wait for 50 seconds before declaring the tracked object as up (reachable) and 60 seconds before declaring it as down (unreachable). This delay can help prevent the router from switching routes too quickly, thus reducing the chances of reachability flapping and packet loss.

Reference:

[Configure ISP Failover with Default Routes Using IP SLA Tracking - Cisco](#)

[Using IPSLA to change routing - Cisco Community](#)

[How to Use IP SLA Technology to Assess WAN Performance](#)

[Reliable Static Routing with IP SLA - NetworkLessons.com](#)

[Configuring Static Route Tracking using IP SLA \(Basic\)](#)

Question: 567

The network administrator configured CoPP so that all SNMP traffic from Cisco Prime located at 192.168.1.11 toward the router CPU is limited to 1000 kbps. Any traffic that exceeds this limit must be dropped.

```
access-list 100 permit udp any any eq 161
```

```
class-map CM-SNMP
```

```
match access-group 100
```

```
!
```

```
policy-map PM-COPP
```

```
class CM-SNMP
```

```
police 1000 conform-action transmit
```

```
!
```

```
control-plane
```

```
service-policy input PM-COPP
```

The network administrator is not getting the desired result for the SNMP traffic and SNMP traffic is getting dropped frequently. Which set of configurations resolves the issue?

A. no access-list 100
access-list 100 permit tcp host 192.168.1.11 any eq 161

B. no access-list 100
access-list 100 permit udp host 192.168.1.11 any eq 161
!
policy-map PM- COPP
class CM-SNMP
no
police 1000 conform-action transmit
police 1000000 conform-action transmit!
control-plane
no service-policy input PM-
COPP!
interface E 0/0
service-policy input PM- COPP!
interface E 0/1
service-policy input PM-COPP

C. no access-list 100
access-list 100 permit udp host 192.168.1.11 any eq 161!
policy-map PM- COPP
class CM-SNMP
no
police 1000 conform-action transmit
police 1000000 conform-action transmit

D. policy-map PM-COPP
class CM-SNMP
no police 1000 conform-action transmit
police 1000000 conform-action
transmit

Answer: C

Explanation:

In the context of Control Plane Policing (CoPP) in Cisco devices, the rate limit is specified in bits per second (bps), not kilobits per second (kbps). Therefore, a limit of 1000 kbps should indeed be entered as 1,000,000 bps in the CoPP configuration.

Also, the access list should be configured to match the specific SNMP traffic from the Cisco Prime IP address (192.168.1.11), as you correctly pointed out.

Here's the corrected configuration:

```
no access-list 100
access-list 100 permit udp host 192.168.1.11 any eq 161
!
policy-map PM-COPP
```

```
class CM-SNMP
```

```
no police 1000 conform-action transmit
police 1000000 conform-action transmit
```

This configuration ensures that only the SNMP traffic from Cisco Prime is policed and any excess traffic is dropped, preventing the router's CPU from being overwhelmed.

Question: 568

Which Layer 3 VPN attribute allows different customers to connect to the same MPLS network wrth overlapping IP ranges?

- A. VRF
- B. RT
- C. MP-BGP

D. RD

Answer: D

Explanation:

In a Layer 3 VPN (Virtual Private Network) over an MPLS (Multiprotocol Label Switching) network, the attribute that allows different customers to connect with overlapping IP ranges is the Route Distinguisher (RD)123.

RD is a unique identifier that is prepended to each IP address in a customer's VPN to create a unique VPNv4 address. This allows customers to use overlapping IP addresses without conflict123. The RD makes it possible for the same IP prefix to exist in different VPNs, which is crucial when customers have overlapping IP ranges123.

Reference:

MPLS Layer 3 VPN Explained - NetworkLessons.com

MPLS: Layer 3 VPNs Configuration Guide - Cisco

Understanding Using MPLS-Based Layer 3 VPNs on Switches - Juniper

Question: 569

Network engineer must configure an EIGRP stub router at a site that advertises only connected and summary routes.

Which configuration performs this task?

A)

router eigrp 100 eigrp stub

B)

**router eigrp 100
eigrp stub summary**

C)

router eigrp 100 eigrp stub connected

D)

router eigrp 100 eigrp stub redistribute

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

In EIGRP (Enhanced Interior Gateway Routing Protocol), a stub router is one that is connected to one or more neighbors and should not be a transit router¹. The EIGRP stub routing feature improves network stability, reduces resource utilization, and simplifies the stub router's configuration¹.

The network engineer wants to configure an EIGRP stub router that advertises only connected and

summary routes. This can be achieved with the configuration provided in Option B:

router eigrp 100

eigrp stub summary

This configuration ensures that the EIGRP stub router advertises only connected and summary routes. The summary keyword after the eigrp stub command indicates that the router should advertise only auto-summarized or statically configured summary routes.

Reference:

EIGRP STUB and Configuration - Cisco Community

EIGRP Stub And Summary Routes Explained - Networkel

Question: 570

Refer to the exhibit.

R1#debug ip ospf ad


23:42:08.259: OSPF; Send DBD to 2.2.2.2 on Ethernet0/0 seq 0x52 flag 0x7 Ion 32 23:42:08.339: OSPF; Rev DBD from 2.2.2.2 on Ethernet0/0 seq 0x836 opt 0x52 flag 0x7 Ion 32 mtu 1532 atato EXSTART

R2#debug ip ospf adj

23:42:08.423: OSPF; Send DSD to 1.1.1.1 on Ethernet0/0 seq 0x834 opt 0x52 flag 0x7 lan 32

23:42:08.423: OSPF; First DBD and we are not SLAVE

23:42:08.511: OSPF; Rev DBD from 1.1.1.1 on Ethernet0/0 seq 0x836 opt 0x52 flag 0x2 Ion 52 ntu 1300 atato EXSTART



Refer to the exhibit R1 cannot establish a neighbor relationship with R2 Which action resolves the issue?"

- A. Configure the mtu ignore command on the Interfaces of R1 and R2
- B. Configure the ip ospf network point-to-point command on the interfaces of R1 and R2
- C. Configure the ip ospf network broadcast command on the interfaces of R1 and R2
- D. Configure the neighbor 2.2.2.2 command on R1 under the OSPF process

Answer: A

Explanation:

In OSPF (Open Shortest Path First), the MTU (Maximum Transmission Unit) size must match on both sides of a link for OSPF neighbors to form an adjacency¹². If the MTU sizes do not match, the OSPF adjacency will not form, and the routers will not become neighbors¹².

In this case, R1 cannot establish a neighbor relationship with R2. One possible reason for this could be a mismatch in the MTU size on the interfaces of R1 and R2¹².

To resolve this issue, you can configure the `mtu ignore` command on the interfaces of R1 and R2 (Option A). This command allows OSPF to ignore the MTU size when determining if it can form an adjacency with a neighbor¹². This means that even if the MTU sizes do not match, OSPF will still form an adjacency, and R1 and R2 will become neighbors¹².

Reference:

[OSPF MTU Mismatch - NetworkLessons.com](#)

[Understanding OSPF MTU Mismatch Conditions and Solutions - Cisco](#)

Question: 571

Refer to the exhibit.

```

Router# show ip interface brief
Interface      IP-Address      OK? Method Status  Protocol
Ethernet0/0    90.0.1.1        YES NVRAM  up      up
Ethernet1/0    118.8.10.1     YES NVRAM  up      up
Loopback6      10.1.7.6        YES NVRAM  up      up
Loopback7      10.1.7.7        YES NVRAM  up      up

```

```

Router# show ip bgp
BGP table version is 10, local router ID is 10.1.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 10.0.0.0       0.0.0.0         0 32768 ?
*> 10.1.7.7/32    0.0.0.0         0 32768 i
*> 90.0.0.0       0.0.0.0         0   32768 ?
r> 90.0.1.0/24    100.0.1.254     0     0 6 ?
*> 118.0.0.0     0.0.0.0         0   32768 ?
r>118.8.10.0/24   118.8.10.254    0  100  0 ?

```

Which action adds the 10.1.7.6-32 route to the BGP table?

- A. Add a static route for the 10.1.7.6/32 network
- B. Add the network 10.1.7.6 mask 255.255.255.255 backdoor command
- C. Add summary-address 10.1.7.6.255.255.255
- D. Add the network 10.1.7.6 mask 255.255.255.255 command

Answer:

D

Explanation:

Question:
572

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


- A. extended-community
- B. MPLS label
- C. RD
- D. RT

Answer: C

Explanation:

Question: 573

Refer to the exhibit.



Refer to the exhibit. R1 and R2 have been configured where the neighbor relationship must be authenticated using MD5:

```
R1
router bgp 100
neighbor 172.16.1.2 remote-as 200
neighbor 172.16.1.2 password cisco123
```

```
R2
router bgp 200
neighbor 172.16.1.1 remote-as 100
neighbor 172.16.1.1 password cisco
```

The neighbor relationship is not coming up. Which configuration resolves the issue?

R1

```
router bgp 100
neighbor 172.16.1.2 password cisco
```

R2

```
router bgp 200
neighbor 172.16.1.1 password cisco123
```

R1

```
router bgp 100
neighbor 172.16.1.2 password MDS cisco123
```

R2

```
router bgp 200
neighbor 172.16.1.1 password MDS cisco123
```

R1

```
router bgp 100
neighbor 172.16.1.2 password MDS cisco
```

R2

```
router bgp 100
neighbor 172.16.1.1 password MDS cisco
```

A. Option A

B. Option B

C. Option C

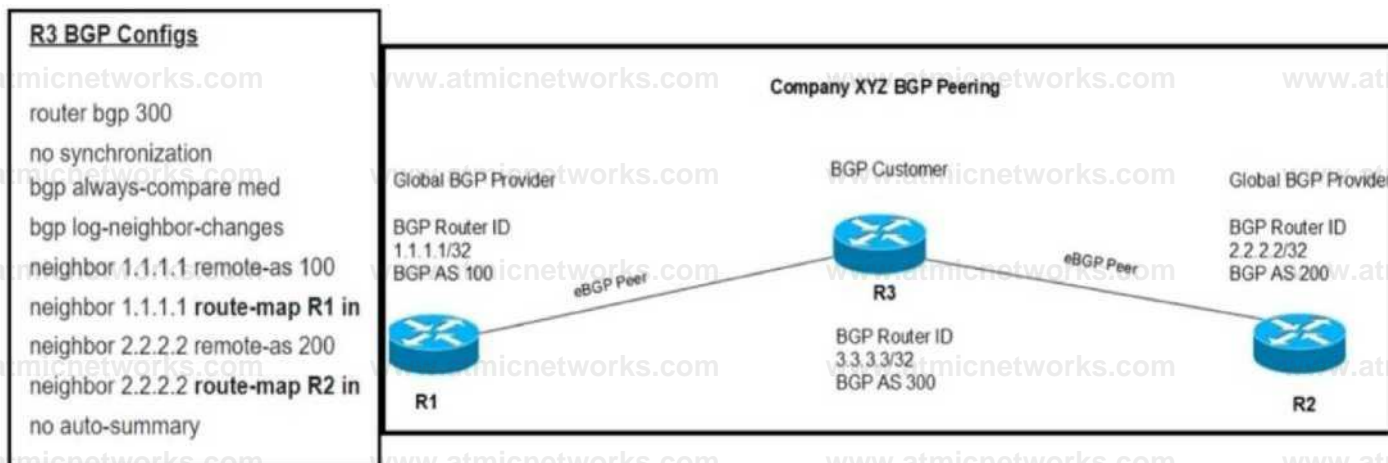
D. Option D

Answer: B

Explanation:

Question: 574

Refer to the exhibit.



Refer to the exhibit. R3 is dual-homed to two service providers for traffic redundancy. R3 prefers its outbound traffic via R2. Which set of configurations achieves this goal?

route-map R1 permit 10 set metric 200
route-map R2 permit 10 set metric 100

route-map R1 permit 10 set metric 100
route-map R2 permit 10 set metric 200

route-map R1 permit 10
set AS-Path Prepend 200
route-map R2 permit 10
set AS-Path Prepend 100

route-map R1 permit 10
set AS-Path Prepend 100

route-map R2 permit 10
set AS-Path Prepend 200

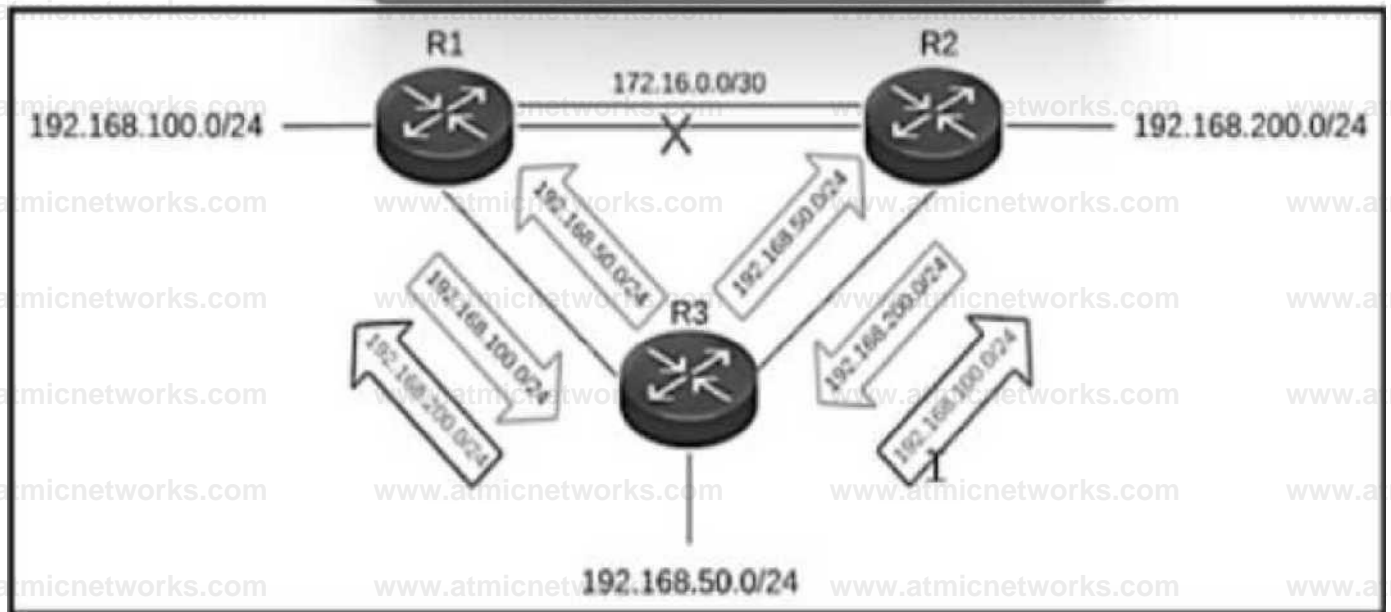
- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 575

Refer to the exhibit.



Refer to the exhibit. The primary link between R1 and R2 went down but R3 is still advertising the 192.166.200.0/24 network to R1 and the 192.166.100.0/24 network to R2, which creates a loop. Which action resolves the issue?

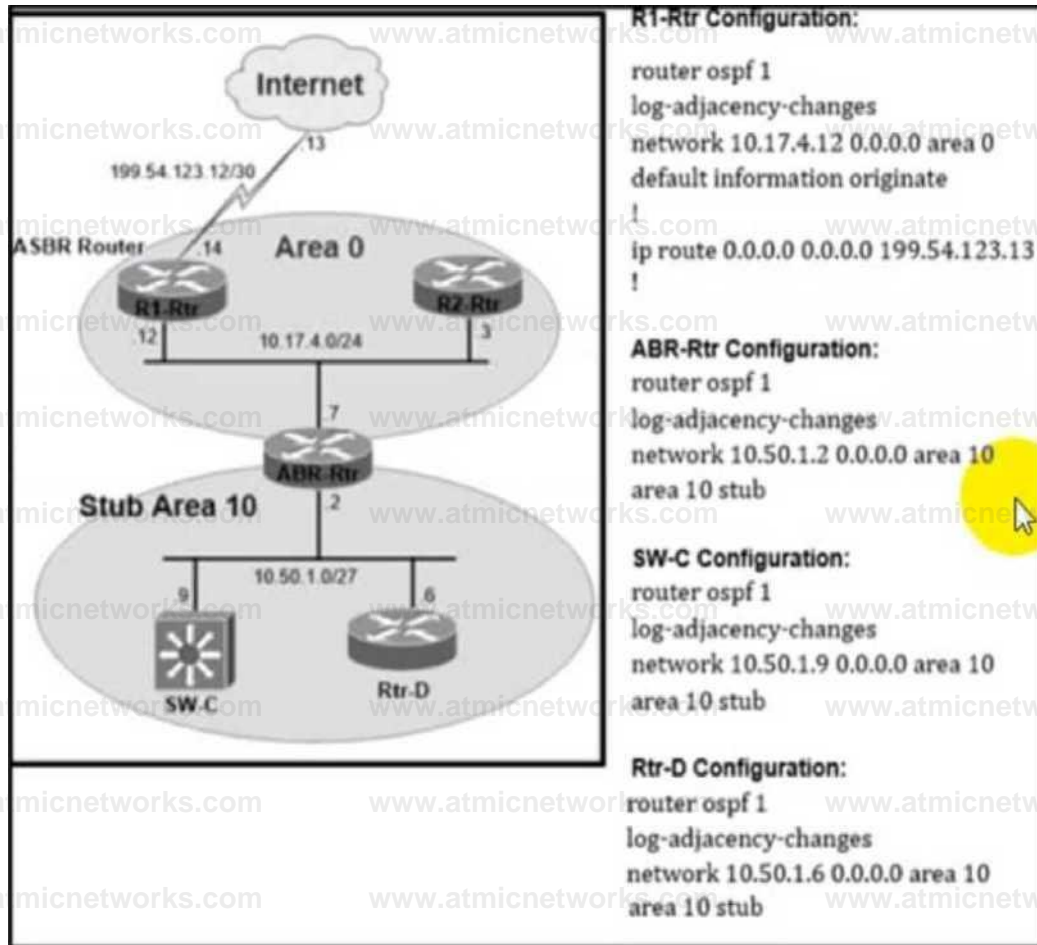
- A. Configure the eigrp stub command under the EIGRP process on R3
- B. Configure the summary address 192.168.0.0 255.255.0.0 100 command on R3
- C. Configure the eigrp stub leak-map command under the EIGRP process on R1
- D. Configure the eigrp stub command under the EIGRP process on R2

Answer: A

Explanation:

Question: 576

Refer to the exhibit.



Refer to the exhibit. Router ABR-Rtr is not propagating the internet routes in OSPF area 10, which causes internet reachability problems in the area.

a. Which action resolves the issue?

- A. ABR-Rtr must configure the default-information originate always command.
- B. ABR-Rtr must configure the area 10 stub no-summary command.
- C. ABR-Rtr network type must be broadcast network.
- D. ABR-Rtr must advertise the 0.0.0.0/0 default route in area 10.

Answer: D

Explanation:

Question: 577

An engineer must configure encrypted packets for a single router OSPF neighborsMp Which configuration meets this requirement?

```
interface Ethernet0/2
```

```
  ip ospf authentication message-digest  
  ip ospf message -digest-key 1 md5 exam
```

```
router ospf 100
```

```
  area 0 authentication message-digest-key 1 md5 exam
```

```
interface Ethernet0/2
```

```
  ip ospf message-digest-key 1 md5 exam
```

```
router ospf 100
```

```
  area 0 authentication
```

```
interface Ethernet0/2
```

```
  ip ospf authentication-key exam
```

```
interface Ethernet0/2
```

```
  ip ospf authentication-key exam
```

```
router ospf 100
```

```
  area 0 authentication message-digest-key 1 md5  
  exam
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

Question: 578

The network administrator deployed the Binding Table Recovery feature. Which two devices recover the missing binding table entries? (Choose two.)

A. DHCP client

B. DHCP server

C. destination host

D. source host

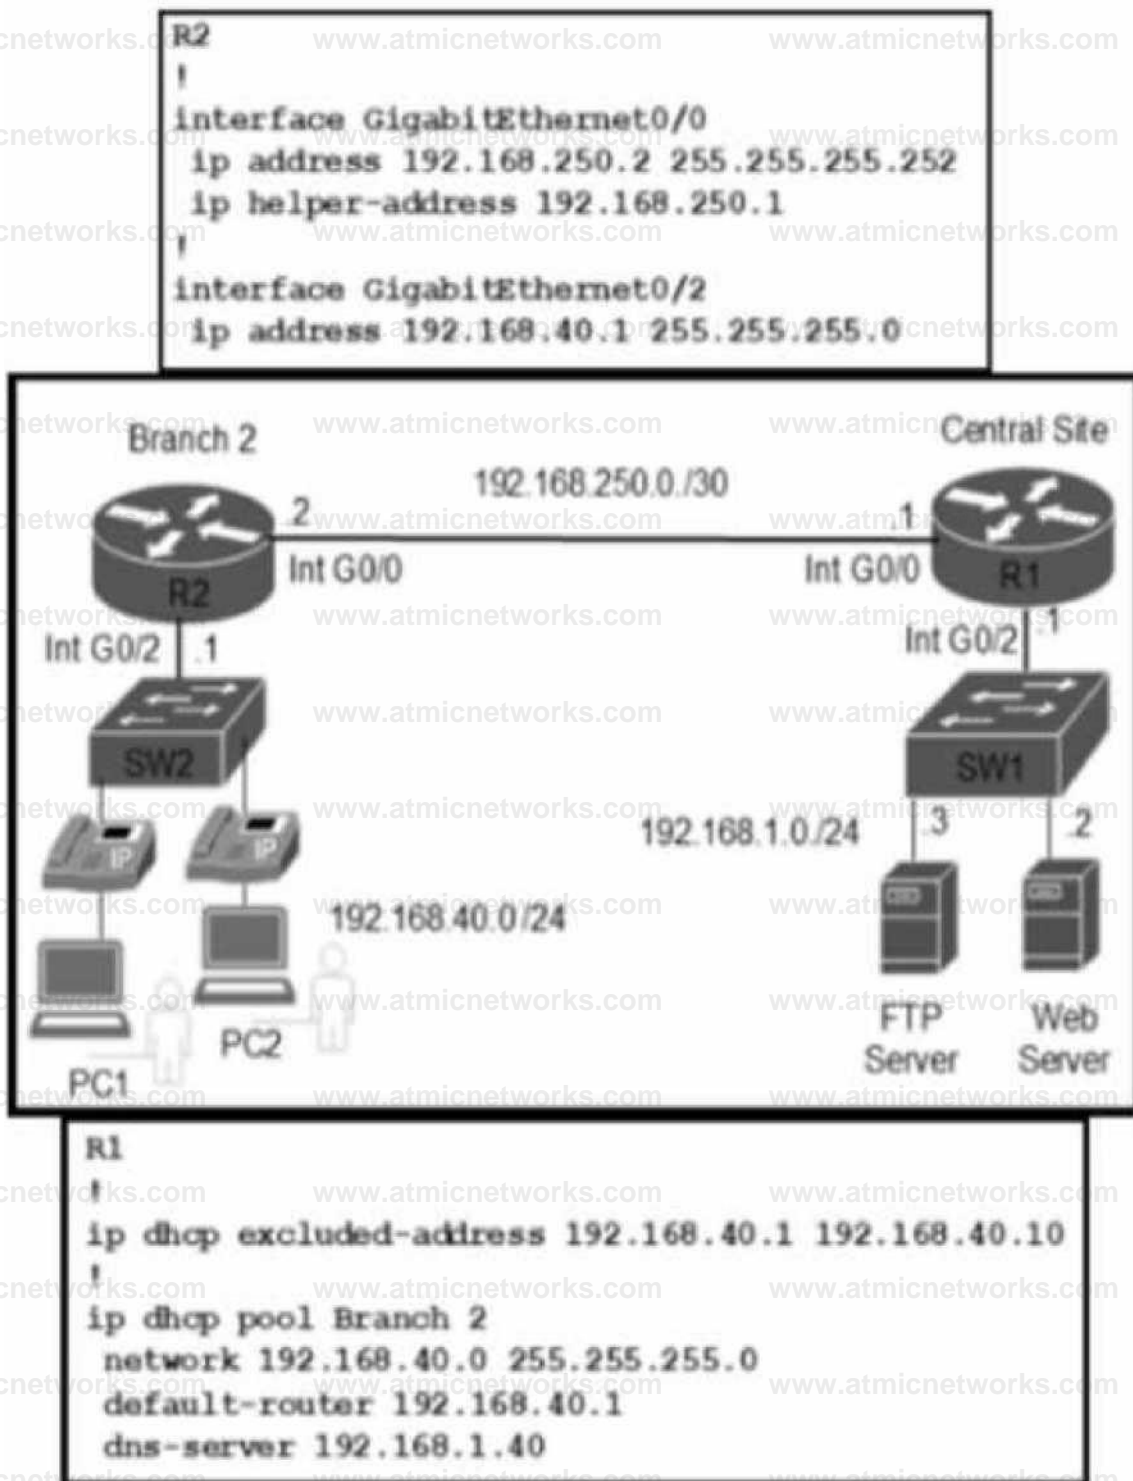
E. DHCP relay agent

Answer: A, E

Explanation:

Question: 579

Refer to the exhibit.



Refer to the exhibit. Branch 2 hosts cannot receive dynamic IP addresses. Which action resolves the issue?

- A. Configure the help command on the interface GigabitEthernet 0/2 of the R2 router
- B. Configure the help command on the Layer 2 switch SW2 interfaces
- C. Configure the help command on the interface GigabitEthernet 0/2 of the DHCProuter.
- D. Configure the help command on the interface GigabitEthernet 0/0 of the DHCProuter.

Answer: A

Explanation:

Question: 580

Refer to the exhibit.

```
R1# show ip route static
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
S       10.100.13.0/24 is directly connected, Loopback10
R1# show run | begin router ospf 1
router ospf 1
 redistribute static
 network 10.100.11.0 0.0.0.255 area 1
 network 10.100.12.0 0.0.0.255 area 1
 network 10.100.14.0 0.0.0.255 area 1

R2# show ip route 10.100.13.0
% Subnet not in table
R2# show ip ospf neighbor | include 10.100.14.1
10.100.14.1      1    FULL/DR           00:00:39      10.100.14.1
GigabitEthernet0/0
```

Refer to the exhibit. R1 is directly connected to R2 over network 10.100.14.0/24. An engineer configures R1 to advertise a static route that is connected to a local loopback for network 10.100.13.0/24. The network is not in the routing table of R2.

Which action resolves the issue?

- A. The redistribution command is incorrect on R1 The default metric metric 200 should be inducted with the redistribution command.

B. The Loopback interface on R1 is administratively down The interface should be enabled with the `no shutdown` command

C. R2 must use a different OSPF process number and should be changed to `ospf 1` to match R1

D. The redistribution command is incorrect on R1 The keyword `subnets` should be included with the `redistribution` command

Answer: D


Explanation:

Question: 581

Refer to the exhibit.

```
RT1#sh ip int br
Interface          IP-Address      OK? Method
-----
GigabitEthernet0/0 unassigned     YES N/A    administratively
GigabitEthernet0/1 172.16.255.1   YES manual  up
GigabitEthernet0/2 172.16.255.14  YES manual  up
GigabitEthernet0/3 172.16.2.17   YES manual  up
RT1#
RT1#sh run | begin router eigrp
router eigrp 100
 network 172.16.2.0 0.0.0.3
 network 172.16.2.16 0.0.0.15
!
no ip http server
no ip http secure-server
ip http access-interface GigabitEthernet0/3
!
line con 0
line vty 0 4
login
transport input none
!
```

```
RT2#sh run
!
hostname R1
!
ip conf
!
interface GigabitEthernet0/0
 ip address 172.16.2.2 255.255.255.252
 ip access-group 120 in
!
interface GigabitEthernet0/1
 ip address 172.16.2.17 255.255.255.240
!
router eigrp 100
 network 172.16.2.0 0.0.0.3
 network 172.16.2.16 0.0.0.15
!
access-list 120 permit udp host 172.16.1.2 host 172.16.2.17 eq tftp
access-list 120 deny   udp any any eq tftp
access-list 120 permit tcp any any
```



Refer to the exhibit. The engineer is trying to transfer the new IOS file to the router R3 but is getting an error. Which configuration achieves the file transfer?

R4 configure access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
 R3 configure access-list 120 permit tcp any any

R4 configure no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
 R3 configure access-list 120 permit udp host 172.16.1.17 host 172.16.2.19 eq 69
 R4 configure access-list 120 permit tcp any any

R4 configure no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
 R4 configure no access-list 120 deny udp any any eq tftp
 R4 configure access-list 120 permit tcp any any

R4 configure no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
 R4 configure access-list 120 permit tcp host 172.16.1.17 host 172.16.2.19 eq 69
 R4 configure access-list 120 permit tcp any any

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

Question: 582

Which type of ports are protected by IPv6 Source Guard?

- A. Layer 2 ports
- B. access ports
- C. Layer 3 ports
- D. trunk ports

Answer: B

Explanation:

Question: 583

While BGP internet routes are redistributed to a lower class of router via RIP, packets are being dropped and routes are failing to be distributed in RIP. Which action resolves the issue?

- A. Use OSPF instead of RIP to accept all BGP routes.
- B. Use the input-queue command to prevent the loss of packets.
- C. Use WFQ in the output queue of the high-performance router.
- D. Use RIP V2 to be able to use classless networks from BGP

Answer: B

Explanation:

Question: 584

Which IP precedence value does BFD use to prioritize traffic within an infrastructure device?

- A. 4
- B. 5
- C. 6
- D. 7

Answer: D

Explanation:

Question: 585

Refer to the exhibit.

R1

```
ip as-path access-list 1 deny 65412_$
ip as-path access-list 1 permit .*

router bgp 64560
  neighbor 10.10.10.10 remote-as 64570
  neighbor 10.10.10.10 route-map FILTER in
```

```
route-map FILTER permit 10
  match as-path 1
```

R1#show ip bgp

Network	Next Hop	1 Metric	LocPrf	Weight	Path
*> 10.0.0.0/8	10.10.10.10	0	100	0	64570 i
*> 10.1.0.0/16	10.10.10.10	0	100	0	64570 i
*> 10.1.1.0/24	10.10.10.10	0	100	0	64570 65412 i
*> 10.1.2.0/24	10.10.10.10	0	100	0	64570 65412 i

*> 10.1.3.0/24	10.10.10.10	0	100	0	64570	65412	i
*> 10.1.4.0/24	10.10.10.10	0	100	0	64570	65412	i
*> 10.1.5.0/24	10.10.10.10	0	100	0	64570	65412	i
*> 10.1.6.0/24	10.10.10.10	0	100	0	64570	65412	i

Refer to the exhibit. An engineer must filter prefixes that originate from AS65412. but it is not working correctly. Which configuration must the engineer apply to R1 to resolve the issue?

```

route-map FILTER permit 10
  match as-path 1
route-map FILTER permit 20
  no ip as-path access-list 1
  ip as-path access-list 1 deny 65412_
  ip as-path access-list 1 permit .*

router bgp 64560
  neighbor 10.10.10.10 route map FILTER out

  no ip as-path access-list 1
  ip as-path access-list 1 deny 65412S
  ip as-path access-list 1 permit .*

```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

Question: 586

A switch has been configured to provide DHCP relay on VLAN 100 to a server with an IP address of 10.1.1.1. The DHCP

server is sending syslog reports of multiple TFTP requests that also originate from the switch. As a result, the server CPU exceeded a configured threshold. Which action does the network administrator recommend to bring the server CPU threshold down?

A. Configure the switch with an access list on VLAN100 to deny TFTP.

B. Configure the switch with a VACL on VLAN100 to deny TFTP.

C. Configure the switch with ip forward-protocol udp 67 globally.

D. Configure the switch with no ip forward-protocol udp 69 on VLAN100.

Answer:

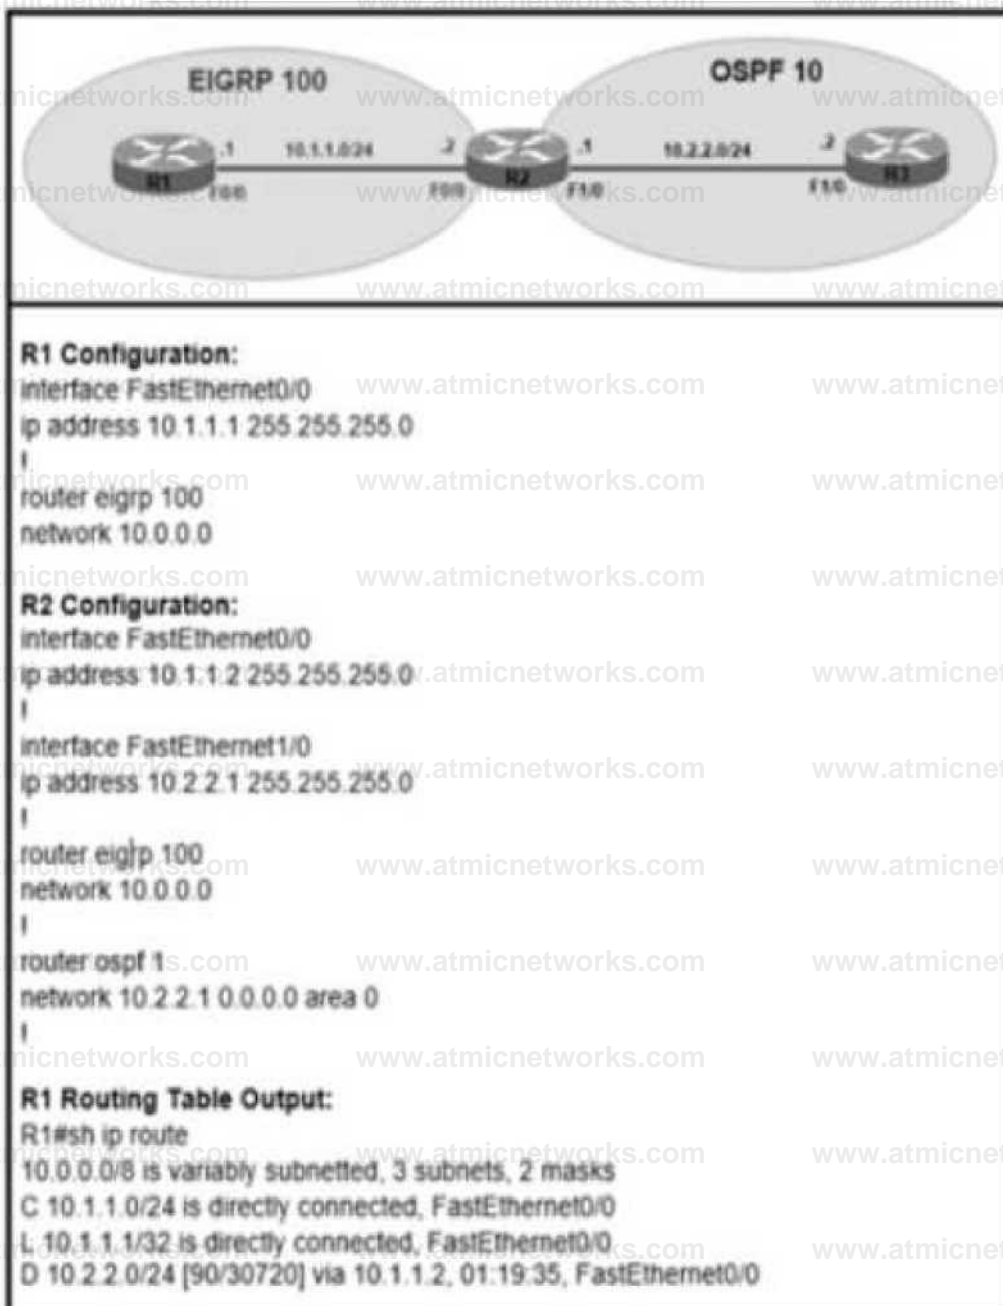
D

Explanation:

Question:

587

Refer to the exhibit.



Refer to the exhibit The R2 OSPF route 10 2 2 0/24 shows in the R1 EIGRP routing table without route redistribution performed between OSPF and EIGRP routing protocols Which configuration is required ON router R2 to resolve the issue?

- A. passive-interface FastEthernet 0/0 command in OSPF1.
- B. Add the no auto-summary command in EIGRP 100.
- C. Replace the network 10.0.0.0 command with FastEthernet0/0 network in EIGRP 100.

D. Add the passive-interface FastEthernet 1/0 command in EIGRP 100

Answer: C

Explanation:

Question: 588

Refer to the exhibit.

```
router eigrp 1
 redistribute ospf 100 route-map ospf-to-eigrp default-metric 20000 2000 255 1 1500
```

1— Output suppressed

```
route-map ospf-to-eigrp deny 10 match tag 6
 match route type external type-2
```

```
route-map ospf-to-eigrp permit 20 match ip address prefix-list pfx 1
 route-map ospf-to-eigrp permit 30 set tag 8
```

Refer to the exhibit Which action fixes the OSPF routes redistribution into EIGRP?

- A. Match external type to type-1
- B. Set tags before matching into EIGRP
- C. Set a default metric in the route map
- D. Match OSPF and EIGR IDs

Answer: B

Explanation:

Question: 589

Refer to the exhibit.

```
R1#ping ipv6 ff02::a
Output Interface: fastethernet1/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FF02::A, timeout is 2 seconds:
Packet sent with a source address of FE80::C801:FFF:FE94:1C%FastEthernet1/0
Reply to request 0 received from FE80::C004:22FF:FE78:1, 40 ms
Request 1 timed out
Request 2 timed out
Request 3 timed out
Reply to request 4 received from FE80::C004:22FF:FE78:1, 56 ms
Success rate is 40 percent (2/5), round-trip min/avg/max = 40/48/56 ms
```

Refer to the exhibit. An engineer investigates an IPv6 EIGRP neighbor adjacency issue that sees the neighbors flapping and issued a ping from R1 to its directly connected neighbor. The link between the switches is stable at Layer 2, and other connected devices are also functioning. Which action resolves the issue?

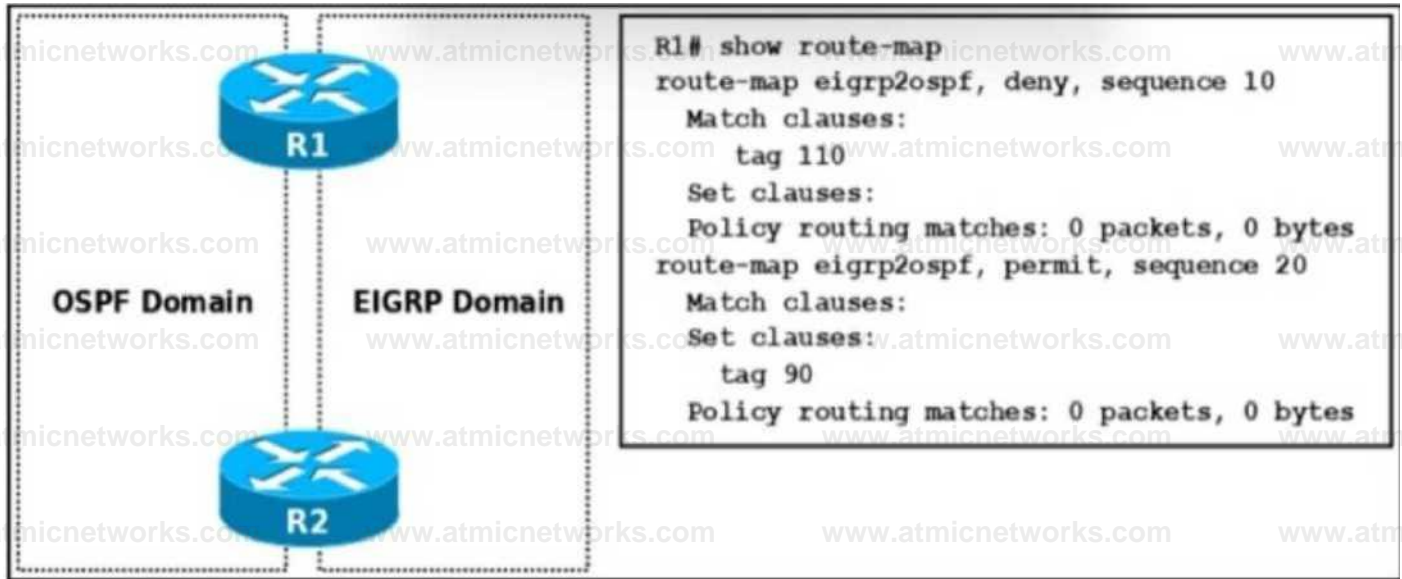
- A. The switch between the two neighbors is not IPv6 compatible and must be replaced with an IPv6 compatible device.
- B. Data is not reliably transmitted between the dynamically assigned link-local addresses of the routers and must be manually assigned.
- C. Multicast packets are not reliably transmitted over the link, and the switch must be replaced.
- D. The switch between the two neighbors is not configured for IPv6 multicast and must be configured for IPv6 multicast.

Answer: C

Explanation:

Question: 590

Refer to the exhibit.



Refer to the exhibit. Which two configurations are required on R2 to prevent a routing loop caused by the redistribution from OSPF back into EIGRP? (Choose two.)


```
route-map eigrp2ospf permit 10
  set tag 90
route-map eigrp2ospf deny 20
  match tag 110
```

```
route-map ospf2eigrp deny 10
  match tag 90
route-map ospf2eigrp permit 20
  set tag 110
```

```
O router ospf 1
  redistribute eigrp 1 route-map ospf2eigrp
```

```
router eigrp 1
  redistribute ospf 1 metric 100000 1 255 1 1500 route-map
  eigrp2ospf
```

```
router eigrp 1
  redistribute ospf 1 metric 100000 1 255 1 1500 route-map
  ospf2eigrp
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

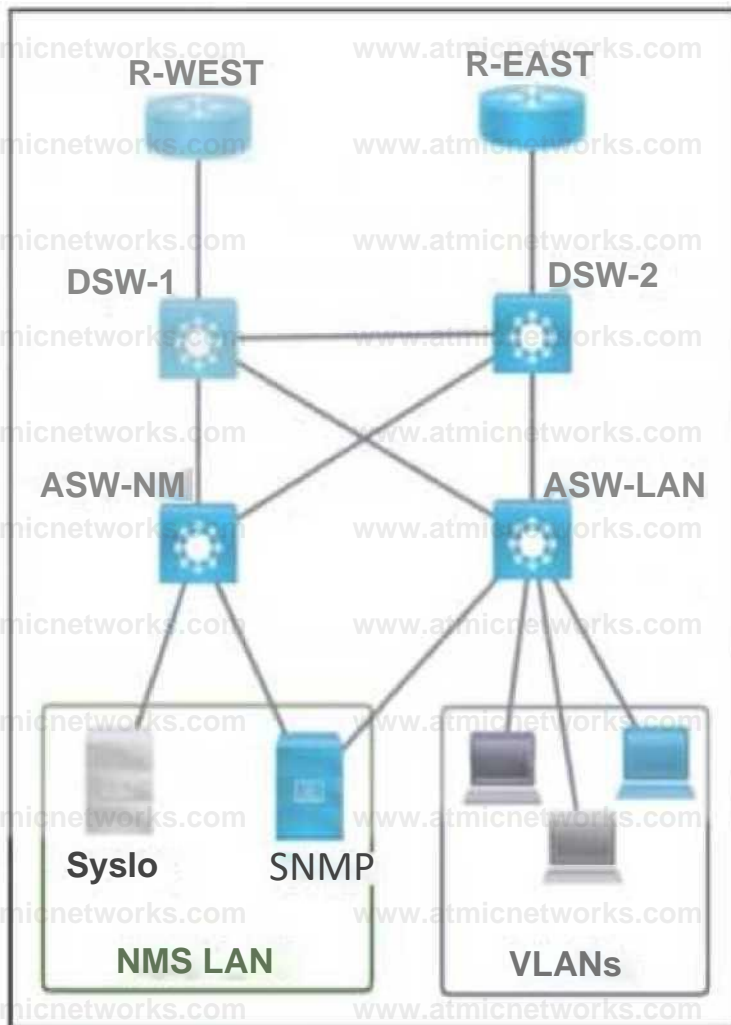
Answer: C, D

Explanation:

Question: 591

Refer to the exhibit.

SIMULATION



Troubleshoot R-WEST to achieve the desired results:

1. The locally generated logs should have sequence numbers, date and time.
2. The SNMP traps related to OSPF and participating interface state changes utilizing RFC1253-MIB OSPFv2 should be sent to SNMP server.

**Answer: See the
solution in**

**Explanation
below.**

Explanation:

Solution:

```
R-WEST(config)tservice sequence-numbers
R-WEST(config)tserv
R-WEST(config)Iservice timestamps 10
R-WEST(config)tsn
R-WEST(config)fsnm
R-WEST(config)Isnrop-se
R-WEST(config)tsnmp-server ena
R WEST(config)tsnmp server enable tr
R-WEST(config)tsnmp-server enable traps ospf
```

```
R-WEST(config)tarchive
R-WEST(config archive)flog co
R-WEST(config-archive)flog config
R-WEST(config-archive-log-cfg)l
R-WEST(config-archive-log cfg)t
R-WEST(config-archive-log-cfg)Hogg
R-WEST(config-archive-log-cfg)Hogging ena
R-WEST(config-archive-log-cfg)flogging enable
R-WEST(config-archive-log cfg)thidek
R WEST(config-archive-log-cfg)thidekeys
R-WEST(config-archive-log-cfg)Inoti
R-WEST(config-archive-log-cfg)Inotify sys
R-WEST(config-archive-log-cfg)fnotify syslog
R-WEST(config-archive-log cfg)tex T
```

```
R-WEST (config) tsnmp-server enable traps ospflsa^^^^^B
```

```
R-WEST(config)tsnmp-server enable traps ospf cisco-specific
```

R-West#

service sequence-numbers

service timestamps log datetime msec

snmp-server enable traps ospf

archieve

log config

logging enable

hidekeys

notify syslog

exit

snmp-server enable traps ospf lsa

snmp-server enable traps ospf cisco-specific lsa

wr

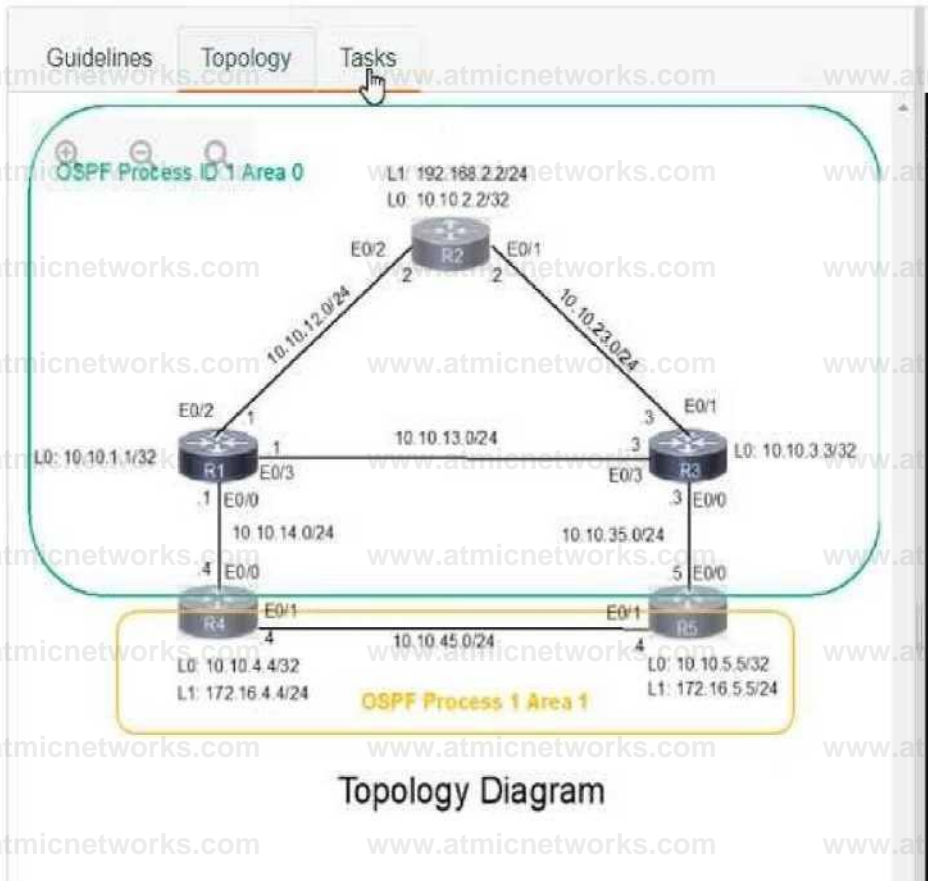
end

wr

Question: 592

SIMULATION

A network is configured with IP connectivity, and the routing protocol between devices started having problems right after the maintenance window to implement network changes. Troubleshoot and resolve to a fully functional network to ensure that:



Guidelines Topology Tasks

A network is configured with IP connectivity, and the routing protocol between devices started having problems right after the maintenance window to implement network changes. Troubleshoot and resolve to a fully functional network to ensure that:

1. Inter-area links have link authentication (not area authentication) using MD5 with the key 1 string CCNP.
2. R3 is a DR regardless of R2 status while R1 and R2 establish a DR/BDR relationship.
3. OSPF uses the default cost on all interfaces. Network reachability must follow OSPF default behavior for traffic within an area over intra-area VS inter-area links.
4. The OSPF external route generated on R4 adds link cost when traversing through the network to reach R2. A network command to advertise routes is not allowed.

R2 R4 R5

```
R2>en | $J I
R2t R2I R2# r
R2f 1
R2# R2i R2#sh run Building configuration

Current configuration : 1279 bytes 1 version 15.8 service
timestamps debug datetime msec service timestamps log datetime msec
no service password-encryption i hostname R2 i
boot-start-marker boot-end-marker 1 I I no aaa new-model I i
I
clock timezone PST -8 0 Go to Settings to 2
mmi polling-interval 60 no mmi auto-configure
```

R2 R4 R5

```
interface Loopback0
ip address 10.10.2.2 255.255.255.255
ip ospf 1 area 0 1
interface Loepback1 ip address 192.168.2.2 255.255.255.0 ip ospf 1
area 0

interface Ethernet0/0 no ip address shutdown duplex auto
I
interface Ethernet0/1
ip address 10.10.23.2 255.255.255.0
ip ospf 1 area 0 duplex auto
j
interface Ethernet0/2
ip address 10.10.12.2 255.255.255.0
ip ospf 1 area 0
duplex auto

interface Ethernet0/3 no ip address shutdown duplex auto j
router ospf i
passive-interface default
no passive-interface Ethemet0/1 no passive-interface Ethernet0/2
```

Activate Windows

Go to Settings to activate

R2 R4 R5

```
interface Ethernet0/3
no ip address
shutdown
duplex auto
router ospf 1
passive-interface default
no passive-interface Ethernet0/1
no passive-interface Ethernet0/2

ip forward-protocol nd
```

```
no ip http server
no ip http secure-server i
ipv6 ioam timestamp f
!
control-plane
i
```

```
!
!
line con 0
Activate
GotoSetti
```

R4

```
R2 R4 R5
R4>
R4>
R4>
R4>
R4>en
R4#sh run
Building configuration...

Current configuration : 1479 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R4
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
!
!
!
clock timezone PST -8 0
mmi polling-interval 60
no mmi auto-configure
no mmi pvc
--More--
```

Activate V
Go to Settin

R2 RW R5

```
key chain CCNP
key 1
  key-string ccnp
  cryptographic-algorithm rod 5
i i
r
i
i
ip address 172.16.4.4 255.255.255.0
1
interface Ethernet0/0
ip address 10.10.14.4 255.255.255.0
ip ospf authentication key-chain CCNP
ip ospf 1 area 0
duplex auto j
interface Ethernet0/1
ip address 172.16.45.4 255.255.255.0
ip ospf 1 area 1
duplex auto
1
interface Ethernet0/2
no ip address shutdown duplex auto
r
interface Ethernet0/3
no ip address
Shutdown
Sett
shutdown
duplex auto
```

Activate

Go to

R2 R4 R5

```
router ospf 1
 redistribute connected subnets route-map to-ospf
 passive-interface default
 no passive-interface Ethernet0/0
 no passive-interface Ethemet0/1
```

```
forward-protocol nd
```

```
no http server
no http secure-server
```

```
ipv6 ioam timestamp
```

```
route-map to-ospf permit 10
 match interface Loopback1
```

```
control-plane
```

```
line con 0
 logging synchronous
```

```
line aux 0
```

Activate Wi
Go to Settings

R5

R2 R4 R5

```
R5>
R5>
R5>en
R5*
R5*
R5*sh run
Building configuration ____
Current configuration : 1496 bytes

version 15.8
service timestamps debug datetime msec service timestamps
log datetime msec no service password-encryption I
hostname R5 i
boot-start-marker boot-end-marker

1
I
no aaa new-model

clock timezone PST -8 0
moi polling-interval 60
no moi auto-configure
no rami pvc
-More- |
```

Activate W
Goro Settings

```
R2  R4  R5
no ip domain lookup
ip cef ipv6 cef
no
raulink bundle-name authenticated
key chain CCNP
key 1
  key-string CCNP
  cryptographic-algorithm md5
Ac
Go
```

R2 R4 R5

```
||i
```

```
»
```

```
interface Loopback0-  
  ip address 10.10.5.5 255.255.255.255  
  ip ospf 1 area 1 i
```

```
interface Loopback1  
  ip address 172.16.5.5 255.255.255.0 1  
interface Ethernet0/0
```

```
  ip address 10.10.35.5 255.255.255.0  
  ip ospf authentication key-chain CCNP  
  ip ospf 1 area 0 duplex auto
```

```
i  
interface Ethernet0/1  
  ip address 172.16.45.5 255.255.255.0
```

```
  ip ospf 1 area 1  
  ip ospf cost 60 duplex auto
```

```
interface Ethernet0/2 no ip address shutdown  
  duplex auto
```

```
A
```

```
interface Ethernet0/3  
  no ip address
```

```
$
```

R2 R4 R5

```
i
router ospf 1
 redistribute connected subnets route-map to-ospf
 passive-interface default
 no passive-interface Ethernet0/0
 no passive-interface Ethernet0/1

ip forward-protocol nd

]
no ip http server
no ip http secure-server !
ipv6 ioam timestamp
I
route-map to-ospf permit 10 match interface Loopback1 • i i
control-plane
I
```

```
1
|
|
I
I
```

```
!
line con 0
 logging synchronous line aux 0
```

Activate Windot
Go to Settings to act

**Answer: See the
solution in
Explanation**

Explanation:

SOLUTION:-

R4

Int range et0/0 – 1

Ip ospf authentication message-digest

Ip ospf message-digest-key 1 md5 CCNP

Router ospf 1

Redistribute connected subnets route-map to-ospf metric-type 1

Copy run start

R5

Int range et0/0 – 1

Ip ospf authentication message-digest

Ip ospf message-digest-key 1 md5 CCNP

Interface eth 0/1

Ip ospf cost 10

Copy run start

VERIFICATION:-

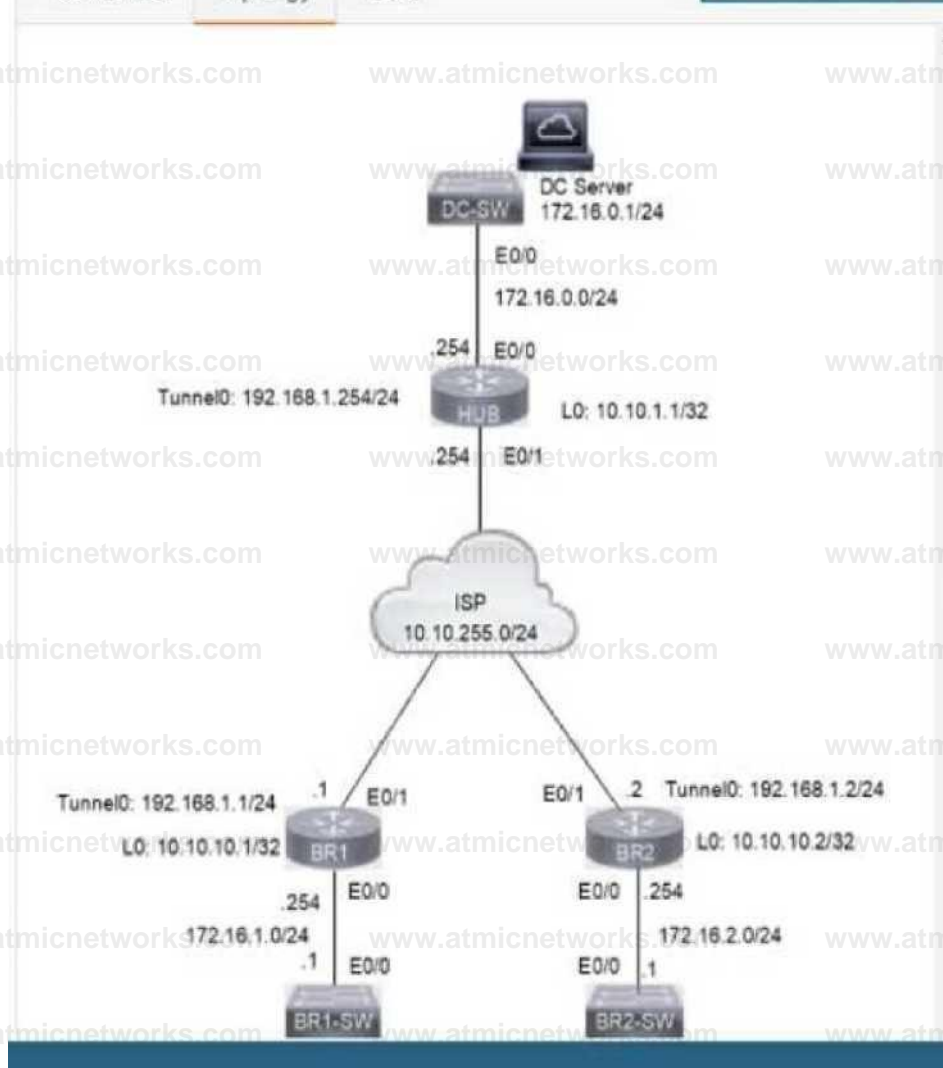
```
R2tshow ip ospf nei
R2tshow ip ospf neighbor

Neighbor ID      Pri  State           Dead Time Address      I
nterface
10.10.1.1        1   FULL/BDR        00:00:38    10.10.12.1   E
thernet0/2
10.10.3.3        i   FULL/BDR        ootoot&ctivateAMtadoavs  E
thernet0/1
R2#|
```

Question: 593

SIMULATION

A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS 100. The NHRP password is ccnp123, and the network ID and tunnel key is EIGRP ASN Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical interface as the tunnel source to achieve these goals:



Guidelines Topology Tasks

A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254

on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS 100.

The NHRP password is ccnp123, and the network ID and tunnel key is

EIGRP-ASN Do not introduce a static route. Configure DMVPN connectivity

between routers BR1 and BR2 to the HUB router using physical interface

as the tunnel source to achieve these goals:


1. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR1.

2. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR2.

3. Ensure that packet fragmentation is done before encryption to account for GRE and IPsec header and allow a maximum TCP

segment size of 1360 on an IP MTU of 1400 on the tunnel interfaces of both branch routers.

4. Apply an IPsec profile to the tunnel. Verify that direct spoke-to-spoke tunnel is functional between branch routers BR1



ADMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS 100. The NHRP password is ccnp123, and the network ID and tunnel key is EIGRP ASN. Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical interface as the tunnel source to achieve these goals:

1. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR1.
2. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR2.
3. Ensure that packet fragmentation is done before encryption to account for GRE and IPsec header and allow a maximum TCP segment size of 1360 on an IP MTU of 1400 on the tunnel interfaces of both branch routers.
4. Apply an IPsec profile to the tunnel. Verify that direct spoke-to-spoke tunnel is functional between branch routers BR1 and BR2 by using traceroute to Ethernet 0/0 IP address to get a full score.

| a Subir# feelbKk about Inn deni I

Answer: See the solution in Explanation

Explanation:

SOLUTION:-

ON BR1

```
Current configuration ; 405 bytes
interface Tunnel0
ip address 192.168.1.1 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccnp123
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp map multi 10.10.255.254
ip nhrp network-id 100
ip nhrp holdtime 5
ip nhrp shortcut 192.168.1.254
ip nhrp shortcut
ip tcp adjust-mss 1360
delay 1000
tunnel source 10.10.255.1
tunnel destination 10.10.255.254
tunnel key 100
end
```

```
BRI(config)#
EEL(config)#
```

ON BR2

```

DC-SW  HUB  BR1  BR1-SW  BR2  BR2-SW
UpDn Time --> Up or Down Time for a Tunnel

Interface: Tunnel0, IPv4 NHRP Details
Type:Spoke, NHRP Peers:1,

# Ent Peer NEMA Addr Peer Tunnel Add State UpDn Tm Attrb
1 10.10.255.254 192.168.1.254 NHRP 00:17:20 S

BR2(config)#do show run int tu 0
Building configuration...

Current configuration : 404 bytes
!
interface Tunnel0
ip address 192.168.1.2 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccmpl23
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp map multicast 10.10.255.254
ip nhrp network-id 100
ip nhrp holdtime 5
ip nhrp nba 192.168.1.254
ip nhrp shortcut
ip tcp adjust-mss 1360
delay 1000
tunnel source 10.10.10.2
tunnel destination 10.10.255.254
tunnel key 100
end
    
```

Verification:-

```

BM!sh dw
aailih l!BVpil
Legend: Rtrb --> S - Static, D - Dynamic, I - Tncwslata
N - HATed, L - bocal, X - No Socket
T1 - Route Installed, T2 - Nextrrp-overt ids
C * CT8 Capable, 12 - Teaporary
♦ Ent --> Hunter of NHRP entries with wise NIKa pee:
NHS Status: E --> Expecting Replies, R --> Responding, N --
> Waiting
UpDn Tim --> Up ox Down Tmc for a Tunnel

Interface: Tunnel0, IPv4 NHRP Details
TypesSpoke, NHRP Peers:1,

I Ent Peer NEMA Addr Feer Tunnel Md State Opta Ta Attrb
1 10.10.255.254 192.168.1.254 UP 00:00:04 S

R»sh <tav
BiClsh dsvpn
Legend: Attrb --> S - Static, D - Dynamic, 1 - Inccoplete
S - HATod, L - Local, X - Me Socket
T1 - Route Installed, T2 - Nexthop-ewerride
C - CTS Capable, 12 - Temporary
♦ Ent --> Hunter of HHRP entries with sama MBNA peer
NH3 status: E --> Expecting Replies, R --> Responding, w --> waiting
OpCn Time --> Up or Down Time for a Tunnel

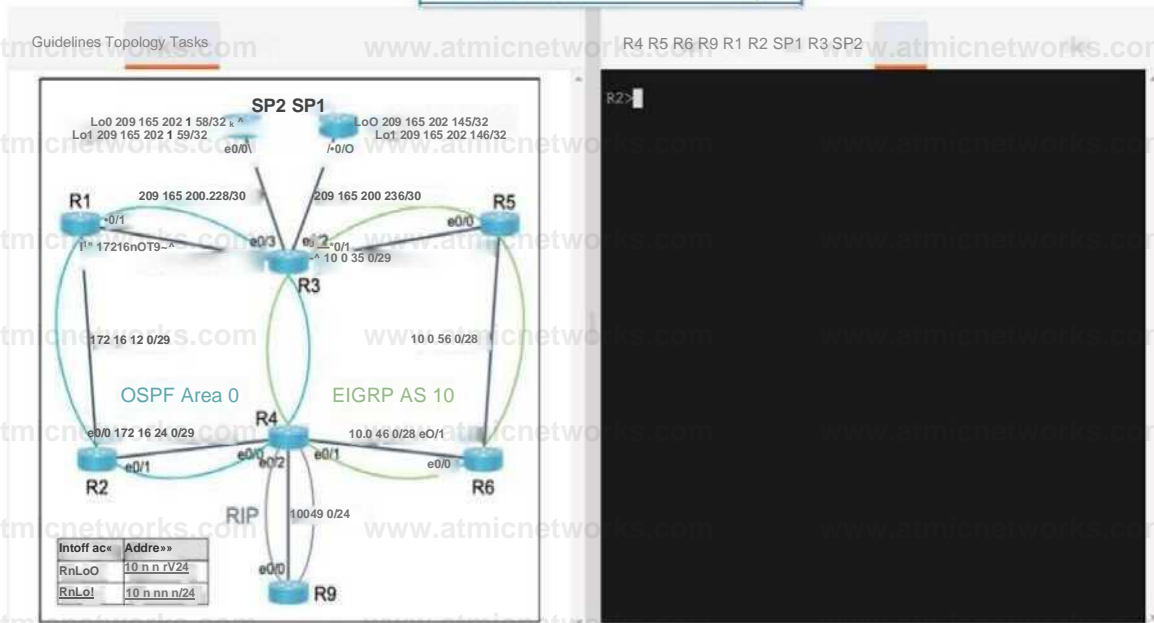
Interface: Tunnel0, IPv4 KHAT Details
Type:Spoke, HHRP Peats:1,

I Ent Pear MBNA Adds Peer Tunnel Add State UpDn lte Attrb
1 10.10.255.254 192.168.1.254 NHRP 00:01:01 8
    
```

auf

Question: 594

SIMULATION



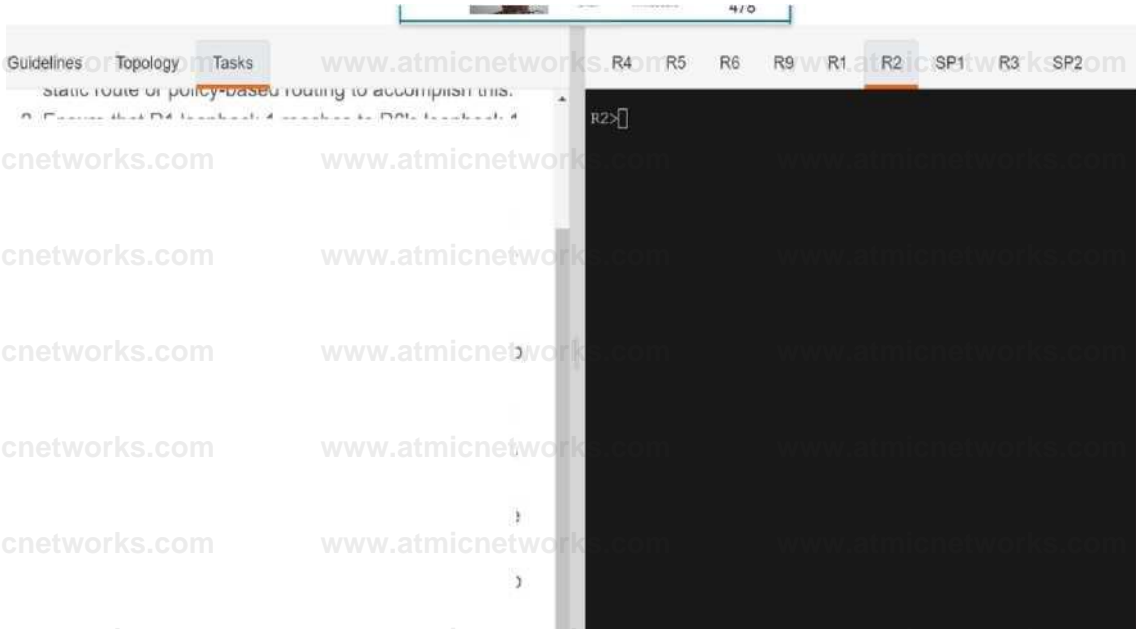
Guidelines Topology Tasks
R3 SP2

R4 R5 R6 R9 R1 R2 SP1

R2>

Troubleshoot and resolve the issues to achieve these goals:

1. Ensure that R1 reaches R5 and R6 loopback's without any single point of failure in the path. Do not use a static route or policy-based routing to accomplish this.
2. Ensure that R1 loopback 1 reaches to R6's loopback 1 by following the path through R1, R3, R5 to R6 and vice versa. Use metric values K1 = 100000, K2=1, K3=255, K4=10, K5=1500 to modify the default metric in EIGRP if required. Do not add or modify the defaultmetric command under router eigrp 10. Do not use a route-map to set metrics.
3. Ensure that on R3, prefix 10.0.56.6/32 uses the SP1 to route to the Internet, whereas prefix 172.16.12.2/32 uses the SP2 to route to the Internet. Do not use BGP to accomplish this. Use the pre-configured route-maps INTERNET1 and INTERNET2 and modify to



2. Ensure that R1 loopback 1 reaches to R6 s loopback 1 by following the path through R1, R3, R5 to R6 and vice versa. Use metric values K1 = 100000, K2=1, K3=255, K4=10, K5=1500 to modify the default metric in EIGRP if required. Do not add or modify the defaultmetric command under router eigrp 10, Do not use a route-map to set metrics.
3. Ensure that on R3, prefix 10.0.56.6/32 uses the SP1 to route to the Internet, whereas prefix 172.16.12.2/32 uses the SP2 to route to the Internet. Do not use BGP to accomplish this. Use the pre-configured route-maps INTERNET1 and INTERNET2 and modify to accomplish the task if required. Use the ping and trace commands from R6 and R2 to prefixes 209.165.202.146 and 209.165.202.158, respectively to verify the results.

**Answer: See the
solution in
Explanation below.**

Explanation:

R3#

```
router eigrp 10
no distance 255 0.0.0.0 255.255.255.255
```

```
exit
```

```
route-map INTERNET1 permit 10
set ip next-hop 209.165.200.238
```

```
exit
```

```
route-map INTERNET2 permit 10
```

```
set ip next-hop 209.165.200.230
```

```
exit
```

```
router eigrp 10
```

```
red ospf 10 metric 100000 1 255 10 1500
```

```
exit
```

```
int et0/0
```

```
ip policy route-map INTERNET1
```

```
exit
```

```
int et0/1
```

```
ip policy route-map INTERNET2
```

```
wr
```


R4

config t

router eigrp 10

no distance 255 0.0.0.0 255.255.255.255 66

red ospf 10 metr 100000 1 255 10 1500

exit

router ospf 10

red eigrp 10 metric 10

wr

Question: 595

What are the two benefits of using BFD? (Choose two.)

- A. forwarding path failure detection
- B. supports all routing protocols
- C. synchronous path determination
- D. supports UDL failure
- E. subsecond failure detection

Answer: A,E

Explanation:

Question: 596

Refer to the exhibit.

```
RI#  
interface Loopback0  
ip address 100 1 1 1 255 255 255 255  
  
interface Loopback 10  
ip address 10 100 1 10 255 255 255 255  
  
interface Loopback20  
ip address 10 100 1 20 255 255.255 255  
  
interface Loopback30  
ip address 10 100 1 30 255 255 255 255  
  
interface Loopback40  
ip address 10 100 1 40 255 255 255 255?
```

```

interface Loopback50
ip address 10 100 1 50 255 255 255 255

interface FastEthernet0/0
ip address 10 1 1 1 255 255 255 252
ip authentication mode eigrp 100 md5
ip authentication key chain eigrp 100 EIGRPKEY
ip summary-address eigrp 100 10100 1 0 255 255 255 0 5 leak-map LOOPBACK50
!

router eigrp 100
network 10 1 100003
network 10 1 1 12 00 0 3
neighbor 101 12 FastEthernet0/0
neighbor 10 1 1 14 FastEthernet0/0 default metric 1000000 10 255 1 1500 no auto-summary
ip prefix list LOOPBACK50 seq 5 permit 10 100 1 50/32

route-map LOOPBACK50 permit 10
match ip address prefix list LOOPBACK50

```

Refer to the exhibit. An engineer configured summarization for the R1 loopback addresses and failed. Which action permits the successful advertisement of the R1 loopback addresses?

- A. Configure EIGRP 100 with a network statement for loopback 0 and configure and redistribute the static route for loopback 50.
- B. Configure EIGRP 100 with a network statement for loopback 0 and redistribute the connected route for loopback 50.
- C. Configure EIGRP 100 with a network statement for loopback 0 and remove the leak map for loopback 50.
- D. Configure 100.1.1.1 permit statement in the prefix list LOOPBACK50 and redistribute the connected route for loopback 50.

Answer: B

Explanation:

Question: 597

Refer to the exhibit.

```

S==
*Mar 10 20:13:58.156: AAA/BIN0(00000055): Bind i/f
•Mar 10 20:13:58.156: AAA/AUTHEN/LOGIN (00000055): Pick method list
'default'
'Mar 10 20:13:58.156: TAC+: Queuing AAA Authentication request 85 for processing
'Mar 10 20:13:58.156: TAC+: (00000055) login timer started 1020 sec
timeout
•Mar 10 20:13:58.156: TAC+: processing authentication start request
id 85
'Mar 10 20:13:58.156: TAC+: Authentication stall packet created for
85{}

```

```

•Mar 10 20:13:58.156: TAC+: Using server 10.106.60.182
`Mar 10 20:13:58.156: TAC+: (00000055)/0/NB_WAIT/225FE2DC; Started 5
sec timeout
•Mar 10 20:13:58.156; TAC+: (00000055)/0/NB_WAIT; socket event 2
`Mar 10 20:13:58.156: TAC+: (00000055)/0/NB_WAIT: wrote entire 38
bytes request
`Mar 10 20:13:58.156: TAC+: (00000055)/0/READ socket event 1
`Mar 10 20:13:58.156; TAC+: (00000055)/0/READ. Would block while
reading
•Mar 10 20:13:58.156: TAC+: (00000055)/0/READ: socket event 1
`Mar 10 20:13:58.156: TAC+: (00000055)/0/READ: read entire 12 header
bytes (expect 6 bytes data)
•Mar 10 20:13:58.156: TAC+: (00000055)/0/READ: socket event 1
•Mar 10 20:13:58.156: TAC+: (00000055)/0/READ: read entire 18 bytes
response
•Mar 10 20:13:58.156: TAO: (00000055)/0/225FE2DC: Processing the reply
packet
`Mar 10 20:13:58.156: TAC+:: received bad AUTHEN packet: length = 6,
expected 43974
•Mar 10 20:13:58.156: TAC+:: Invalid AUTHEN packet (check keys).

```

Refer to the exhibit. An engineer must troubleshoot an issue with the aaa authentication that affected the user's login to router R1. Which command allows the configured user to authenticate?

- A. aaa authentication login default group radius local
- B. aaa authentication login default group radius tacacs+
- C. aaa authentication login default group tacacs+
- D. aaa authentication login default group radius

Answer: C

Explanation:

Question: 598

Which feature is required for IPv6 Source Guard to block traffic arriving on the server interface from unknown sources?

- A. Dynamic ARP Inspection
- B. IPv6 RA Guard
- C. IPv6 ND Inspection
- D. DHCPv6 Guard

Answer: C

Explanation:

Question: 599

How many labels are present in an MPLS Layer 3 packet traversing through the network without traffic engineering?


- A. 1
- B. 2
- C. 3
- D. 4

Answer: A

Explanation:

Question: 600

Refer to the exhibit.



R2» show ip bgp 130.130.130.0/24
 BGP table version 4, local router ID 10.10.20.1

```
Network: 130.130.130.0/24, Local: 10.10.20.1, Path: 0 100 O i
```

R2* show ip protocol

Routing Protocol is "bgp 1"

Outgoing update filter list for all interfaces is not set

Incoming update filter list for all interfaces is not set IGP synchronization is enabled

Automatic route summarization is disabled

Neighbor(s)

Address 10.10.20.1, 10.10.20.3

Maximum path 1

Routing for Networks

Routing Information Sources

Gateway of Last Resort

10.10.20.3, 200.0.1.48, 24

Distance: external 20, internal 200, local 700

Refer to the exhibit. The 130.130.130.0/24 route shows in the R2 routing table but is getting filtering toward R3. Which action resolves the issue?

- A. Automatic route summarization must be enabled on R2.
- B. The incoming filter list for all interfaces must be set on R2.
- C. IGP synchronization must be disabled on R2.
- D. The outgoing filter list for all interfaces must be set on R2.

Answer: A

Explanation:

Question: 601

Refer to the exhibit.

```
Nov 30 10:19:29.595: IP: tableid=0, s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.595: IP: s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 92, rcvd 3
Nov 30 10:19:29.603: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.603: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 92, sending
Nov 30 10:19:29.607: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.611: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 92, sending
Nov 30 10:19:29.615: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.615: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 92, sending
Nov 30 10:19:29.723: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.727: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 108, sending
Nov 30 10:19:29.739: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.743: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 40, sending
Nov 30 10:19:29.807: IP: tableid=0, s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.807: IP: s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.815: IP: tableid=0, s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.815: IP: s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.819: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.819: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 40, sending
Nov 30 10:19:29.827: IP: tableid=0, s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.831: IP: s=172.17.132.95 (GigabitEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.831: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), routed via FIB
Nov 30 10:19:29.835: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigabitEthernet0/0), len 40, sending
```

Refer to the exhibit. Users are experiencing slow response times from critical application servers connected to the network. The engineer determined from the console log that there is a DOS attack on the interface GigabitEthernet0/0. Which action must the engineer take to resolve the issue?

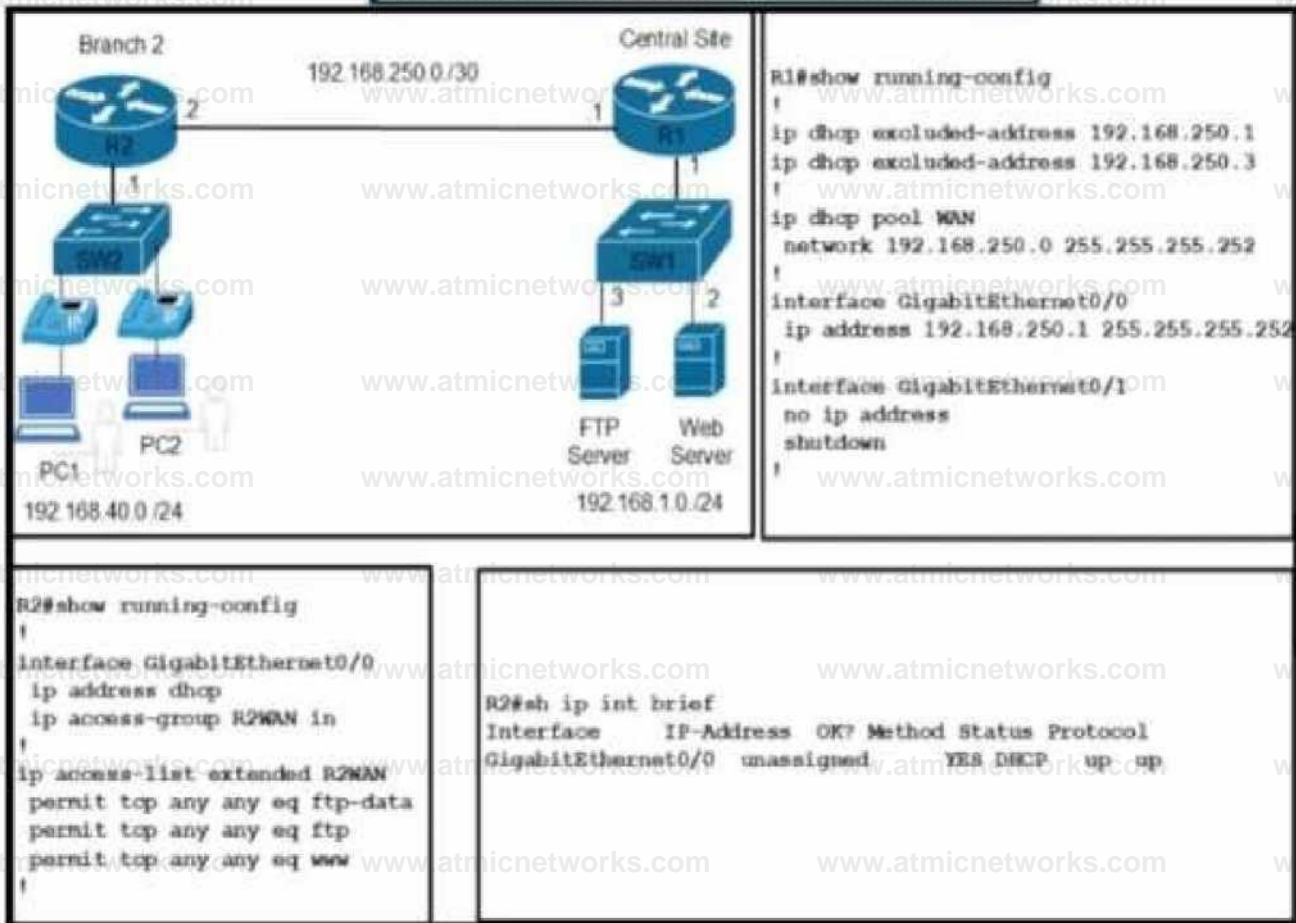
- A. Configure an access group on GigabitEthernet0/0 to block traffic from 192.168.1.1.
- B. Configure an access group on GigabitEthernet0/0 to block traffic from 172.17.132.95.
- C. Restrict CoPP traffic on the management interface from 172.17.132.95 and 192.168.1.1.
- D. Configure QoS on FastEthernet0/0 to restrict bandwidth usage for 172.17.132.95 and 192.168.1.1.

Answer: D

Explanation:

Question: 602

Refer to the exhibit.



Refer to the exhibit. Which configuration is required for R2 to get the IP address from the DHCP server?

- A. ip access-list extended R2WAN permit tcp any any eq 68
- B. ip access-list extended R2WAN permit udp any any eq 68
- C. ip access-list extended R2WAN permit udp any any eq 67
- D. interface GigabitEthernet0/0 ip access-group R2WAN out

Answer: B

Explanation:

Question: 603

Refer to the exhibit.

<pre>R1#show run begin router eigrp 100 router eigrp 100 network 172.16.250.0 0.0.0.3 redistribute ospf 10 metric 1 1 1 1 ! router ospf 10 network 192.168.1.0 0.0.0.3 area 0 ! ip forward-protocol nd ! no ip http server</pre>	<pre>R3#show ip route Gateway of last resort is not set 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks C 192.168.1.0/30 is directly connected, GigabitEthernet0/1 L 192.168.1.2/32 is directly connected, GigabitEthernet0/1 C 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks C 192.168.2.0/24 is directly connected, Loopback2 L 192.168.2.33/32 is directly connected, Loopback2 C 192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks C 192.168.3.0/24 is directly connected, Loopback1 L 192.168.3.17/32 is directly connected, Loopback1 R3#</pre>
<pre>R3#traceroute 172.16.2.48 Type escape sequence to abort. Tracing the route to 172.16.2.48 VRF info: (vrf in name/id, vrf out name/id) 1 * * * 2 * * * 3 * * *</pre>	<pre>R4#show running-config begin router eigrp router eigrp 100 network 172.16.2.0 0.0.0.3 network 172.16.2.16 0.0.0.15 network 172.16.2.32 0.0.0.15 redistribute static metric 100 1 1 1 1 route-map CCNP ! ip forward-protocol nd ! no ip http server no ip http secure-server ip route 172.16.2.48 255.255.255.240 172.16.2.34 ! route-map CCNP permit 10 match ip address 10 set tag 200 ! access-list 10 permit 172.16.2.48 0.0.0.15</pre>

Refer to the exhibit. An engineer must troubleshoot a connectivity issue impacting the redistribution of the subnet 172.16.2.48/28 into the OSPF domain. Which configuration on router R1 advertises this subnet into the OSPF domain?

- A. R1(config)#route-map CCNP permit 10R1(config-route-map)#match route-type internalR1(config-router)#router ospf 10R1(config-router)#redistribute eigrp 100 subnets route-map CCNP
- B. R1(config)#route-map CCNP permit 10R1(config-route-map)#match route-type level-2R1(config-router)#router ospf 10R1(config-router)#redistribute eigrp 100 subnets route-map CCNP
- C. R1(config)#route-map CCNP deny 10R1(config-route-map)#match tag 200R1(config)#route-map CCNP permit 10R1(config-router)#router ospf 10R1(config-router)#redistribute eigrp 100 subnets route-map CCNP
- D. R1(config)#route-map CCNP permit 10R1(config-route-map)#match tag 200R1(config-route-map)#exitR1(config-router)#router ospf 10R1(config-router)#redistribute eigrp 100 subnets routemap CCNP

Answer: D

Explanation:

Question: 604

What are two features of BFD? (Choose two.)

- A. scalable
- B. replaces hello messages

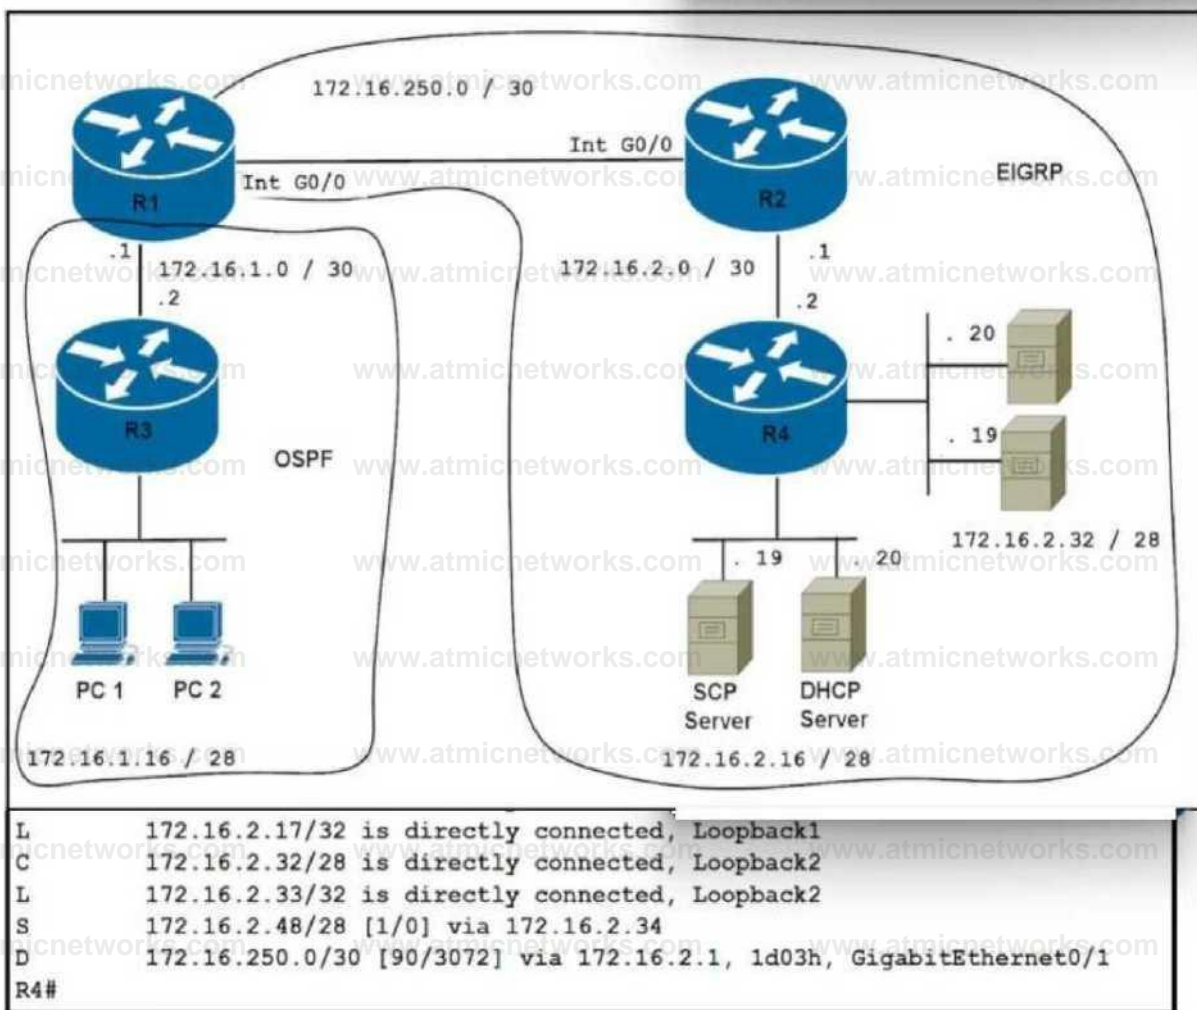
- C. reliable
- D. intensive on CPU for Layer 2 links
- E. requires routing protocols

Answer: A, C

Explanation:

Question: 605

Refer to the exhibit.



Refer to the exhibit. The users from subnet 172.16.1.16/28 are experiencing connectivity issues to subnet 172.16.2.32/48. Which configuration resolves the issue?

- A. R1(config)#route-map REDIST permit 10
R1(config-route-map) #match ip address 15
R1 (config-route-map) #exit
R1(config)# access-list 15 permit 172.16.0.0.0.255.255
R1(config-router) #router eigrp 100
R1(config-router) #redistribute ospf 1 route-map REDIST

- B. R1(config-router) #router ospf 1
R2(config-router) #redistribute eigrp 100 subnets metric 100C.
- C. R1(config-router) #router eigrp 100
R1(config-router) #redistribute ospf 1 metric 1000000 1111
- D. R1(config)#route-map REDIST permit 10
R1(config-route-map) #match ip address 15
R1 (config-route-map) #exit
R1(config)# access-list 15 permit 172.16.0.0 0.0.255.255
R1(config-router) #router ospf 1
R1(config-router) #redistribute eigrp 100 subnets route-map REDIST

Answer: D

Explanation:

Question: 606

Refer to the exhibit.

```

B# show run | section router ospfv3
router ospfv3 20
 area 11 stub
 |
 address-family ipv4 unicast
 passive-interface default
 no passive-interface GigabitEthernet0/0
 no passive-interface GigabitEthernet1/0
 default-information originate
 router-id 10.10.10.10
 exit-address-family
 |
 address-family ipv6 unicast
 passive-interface default
 no passive-interface GigabitEthernet0/0
 no passive-interface GigabitEthernet1/0
 router-id 10.10.20.20
 exit-address-family

```

```

D# traceroute 2001:db8:e::e
Type escape sequence to abort.
Tracing the route to 2001:DB8:e::e
 1 2001:DB8:0:22::1 1U 1U 1U

```

```

C# show ipv6 route 2001:db8:e::e
% Route not found

```

Refer to the exhibit. An enterprise user reports an access issue with IPv6 content on the Internet. The user can access IPv4 content that is at the data center. Which action resolves the issue with IPv6 content?

- A. Change Area 11 to Area 0 between routers B and C.
- B. Add a static route for 2001: db8:0:6::1/64 on router B RIB.
- C. Advertise the 2001:db8: e::e/64 route on the router D OSPF process.
- D. Enable default information to originate in IPv6 AF on router B.

Answer: D

Explanation:

Question: 607

Refer to the exhibit.

R1#sh ip route

C 192.168.10.0/24 is directly connected, Serial1/0
172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
C 172.16.160.0/19 is directly connected, Loopback1
C 172.16.128.0/19 is directly connected, Loopback0
C 172.16.224.0/19 is directly connected, Loopback1
C 172.16.192.0/19 is directly connected, Loopback2
B 172.16.0.0/16 is a summary, 00:01:27, Null0

Refer to the exhibit. Which configuration advertises more specific routes to R1 without sending a BGP summary route?

- A. R1#configure terminal
R1 (config)#router BGP 100
R1 (config-router)#no auto-summary
- B. R1#configure terminal
R1 (config)#router BGP 100
R1 (config-router)#auto-summary
- C. R2#configure terminal
R2 (config)#router BGP 100
R2 (config-router)#no auto-summary
- D. R2#configure terminal
R2 (config)#router BGP 100
R2 (config-router)#auto-summary

Answer: C

Explanation:

Question: 608

Refer to the exhibit.

R Wing 10.1.1.1

Type escape sequence to abort

Sending 5,100-byte ICMP Echoes to 10.11.1.1, timeout is 2 seconds
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/99/152 ms

R10#copy tftp://10.1.1.1/c890-universal-mz-154-2-M3.bin flash:

Destination filename [c890-universal-mz-154-2-M3.bin]:

Accessing tftp://10.1.1.1/c890-universal-mz-154-2-M3.bin

%Error opening tftp://10.1.1.1/c890-universal-mz-154-2-M3.bin (Timed out)

```
R11#sh run interface GigabitEthernet 0/0
```

```
interface GigabitEthernet0/0
ip vrf forwarding TFTP
ip address 10.1.1.1 255.255.255.0
speed auto
duplex auto end
```

Refer to the exhibit. R10 attempts to copy an image file from R11 using TFTP, and the operation is timing out. Which action resolves the issue?

- A. Change R10 duplex to auto to resolve the duplex mismatch between routers R10 and R11.
- B. R10 must be configured to belong to the same VRF TFTP.
- C. R11 requires the TFTP server source-interface to be set to GigabitEthernet0/0.
- D. R11 requires the TFTP server command for the image to be set to the VRF TFTP.

Answer: C

Explanation:

Question: 609

How are LDP neighbors discovered?

- A. Unicast hellos are sent to directly connected neighbors IP addresses.
- B. Multicast hellos are sent to the 224.0.0.2 group address.
- C. Broadcasts hellos are sent to the 255.255.255.255 broadcast address.
- D. Multicast hellos are sent to the 224.0.0.5 group address.

Answer: B

Explanation:

Question: 610

Refer to the exhibit.

The screenshot shows a network device log for 'Device 360'. The log entry is highlighted in blue and reads: 'Notice DUAL:NBRCHANGE Syslog 8:46:37.207 PM'. Below the log entry, a 'Detailed Information' table is displayed:

Severity	Notice
Mnemonic	NBRCHANGE
Facility	DUAL
Message Text	682: *Jan 11 15:41:03.036: EIGRP-IPv4 88: Neighbor 172.16.33.2 (GigabitEthernet2.10) is down: authentication mode changed
Message Type	Syslog

Refer to the exhibit. R1 lost its directly connected EIGRP peer 172.16.33.2 (SW1). Which configuration resolves the issue?

A. key chain EIGRP

keyl

key-string Cisco !

interface GigabitEthernet 2.10

ip authentication mode elgrp 88 md5

ip authentication key-chain elgrp 88 EIGRP

B. key chain EIGRP

keyl

key-string Cisco !

interface GigabitEthernet 2

ip authentication mode elgrp 88 md5

ip authentication key-chain elgrp 88 Cisco

C. key chain EIGRP

keyl

key-string Cisco !

interface GigabitEthernet 2

```
ip authentication mode eigrp 88 md5
ip authentication key-chain eigrp 88 EIGRP
D.key chain EIGRP
keyl
key-string Cisco !
interface GigabitEthernet 2.10
ip authentication mode eigrp 88 md5
ip authentication key-chain eigrp 88 Cisco
```

Answer: A